

The Media Streaming Journal

November 2019



Covering Audio and Video Internet Broadcasting

Brought To You By

RADIOSOLUTION

www.radiosolution.info



The Media Streaming Journal Staff

Derek Bullard
Publication Director
info@radiosolution.info

David Childers
Editor In Chief
editor@radiosolution.info

Advertising
advertising@radiosolution.info

www.radiosolution.info

publicdomainvectors.org/en/free-clipart/Vintage-microphone-vector-graphics/6111.html

Welcome to The Media Streaming Journal

Greetings,

The world of original video content is rapidly growing among both the large media production companies as well as smaller independent businesses. The ability to create unique video scenes can be accomplished without the use of advanced CGI animation processes. Shading, lighting, background, and optical characteristics can be worked to achieve genuinely remarkable video footage.

The deciding factor is the ability to compile and edit these videos to achieve noteworthy results. It is essential to learn beyond button pushing and understand the processes that occur when editing and mixing video content.

Please feel free to contact either the Publication Director (Derek Bullard) or myself if you have any questions or comments regarding The Media Streaming Journal.

Namaste

David Childers

The Grand Master of Digital Disaster
(Editor In Chief)



www.linkedin.com/pub/david-childers/4/736/72a

David Childers

The Grand Master of Digital Disaster

Current Member: International Association Of Internet Broadcasters

Former Member: Society of Motion Picture and Television Engineers

Published Author

Introduction To Internet Broadcasting
Amazon Publishing

Over 25 Creative Commons Computer, Technical and Internet Broadcasting Guides

Newspaper Interviews

New York Times

Lagniappe - "Something Extra for Mobile"

Internet TV: Don't Touch That Mouse!
Tim Gnatek
July 1, 2004

Mobile Gets Hoaxed
Rob Holbert
Mar 16, 2016

Cited By

Five Essays on Copyright In the Digital Era
Ville Oksanen
2009

Turre Publishing
Helsinki Finland

Open Source Developer

Developed software architecture to continuously source multimedia content to Youtube Live servers.
Scenic Television - The sights and sounds of nature on the Internet.

<http://www.ScenicTelevision.com>

Projects

Researched and developed documentation for Peercast P2P multimedia streaming project.

<http://en.wikipedia.org/wiki/PeerCast>

Researched and developed technical documentation for NSV / Winamp Television.

http://web.archive.org/web/20080601000000*/http://www.scvi.net

MidSummer Eve Webfest

A virtual International festival focusing on Digital art and Free Software that was coordinated by OrganicaDTM Design Studio.

Presentation and discussion regarding Internet multimedia content distribution.

<http://web.archive.org/web/20061104230522/http://www.organicdtm.com/index.php?module=articles&func=display&catid=37&aid=61>

LinkedIn Contact Information

<http://www.linkedin.com/pub/david-childers/4/736/72a>

The Media Streaming Journal

What is in this edition of the Media Streaming Journal

Cinelerra-GG Version Infinity
User Manual



Join our technical discussion on Facebook

<http://www.facebook.com/groups/internetradiosupport/>

Magazine cover:

https://commons.wikimedia.org/wiki/File:Fotothek_df_pk_0000165_012.jpg

**The Media Streaming Journal is licensed under the
Attribution-ShareAlike 4.0 International
(CC BY-SA 4.0)
Creative Commons License.**

www.creativecommons.org/licenses/by-sa/4.0/



RADIO SOLUTION

www.radiosolution.info

Our Mission

Let our friendly, knowledgeable staff assist you to build your project, such as an online radio station using our high end reliable video and audio streaming technologies. We want to become your partner for all your hosting needs, as well as your one stop shop for radio products such as custom DJ drops and radio ID's.

Start An Internet Radio Station

Whatever you need to start Internet radio station, we will deliver! We provide high quality Internet Radio services to make your music radio project a success. We can provide Wowza, Icecast, SHOUTcast hosting and internet radio services to hobbyists, deejays, amateurs and established professionals. No radio station client is too big or too small for Radiosolution.

Choose between complete hassle-free service packages or new features to add to start internet radio station. Benefit from customized services and the latest in internet radio technology. You will receive professional, personalized and better Internet Radio Station services than you have received up till now. If you already have an Icecast or SHOUTcast hosting provider, we can still help you transfer your radio server over to us with no hassle and at no charge.

Internet Radio Station Services

Launch your internet, digital, satellite or AM/FM radio station anywhere in the world with all of the right tools. A broadcasting specialist is on standby to help you get started with an SHOUTcast or Icecast hosting package. We have servers ready for reliable streaming in North America and Europe. Our hosting packages have all the features you need to make your radio station project a success.

If you stream live or with an Auto DJ, we can provide you with the latest in web-based Cloud technology. You will love the simple to use control panel. Discover how easy it is to manage live deejays, upload fresh music and create custom scheduled programming. You will be able to track your listeners by getting real time statistics.

Starting your own Internet radio has never been easier. Get in touch with us anytime to start your Internet radio station.

Radiosolution is a SHOUTcast hosting provider located in Quebec Canada. We also offer Icecast, Wowza and Web Hosting services. Contact us to discuss the best option available as you start internet radio station. Radiosolution can provide personalized service in English, Dutch, and French. Starting an internet radio station can be intimidating, many people want to start one, but have no idea where to start. Radiosolution will be there for you every step of the way. Everyday people are searching the internet for free SHOUTcast servers. With Radiosolution SHOUTcast hosting we will allow you to try our services for FREE. By trying our services, you can be confident that you have chosen the best radio server hosting provider. You have nothing to loose because we offer a 30 day satisfaction guarantee. What are you waiting for? Contact us now! Radiosolution offers everything you need to start internet radio station. You will not need to go anywhere else. We can create your website, market your station and help you submit your station to online directories. We also feature the voice of Derek Bullard aka Dibblebee He can create affordable commercials, DJ intros, sweepers, jingles, ids and so much more.



Relax With The Sights And Sounds Of Nature

Scenic Television

Your Window To The World

Scenic Television is an Internet television station that broadcasts the sights and sounds of nature 24 hours a day. Savor exotic tropical beaches, or relax in a remote rain forest. Meditate at a bubbling stream, or relish the view of soft rolling waves at a lake. We have beautiful nature video from locations all around the world.

Scenic Television originates from the Gulf coast of South Alabama and broadcasts to a global audience. The television broadcast is accessible on any device with an Internet connection. Such electronic devices include desktop computers, laptops, tablets, smartphones, game platforms, and Internet-connected televisions.

<http://www.scenictelevision.com>



all-free-download.com/free-vector/download/magnifying_glass_clip_art_23181.html

We Are Your Information Resource

Are you looking for specialized data?

Are you swamped with information overload?

Do you need help finding the right information?

We Can Help You
Find The Information
That You Need

Our experienced data research analysts can wade through the vast information wasteland and find the information that you need.

We can save you both time and money.

We can streamline data requirement planning.

We can provide business critical information acquisition.

Contact us today

info@radiosolution.info

Cinelerra-GG Version Infinity User Manual

Cinelerra is a software program NLE, Non-Linear Editor, that provides a way to edit, record, and play audio or video media. It can also be used to retouch photos.

Hey You! Yes, You! Why Should Anyone Listen to You?!

Do you need compelling, clever copy or catchphrases for your Internet station? If you do, please visit and lets talk!

<http://www.ielectrify.com/work-with-me/>

I am a professional writer with 15+ years of experience creating high-converting copy, for a variety of radio, broadcasting and marketing applications.



https://www.wpclipart.com/people/professions/professions_3/radio_announcer.png.html

Cinelerra-GG Version Infinity



User Manual

Creation Date 2019

<https://www.cinelerra-gg.org>



This page intentionally left blank so that chapters begin on a right facing page.

1 - Introduction

Cinelerra is a software program NLE, Non-Linear Editor, that provides a way to edit, record, and play audio or video media. It can also be used to retouch photos.

This manual covers Cinelerra-GG Infinity version. The author of the original Cinelerra, Adam Williams, as well as many different people worked on Cinelerra over the years. The software and this manual were merged in from various sources and each person is to be thanked and commended for their efforts. Numerous software modifications were made by William Morrow. These are all under GPLv2+ license. Refer to: <https://www.gnu.org/licenses/old-licenses/gpl-2.0-standalone.html>

This is a copy of the header from the original source code.

```
/*
 * CINELEERRA
 * Copyright (C) 1997-2012 Adam Williams <broadcast at earthling dot net>
 *
 * This program is free software; you can redistribute it and/or modify
 * it under the terms of the GNU General Public License as published by
 * the Free Software Foundation; either version 2 of the License, or
 * (at your option) any later version.
 *
 * This program is distributed in the hope that it will be useful,
 * but WITHOUT ANY WARRANTY; without even the implied warranty of
 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
 * GNU General Public License for more details.
 *
 * You should have received a copy of the GNU General Public License
 * along with this program; if not, write to the Free Software
 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
 */
```

This is a copy of the information in the Cinelerra-CV manual.

Copyright c 2003, 2004, 2005, 2006 Adam Williams - Heroine Virtual Ltd.
Copyright c 2003, 2004, 2005, 2006, 2007 Cinelerra CV Team.

This manual is free; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This document is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA 02110, USA.

1.1 Cinelerra Overview

Presented briefly here is an overview of Cinelerra-GG Infinity and information provided in this manual. The GG version of Cinelerra has been improved for *stability*, *modernized* to accommodate the *current state* of Linux software, enhanced with additional *basic features*, and enriched with *new features* imagined by dedicated users and then implemented by professional programmers.

Website <https://www.cinelerra-gg.org>

The website for the cinelerra-gg software is a good place to start for information, help, and documentation. It is professionally maintained and continuously updated with more language translations ongoing.

Stability

Software programs that “just work” are a #1 priority in order to be of use for producing quality videos. A large amount of time has been invested in debugging problems and resolving crashes. And in a continuous process to do so, a chapter on Troubleshooting is included in order to easily provide sufficient information for users to capture issues and crashes so that they can be analyzed and quickly fixed to avoid repeat problems.

Modernization

Artistic creativity has been applied to modernize the Cinelerra-GG plugin icons for video and audio to even include the ffmpeg plugins. The CInfinity set of plugin icons come in square or roundish versions – your choice. In keeping up with the current expectation of users for a certain “look and feel”, a very modern Neophyte theme recent addition provides an alternative to the already existing 9 themes. These 10 themes give the user the choice to get the look they like best for their own eyes.

Current and up-to-date

For today’s software, included thirdparty libraries are kept up to date in a timely manner and an effort is made to always used a relatively recent version of FFmpeg, if not the latest. This is a big deal because there is a whole set of separate programmers continuously working diligently to cover all of the old and newly created formats. Thus Cinelerra programmers can be dedicated to working solely on Cinelerra rather than just trying to keep up with many formats.

FFmpeg usage integration

By using FFmpegwith Cinelerra there is the advantage that users can directly convert videos via pre- and post-processing, without the need for command lines to be executed manually before or afterwards.

Import and Export formats

Listed here are only a few of the supported Import and Export formats:

- several standard native formats, such as mpeg, ac3, flac, exr and jpeg/png/ppm/tiff sequences
- FFmpeg’s over 400 decoders and 150 encoders accessible from within Cinelerra to include:
 - mp4, mkv, mpeg, mov, m2ts, mp3, dvd, ogg, theora, prores, tiff, webm, flac, opus, vorbis, quicktime (div, dnxhd, jpeg, mjpeg, mp4, rle, v308, v410), h264 & h265 usage, avc, hevc, and recently released AV1 and WebP
- raw image format for over 700 supported cameras, courtesy Dave Coffin DCraw

Standard Features

Program window for video and audio tracks, navigation, popups, playing and seeking functions.
Editing via track manipulation with either drag and drop editing or cut and paste editing.
Patchbay tools: fader, gang, arm, enable output, expanders, mixers, and blending operations.
Undo and Redo capability for many editing functions.
Drag handle functionality for Ripple, Roll, Slip, Slide, Edge, and No Effect.
Dynamic Keyframe support with curve, toggle, automatic, compositor, and editing capabilities.
Proxy editing support to speed up editing for large formatted files or slower computers.
Compositor window with masking, zooming, cropping, projector and camera capabilities.
Viewing window for quickly viewing/playing audio, video, clips, or proxies.
Resources window with Media, Proxy, Clips, Video and Audio Effects/Transitions.
User defined bins/folders for a project.
Expandable Video/Audio effects for grouping selections such as Color Correction or Motion plugins.
Included are over 400 Plugins – audio and video native, FFmpeg, plus expandable to include others.
LV2 and Ladspa plugins audio support.
Mouse over Media or Proxy folders to get a 5 second real time preview of the media.
Single or multiple files rendering; background rendering, batch rendering, command line rendering,
and use of a Render Farm to take advantage of multiple computer CPUs.
Numerous pre-defined output formats automatically available and allowance for user-defined formats.
Capture and Recording capability to include Broadcast TV recording, editing, and viewing.
Hundreds of Shortcuts are defined which are easily viewed using the shell commands pulldown.
PorterDuff operations are available in the patchbay of the main timeline window for alpha blending.
Color correction + 8-bit and 10-bit color space.
Up to 8K video supported.
ShuttlePro and ShuttleXpress functionality.
International language support.
Multi-screen, Multi-viewers, Multi-session, Multi-pane capability.
Menu bar shell commands with reference to some manuals and user-definable additions.
Video stabilization, Motion tracking, Motion Graphics.
Blade Cut / Hard Edges, Snap/Cut editing and Snap Dragging.
OpenGL support as well as Direct X11 speedup.
Advanced Trim Features and Grouping.
Asset and Title colors selection including Alpha slider adjustment.

Innovative New Features

Inter-View mode as invented through a collaborative effort by a user and the developer.
Proxy settings for large formatted media with the coveted Scalar option and proxy quick switch.
Multiple cameras / Mixer Viewer with the number of cameras only limited by the resources on the
user's computer (on a large system, 50 are known to work).
DVD/bluray creation, editing, and copying for non-commercial media greatly enhanced for usability.
Title plugin virtually unlimited script size with many changeable attributes such as size, blink, color.
Motion Graphics using the Sketcher plugin to create elements such as ellipses, rectangles and shapes
for simpler motion graphics.

1.2 Chapters Overview

Below is listed a brief overview of each chapter to help you to decide which chapters you should pay the most attention to and read more thoroughly. And which chapters are important for beginning to learn to use Cinelerra-GG. At the end of the descriptions is a list of the sections to read for beginners.

Chapter 2 – Installation

If you just want to get started using the program, you can safely skip this chapter and instead go to:

<https://www.cinelerra-gg.org>

and simply download a pre-built linux version for your Operating System. If you would like to do your own builds so that you always have the latest, refer to this chapter to learn how. But if you are already familiar with doing your own builds, you can just refer to this chapter if you encounter issues.

Chapter 3 – The 4+ Windows

It is important to understand the window setup in Cinelerra-GG and what each of them is used for. In addition this chapter covers some basic navigation functions that you will need to know. However, if you have used an NLE before and don't mind experimenting to learn how to execute certain functions, probably there is no need to peruse this chapter.

Chapter 4 – Project and Media Attributes

This chapter is helpful for basic understanding of setting up your session but it can easily be skipped.

Chapter 5 – Load, Save, and the EDL

Since this is important to not losing your work, you should read this chapter for some basic usage concepts and for some lesser used functions that may come in handy. Besides how to Load and Save files, there is also information on using raw camera formats. Helpful hints on working with image sequences, such as a bunch of pictures from your camera all loaded at once, is a time saver.

Chapter 6 – Editing

New and occasional users will find it necessary to read this chapter. However, you can decide which editing mode you prefer and concentrate on reading that section. Even seasoned users will want to read this chapter in order to take advantage of features that have been newly implemented or may be less familiar due to infrequent usage. Anyone who wants the editing capabilities that comes with use of the Shuttle Jog Wheel, will initially need to read that section to become familiar with its usage.

Chapter 7 – Rendering

Minimally most users will need to read the section on single file rendering to get started; since after all the whole purpose of using an NLE is to create your own media. The end of this chapter includes some helpful specific rendering scenarios which could prove quite useful. You might also want to consider using batch rendering and background rendering for performance reasons. And, of course, render farm usage to take advantage of multiple computers can be a real boost.

Chapter 8 – Keyframes

This is a more advanced topic but extremely useful to know. Keyframes are data values that have been associated to the timeline that affect the media presentation. You can go for years without knowing all of the nuances but you will most likely have a need to use and understand them sooner or later.

Chapter 9 – Ffmpeg Interactions

Skip this chapter if you are just getting started. Skip this chapter if you just want to use common formats that are already set up for you. Experts may want to refer to this chapter in order to set up their own option files.

Chapter 10 – Effect Plugins

Use this chapter as a reference to add an audio or video plugin in order to correct color or add some fancy effect. You will not want to read about each and every plugin, most of which you will never use. On the other hand, going over the section on “some specific details concerning plugins” might come in handy.

Chapter 11 – Transition Plugins

Everyone who does not know about using transitions should read this chapter because you will want to use transitions between cuts in your video for smoothly changing scenes. It is short and easy reading.

Chapter 12 – Overlay Modes, Alpha Blending, and Porter Duff

Experts can really use the information in this chapter to spiff up their editing work through some blending type techniques. A lot of information and possibilities are described here.

Chapter 13 – Capturing and Recording Media + DVD/Bluray Creation

If this is what you want to do, then read this chapter. In the case of DVD media creation it includes many examples and lots of tips for checking your methods. A section on using cinelerra with broadcast TV is especially useful if that is something you want to do.

Chapter 14 – The 5 Multi-s

For using multiple cameras in your video creation, this chapter is a must to read. Otherwise, just refer to the section that becomes pertinent as you familiarize yourself with using Cinelerra.

Chapter 15 – Shortcuts with Keyboard and the Mouse

Everyone will want to use this chapter to speed up their editing sessions. Once you find yourself doing the same thing over and over again, you will want to know if there is a shortcut. Then look here!

Chapter 16 – Configuration, Settings, and Preferences

Refer to this chapter when you want to change some setting or preference. Otherwise if a new user, you should make sure to read the first page which includes some important basic information.

Chapter 17 – How Some Stuff Works

Only read this chapter if you are confused about a specific covered topic and how it works.

Chapter 18 – Troubleshooting

Use this chapter for diagnosing a problem and find out what to report to get the best resolution or help.

Chapter 19 – Performance Tips and Other Tips

There are a few performance tips in this chapter that may eventually prove useful.

Chapter 20 – Language Translations

If you want to help by providing translations for a specific language, there is a complete description of how to do this and what to do.

Chapter 21 – Licenses

This chapter covers the GPL and other license information, where to read about them online and which licenses cover the Cinelerra software.

Appendix A – QuickStarts

Here you go! If you hate to read, just go over the quick start guide or the youtube video creation and simply be on your way.

Appendix B – Developer’s Section

Some extra information that developers may or may not find useful.

Appendix C – Auxiliary Programs

This is a catch-all for any useful programs that may need to be included, mostly for analysis.

In summary, “must” reads for a new user would be these chapters or sections:

Chapter 3 - The 4+ Windows

Chapter 5 - Load, Save, and the EDL

Chapter 6 - Editing; read all sections except emphasize only either Cut&Paste or Drag&Drop editing to suit your purpose and then skim the editing mode that is not your preference as some operations work in either mode. Skip the Shuttle section unless you have this jog wheel in hand.

Chapter 7 – Rendering; minimally read the “single file rendering” section.

Chapter 10 – Effect Plugins; read the section on “how to use plugins”.

Chapter 11 – Transition Plugins

Chapter 16 – Configuration, Settings and Preferences; read at least the first couple of paragraphs.

Table of Contents

1 Introduction

1.1 Cinelerra Overview

1.2 Chapters Overview

2 Installation

2.1 How to Build Cinelerra-GG from Developer's Git Repository

2.1.1 Notable Options and Caveats

2.1.2 Notes about Building from Git in your Customized Environment

2.1.3 Cloning the Repository for Faster Updates

2.1.4 How to Build from a Previous GIT Version

2.1.5 Debuggable Single User Building

2.1.6 Unbundled Builds

2.2 Download Already Built Cinelerra-GG

2.3 Distribution Systems with Cinelerra Included

2.4 Cinx and a "Bit" of Confusion

3 The 4+ Windows

3.1 Program Window

3.1.1 Video and Audio Tracks and Navigation

3.1.2 Zoom Panel

3.1.3 Track Popup Menu

3.1.4 Insertion Point

3.1.5 Editing Modes

3.1.6 In/Out Points

3.1.7 Labels

3.1.8 Color Title Bars and Assets

3.1.9 More about Pulldowns

3.1.10 Window Layouts

3.1.11 Just Playing!

3.2 Compositor Window

3.2.1 Compositor Controls

3.2.2 Compositing

3.2.3 Camera and Projector

3.2.4 Masks

3.2.5 Cropping

3.2.6 Safe Regions

3.2.7 Track and Output Sizes

3.3 Viewer Window

3.4 Options in both the Compositor and Viewer Windows

3.4.1 Click to Play in Compositor and Viewer

3.4.2 Timebar + Preview Region Usage in the Compositor and Viewer

3.5 Resources Window

3.5.1 Info Asset Details

3.5.2 User Folders/Bins

3.5.3 Vicons & Aicons – aka Video Icons / Audio Icons

3.5.4 Resources Window Preview Mode

3.5.5 Moving clips/media from/to Resources window

3.5.6 Snapshot / Grabshot

3.6 Other Options and Other Windows

3.6.1 Transport Controls

3.6.2 Zoombar

3.6.3 Show Overlays

3.6.4 Sound Level Meters

4 Project and Media Attributes

5 Load, Save, and the EDL

5.1 EDL – Edit Decision List

5.2 Supported File Formats

5.2.1 Working with Still Images

5.2.2 Timelapse Sequences of Images, and Animation

5.2.3 Raw Image Format of Some Digital Cameras & Probe Order

5.3 Loading Files

5.3.1 Sort within Sort in File Load Dialog

5.3.2 Size Numeric Format Displayed in File Load

5.3.3 Probe Order when Loading Media

5.3.4 Program Selection Support after Load

5.4 Saving Your Work

5.4.1 Saving Project Files

5.4.2 Export Project – Save or Moving Project to another Computer

5.4.3 Information about Backups and Perpetual Session

6 Editing

6.1 The Patchbay

6.2 Manipulating Tracks

6.3 Two Screen Editing

6.4 Cut and Paste Editing

6.5 Drag and Drop Editing

6.5.1 Cut/Paste Behavior

6.5.2 Groups

6.5.3 Dragging Groups

6.6 Inter-View Mode / Identifying Source Targets Editing

6.7 Some Specific Editing Tools

6.7.1 Edit Length

6.7.2 Align Edits

6.7.3 Reverse Edits

6.7.4 Shuffle Edits

6.7.5 Drag Handle Management / Trimming

6.7.6 Split View in Compositor Using the Drag Handle with Trim

6.7.7 Snapping while Cutting and Dragging

6.7.8 Nesting

6.7.9 Copy/Paste clips/media across Multiple Instances

6.8 ShuttlePROv2 and ShuttleXpress Jog Wheels for Editing

7 Rendering

7.1 Single File Rendering

7.2 Batch Rendering

7.3 Background Rendering

7.4 Render Farm Usage

7.4.1 Basic Steps to Start a Render Farm

7.4.2 Render Farm Menu and Parameter Description

- [7.4.3 Detailed Setup Descriptions](#)
- [7.4.4 Quick and Easy Render Farm Setup – The Buddy System Way](#)
- [7.4.5 Multi-core Computers Render Farm Setup \(the “Epyc” method!\)](#)
- [7.4.6 Troubleshooting Tips and Warnings](#)

[7.5 Some Specific Rendering](#)

- [7.5.1 FFmpeg Common H.264 Rendering](#)
- [7.5.2 Lossless Rendering](#)
- [7.5.3 Extra “cin ” Options for Render with FFmpeg](#)
- [7.5.4 Two-Pass Encoding with FFmpeg](#)
- [7.5.5 Piping Video to a Command Line](#)
- [7.5.6 Faststart Option for MOV type files](#)

[8 Keyframes](#)

[8.1 Automation Keyframes / Autos](#)

- [8.1.1 Using Autos](#)
- [8.1.2 Speed / Fade Automation Usage and Auto Gang](#)

[8.2 Plugin Keyframes](#)

- [8.2.1 Default Keyframe / Typeless Keyframe](#)
- [8.2.2 Keyframes “Edit Params” for Plugins](#)

[8.3 Generate Keyframes while Tweaking / Automatic Keyframe Mode](#)

[8.4 Compositor Keyframes](#)

[8.5 More about Editing Keyframes](#)

[9 FFmpeg Interactions](#)

[9.1 FFmpeg Early Probe Explanation](#)

[9.2 How to Create Ffmpeg Options Files](#)

- [9.2.1 File naming convention](#)
- [9.2.2 Option File Format / Content](#)
- [9.2.3 Complete Options File Example](#)
- [9.2.4 Viewing and Modifying FFmpeg format options inside Cinelerra](#)

[9.3 The FFmpeg Image2 Streams](#)

[9.4 Raw Input Opts File for Video/Audio](#)

[9.5 FFmpeg Items of Note](#)

[10 Effect Plugins](#)

[10.1 How to Use Plugins](#)

[10.2 Editing Effects](#)

[10.3 Shared Effects and Shared Tracks](#)

[10.4 Saved Plugin Presets](#)

[10.5 Some specific details concerning Plugins](#)

- [10.5.1 How to see short Description of a Plugin](#)
- [10.5.2 Delete Plugins to save Resources Space or make them Unavailable](#)
- [10.5.3 Updatable Icon Image Support](#)
- [10.5.4 Details on where to put your own Plugin Icons](#)
- [10.5.5 Example of new Plugin Icon Testing](#)
- [10.5.6 Plugins/Effects Visibility](#)
- [10.5.7 Expanders for Plugin Subtrees in the Resources Window](#)
- [10.5.8 Speed-up of Ffmpeg plugin usage with OPTS files](#)

[10.6 Audio Effects – Native](#)

[10.7 Audio Ladspa Effects](#)

[10.8 Audio LV2 / Calf Plugins \(Optional Feature – OS dependent\)](#)

[10.9 Video Effects – Native](#)

[10.10 Video Find Object and other plugins from OpenCV – Stylize, Puzzle, Flow, Gabor, Move](#)

[10.10.1 How to Build OpenCV Plugins](#)

[10.10.2 Using OpenCV Plugins from the Automatic Builds](#)

[10.10.3 Description of Find Object Plugin](#)

[10.10.4 An example of how to use Find Object](#)

[10.10.5 Additional OpenCV – MoveObj, FlowObj, GaborObj, StylizeObj, PuzzleObj](#)

[10.11 FFmpeg Audio and Video Plugins](#)

[10.11.1 FFmpeg Icons and how to Replace](#)

[10.11.2 How to use FFmpeg Audio/Video Plugins](#)

[10.11.3 FFmpeg Audio Plugins](#)

[10.11.4 FFmpeg Video Plugins](#)

[10.12 Rendered Effects](#)

[10.12.1 Rendered Audio Effects](#)

[10.12.2 Rendered Video Effects](#)

[11 Transition Plugins](#)

[11.1 Audio Transitions](#)

[11.2 Video Transitions](#)

[12 Overlay Modes, Alpha Blending, and Porter Duff](#)

[13 Capturing and Recording Media + DVD/Bluray Creation](#)

[13.1 Record Web Media by real-time Audio Capture/Screen Capture](#)

[13.2 Digital Video Broadcasting \(DVB\) Capture Configuration](#)

[13.2.1 Overview](#)

[13.2.2 Initial Setup for Digital Video Broadcast Capture](#)

[13.2.3 Watching TV using Cinelerra](#)

[13.2.4 Recording TV to Capture at some Future Time via Batch](#)

[13.2.5 Recording Immediately to Capture a current TV program](#)

[13.2.6 Record “on the fly” Table of Contents \(TOC\)](#)

[13.2.7 Watching a Previously Recorded TV program](#)

[13.2.8 Dual Screen Mode](#)

[13.2.9 Remote Control for DVB](#)

[13.2.10 Android Remote Control for DVB](#)

[13.2.11 How to Setup a Remote Control Device to work with Cinelerra](#)

[13.3 The commercial DB... What is that? A TDB – Traveling Data Base](#)

[13.4 DVD and Bluray Creation](#)

[13.4.1 Output Terminal Messages from Creating DVDs](#)

[13.4.2 Helpful Hints for Problems/Debugging DVDs Creation](#)

[13.4.3 Subtitles](#)

[13.4.4 Dvd Interlaced Chroma](#)

[13.4.5 MPEG utility programs](#)

[13.4.6 HDV-sourced Video Recording on Blu-ray Disc Without Re-encoding](#)

[13.4.7 User-mode Blu-ray Render Workaround for Mount/Umount](#)

[13.4.8 How to Create Bluray Media from Multiple Cinelerra Output](#)

[13.5 Complete Example of Video Source with 4:3 Aspect Ratio, Being Transcribed to 16:9 and Creating a DVD to be Displayed on a Digital TV](#)

[14 The 5 Multi-s](#)

[14.1 Multi-Camera / Mixer](#)

[14.1.1 Mixer Align by Audio](#)

[14.1.2 Recover Mixer Windows](#)

[14.2 Multi-Pane](#)

- [14.3 Multi-Screen / Playback Configuration](#)
 - [14.4 Multi-Session](#)
 - [14.5 Multi-Viewer Window Support](#)
 - [15 Shortcuts with Keyboard and Mouse](#)
 - [16 Configuration, Settings and Preferences](#)
 - [16.1 Playback A / Playback B](#)
 - [16.2 Recording](#)
 - [16.3 Performance](#)
 - [16.4 Interface](#)
 - [16.5 Appearance](#)
 - [16.6 About](#)
 - [16.7 Environment Variables for Customization](#)
 - [17 How Some Stuff Works](#)
 - [17.1 Copy/Paste and Highlight Usage with Cinelerra](#)
 - [17.2 Playing is Different than Seeking/Positioning!](#)
 - [17.2.1 Always Show Next Frame](#)
 - [17.2.2 Examples](#)
 - [17.2.3 Seeking Issues](#)
 - [17.3 Color Space and Color Range Affecting Playback](#)
 - [17.4 Simple Animation – Filescene using Festival](#)
 - [17.5 Textbox Non-standard Character / Unicode Insertion](#)
 - [18 Troubleshooting and Help](#)
 - [18.1 What to Include in Problem Reports](#)
 - [18.2 Crash Dumps for Analysis](#)
 - [18.3 When Things go Wrong](#)
 - [18.4 Menu Bar Shell Cmds](#)
 - [19 Performance Tips and Other Tips](#)
 - [19.1 Optimized Playback of Large Format Video – X11 Direct](#)
 - [19.2 Proxy Settings](#)
 - [19.3 Some Settings Parameter Values](#)
 - [19.4 Tips for Improving Smaller Computer Use](#)
 - [19.5 General Crash Handling Tips](#)
 - [19.6 Usage Tips for specific Operations](#)
 - [20 Language Translations](#)
 - [21 Licenses](#)
- [Appendix](#)
- [Appendix A – QuickStarts](#)
 - [A.1 Cinelerra-GG Quick Start Guide](#)
 - [A.2 YouTube with Cinelerra + Basic Video Creation in 8 Steps](#)
 - [Appendix B - Developer's Section](#)
 - [B.1 How builds really work and even more options](#)
 - [B.2 Thirdparty Parallel Build](#)
 - [B.3 Find Lock Problem with Booby Trap](#)
 - [B.4 Valgrind Support Level](#)
 - [B.5 CFLAGS has -Wall](#)
 - [Appendix C - Auxiliary Programs / Information](#)
 - [C.1 Using ydiff to check results](#)
 - [C.2 Image Sequence Creation](#)

C.3 Webm / Vp9 Usage and Example File

2 – Installation

2.1 How to Build Cinelerra-GG from Developer's Git Repository

These are generic build instructions for building Cinelerra-GG Infinity. Known to work on ubuntu, mint, suse/leap, fedora, debian, centos, arch, and slackware. It has not been tested on every single possible distro yet so you might expect to have to make some minor changes. Patches have been created to build on FreeBSD through the work of another programmer and a Gentoo version is being maintained elsewhere by another programmer.

Alternatively, there are some pre-built dynamic or static binaries which are updated on a fairly regular basis (as long as code changes have been made) available at the link below.

<https://cinelerra-gg.org/download/>

There are 2 kinds of builds, the default system-build and a single-user build. A system build has results which are installed to the system. The majority of the files are installed in the standard system paths, but some customization is possible. The single user build allows for running completely out of a local user directory so it doesn't affect the system.

We recommend the single-user version when possible. It makes it very easy to install a new version without having to delete the older version in case you want it for backup – once you are happy with the new version, all you have to do is delete the entire old directory path. Also, if you install a new Operating System version and if you have Cinelerra on separate disk space that is preserved, you won't have to reinstall Cinelerra. In addition for purposes of having the ability to interrupt or to see any possible error messages, if you start the application from a terminal window command line you will have more control to catch problems. However, the system builds can be useful in a university lab setting where there are possibly multiple users, or multiple versions.

There are two notable differences between “standard” views of Cinelerra and this implementation for the system builds. Both of these can be configured during installation. These differences make it possible to have several different versions installed without having them “walk” on each other.

- 1) application name can be set during installation and defaults to: “cin”
- 2) the home configuration directory can also be set and defaults to: “\$HOME/.bcast5”

To do a system build, you should read the README that is at the top level after you get the source.

- 1) You need at least 2.5GB of disk storage to operate a build + you need to have “git” installed.
- 2) Obviously in order to install into the system, you must run as root.
- 3) The "git" step has to download many files (approx 100MB) so allow time.
- 4) Run the following commands (this takes awhile):

```
cd /<build_path>/          # this is where you need the 2.5GB of disk space
git clone --depth 1 "git://git.cinelerra-gg.org/goodguy/cinelerra.git" cinelerra5
cd cinelerra5/cinelerra-5.1 # toplevel directory
```

NOTE: if your system has never had Cinelerra-GG Infinity installed, you will have to make sure you have all of the compilers and libraries necessary. So on the very first build you should run:

```
./blds/bld_prepare.sh <os> # where <os> represents the Operating System of
centos, fedora, suse, leap, ubuntu, debian.
./autogen.sh
./configure --prefix=/usr # optional parameters can be added here
make 2>&1 | tee log # make and log the build
```

5) Check for obvious build errors:

```
grep "\*\*\.*error" -ai log
```

If this reports errors and you need assistance or you think improvements can be made to the builds, email the log, filename is listed on the next line, to cin@lists.cinelerra-gg.org:

```
/<build_path>/cinelerra5/cinelerra-5.1/log
```

6) If there are no build errors, finally just run:

```
make install
```

7) If it all worked, you are all setup. Just click on the cinelerra desktop icon.

To do a single-user build, read the README that is at the top level after you get the source.

- 1) You need at least 2.5GB of disk storage to operate a build + you need to have “git” installed.
- 2) Recommend you build and run as root, just to avoid permission issues initially.
- 3) The "git" step has to download many files (approx 100MB) so allow time.
- 4) Run the following commands (this takes awhile):

```
cd /<build_path>/ # this is where you need the 2.5GB of disk space
git clone --depth 1 "git://git.cinelerra-gg.org/goodguy/cinelerra.git" cinelerra5
cd cinelerra5/cinelerra-5.1
```

NOTE: if your system has never had Cinelerra-GG Infinity installed, you will have to make sure all of the compilers and libraries necessary are installed. So on the very first build you should run as root:

```
./blds/bld_prepare.sh <os> # where <os> represents the Operating System of
centos, fedora, suse, leap, ubuntu, debian.
./autogen.sh
./configure --with-single-user # the “with-single-user” parameter makes it so
make 2>&1 | tee log # make and log build (check for errors before proceeding)
make install
```

Then just start the application by keying in: `./cin` in the bin subdirectory OR add a desktop icon by using the appropriate directory to copy the files to, run as root, and edit to correct the directory path.

```
cd /cinelerra_directory_path
cp -a image/cin.{svg,xpm} /usr/share/pixmaps/
cp -a image/cin.desktop /usr/share/applications/cin.desktop
change the “Exec=cin” line to be “Exec=<your_directory_path>/bin/cin”
```

The preceding directions for doing a single-user build has been meticulously followed to build and run on a newly installed ubuntu 15 system WITHOUT BEING ROOT except for the `bld_prepare.sh` and creating the desktop icon.

2.1.1 Notable Options and Caveats

These procedures and the Cinelerra-GG Infinity software have all been run as “root” on various home laptops and desktops. This provides the best chance to ensure all works correctly and also allows for handling errors, other problems and potential crashes with the most success. Included in this section are some of the build variations easily available for normal builds.

To see the full list of features use: `./configure --help`

The default build is a system build which uses: `./configure --without-single-user`

In the single-user build, the target directory is always “cin”. Because this is also the developer build, constant names are used throughout. However, you can rename files after the install is complete.

If your distro/operating system has issues with the default install to `/usr/local`, you might have to change the location to `/usr` for a system build. Then you will have to use:

```
./configure --prefix=/usr
```

If you wish to change the default directory for a system build you will have to add the destination directory path on the “make install” line. For example:

```
make install DESTDIR=<your selected target directory path>
```

The application name can be set during installation, but defaults to cin so that the GG/Infinity build can coexist with other Cinelerra builds if necessary. To override the default cin name, use:

```
./configure --with-exec-name=cinelerra
```

The home configuration directory can also be set, but defaults to `$$HOME/.bcast5`. For example:

```
./configure --with-config-dir=/myusername/.bcast5
```

NOTE: when you specify parameters to the configure program, it will create a make file as a consequence. Since in a make file, the \$ is a special character, it must be escaped so in order to represent a \$ as part of an input parameter, it has to be stuttered. That is, you will need \$\$ (2 dollar signs) to represent a single dollar sign.

It may be necessary on some distros which have missing or incomplete up-to-date libraries, to build cinelerra without Ladspa. To do so, use:

```
./configure --prefix=/usr --without-ladspa-build
```

Note that the `with-ladspa-dir` is the ladspa search path, and exists even if the ladspa build is not selected. This gives you the ability to specify an alternate ladspa system path by utilizing the `LADSPA_PATH` environment variable (that is, the default ladspa build is deselected).

Note for 32-bit 14.2 Slackware, Debian, Gentoo, Arch, FreeBSD, before running the configure, you will need to set up the following:

```
export ac_cv_header_xmmintrin_h=no
```

```
export FFMPEG_EXTRA_CFG="--disable-udpau"
```

2.1.2 Notes about Building from Git in your Customized Environment

Getting a build to work in a custom environment is not easy. If you have already installed libraries which are normally in the thirdparty build, getting them to be recognized means you have to install the "devel" version so the header files which match the library interfaces exist. Below is the list of thirdparty builds, but this list may have changed over time:

a52dec	yes
djbfft	yes
fdk	auto
ffmpeg	yes
fftw	auto
flac	auto
giflib	yes
ilmbase	auto
lame	auto
libavc1394	auto
libraw1394	auto
libiec61883	auto
libdv	auto
libjpeg	auto
openjpeg	auto
libogg	auto
libsndfile	auto
libtheora	auto
libuuid	yes
libvorbis	auto
mjpegtools	yes
openxr	auto
tiff	auto
twolame	auto
x264	auto
x265	auto
libvpx	auto
libwebp	auto
libaom	auto

The "yes" means force build and "auto" means probe and use the system version if the build operation is not static. To get your customized build to work, you need to change the probe options for the conflicting libraries from "yes" to "auto", or even rework the configure.ac script. There may be several libraries which need special treatment.

An example of a problem you might encounter with your customized installation is with "a52dec" which has probes (CHECK_LIB/CHECK_HEADER) in configure.ac, but djbfft does not. In this case, djbfft is only built because a52dec is built, so if your system has a52dec, set a52dec to auto and see if that problem is solved by retrying the build with:

```
./configure --with-single-user --enable-a52dec=auto .
```

With persistence, you can get results, but it may take several tries to stabilize the build. If you need help, email the "log" and "config.log", which is usually sufficient to determine why a build failed.

If you have already installed the libfdk_aac development package on your computer because you prefer this version over the default aac, you will have to do the following to get this alternative operational.

```
export FFmpeg_EXTRA_CFG="--enable-libfdk-aac --enable-nonfree"
export EXTRA_LIBS="-ldk-aac"
```

```
for f in `grep -lw aac cinelerra-5.1/ffmpeg/audio/*`; do
  sed -e 's/<aac>/libfdk_aac/' -i $f
done
```

2.1.3 Cloning the Repository for Faster Updates

If you want to avoid downloading the software every time an update is available you need to create a local "repository" or repo. The repo is a directory where you first do a "git clone". For the initial git clone, setup a local area for the repository storage, referred to as <repo_path>. The "git clone" creates a repo named "cin5" in the /<repo_path>/ directory. This accesses over 300MB of repo data, so the device has to have at least that available. The repo path is always a perfect clone of the main repo.

Setting up the initial clone: (add "--depth 1" before cin5 which is faster/smaller, but has no history)

```
# cd /<repo_path>/
# git clone "git://git.cinelerra-gg.org/goodguy/cinelerra" cin5
Cloning into 'cin5'...
remote: Counting objects: 20032, done.
remote: Compressing objects: 100% (11647/11647), done.
remote: Total 20032 (delta 11333), reused 16632 (delta 8189)
Receiving objects: 100% (20032/20032), 395.29 MiB | 3.26 MiB/s, done.
Resolving deltas: 100% (11333/11333), done.
Checking connectivity... done.
```

Update an existing repo:

```
# cd /<repo home>/cin5
# git pull
```

Useful git commands:

```
git clone "git://git.cinelerra-gg.org/goodguy/cinelerra.git" cin5
git pull      - pull remote changes to the local version
git status    - shows changed files
git clean -i  - interactive clean, use answer 1 to "clean"
```

2.1.4 How to Build from a Previous GIT Version

```
# cd /<path>/cin5_repo  
# git log  
# git checkout <version>
```

The “git log” command produces a log file with hash values for commit keys. The hash ids are the commit names to use when you use git checkout. Next is displayed sample output:

```
...  
    delete stray line in last checkin  
  
commit 4a90ef3ae46465c0634f81916b79e279e4bd9961  
Author: Good Guy <good1.2guy@gmail.com>  
Date: Thu Feb 22 14:56:45 2018 -0700  
  
    nested clips, big rework and cleanup, sams new icons, leaks and tweaks  
  
commit f87479bd556ea7db4afdd02297fc00977412b873  
Author: Good Guy <good1.2guy@gmail.com>  
Date: Sat Feb 17 18:09:22 2018 -0700  
...
```

For the “git checkout <version>”, you would then keyin the line below for the following results:

```
# git checkout f87479bd556ea7db4afdd02297fc00977412b873
```

Note: checking out 'f87479bd556ea7db4afdd02297fc00977412b873'.

You are in 'detached HEAD' state. You can look around, make experimental changes and commit them, and you can discard any commits you make in this state without impacting any branches by performing another checkout.

If you want to create a new branch to retain commits you create, you may do so (now or later) by using -b with the checkout command again. Example:

```
git checkout -b <new-branch-name>
```

HEAD is now at f87479bd... more file size icon updates, and more to follow

Later to get the repo back to current, use: # git checkout master

2.1.5 Debuggable Single User Build

To build from source with full debugging symbols, first build a full static (non_debug) build as follows but instead /tmp substituted with a permanent disk path if you want to keep it.

```
# git clone ...
# cp -a /path/cinelerra-5.1 /tmp/.
# cd /tmp/cinelerra-5.1
# ./bld.sh
```

Then, to run as a developer in the debugger:

```
# CFLAGS="-O2 -ggdb" make -j8 rebuild_all
# cd cinelerra
# gdb ./ci
```

2.1.6 Unbundled Builds

There are some generic build scripts included in the Cinelerra-GG GIT repository for users who want to do unbundled builds with ffmpeg already available on their system. This has been tested on Arch, Ubuntu 18, FreeBSD, and Leap 15 (rpm) at the time this was documented. The names of the build scripts are: arch.bld , bsd.bld , deb.bld , and rpm.bld . These scripts are in the “blds” subdirectory. The bsd.bld should be used with the bsd.patch file in that same directory.

The reason that Cin Infinity traditionally uses thirdparty builds (bundled builds) is because there are a lot of different distros with varying levels of ffmpeg and other needed thirdparty libraries. However, some users prefer using their current system baseline without another/different copy of ffmpeg. With different levels of the user’s libraries, uncertainty, potential instability, and unknown issues may come up while running Cinelerra and this will make it, for all practical purposes, impossible to diagnose and debug problems or crashes. There may be no help in these cases. You are encouraged to report any errors which potentially originate from Cin Infinity, but if the data indicates alternate library sources, please report the problems to the appropriate maintainers.

With the unbundled builds, some features may not be available and no attempt to comment them out has been made. So if you use a pulldown, or pick a render option, or choose something that is not available, it just will not work.. For example, unless special options were set up by you, the LV2 audio plugins will not be available. Nor will the codec libzmpeg, the file codec ac3, or DVD creation. The old school file classes will all work, but some of the formats that come with ffmpeg may not because of the way that ffmpeg was installed on your operating system. That is because the Cinelerra ffmpeg is a known static build and is usually the latest stable/released version. In the current case of Leap 15, libx264 and libx265 are not built in and this can be debilitating; You can always run “ffmpeg -formats” and “ffmpeg -codecs” to see what is available on your system.

2.2 Download Already Built Cinelerra-GG

If you prefer to not have to take the time to build Cinelerra-GG Infinity yourself, there are pre-built dynamic or static binaries for various versions of ubuntu, mint, suse, fedora, debian, centos, arch, and slackware linux as well as Gentoo and FreeBSD. There are also 32-bit i686 ubuntu, debian, and slackware versions available. These are updated on a fairly regular basis as long as significant code changes have been made. They are in subdirectories of:

<https://cinelerra-gg.org/download/tars>
<https://cinelerra-gg.org/download/pkgs>

The “**tars**” directory contains single-user static builds for different distros. This is the recommended usage of Cinelerra-GG because all of the files will exist in a single directory. To install the single user builds, download the designated tarball from the ./tars subdirectory and unpack as indicated below:

```
cd /path
mkdir cin
cd cin
tar -xJf /src/path/cinelerra-5.1-*.txz # for the *, substitute your distro tarball name
```

Do NOT download the LEAP 10-bit version unless you use h265 (it can't render 8-bit h265).

The “**pkgs**” directory contains the standard packaged application for various distros. This will install a dynamic system version for users who prefer to have the binaries in the system area and for multi-user systems. In addition, performing the package install checks the md5sum in the file md5sum.txt to ensure the channel correctly transmits the package. There is a README.pkgs file in the “pkgs” directory with instructions so you can “cut and paste” and avoid typos; it is also shown next.

README.pkgs

Depending on the distro, use the instructions below and select the appropriate setup operations to install, update or remove cinelerra-gg infinity. (03/04/2019)
To upgrade, refresh repo, then replace "install" with "update", or whatever.

Email problems to cin@lists.cinelerra-gg.org

If repository problems, usually you can manually do an install by using:

```
wget https://cinelerra-gg.org/download/pkgs/{substitute_name}/cin_5.1.<sub_name>.deb
and install it manually, for example: dpkg -i cin_5.1.{substitute_filename}.deb
```

GENTOO – courtesy Dominique Michel

There is an ebuild package at this time as of 01/03/2019 at:

```
# https://svnweb.tuxfamily.org/listing.php?repname=proaudio%2Fproaudio&path=
# %2Ftrunk%2Foverlays%2Fproaudio%2Fmedia-video%2Fcinelerra
# %2F&#ab000caf7024d83112f42a7e8285f2f29
```

FREEBSD – courtesy Yuri

There is a port available at: <https://www.freshports.org/multimedia/cinelerra-gg/>

To use this port: `cd /usr/ports/multimedia/cinelerra-gg && make install clean`

and then install this precompiled package via: `pkg install cinelerra-gg`

```

# FEDORA
# Replace the XX in fedoraXX in the next line with your current O/S version number
dnf install cinelerra --nogpgcheck --repofrompath
cingg,https://cinelerra-gg.org/download/pkgs/fedoraXX/
##dnf erase cinelerra

# CENTOS
# Python 2 has been updated for other distros to Python 3 so you might have to create a soft link
# to get the correct version. For help, send email to cin@lists.cinelerra-gg.org
# first create the file /etc/yum.repos.d/cin_gg, with the following contents:
[cin_gg]
name=cingg
baseurl=https://cinelerra-gg.org/download/pkgs/centos7
gpgcheck=0
# end of cin_gg
yum install cinelerra
##yum erase cinelerra

# UBUNTU, replace ub14 with your distro id: ub16,ub17,ub18
# Some ubuntu apt downloads register status as working 0% constantly while running the package
# download, like ubuntu 14. It may take a few minutes for this step so be patient.
apt install software-properties-common apt-transport-https
apt-add-repository https://cinelerra-gg.org/download/pkgs/ub14
# UBUNTU 16/17/18 note - This has been known to work, but things change quickly:
# VIP - for the first install, the above line adds cinelerra to /etc/apt/sources.list but...
# Version 16/17/18 of Ubuntu are more strict for licensing so you will have to edit
# the file /etc/apt/sources.list to add [trusted=yes] after deb and before https...cin...
# For example the line should be: deb [trusted=yes] https://cinelerra-gg.org/download/pkgs/ub16 xenial
main
# Or for ub17: deb [trusted=yes] https://cinelerra-gg.org/download/pkgs/ub17 zesty main
# Or for ub18: deb [trusted=yes] https://cinelerra-gg.org/download/pkgs/ub18 bionic main
# Also, on the install you will get an error message that you can either ignore as cinelerra
# will run anyway, or else (the first time only) on the command line keyin:
# echo > /etc/sysctl.d/50-cin.conf "kernel.shmmax=0x7ffffff"
apt update
apt install cin
#to update a previous install (ignore any i386 errors as only 64 bit version available):
apt update
apt upgrade cin
##apt remove cin

# MINT should use the same procedure as Ubuntu, but:
# Note: apt-add-repository did not work for me, I had to use the gui version:
# System_OR_Administration->Software Sources->Additional Repositories->Add a new repository
# For Mint18,add: deb [trusted=yes] https://cinelerra-gg.org/download/pkgs/mint18 xenial main
# For Mint19,add: deb [trusted=yes] https://cinelerra-gg.org/download/pkgs/mint19 bionic main
apt update
apt install cin

```

```
#to update a previous install
apt update
apt upgrade cin
##apt remove cin
```

```
# DEBIAN uses the same basic procedure as Ubuntu.
# The apt-add-repository varies per system so you will have to use your best judgement
apt install software-properties-common apt-transport-https
apt-add-repository https://cinelerra-gg.org/download/pkgs/debian8
OR apt-add-repository https://cinelerra-gg.org/download/pkgs/debian9
# VIP - for the first install, the above line adds cinelerra to /etc/apt/sources.list but...
# Debian stretch and jessie are more strict for licensing so you will have to edit
# the file /etc/apt/sources.list to add [trusted=yes] after deb and before https...cin...
# For example for debian8: deb [trusted=yes] https://cinelerra-gg.org/download/pkgs/debian8 jessie
main
# For example for debian9: deb [trusted=yes] https://cinelerra-gg.org/download/pkgs/debian9 stretch
main
apt update
apt install cin
#to update a previous install
apt update
apt upgrade cin
##apt remove cin
```

```
# SUSE/LEAP
# (Note: you may have to zypper libavc and libiec versions if not already installed)
# cinelerra packages are unsigned so you will have to ignore: Package is not signed!
# openSUSE LEAP 15
zypper ar -f https://cinelerra-gg.org/download/pkgs/leap15/ cingg
zypper install -r cingg cinelerra # or cinelerra10bit for 10 bit
# openSUSE LEAP 42
zypper ar -f https://cinelerra-gg.org/download/pkgs/leap42/ cingg
# as of 42.3 SUSE there is a new requirement, so you will need to add:
zypper mr -G cingg
zypper install -r cingg cinelerra # or cinelerra10bit for 10 bit
##zypper remove cinelerra # or cinelerra10bit for 10 bit
#to update a previous install (assuming you enabled autorefresh as above)
zypper refresh cingg
zypper up cinelerra # or cinelerra10bit for 10 bit
```

```
# SLACKWARE, substitute slk32 for slk64 and i486-1 for x86_64-1
wget -P /tmp https://cinelerra-gg.org/download/pkgs/sl64/cin-  
{date}-slk64-x86_64.txz
installpkg /tmp/cin... # name you used in the above line
#to update a previous install
upgradepkg /tmp/cin... # name you used in the above line
##removepkg cin
```

```
# ARCH linux
# first edit the file /etc/pacman.conf, to include the following:
```

```
[cingg]
SigLevel = Optional TrustAll
Server = https://cinelerra-gg.org/download/pkgs/arch
# end of cingg
pacman -Sy
pacman -S cin
##pacman -R cin
```

2.3 Distribution Systems with Cinelerra Included

There are also some special complete distribution systems available that include Cinelerra-GG for audio and video production capabilities.

AV Linux is a downloadable/installable shared snapshot ISO image based on Debian. It provides the user an easy method to get an Audio and Video production workstation without the hassle of trying to find and install all of the usual components themselves. Of course, it includes Cinelerra-GG! It is at:

<http://www.bandshed.net/avlinux/>

Bodhi Linux is a free and open source distribution that comes with a curated list of open source software for digital artists who work with audio, video, includes Cinelerra GG, games, graphics, animations, physical computing, etc. It is at:

<https://gitlab.com/giuseppetorre/bodhilinuxmedia>

2.4 Cinx and a “Bit” of Confusion

Cinx is the exact same program as Cin. The X (x) represents the roman numeral 10 for 10-bit as opposed to 8-bit standard. The third-party library used for x265 must be specially compiled with `--bit-depth=10` in order to produce 10-bit rendered output. This build will not be able to output 8-bit depth which means you have to retain the Cin version also. Whatever build ffmpeg is linked to will determine what bit depth it can output. This is why there have to be separate builds. If you install both packages, Cin and CinX, you may get “file conflicts of same file name” - just continue.

Keep in mind that the regular 8-bit version works on 8-bit bytes – the standard word size for computers, but the 10-bit version has to use 2 words to contain all 10 bits so you can expect rendering to be as much as twice as slow. There is also a 12-bit version for consideration but currently the results are simply the same as 10-bit with padding to make 12-bit so it is of no value.

This page intentionally left blank so that chapters begin on a right facing page.

3 – The 4+ Windows

3.1 Program Window

The main window is called the Program window and contains the timeline as well as the entry point for all menu driven operations. It is often just called the “timeline”. The timeline consists of a vertical stack of tracks with a horizontal representation of time. This defines the output of rendering operations and what is saved when you save files. To the left of the timeline is the patchbay which contains options affecting each track. The patchbay is described in detail in the Editing section.

The *Window* pulldown on this main window contains options that affect the 4 main windows. *Default positions* repositions all the windows to a 4 screen editing configuration. On dual headed displays, the *Default positions* operation fills only one monitor with windows.

3.1.1 Video and Audio Tracks and Navigation

The program window contains many features for navigation and displays the timeline as it is structured in memory: tracks stacked vertically and extending across time horizontally. The horizontal scroll bar allows you to scan across time. The vertical scroll bar allows you to scan across tracks.



Figure 1: Patchbay | Timeline with pulldowns & navigation icons, Video/Audio tracks & bottom Zoom

Video tracks represent the duration of your videos and clips, just as if you placed real photographic film stock end-to-end on a table. The individual images you see on the track are samples of what is located at that particular instant on the timeline.

Audio tracks represent your sound media as an audio waveform. Following the film analogy, it would be as if you "viewed" magnetic tape horizontally on your table. You can adjust the horizontal and vertical magnification of the tracks and the magnification of the audio "waveform" display using the zoom panel controls. Every track on the timeline has a set of attributes on the left, called the patch-bay. It is used to control some of the behavior of the tracks.

Track Navigation involves both selecting a specific audio or video track and moving to a certain time in the track. The vertical scroll bar allows you to scan across tracks. For vertical scrolling you can also use the mouse wheel. The horizontal scroll bar allows you to scan across time. For horizontal scrolling you can use the mouse wheel with the Ctrl key.

In addition to the graphical tools, you can use the keyboard to navigate. There is a shortcuts document for keyboard navigation; it includes, for example, shortcuts like use the Home and End keys to instantly go to the beginning or end of the timeline. Or in the default cut and paste mode, hold down Shift while pressing Home or End in order to select the region of the timeline between the insertion point and the key pressed.

3.1.2 Zoom Panel

Below the timeline, you will find the zoom panel. The zoom panel contains values for sample zoom (duration visible on the timeline), amplitude (audio waveform scale), track zoom (height of tracks in the timeline), and curve zoom (automation range). In addition to the scrollbars, these zooms are the main tools for positioning the timeline. Also on the zoom panel is selection change and alpha slider.



Figure 2: Zoom panel on the bottom of the main program window

Changing the *sample zoom* causes the unit of time displayed in the timeline to change size. It allows you to view your media all the way from individual frames to the entire length of your project. The higher the setting, the more frames you can see per screen. The sample zoom value is not an absolute reference for the unit of time since it refers to the duration visible on the timeline and thus changes also as you modify the length of the program window horizontally. Use the Up and Down arrows to change the sample zoom by a power of two. Or if your mouse has a wheel, mouse over the tumblers and use the wheel to zoom in and out.

The *amplitude* only affects audio which determines how large the waveform appears. Ctrl-up and Ctrl-down cause the amplitude zoom to change.

The *track zoom* affects all tracks. It determines the height of each track. If you change the track zoom, the amplitude zoom compensates so that the audio waveforms look proportional. Ctrl-pgup and Ctrl-pgdown cause the track zoom to change.

The *curve zoom* affects the curves in all the tracks of the same type. It determines the value range for curves. First select the automation type (audio fade, video fade, zoom, X,Y) then use the left tumblers for the minimum value and the right tumblers for the maximum value or manually enter the values in the text box. Normally you will use -40.0 to 6.0 for audio fade and 0.0 to 100.0 for video fade. The tumblers change curve amplitude, but the only way to curve offset is to use the fit curves button.

The *selection start time*, *selection length*, and *selection end time* display the current selected timeline values. The *alpha slider* allows for varying the alpha value when using colors on the tracks as set in your appearance preferences for Autocolor assets. It has no function without that flag set.

3.1.3 Track Popup Menu

Each Track has a popup menu. To activate the track popup menu, Right mouse click on the track. The popup menu affects the track whether the track is armed on the patchbay or not. The Track Menu contains a number of options:

Attach Effect - opens a dialog box of effects applicable to the type of track of audio or video.

Move up - moves the selected track one step up in the stack.

Move down - moves the selected track one step down in the stack.

Delete track - removes the track from the timeline.

Add Track - adds a track of the same media type, audio or video, as the one selected above that track.

Find in Resources - that media file will be highlighted in the media folder in the Resources window.

Show edit - will point out the exact start and stop points along with the length of the current edit on that track as well as the media name.

User title - is used to change the title name. This is really handy for files that have very long and similar names that would get cut off during edits. You can use short names to better differentiate the media. If you select multiple, all those clips will have title name changed.

Bar color – allows the user to select a specific color for the title bar. This helps ease of locating.

Resize Track - resizes the track.

Match Output Size - resizes the track to match the current output size.

3.1.4 Insertion Point

The insertion point is the flashing hairline mark that vertically spans the timeline in the program window. Analogous to the cursor on your word processor, the insertion point marks the place on the timeline where the next activity will begin. It is the point where a paste operation takes place. When rendering, it defines the beginning of the region of the timeline to be rendered. It is also the starting point of all playback operations.

Normally, the insertion point is moved by clicking inside the main timebar. Any region of the timebar not obscured by labels and in or out points is a hotspot for repositioning the insertion point. In cut and paste editing mode only, the insertion point can be moved also by clicking in the timeline itself. When moving the insertion point the position is either aligned to frames or aligned to samples. When editing video, you will want to align to frames. When editing audio you will want to align to samples. Select your preference by using Settings->Align cursor on frames.

Pulldowns ->
Transport & Buttons Bar ->

Timebar ->
Media Title ->

Video Track ->

Audio Track ->

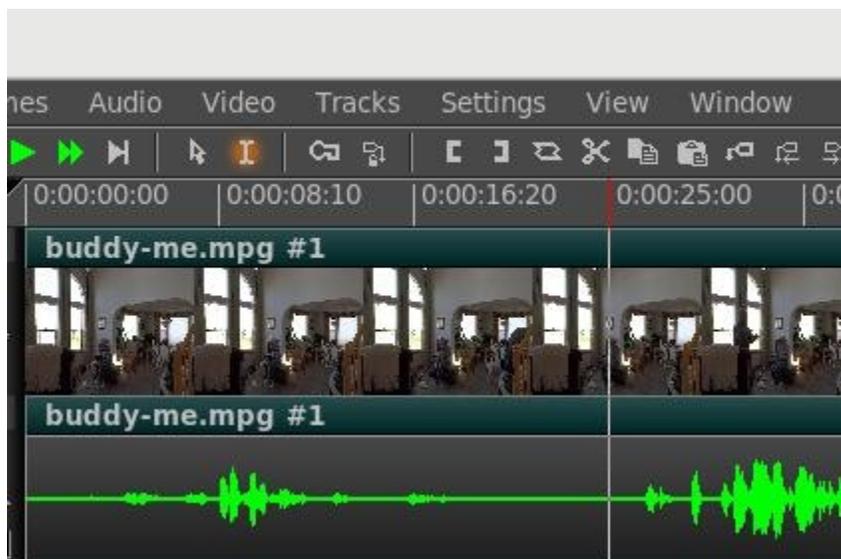


Figure 3: Insertion point is at 0:00:25:10 in Hr:Mn:Sec:Frames

3.1.5 Editing Modes

There are 2 different editing methods of operation that affect the insertion point and the editing on the timeline. There is: *drag and drop mode* and *cut and paste mode*. The editing mode is determined by selecting the arrow or the I-beam in the Transport and Buttons bar.

If the arrow is highlighted, it enables *drag and drop mode*. In drag and drop mode, clicking in the timeline does not reposition the insertion point. Double-clicking in the timeline selects the entire edit the mouse pointer is over. Dragging in the timeline repositions the edit the mouse pointer is over. This is useful for reordering audio playlists, sorting movie scenes, or moving effects around. To cut and paste in drag and drop mode you need to set in/out points to define an affected region.

If the I-beam is highlighted it enables *cut and paste mode*. In cut and paste mode, clicking in the timeline repositions the insertion point. Double-clicking in the timeline selects the entire edit the cursor is over. Dragging in the timeline highlights a region. The highlighted region becomes the region affected by cut and paste operations and the playback range during the next playback operation. Shift-clicking in the timeline extends the highlighted region.

When highlighting a region, the start and end points are either aligned to frames or aligned to samples. When editing video, you will want to align to frames. When editing audio you will want to align to samples. Select your preference by using settings->align cursor on frames.

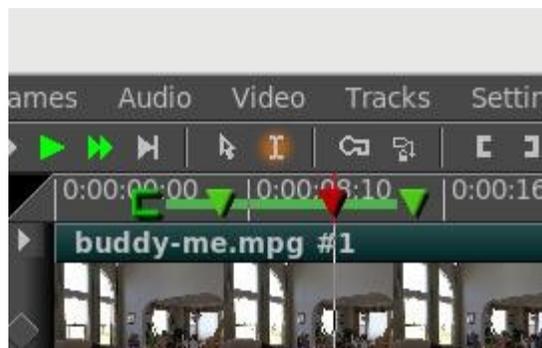


Figure 4: I-beam + in/out + labels

3.1.6 In/Out Points

In both editing modes, you can set one In point and one Out point. The in/out points define the affected region. In drag and drop mode, they are the only way to define an affected region. In both cut and paste mode and drag and drop mode, the highlighted area overrides the In/Out points. If a highlighted area and In/Out points are set, the highlighted area is affected by editing operations and the In/Out points are ignored. If no region is highlighted, the In/Out points are used. To avoid confusion, it is better to use either highlighting or In/Out points but not both simultaneously.

To set in/out points, go to the timebar and position the insertion point somewhere. Select the In point button. Move the insertion point to a position after the In point and click the Out point button. Instead of using the button bar, you can use the [or < and] or > keys to toggle in/out points.

If you set the insertion point somewhere else while In/Out points already exist, when you click the In/Out buttons the existing points will be repositioned. If you click on in/out points while a region is highlighted, the insertion point will be ignored and In/Out points will be set at the beginning and at the end of the highlighted area.

If you select either the In point or the Out point, the insertion point will jump to that location. After selecting an In point, if you click the In point button the In point will be deleted. After selecting an Out point, if you click the Out point button the Out point will be deleted. Shift-clicking on an In/Out point highlights the region between the insertion point and that In/Out point. If a region is already highlighted, it extends the highlighted region up to that In/Out point.

To quickly get rid of In/Out points, without caring about where they are or if they are set or not, just double click on [and] buttons. The first click will set a new point or reposition an old one at the insertion point; the second click will delete it. This trick does not work if the In point or the Out point is already set at insertion point.

Some of the useful operations concerning the In/Out pointers are listed next.

- Ctrl-KeyPad# - if in/out set, KP 2,3,5,6 + Enter, play between In/Out point
- Shift-Ctrl – loops play between In/Out points
- Click in/out while holding the left mouse button, drags In/Out pointer elsewhere
- Shift-Ctrl with transport button, loops play between In/Out points
- Ctrl-t – clears both In/Out points

3.1.7 Labels

The insertion point and the In/Out points allow you to define an affected region, but they do not let you jump to exact points on the timeline very easily. Labels are an easy way to set exact locations on the timeline that you want to jump to. When you position the insertion point somewhere and click the label button, a new label appears on the timeline. With label traversal you can quickly seek back and forth on the timeline.

No matter what the zoom settings are, clicking on the label highlights it and positions the insertion point exactly where you set the label. The lower case letter “L” is a shortcut for the label button.

Labels can reposition the insertion point when they are selected but they can also be traversed with the label traversal buttons. When a label is out of view, the label traversal buttons reposition the timeline so the label is visible. Keyboard shortcuts for label traversal are:

Ctrl-left repositions the insertion point on the previous label.

Ctrl-right repositions the insertion point on the next label.

The Label folder in the Resources window lists the timestamp of every label. You can edit the label list and add a title for every item using the popup menu. To open the Label info dialog right click on the label icon in the Resources window or directly on the label symbol on the timebar. With labels you can also select regions:

Shift-Ctrl-left highlights the region between the insertion point and the previous label.

Shift-Ctrl-right highlights the region between the insertion point and the next label.

Double-clicking on the timebar between two labels highlights the region between the labels.

Shift-clicking on a label highlights the region between that label and the insertion point.

If a region is already highlighted, it extends the highlighted region up to that label.

If you hit the label button when a region is highlighted, labels are created at each end of the highlighted region. However, if one end already has a label, then the existing label is deleted. Hitting the label button again when a label is selected deletes it. Manually hitting the label button or L key over and over again to delete a series of labels can get tedious. To delete a set of labels, first highlight a region, then use the Edit->Clear labels function. If in/out points exist, the labels between the in/out points are cleared and the highlighted region is ignored.

In Cut and Paste editing mode only, by enabling *Edit labels* in the settings menu or by disabling the *Lock labels from moving* button on the program toolbar, labels will be cut, copied or pasted along with the selected region of the first armed track. Similarly, if a selected area of a resource is spliced from the viewer to the timeline in a position before labels, these labels will be pushed to the right on the timebar for the length of the selected area. To prevent labels from moving on the timebar, just disable the *Edit labels* option or enable the *Lock labels from moving* button.

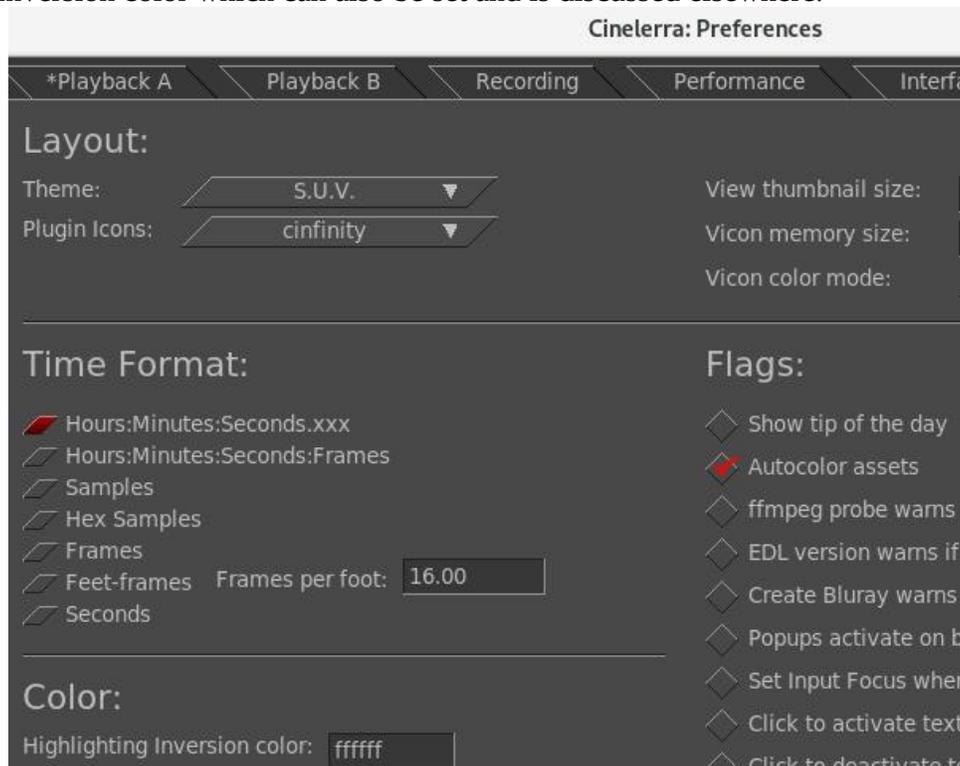
Originally in Drag and Drop editing mode labels will be always locked to the timebar, even with the *Edit labels* option enabled. This may no longer be correct in all cases.

3.1.8 Color Title Bars and Assets

In order to visually aid in locating clips on the timeline that are from the same media file, you can have them auto-colored or self-colored. Use of this feature requires additional memory and cpu on every timeline redraw, therefore it is recommended that smaller computers leave it turned off.

For auto-color the color will be based on a hashed filename so that whenever you load this particular media, it will always have the same color on the title bar even if you use proxy. To enable auto-color, go to Settings → Preferences, Appearance tab and check on “Autocolor assets”. It is disabled by default. Each media will have a random muted color and there could easily be close duplicates as generated by the program algorithm. There will be no total black, but some dark shades are possible.

Screencast shows the red colored checkmark to enable Autocolor assets. In the lower left corner is Highlighting Inversion color which can also be set and is discussed elsewhere.



To change a specific clip to your own chosen color, middle mouse button over that clip and an Edits popup will be displayed. Choose the option *Bar Color* to bring up the color picker and choose a color. You can also change the alpha value in the color picker and this alpha takes precedence over the current alpha slider bar value unless it was set to 1.0. The color will only change after you click on the checkmark. The *Bar Color* option works in either Drag and Drop or Cut and Paste editing mode and also works if “Autocolor assets” is not set. In Drag and Drop editing mode, if you select several clips and then bring up the Edits popup with the middle mouse button over a track, you can use the *Bar Color* option to change all of those selected to the same color.

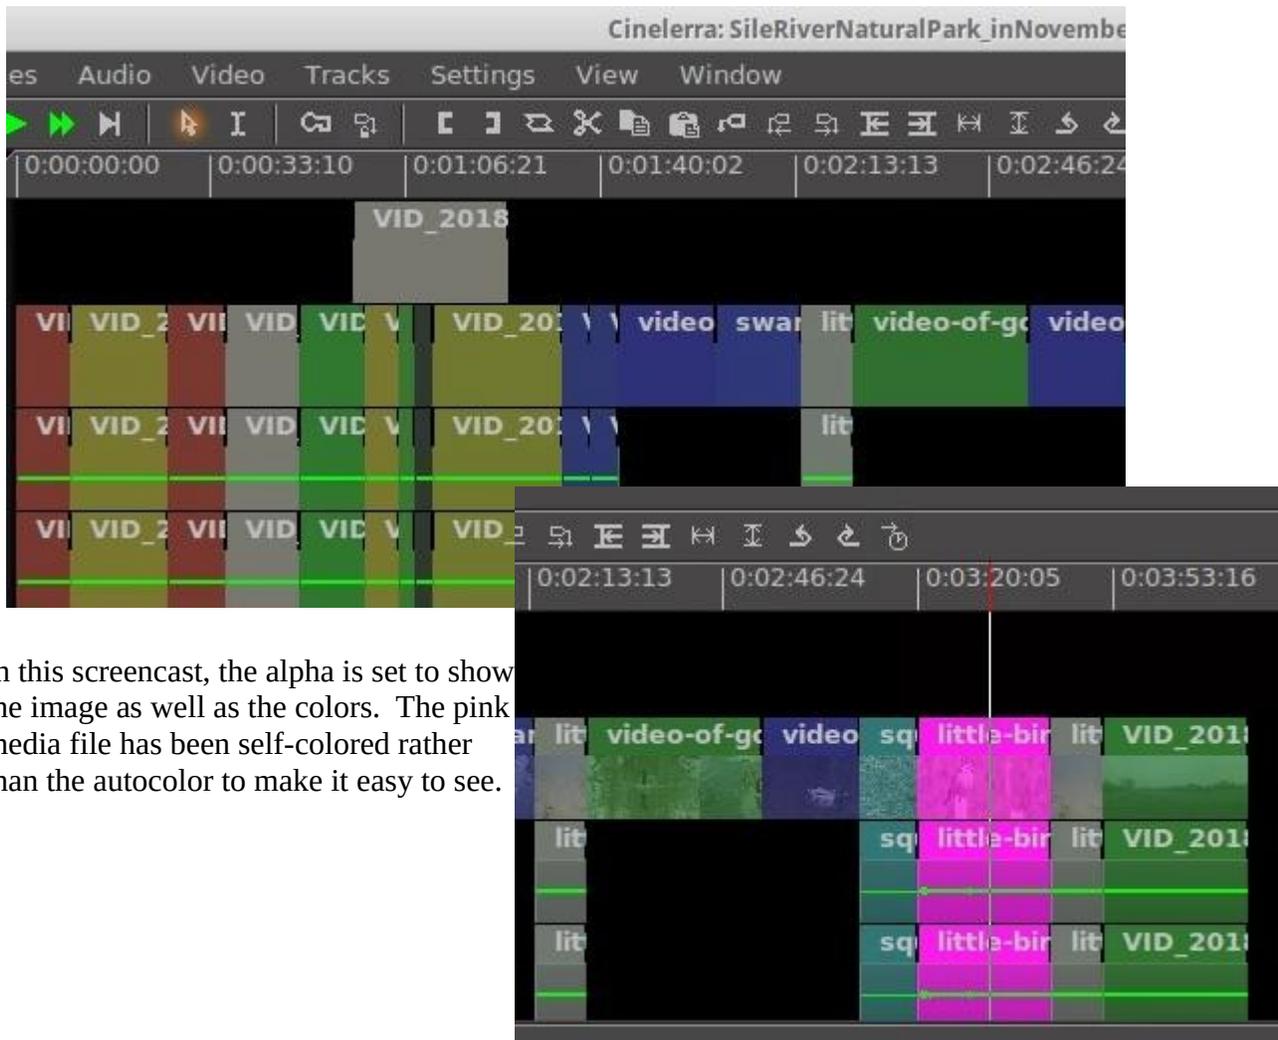
To go back to the the default colors, uncheck “Autocolor assets” in Preferences, but this does not affect the specially chosen self-colored ones as they are preserved. To change these individually or selectively use the Edits popup *Bar Color* option and click on “Default” in the color picker window. Auto-color does not honor armed/disarmed tracks. Self-color does honor armed/disarmed tracks.

And that’s not all! There is an *alpha fader slider bar* on the bottom of the main window on the right hand side of what is referred to as the Zoom Panel. With this alpha slider, you can colorize your video and audio tracks to either see only the color at 0.0 or see only the image/audio waveform at 1.0. This slider bar affects all colored areas of the Autocolor assets and the self-colored ones. In the case when a specifically changed edit alpha value is any value except 1, the slider bar will not affect that. Once you use the slider bar, it is activated so gets first shot at any keystrokes in the main window. You deactivate this by simply clicking in a different part of the main window.

As long as we are on the subject of color, just a reminder that you can also change the “Highlighting Inversion color” in Settings → Preferences, Appearance tab. This is on right left hand side of the menu

more than half the way down and you can see this in the previous screencast. That setting defaults to white (ffffff) but sometimes this is a little bright so you can put any hex value in that suits you.

Screencast which shows an example of the Autocolor assets with alpha set to 0.0.



In this screencast, the alpha is set to show the image as well as the colors. The pink media file has been self-colored rather than the autocolor to make it easy to see.

3.1.9 More about Pulldowns

The main window pulldowns are quite obvious in their meaning and usage, so here is only a summary. Figure 3 shows an example of the pulldowns as displayed in the main window.

File – options for loading, saving, and rendering as described in other sections.

Edit – edit functions; most of which have shortcuts that you will quickly learn.

Keyframes – keyframe options which are described in the Keyframe section.

Audio – audio related functions such as “Add track”, “Attach transition/effect”.

Video – video functions such as “Default/Attach transition”.

Tracks – move or delete tracks are the most often used.

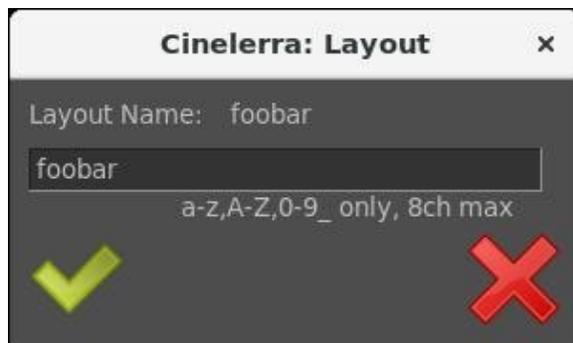
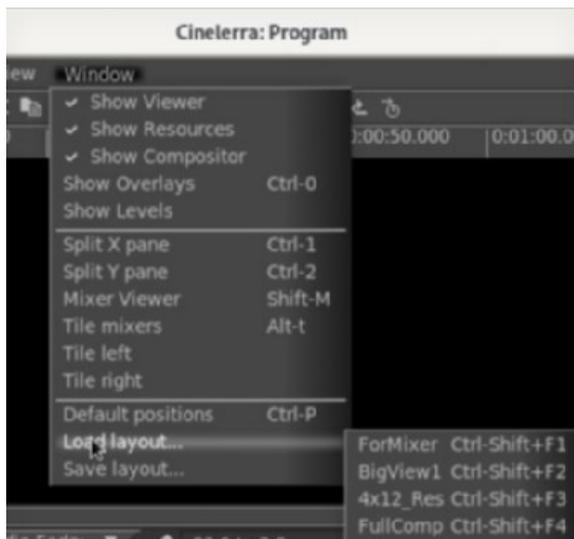
Settings – this is mostly described in other sections. However, typeless keyframes are not described anywhere else. They allow keyframes from any track to be pasted on either audio or video tracks.

View – for display or modifying asset parameters and values to include Fade, Speed, and Cameras.

Window – window manipulation functions.

3.1.10 Window Layouts

If you like to use different window layouts than the default for certain scenarios, you can setup, save, and load 4 options. First position your Cinelerra windows where you want them to be and then use the Window pulldown and choose *Save layout*. To use the default name of Layout #, when the popup comes up, just click the green checkmark OK on the Layout popup menu. If you would like a specific name for your layout so you can remember what it is for, keyin 1-8 english characters that are meaningful to you (english characters mean you can not use the German umlaut or the French accent). Legal characters are a-z, A-z, 0-9, _ (the underscore character) and a limit of 8 total. If you keyin more than 8, only the last 8 characters will be used. To rename a currently existing layout, use the Save layout option again on the one to rename, and keyin a different name into the text box or blank for the default name.



The files containing the coordinates for your layouts will automatically be saved in the \$HOME/.bcast directory as layout#_rc or layout#_8chars_rc.

To use the desired layout, keyin the shortcut or use the Window pulldown and choose *Load layout* and then make your choice.

3.1.11 Just Playing!

What if you are just using Cinelerra to play media and listen to tunes? After loading your media, just hit the space bar to start playing and then again to stop playing. Other than that, use the transport buttons on the top bar of the Program window. Other ways, not previously mentioned to “play around” are described next.

Repeat Play / Looping Method

There are 2 methods for repeat play or looping on the timeline and 1 method for both the Compositor and the Viewer. This works in conjunction with any of the transport buttons or shortcuts in either forward or reverse as usual. The 1 exception is that the Shift button can not be used to either add or subtract audio within the repeat area.

Shift-L on the Timeline, repeats the selection per the algorithm outlined next. When setup, long green lines are displayed across the entire set of tracks which shows the start and end of the loop.

- 1) Highlighted selection repeats loop and takes precedence over all other possibilities. If the cursor is before the highlighted area, it will play up to the area and then repeat the highlighted section. If the cursor is after the highlighted section, play will start at the beginning until you get to the highlighted section and then repeat.
- 2) When both In and Out pointers are set, it repeats the section between [and].
- 3) If only one of the In or Out pointers is set, it loops the whole media.

Ctrl+Shift+transport button on the Timeline, Viewer, and Compositor

- 1) Repeats entire media if no In or Out pointer set.
- 2) In and Out pointer set, repeats area between pointers.
- 3) Only In pointer set, repeats from In to end of media.

Last Play Position Memory

When you play media, the start/end playback positions are saved as if they had been made into temporary labels. They appear on the timeline as purple/yellow hairline markers representing the last start/end labels for the last playback. They can be addressed as if they are label markers using:

- Ctrl ← tab to the label before the cursor, that is “play start”
- Ctrl → tab to the label after the cursor, that is “play stop”

You can use these markers for re-selection. Additionally, the selection region can be expanded by “pushing” the markers using single frame playback. Use frame reverse (keypad 4) to push the start play marker backward, or use frame forward (keypad 1) to push the end play marker forward.

Another handy feature is to use the combination of Ctrl-shift-arrow (left or right) to select the media from the cursor position (red hairline) to the start or end marker by “tabbing” to the label markers. For example, tab to the beginning of the previous play region using Ctrl-left-arrow to move the cursor to the beginning of last play, then press Ctrl-Shift-right-arrow to tab to the end of the playback region. Now you can clip/play/expand or edit the previous playback selection.

- Ctrl SHIFT → tab cursor to label right of cursor position and expand selection
- Ctrl SHIFT ← tab cursor to label left of cursor position and expand selection

Playback Speed Automation Support

The speed automation causes the playback sampling rate to increase or decrease to a period controlled by the speed automation curve. This can make playback speed-up or slow-down according to the scaled sampling rate, as “time is multiplied by speed” (speed X unit_rate).

Alternative to using Numeric Keypad for Playing

For the keyboards without a numeric keypad or if you prefer to use keys closer to where you normally type, there are alternative keys for the play/transport functions. These are listed below.

Alt + m = stop playback

Alt + j = forward single frame

Alt + k = forward slow playback

Alt + l = forward normal playback

Alt + ; = forward fast playback

Alt + u = reverse single frame

Alt + i = reverse slow playback

Alt + o = reverse normal playback

Alt + p = reverse fast playback

+ Shift key with the above, results in the reverse of whether audio is included or not.

+ Ctrl with the above, results in the transport function operating only between the in/out pointers.

3.2 Compositor Window

The Compositor window displays the output of the timeline. It is the interface for most compositing operations or operations that affect the appearance of the timeline output. Operations done in the Compositor affect the timeline but do not affect clips.

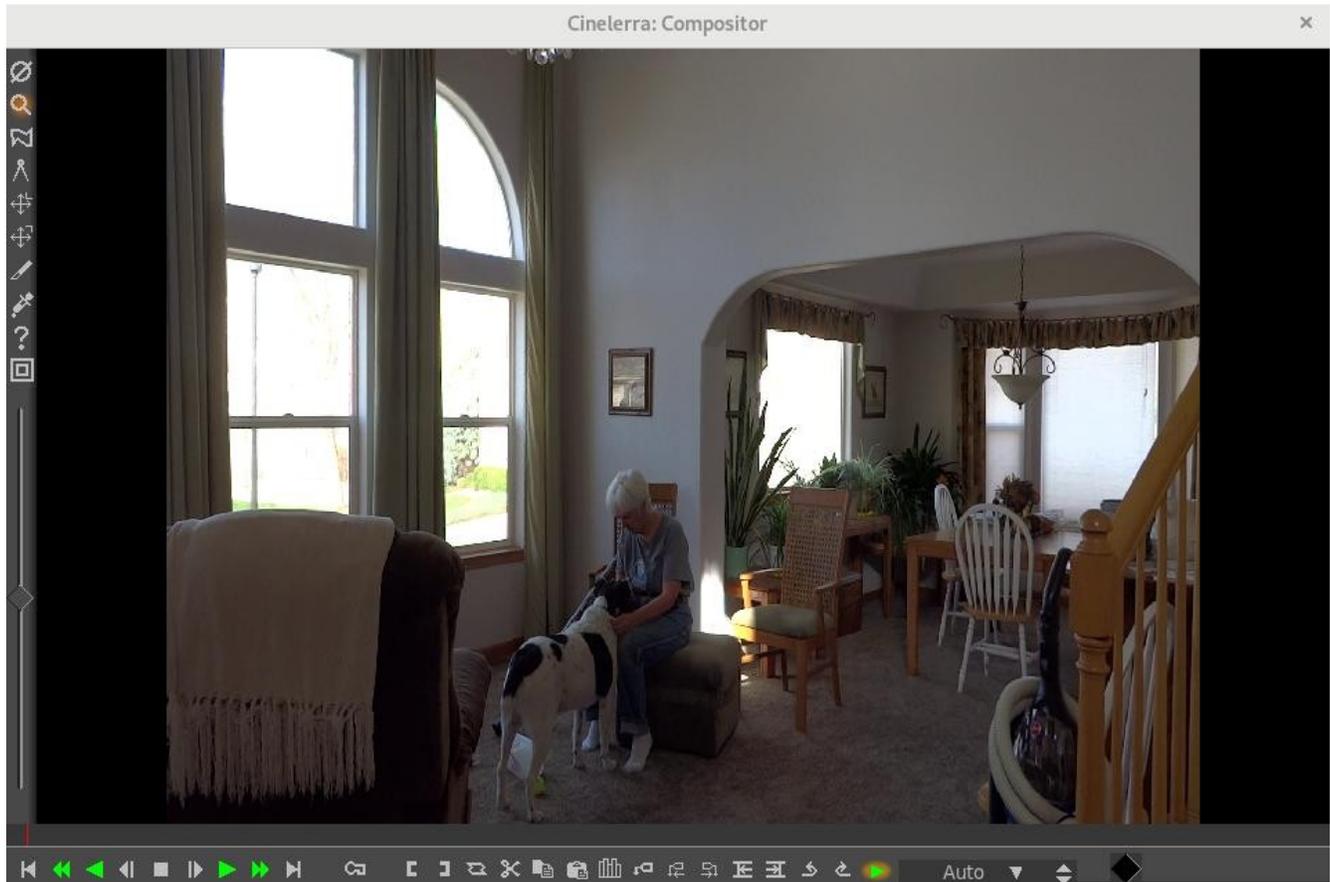


Figure 5: Upper right side contains navigation tools / bottom bar has manu control functions

3.2.1 Compositor controls

Navigating the video output does not affect the rendered output; it just changes the point of view in the compositor window. The video output has several navigation functions. The video output size is either locked to the window size or unlocked with scrollbars for navigation. The video output can be zoomed in and out and panned. If it is unlocked from the window size, middle clicking and dragging anywhere in the video pans the point of view. Hitting the + and - keys zooms in and out of the video output.

Underneath the video output are copies of many of the functions available in the main window. In addition there is a zoom menu and a tally light. The zoom menu jumps to all the possible zoom settings and, through the Auto option, locks the video to the window size. The zoom menu does not affect the window size. The tally light turns red when rendering is happening. This is useful for knowing if the

output is current. Right clicking anywhere in the video output brings up a menu with all the zoom levels, zoom auto mode, and some other options. In this particular case the zoom levels resize the entire window and not just the video. The *Reset camera* and *Reset projector* options center the camera and projector. The *Hide controls* option hides everything except the video.

On the left of the video output is a toolbar specific to the compositor window. The toolbar has the following functions:

Protect video - disables changes to the compositor output from clicks in it. It is an extra layer on top of the track arming toggle to prevent unwanted changes.

Magnifying glass - this tool zooms in and out of the compositor output without resizing the window. If the video output is currently locked to the size of the window, clicking in it with the magnifying glass unlocks it and creates scrollbars for navigation.

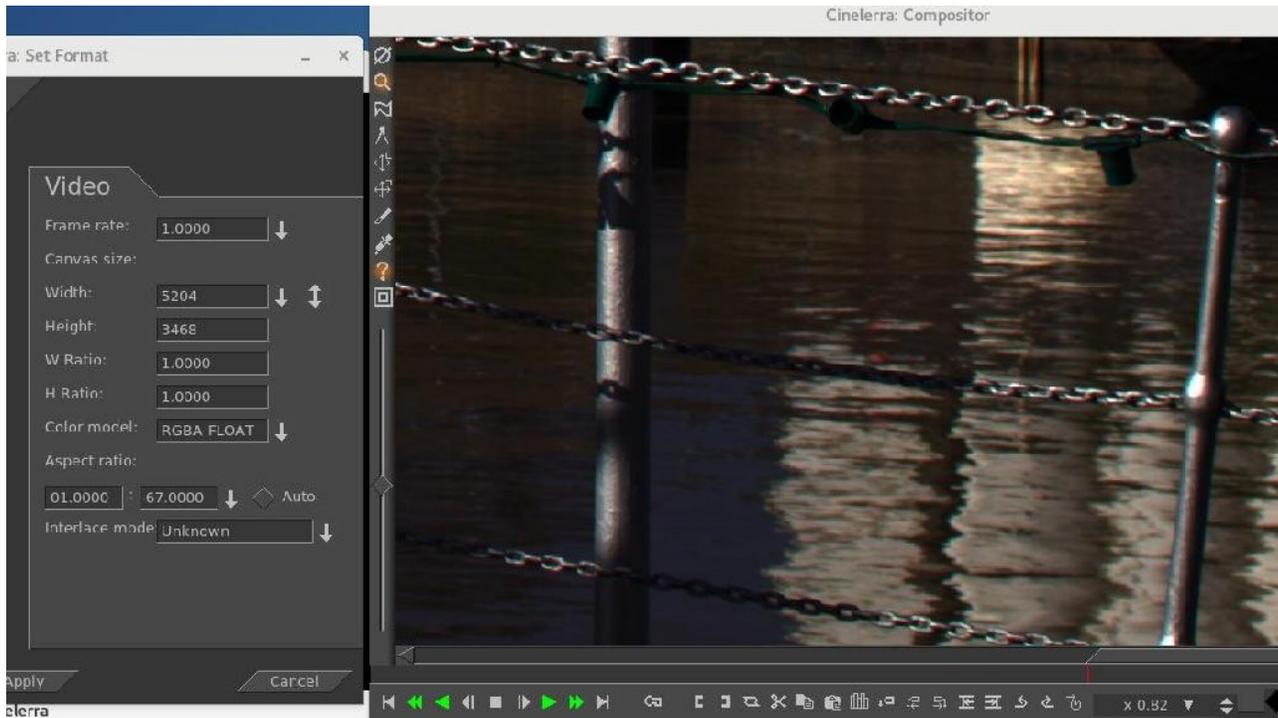
Left clicking in the video zooms in;

Ctrl clicking in the video zooms out;

Rotating the wheel on a wheel mouse zooms in and out.

In addition, if you enable the Magnifying glass, a zoom slider for fine-viewing appears below these tools. It allows you to zoom to most any size. A “zoom slider” will pop-up towards the bottom on the left-hand side of the Compositor when you enable “Zoom view” via the magnifying glass or when you click on the icons for “Adjust camera automation” or “Adjust projector automation”. This will allow for adjusting the amount of zoom at any level between .01 and 100 based on a logarithmic scale. When using the zoom slider, the number by which the view is zoomed can be seen in the textbox where the original-also-working % zoom is located. The zoom slider size is in the form of “times”, such as x 0.82 which indicates that the picture is zoomed to 82/100th of the original size as seen in Settings → Format. Once you have set the zoom to the desired size, use the vertical and horizontal scroll bars to position the view as needed.

Screencast shows below at  a zoom slider bar with the diamond shaped slider in the middle. Note that the magnifying glass is  enabled which automatically pops-up the slider.



The Format shows a large 5204x3468 video and the box at the arrow shows x 0.82 size. 

Masks tool - this tool brings up the mask editing tool. Enable “Show tool info” to see the options.

Camera – the camera brings up the camera editing tool. Enable “Show tool info” to see options.

Projector – the projector brings up the projector editing tool. Enable “Show tool info” for options.

Crop tool - this tool brings up the cropping tool. “Show tool info” must be enabled to use this tool.

Eyedropper - brings up the eyedropper. The eyedropper detects whatever color is under it and stores it in a temporary area. Enabling the “Show tool info” shows the currently selected color. Click anywhere in the video output to select the color at that point. The eyedropper not only lets you see areas which are clipped, but its value can be applied to many effects. Different effects handle the eyedropper differently.

Show tool info - this tool button works only in conjunction with the other controls on the compositor.

Based on what compositing control is active, the toggle button will activate or deactivate the appropriate control dialog box. Controls with dialog boxes are: Edit mask, Camera and Projector automation, Crop control, and Get color.

Safe regions tool - draws the safe regions in the video output. This does not affect the rendered output

3.2.2 Compositing

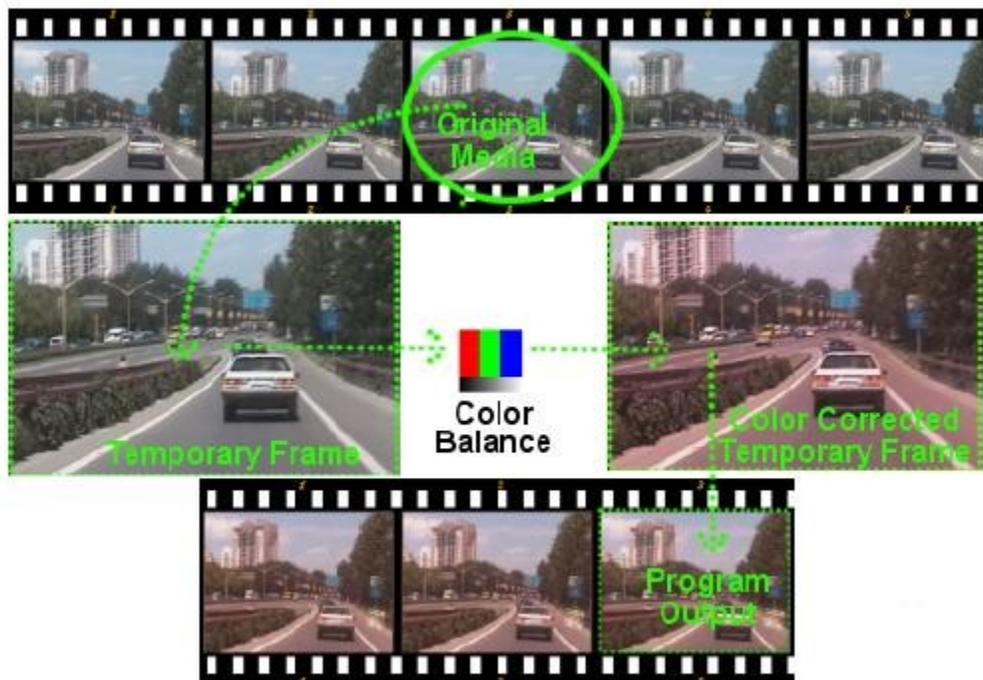
A large amount of Cinelerra's editing is directed towards compositing. Changing the resolution of a show, making a split screen, and fading in and out among other things are all compositing operations in Cinelerra. Cinelerra detects when it is in a compositing operation and plays back through the compositing engine only then. Otherwise, it uses the fastest decoder available in the hardware.

Compositing operations are done on the timeline and in the Compositor window. Shortcuts exist in the Resource window for changing some compositing attributes. Once some video files are on the timeline, the compositor window is a good place to try compositing.

3.2.3 Camera and Projector

In the compositor window, two of the more important functions are the adjust camera automation and the adjust projector automation which control operation of the camera and projector. Cinelerra's compositing routines use a "temporary", a frame of video in memory where all graphics processing is performed. Inside Cinelerra's compositing pipeline, the camera determines where in the source video the "temporary" is copied from. The projector determines where in the output the "temporary" is copied to. Each track has a different "temporary" which is defined by the track size. By resizing the tracks you can create split screens, pans, and zooms.

In compositing, each frame can be digitally altered using various options, such as a color correction plugin. Once the image has been transformed, the finished image is then projected to the compositor thus creating a modified version of the original.



When editing the camera and projector in the compositing window, the first track with record enabled is the track affected. Even if the track is completely transparent, it is still the affected track. If multiple video tracks exist, the easiest way to select one track for editing is to Shift-click on the record icon of the track. This solos the track.

The purpose of the projector is to place the contents of the "temporary" into the project's output. The intent of the projector is to composite several sources from the various tracks into one final output track. The projector alignment frame is identical to the camera's viewport, except that it guides where on the output canvas to put the contents of each temporary.

Compositing projector controls

When the projector button is enabled in the compositor window, you are in projector editing mode. A guide box appears in the video window. Dragging anywhere in the video window causes the guide box to move along with the video. Shift-dragging anywhere in the video window causes the guide box to shrink and grow along with the video. Once you have positioned the video with the projector, you may want to work with adjusting the camera automation.

Compositing camera controls

Select the camera button to enable camera editing mode. In this mode, the guide box shows where the camera position is in relation to past and future camera positions but not where it is in relation to the source video. Dragging the camera box in the compositor window does not move the box but instead moves the location of the video inside the box. The viewport is a window on the camera that frames the area of source video to be scanned. The viewport is represented as a red frame with diagonal cross bars.

Viewport sizes

The size of the viewport is defined by the size of the current track. A smaller viewport (640x400) captures a smaller area. A larger viewport (800x200) captures an area larger than the source video and fills the empty spaces with blanks. Once we have our viewport defined, we still need to place the camera right above the area of source video we are interested on. To control the location of the camera:

- 1) Open the compositor window with a track selected.
- 2) Select the camera button to enable camera editing mode.
- 3) Drag over the display window.

When we drag over the viewport in the compositor window, the way it looks is as if you “move the camera with the mouse”. The viewport also moves with it.

In the compositing window, there is a popup menu of options for the camera and projector. Right click over the video portion of the compositing window to bring up the menu.

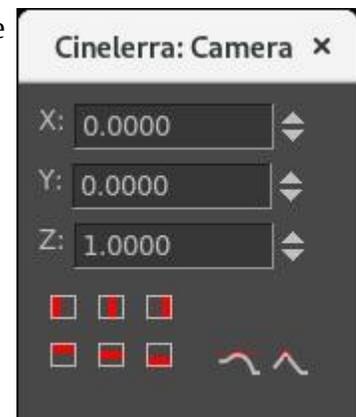
Reset Camera causes the camera to return to the center position.

Reset Projector causes the projector to return to the center.

The camera and projector tool window

The camera and projector have shortcut operations that do not appear in the popup menu and are not represented in video overlays. These are accessed in the *Show tool info* window. Most operations in the Compositor window have a tool window which is enabled by activating the question mark icon.

In the case of the camera and projector, the tool window shows x, y, and z coordinates. By either tumbling or entering text directly, the camera and projector can be precisely positioned. Justification types are also defined for easy access. A popular justification operation is upper left projection after image reduction. This is used when reducing the size of video with aspect ratio adjustment. In the last figure you see the choices for justification as the red stripe in the 6 boxes in the order of left, center horizontal, right, top, center vertical, and bottom.



The translation effect allows simultaneous aspect ratio conversion and reduction but is easier to use if the reduced video is put in the upper left of the “temporary” instead of in the center. The track size is set to the original size of the video and the camera is centered. The output size is set to the reduced size of the video. Without any effects, this produces just the cropped center portion of the video in the output.

The translation effect is dropped onto the video track. The input dimensions of the translation effect are set to the original size and the output dimensions are set to the reduced size. To put the reduced video in the center subsection that the projector shows would require offsetting out x and out y by a complicated calculation. Instead, we leave out x and out y at 0 and use the projector's tool window. By selecting left justify and top justify, the projector displays the reduced image from the top left corner of the “temporary” in the center of the output.

3.2.4 Masks

Masks select a region of the video for either displaying or hiding. Masks are also used in conjunction with another effect to isolate the effect to a certain region of the frame. A copy of one video track may be delayed slightly and unmasked in locations where the one copy has interference but the other copy does not. Color correction may be needed in one subsection of a frame but not another. A mask can be applied to just a subsection of the color corrected track while the plain track shows through. Removal of boom microphones and airplanes are a common kind of mask uses.

The order of the compositing pipeline affects what can be done with masks. Mainly, masks are performed on the temporary after effects and before the projector. This means multiple tracks can be bounced to a masked track and projected with the same mask.

The compositing pipeline graph has a masking stage. There are 8 possible masks per track. Each mask is defined separately, although they each perform the same operation, whether it is addition or subtraction.

Compositing pipeline with masks

- 1) To define a mask, go into the Compositor window and enable the mask toggle.
- 2) Next go over the video and click-drag. Note: You have to select automatic keyframes if you wish to move a mask over time. If you do not, the mask position will be the same even if you edit at different places on the timeline.
- 3) Click-drag again in another part of the image to create each new point of the mask. While it is not the conventional Bezier curve behavior, this masking interface performs in realtime what the effect of the mask is going to be. Creating each point of the mask expands a rubber band curve.

Once points are defined, they can be moved by Ctrl-dragging in the vicinity of the corner. Shift-drag allows you to move existing points to new locations, thus altering the shape of the mask. However, this does not smooth out the curve. The In/Out points of the Bezier curve are accessed by Ctrl-dragging in the vicinity of the corner. Then Ctrl-dragging near the In or Out point causes the point to move. Shift-drag activates bezier handles to create curves between mask points.

- 4) Finally once you have a mask, the mask can be translated in one piece by Alt-dragging the mask. The effect of the mask is always on. Ctrl-Alt-drag translates an entire mask to a new location on the screen.

The masks have many more parameters which could not be represented with video overlays. These are represented in the tool window for masks. Selecting the question mark when the mask toggle is highlighted brings up the mask options window.

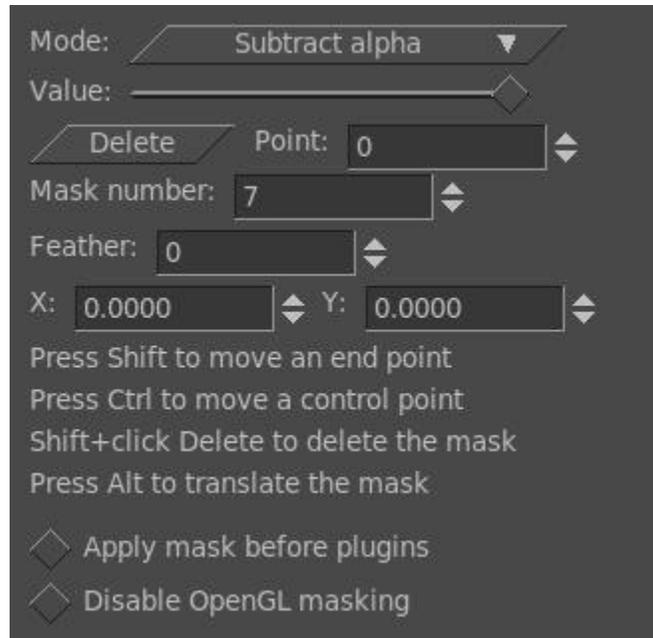


Figure 6: Mask options window

The mode of the mask determines if the mask removes data or makes data visible. If the mode is *Subtract alpha*, the mask causes video to disappear. If the mode is *Multiply alpha*, the mask causes video to appear and everything outside the mask to disappear.

The *Value* of the mask, set by a slider bar, determines how extreme the subtraction or addition is. In the subtractive mode, higher values subtract more alpha. In the additive mode, higher values make the region in the mask brighter while the region outside the mask is always hidden.

The *Mask number* determines which one of the 8 possible masks we are editing. Each track has 8 possible masks. When you click-drag in the compositor window, you are only editing one of the masks. Change the value of mask number to cause another mask to be edited. The previous mask is still active but only the curve overlay for the currently selected mask is visible. When multiple masks are used, their effects are OR-ed together. Every mask in a single track uses the same value and mode.

The edges of a mask are hard by default but this rarely is desired. The *Feather* parameter determines how many pixels to feather the mask. This creates softer edges but takes longer to render.

Two checkbox options are *Apply mask before plugins* and *Disable OpenGL masking*. Note that the OpenGL mask renderer is of low quality and only suitable as a preview for initial work. For fine-tuning of masks (with large feather values) *OpenGL masking* should be switched off so that the software renderer is used instead.

Finally, there are parameters which affect one point on the current mask instead of the whole mask. These are *Delete*, *X*, and *Y*. The active point is defined as the last point dragged in the compositor

window. Any point can be activated merely by Ctrl-clicking near it without moving the pointer. Once a point is activated, *Delete* deletes it and *X, Y* allow repositioning by numeric entry.

3.2.5 Cropping

Cropping reduces the visible picture area of the whole project. It changes the values of the output dimensions (width and height in pixels) and the X, Y values of the projector in a single operation. Since it changes project settings it affects all the tracks for their entire duration and it is not keyframable.

- Enable the crop toggle and the tool window in the compositor window to display the Crop control dialog box.
- Click-drag anywhere in the video to define the crop area. This draws a rectangle over the video.
- Click-drag anywhere in the video to start a new rectangle.
- Click-drag over any corner of the rectangle to reposition the corner.
- Alt-click in the cropping rectangle to translate the rectangle to any position without resizing it.

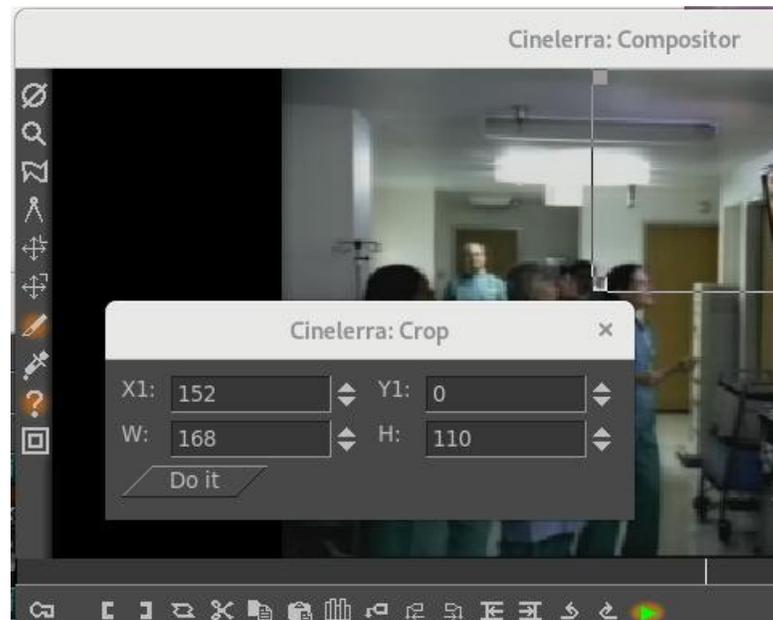


Figure 7: Cropped area is in the top right corner

- The crop control dialog allows text entry of the top left coordinates (X1,Y1) and bottom right coordinates (X2,Y2) that define the crop rectangle. When the rectangle is positioned, hit the *Do it* button in the crop control dialog to execute the cropping operation: the portion of the image outside the rectangle will be cut off and the projector will make the output fit the canvas.
- The Set Format window will show the new project Width and Height values.
- The projector tool window will show the new X, Y values.
- Track size will remain unchanged.

To undo the cropping enter the original project dimensions in the Set Format window and click on *Reset projector* in the popup menu of the compositor.

3.2.6 Safe Regions

On consumer displays the borders of the image are cut off and within the cut-off point is a region which is not always square like it is in the compositor window. The borders are intended for scratch room and vertical blanking data. You can show where these borders are by enabling the safe regions toggle. Keep titles inside the inner rectangle and keep action inside the outer rectangle.

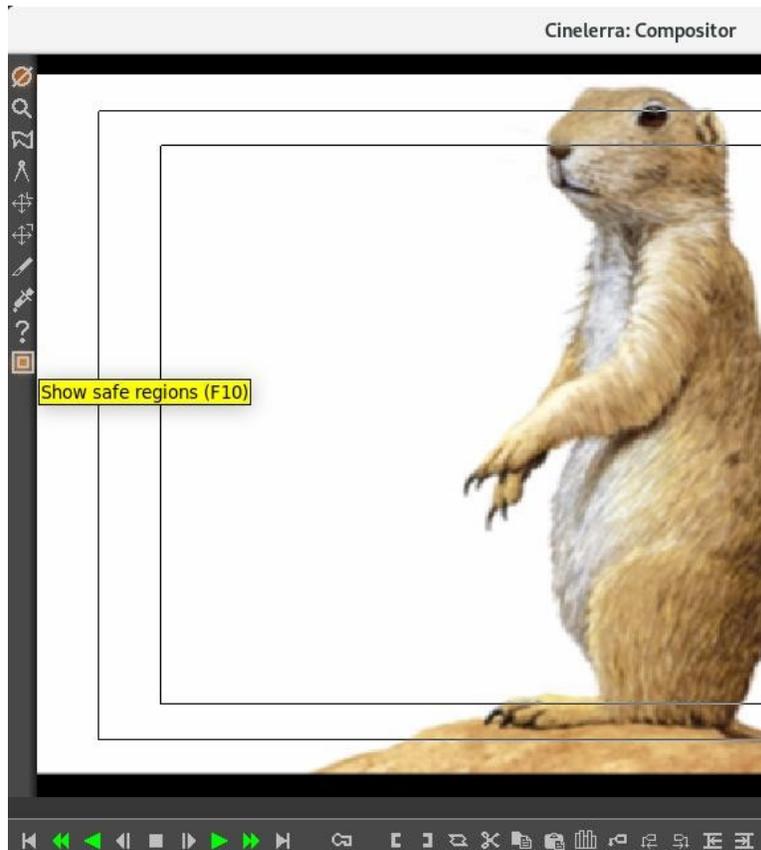


Figure 8: Note the black frames showing the safe regions

3.2.7 Track and Output Sizes

The size of the temporary and the size of the output in our compositing pipeline are independent and variable. The camera's viewport is the temporary size. Effects are processed in the temporary and are affected by the temporary size. Projectors are rendered to the output and are affected by the output size. If the temporary is smaller than the output, the temporary is bordered by blank regions in the output. If the temporary is bigger than the output, the temporary is cropped.

Track size

The temporary size is defined as the track size. Each track has a different size. Right click on a track to bring up the track's menu. Select *Resize Track* to resize the track to any arbitrary size. Alternatively you can select *Match output size* to make the track the same size as the output. If you resize a track, then its appearance on the compositor changes accordingly. Using the relationship between the track and the project's output size you can effectively reduce or magnify the size of a particular track with regards to the final output and therefore create visual effects like split screens, pans, and zooms on the compositor.

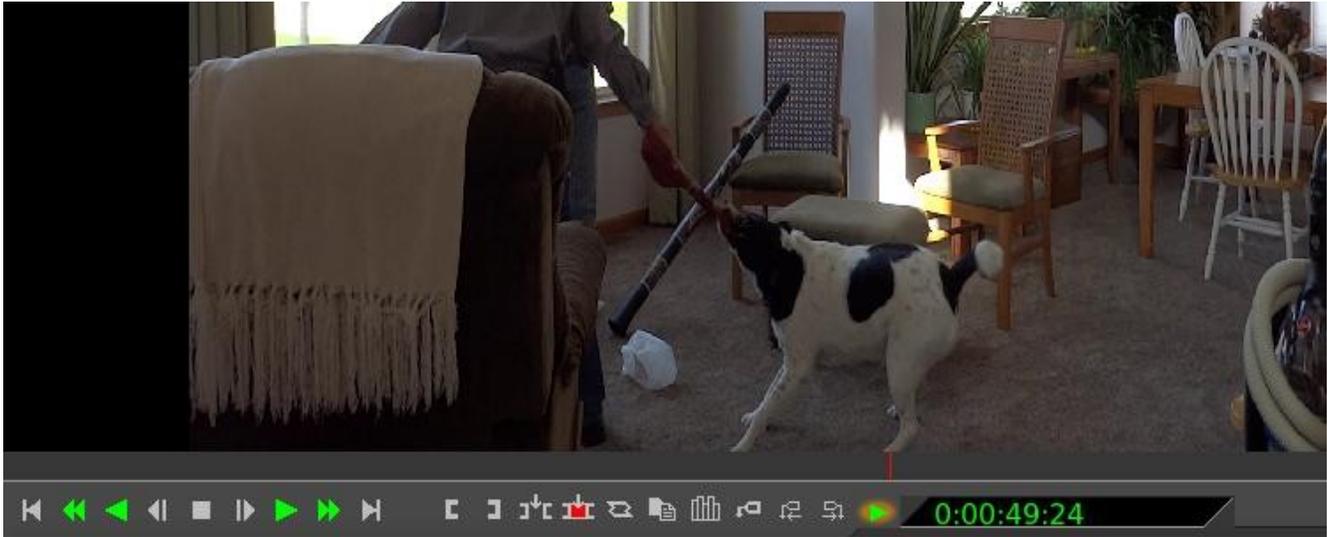
Output size

The output size is set in either File->New when creating a new project or Settings->Format. In the Resources window there is another way to change the output size. Right click on a video asset and select *Match project size* to conform the output to the asset. When new tracks are created, the track size always conforms to the output size specified by these methods.

When rendering, the project's output size is the final video track size where the temporary pipeline is rendered into. If the output size is larger than the temporary then the image transferred from the temporary will fit inside the Output Track. Any space left on the Output is left blank. If the output size is smaller than the temporary then some of the temporary video will be cropped.

3.3 Viewer Window

The viewer window is a place to load and preview your source media and clips. Here you can quickly browse through an asset using the slider control, focus on an area of work with the preview region or you use editing controls to cut and paste segments into the project or create a clip for later use.



To open the viewer window, go to Window->Show Viewer. The display is the area on the viewer where you actually see media playing. Before you can play any media, you must first load it on the viewer. To load media into the viewer:

- 1) Open the resources manager window and select the Media folder or the Clip folder.
- 2) Drag a file from the Media or the Clip folder to the viewer.

You can also load media onto the viewer by right clicking on a file in the asset manager and selecting View from the popup menu or by double clicking on the icon. Once your media loads you will see it appear on the display. To play, rewind or forward through it use the slider control or the transport controls. You can change the media display size by right clicking on the screen to activate the display zoom menu. Select zoom levels of 50%, 100% or 200% of the original media size.

When displaying media, the viewer uses the project's defined output size format settings, not the original assets format. You can change the project's output to match the asset's format using the match project size menu option in the asset manager. In the Viewer you can move around in the source media or clips and select regions to paste into the project. Operations done in the viewer affect a temporary EDL or a clip, but not the timeline.

3.4 Options in both the Compositor and Viewer Windows

The next sections describe capabilities that are available in both the Compositor and Viewer windows.

3.4.1 Click to Play in Viewer and Compositor

In both the Viewer and Compositor windows, there is an arrow on the right hand side of the other buttons in the edit panel. Mouse action can be toggled on/off via this arrow, which has a tooltip of “Click to play” with the letter ‘p’ to be used for a shortcut. When enabled there is a gold-colored shadow around the usual green-colored arrow. The purpose of enabling this capability is to make it really easy to play the media in the window by just using the left mouse button to start or stop the play. The entire main canvas surface becomes a big play button! Although the default is initially off, a good reason to enable this, at least temporarily, is so that you can quickly review your video before a render.

- left click - forward play or stop forward play if already playing
- middle wheel - single frame forward or back
- middle click - reverse play or stop reverse play if already playing. Note that some 3 button mice do not accommodate a middle click for reverse but you can find out by testing from a terminal window with the command `xev`.

3.4.2 Timebar + Preview Region Usage in the Compositor and Viewer

The navigation features of the Viewer and Compositor behave very similarly. Each has a timebar and slider below the video output. The timebar represents the entire time covered by the program. When you have a file loaded in the main window and then slide around it using the compositor slider. The insertion point in the main window follows the compositor.

Labels and in/out points are fully supported in the viewer and compositor. In the viewer and compositor, labels and in/out points are displayed in the timebar. But there is a difference between the viewer and compositor in that the compositor reflects the state of the program while the viewer reflects the state of a clip but not the program. When you hit the label button in the compositor, the label appears both in the compositor timebar and the program timebar. When you select a label or in/out point in the compositor, the insertion point in the program window jumps to that position.

The timebar in the compositor and the viewer can be used to define a region known as the preview region. This preview region is the region of the timeline which the slider affects. By using a preview region inside the entire program and using the slider inside the preview region you can very precisely and relatively quickly seek in the compositor and viewer. The preview region can be especially handy when you have large pieces of media by previewing one section, then move to the next section.

The active preview region is the zone between the edge bars. The full range of the window slider pointer action is down-scaled to the active preview region. To use this, set the preview active region as a media time region of interest. Now addressing the timebar with the mouse only operates as if the timebar is zoomed to the scale of the active preview zone. This has the effect of magnifying the interesting media in terms of the mouse pointer addressing, for fine-tuning.

The arrow above the green-colored “play forward” transport button is on the timebar.



To create and use a preview region, hold down the right mouse button inside the timebar on either end of the timebar close to the edge until you see the resize pointer. While continuously holding the right mouse button down, drag the arrow away from the end towards the middle of the timebar until you have the desired area outlined. The slider will be a light blue color while the selected preview region will remain the same initial black color. There are either a left or right resize pointer and you can click and drag in either direction to expand or shrink the region.

A left-facing arrow on the right side of the blue slider bar is used to drag the bar.



Here you can see the right-facing arrow used to drag the other end of the slider bar. The black area between is the actual preview area.



You can slide the preview zone left or right by holding the right mouse button over the preview zone where you will see a fat double-headed arrow. The selected area will move left or right as you drag and still retains the same size.

Note the double-headed fat arrow in the preview area used to move the selection over.



Settings:

- 1) If no preview region is set, increasing the length of the media on the timeline by inserting media or appending, has no effect on the non-selected preview region. That is, you will not see the blue slider suddenly mysteriously appear.
- 2) If the preview region is set, when you replace the current project or file, the preview region is automatically disabled.
- 3) If the preview region is set, when you append data or change the size of the current project, the preview region may appear to either move, shrink, or grow depending on the new length of the media on the timeline.
- 4) To disable the preview region, you will have to drag both the right and the left blue slider bars completely to their corresponding end so that there is no longer any visible blue slider.

A good method for taking advantage of the preview region is described here. On the main track canvas, scroll to the beginning of the area of interest. When you do that, you will see in the compositor the red indicator line of that location. Now in the compositor window, right mouse drag from the left side of the edge of the timebar to create the blue slider bar line up to the red indicator. Back in the main track canvas, move to the location of the area you want to end looking and again you will see the

red indicator line in the compositor. Use the right mouse drag from the right to stop at that end point. Using this method is often easier than continuous usage of the single frame move which can be tedious.

One last interesting item of note – sometimes you may wish to see just a little more that is outside the preview region and you can do so! You can actually move outside the compositor or viewer window space and view more, at least until you hit the end of the monitor space.

3.5 Resources Window

Effects, transitions, labels, clips, proxies, user bins, and media assets are accessed here. Most of the resources are inserted into the project by dragging them out of the resource window. Management of resource allocation is also performed here.

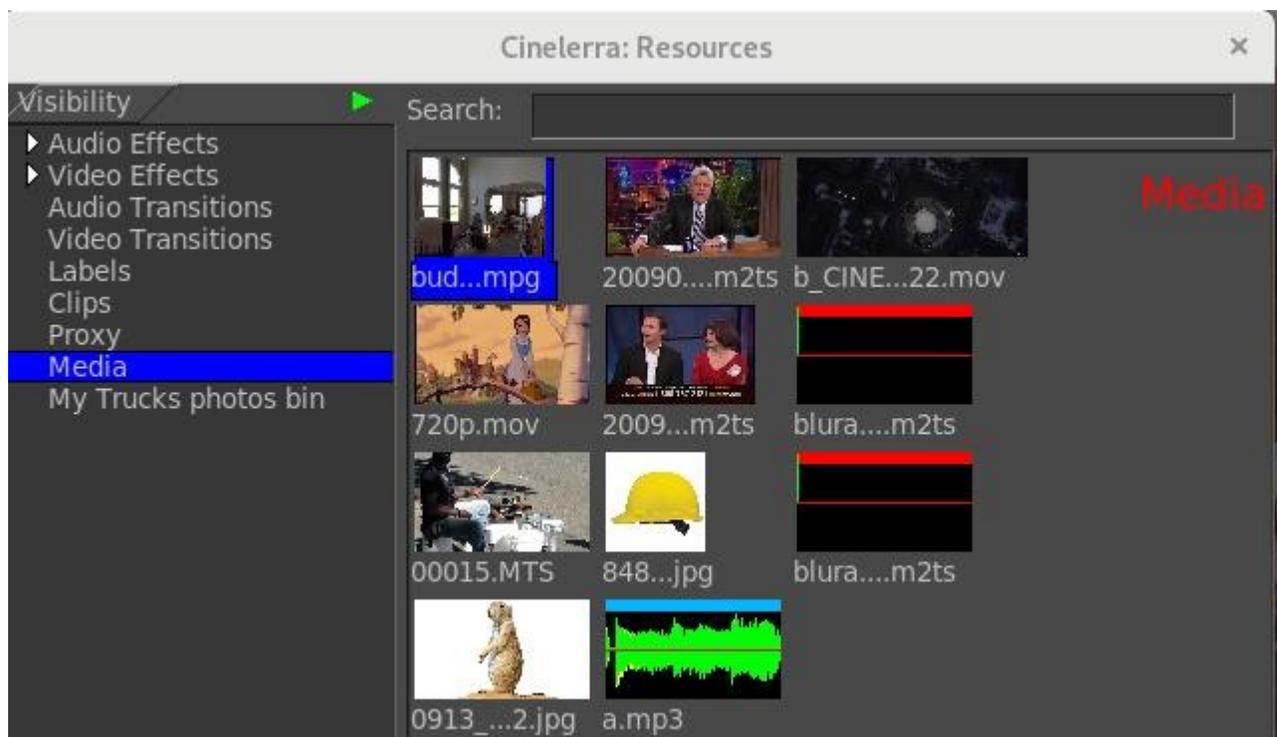


Figure 9: Folders are in the first column with contents of that folder on the right hand side

The resources window is divided into two areas. One area lists folders and another area lists the folder contents. Going into the folder list and clicking on a folder updates the contents area with the contents of that folder. The folders can be displayed as icons or text. There are several variations for displaying the contents; select *Display text*, *Display icons*, *Display icons packed*, *Display icons list* as types of display for the assets or plugins. Use the letter “v” to easily scroll through the choices and see which you prefer. You can also get to these options from the menu by a right mouse click in the window.

A *Search* option is available for any of the folders in the Resources window (and when using “Attach effect” on the main track canvas for the Plugins). As you type in characters a match is made with that substring. Names that do not match are filtered out making it a lot easier to find the item you are looking for. The characters can be any where within the phrase and it does not matter if upper or lower case.

Other options you will see if you **right mouse click in the folder** which brings up the menu are described next.

Load files – for convenience to load files same as from the main window so you do not have to move the mouse so far in case you have multiple monitors.

Display text/icons – as described previously for format variations preference.

Select - options are All, Used, Unused, and None. This gives you the capability to have a set of the contents highlighted for ease of use so you can see what is or is not loaded, or unset the highlight.

Sort items - to sort the contents of the folder alphabetically. Especially helpful if you accidentally did a drag and changed your mind or dropped suddenly so that the assets no longer look nicely aligned.

Copy/Paste file list – use to easily copy a set of files or paste a set of files between Cin and windows.

Snapshot/Grabshot – described elsewhere in more detail.

Using the right mouse click to bring up a menu in the folder area, you can also switch from *Display text* to *Display icons*, *Sort items* and create, delete and manipulate user defined folders/bins. Select *Folder* to create a user *Folder* or modify an existing folder.

If you **right mouse click on a highlighted/selected resource**, several options are available depending on whether the resource is an effect or transition or a piece of media. You can highlight several for some options so that it is applicable to all of them, such as *Info*. Those listed immediately below are the available choices for media assets.

Info – provided basic Asset information; details are described later in this section.

Display text/icons – same as mentioned previously.

Sort – same as mentioned previously.

Rebuild index - if you switch from/to using ffmpeg/native for media loading, you should rebuild indexes. Or if you get hangs on media or strange looking tracks, you might want to rebuild indexes.

View – use this option to bring up the media in the Viewer window.

View in new window – in order to not overwrite your current viewer window, you can open any number of viewer windows to simultaneously view multiple media.

Open mixers – when you record with multiple cameras setup, you can work with them most easily using the mixer mode. This is described in detail

Paste –

Match – if you need to change your media parameters you can choose from the following: Match frame rate, Match project size, Match all

Remove – use to Remove the asset from the project or with caution, to Remove from disk permanently.

In the case of Effects or Transitions, a right mouse click on a highlighted selection leads to an *Info* button which gives a short 1 line description of what the effect/transition can be used for.

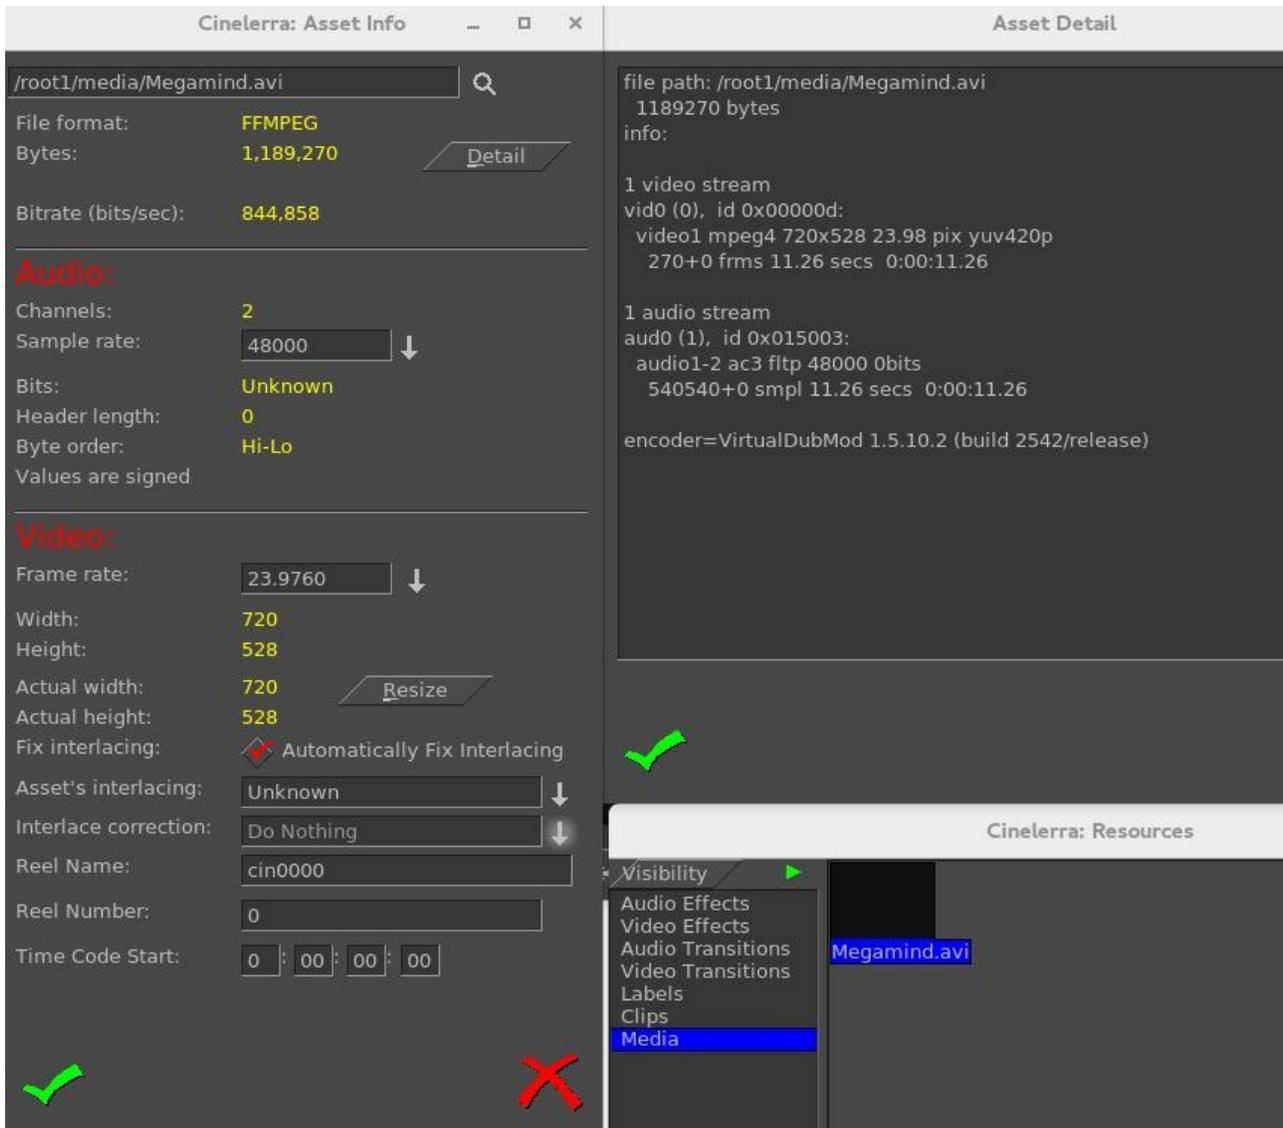
For Labels, choices are *Edit*, *Label*, and *Go to*.

For Clips, *Nest* and *UnNest* as described elsewhere are available.

3.5.1 Info Asset Details

The asset *Info* window also can be used to display detailed information about the selected/highlighted media file - available for any loaded media of type mpeg or ffmpeg. This is extremely helpful in determining what type of media it is, size, resolution, format, and type/number of audio streams. It is especially useful for multiple program streams. You can have the info window popped on several of your assets simultaneously.

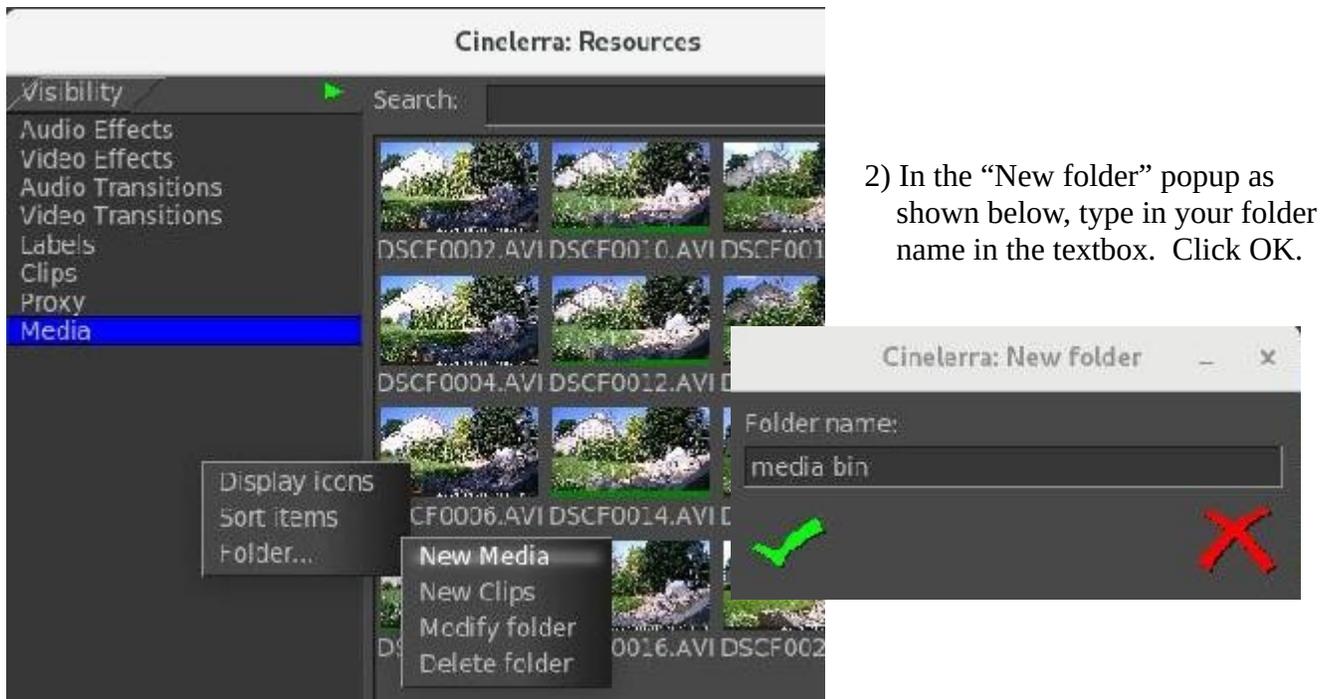
Screenshot below shows the “Detail” box to click on the left side and a simple, typical output in the Asset Detail window on the right side. Also, note the highlighted media in the Resources window.



3.5.2 User Folders/Bins

Creating folders that are more specific to a particular project is helpful in better organizing your work. This can be done by utilizing the files already loaded to the “master” Media or Clips folders in the Resources window. Below are steps illustrating an easy way to set up a folder.

- 1) In the Resources window, in the location of the Video/Audio effects and Media folders, bring up the “Folder...” popup by clicking the right mouse button. Highlight, then click “New Media or Clips”.

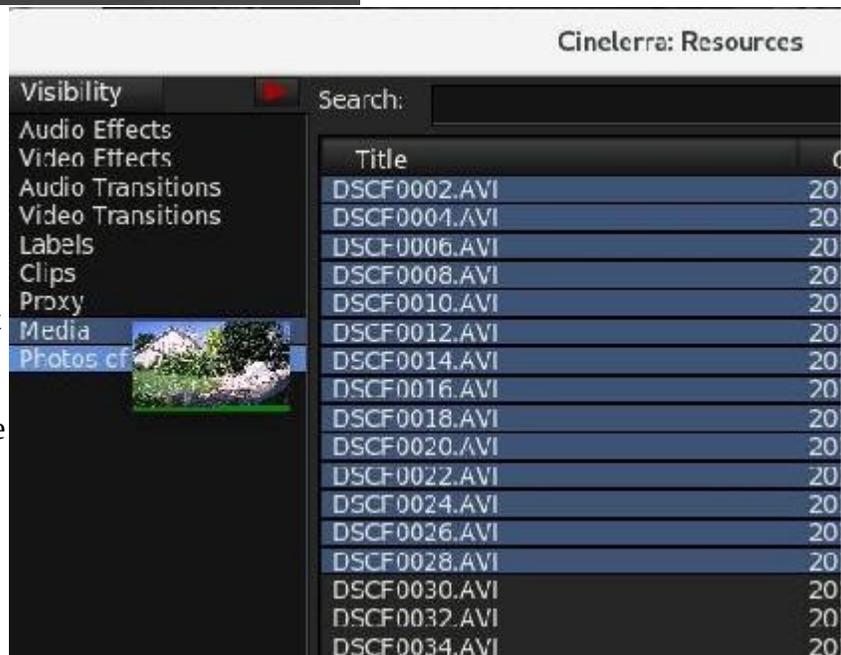


- 2) In the “New folder” popup as shown below, type in your folder name in the textbox. Click OK.

“Modify folder” can be used to || change the name of a folder.

“Delete folder” in the popup can be used to delete a folder.

- 3) Select the “master” Media folder to see which files are currently loaded. Highlight the files there that you want to copy to your new folder (named Photos of Garden seen in this screencast). Drag the files to the left and when you see the Photos of Garden folder become highlighted, then drop there. You can drag and drop any of the media from the “master” Media at any time. It flashes when the drop is successful.

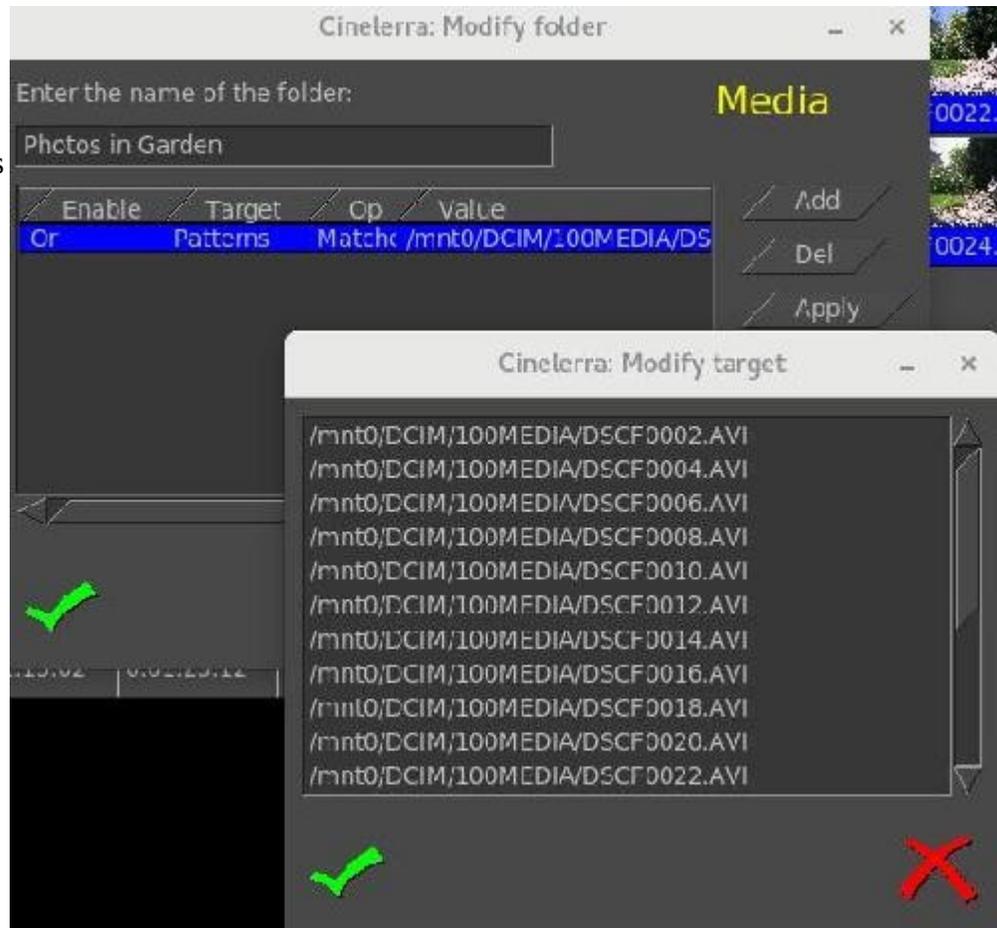


Adding the Shift key before the actual drop, will allow for relative path filenames instead of full path.

But you might want to include or eliminate some of the media that exists in one of the folders that you have set up already. In this case you will want to click on the “Modify folder” in the popup. When you bring up the Modify folder window, if you already have files in that folder, you will see filters that were generated automatically when you did a Drag and Drop.

Here is the filter that was generated with the original drop →

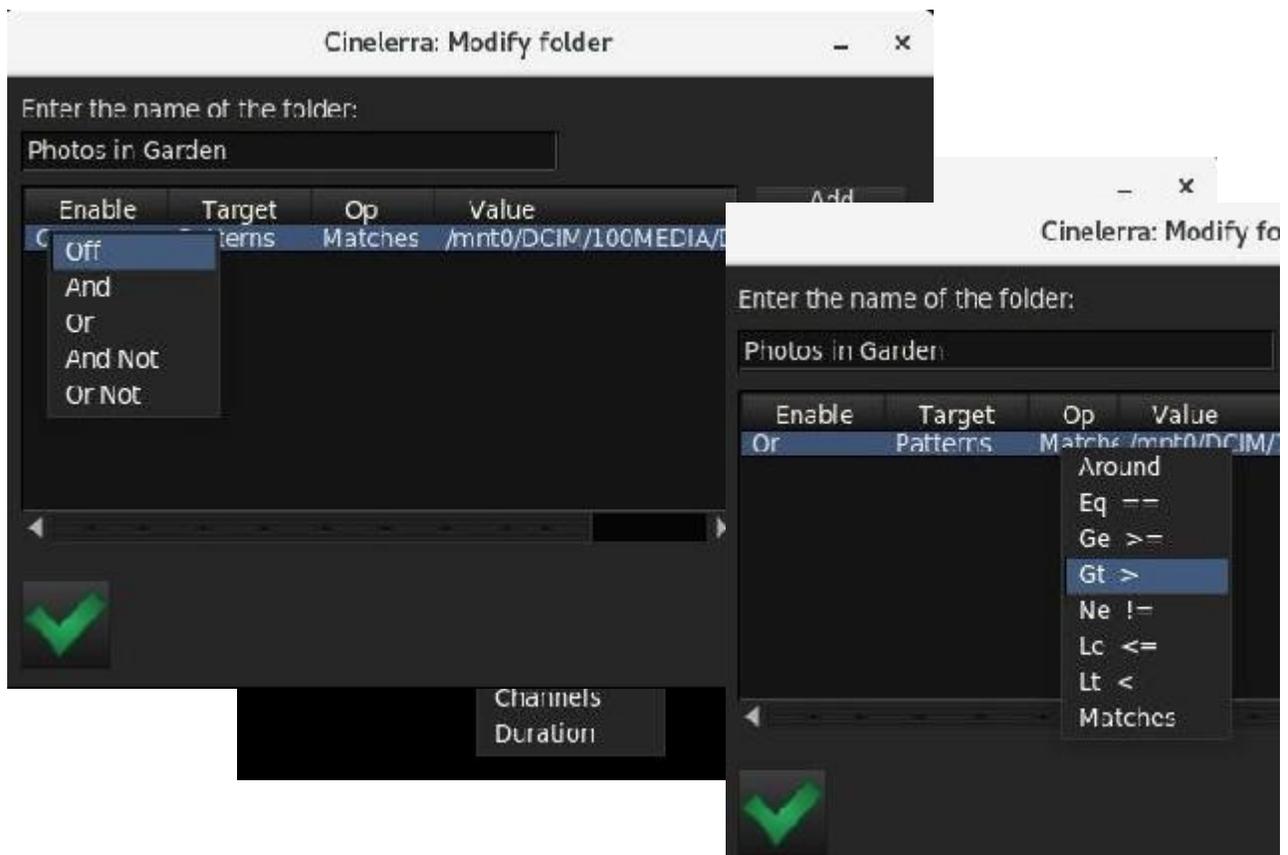
When you click on the Value portion of that filter, the entire set of files that are covered by the filter rules pops up. Now you can highlight a target filename that you would like to remove, and just erase that line and check the green checkmark for OK.



To delete the entire set of files listed on the filter rule, highlight the rule line and hit the “Del” button. To add a new filter rule, click on the “Add” button which will automatically add a default line after any current lines. The default line will be a line that matches everything in the “master” Media folder which is “Or Patterns Matches *”. Click the right mouse button on the current field underneath the column header to see the choices available for each column.

Modifications will not be in effect until you click on the green arrow OK button or click on the Apply button. But once you hit Apply, clicking on the red X button will not undo your changes. The filter/search rules are applied in the order listed in the Modify folder window. You can change the order of the filter rules by highlighting the rule you want to move and then drag and drop to a new location.

The screencast below displays the available choices for each field.



Information about the columns and rules for the search filters in the Modify folder window follows.

Column headers:

Enable – this column is used to designate the state of that filter rule

- Off disable the filter
- And narrow your search; all of your search terms must be present
- Or broaden your search to include more values
- And Not exclude terms that do not contain the given value from your search results
- Or Not include terms that do not contain the given value from your search results

Target – this column designates which media asset attribute to look at

- Patterns each line contains a filename filter, matches the file path
- Filesize number of bytes in a file
- Time date file was created
- Track Type track type of video, audio, or audio video (for both)
- Width Format width
- Height Format height
- Framerate Video framerate
- Samplerate Audio samplerate
- Channels Number of audio channels
- Duration Playback time in seconds – it uses the largest of audio or video if contains both

Op – boolean operators used to narrow or broaden the relationship between your search terms

Around	about this value; use “+radius” for a search range: [target–radius ... target+radius]
Eq	equal to
Ge	greater than or equal to
Gt	greater than
Ne	not equal
Le	less than or equal
Lt	less than
Matches	exactly matches for strings

Value – the characteristic you are looking for with expressions that can be written with the following:

Number (decimal points are allowed and will be converted to a standard form):

inf – representing infinity

#[TtGgMmKk] - where # represents a number and the characters mean:

inf = infinity

T = 1099511627776

t = 1000000000000

G = 1073741824

g = 1000000000

M = 1048576

m = 100000

K = 1024

k = 1000

Scalar:

Number

Number+Number

Date time:

date = year/month

date = year/month/day

time = hour:minute

time = hour:minute:second

date_time = date time

Duration:

day = #day | #days

week = #week | #weeks

month = #month | #months

year = #year | #years

delta = secs

delta = mins:secs

delta = hours:mins:secs

Around time:

date time+duration

Around length:
duration+duration

Table showing the allowed usage:

```
target:      | eq ge gt ne le lt  matches  around
+++++
patterns    | <---- strcmp -----> + filter + nearest
file_size   | <---- arithmetic -----+-----> + radius
mod_time    | <---- arithmetic -----+-----> + radius
track_type  | <---- member test -----+-----+----->
width       | <---- arithmetic -----+-----> + radius
height      | <---- arithmetic -----+-----> + radius
framerate   | <---- arithmetic -----+-----> + radius
samplerate  | <---- arithmetic -----+-----> + radius
channels    | <---- arithmetic -----+-----> + radius
duration    | <---- arithmetic -----+-----> + radius
```

where in the above, the filter can be:

- filter=list
- filter=token
- list=[token]
- list=[token]list
- string=<chars>|<empty>
- token=string
- token=string*token

Examples with some caveats first:

- 1) “Or” generally includes or adds whereas “And” generally excludes or subtracts.
- 2) The filters only work on media in the folder; if there is no media, then there is nothing to search.
- 3) The examples below are not meant to be executed as a list of filters in Modify folder, they are just single line examples to indicate what can work.
- 4) Sort is by filename base name (directory path not included automatically) except when the “Around” operation is used and then it is sorted by that Target distance first and then filename.

<u>Enable</u>	<u>Target</u>	<u>Op</u>	<u>Value</u>	- <u>meaning</u>
Or	Patterns	Matches	*	- all files from the Media folder are included
And Not	Filesize	Lt	160000000	- no files that are less than 160MB in size
Or Not	Time	Ge	2018/07/30 06:13:00	- files not greater than or equal date
And	Duration	Eq	01:00	- files included must have 60 secs. Duration
Off	Samplerate	Ne	44000	- off for now, but may want to include later
And	Framerate	Around	24+1	- files included all have 24 to 25 framerate
Or	Patterns	Matches	[*.mp4]	- all files with the extension of mp4
Or	Time	Around	2018/08/02 06:00:00 + 02:00:00	- files at 4AM to 8 AM

3.5.3 Vicons & Aicons – aka Video Icons / Audio Icons

Vicons are video icons. Aicons are audio icons. By default the Resources window will play the first 5 seconds of video or audio waveform looped in the area occupied by the media icons. This is enabled for the Media/Proxy folders in icon mode when the mouse pointer is inside the Resources window.

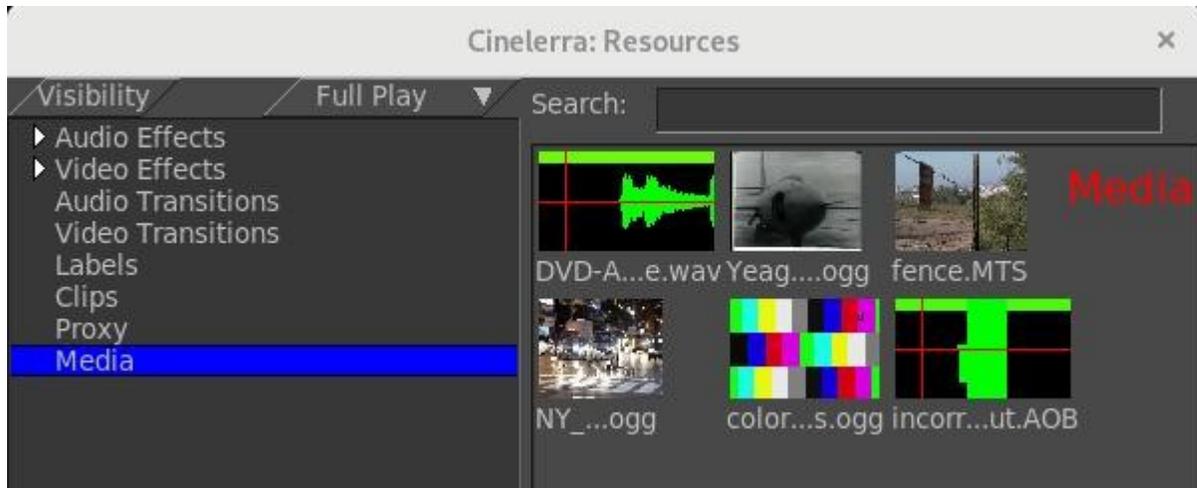
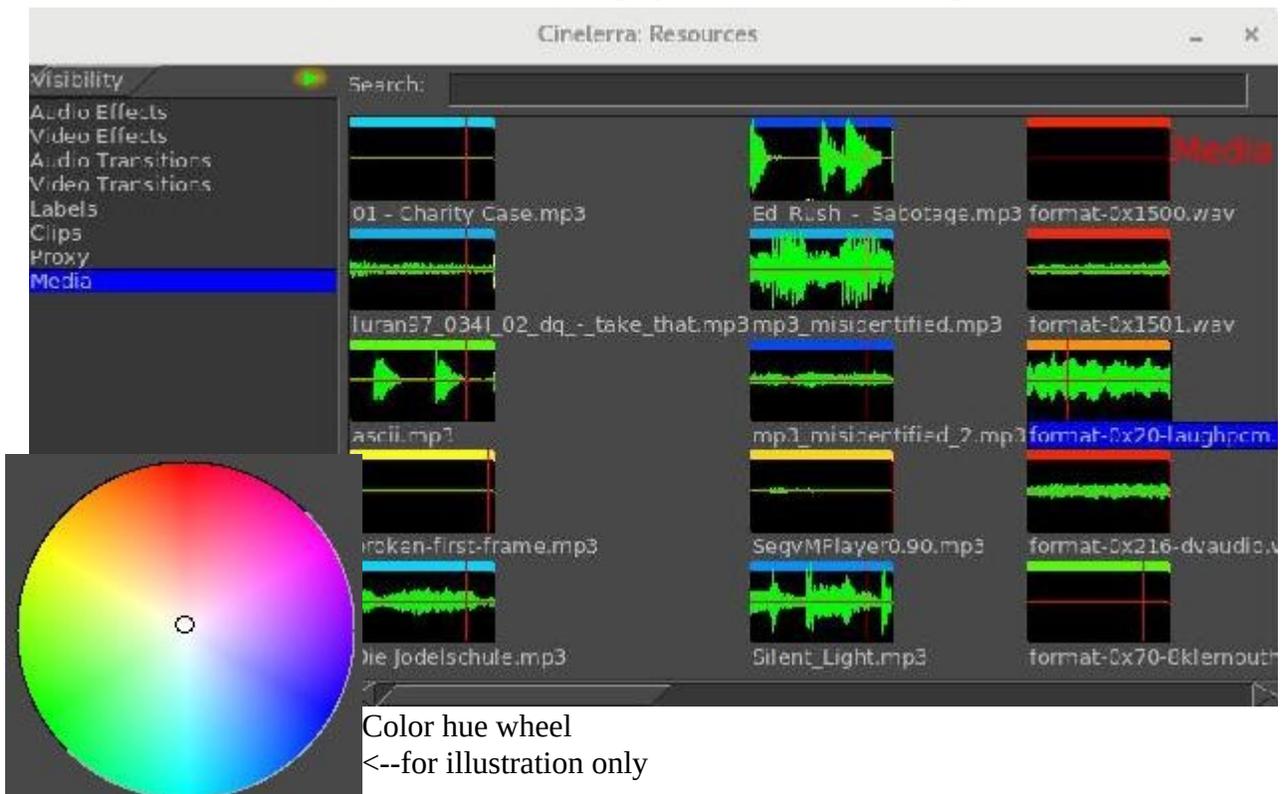


Figure 10: Note "Full Play" mode and Vicons and Aicons in Media folder

The waveform in the last figure is displayed in the Resources window in the color green/yellow for the 2 audio tracks. There is a colored bar on the top of each a-icon where the color is based on the Color Spectrum - the smaller the time duration, the redder the color; then as the time duration goes up, the color goes up so that you will go to green, then yellow, then blue, then really dark blue, then purple for the audio files 1 hour and over. There are various other colors between these colors same as that seen in the color spectrum in the screenshot below. Colors are utilized from the hue wheel in the counter-clockwise direction. Note that the horizontal line in the middle of the a-icon is yellow/red representing the 2 audio tracks and is only red for mono.

Draw Vicons | Screenshot display various audio file lengths; red is shortest.



Note that if in Settings->Preferences under the Appearance tab, you have unchecked “Use thumbnails in resource window” you will only have default icons and none of the above capabilities.

3.5.4 Resources Window Preview Mode

Preview mode can be used to pop up a window which draws the vicons/aicons thumbnails in a larger size. Preview or “draw vicons” mode is a helpful feature of cinelerra that lets you see and/or hear the first 5 seconds of the video for identification purposes. The Preview mode/playback toggle is to the right of the Visibility label as seen in the screenshot above. Preview mode is available for the Media, Proxy, Media User Bins, and Clips but clips are only 1 image.

When “Preview”/“draw vicons” is enabled/active, if you click on one of the video icons or an audio waveform icon, a view pops up that increases the size to 4 times the surface area larger. This makes it easier to see or hear if it is the media you are looking for in case you have many similar media files. To conserve memory, the video is stored 8-bits per pixel which results in low image quality while the audio is 16-bit. The reason for playing 5 seconds of a video for a vicon is that until the first I-frame, the media frequently does not decode properly. In other words, a lot of media does not begin at the “beginning” point and will not be properly rendered until enough data has been read to assemble a picture. You can increase the thumbnail size, clarity of pixels (memory size) and color mode but it takes a lot more memory. Change these values in Settings → Preferences, Appearance tab, right hand side of the Layout section – be aware that when you click OK, your session will re-initialize. You can also temporarily increase the preview mini-window by use of the mouse wheel up or down.

There are 4 options for the preview mode.

1) *Full Play* is the default mode. This means all of the media will automatically play when the mouse is in the Resources window and you can use the left mouse button to click on specific media to see it pop up in a larger view. Audio only files do not play the audio until the icon is clicked on and the waveform aicon pops up into the 4x larger mode. *Full Play* includes the *Mouse Over* capabilities as described below as well as the *Inter-View Src Target* functions.

2) *No Play* mode is especially useful on smaller computers and for users who find the constant loop play to be somewhat distracting.

3) *Mouse Over* mode is activated by a single click on one of the vicons/aicons and deactivated with another single click over any of the icons. Once activated, whenever you just move the mouse over an icon, it automatically pops up the increased size preview. The first time in your session that you enable this feature, it may take a few seconds to load all of the icon previews into memory so be patient and just wait. *Mouse Over* mode makes it quick and easy to preview without having to drag the media to the viewer. You can still drag the media same as without preview enabled.

4) *Src Target* mode gives easy access to the *Inter-View* source target available by using the middle mouse button on media. There are 2 advantages to this mode - there is no 5 second play loop taking up cpu time and the popup allows for the use of the letter “f” on that popup to have it go to fullscreen mode. *Src Target* mode in any scenario never plays sound as that is nonsensical usage. After the initial click to pop media in this mode, you also have the *Mouse over* feature.

For any of the options, but not *No Play*, you can temporarily turn off that option by clicking on the button using the middle mouse button. This helps to avoid having the thumbnail get in the way of

dragging or other functions. When you do, a line will be drawn through the current preview mode so that you are aware that it is in *No Play* mode until click it again.

Note that if in Settings->Preferences under the Appearance tab, you have unchecked “Use thumbnails in resource window” you will only have default icons and no active previews.



Screencast shows location of the Preview/Draw Icons mode.

Note the line through the mode.

3.5.5 Moving clips/media from/to Resources window

If you have several media files loaded into the Resources window of one instance of Cinelerra and want to load some of the same ones into another instance or just want a listing to save in a file for later use, you can do this with these set of steps:

Copy or paste a list of files in the Media Resources window:

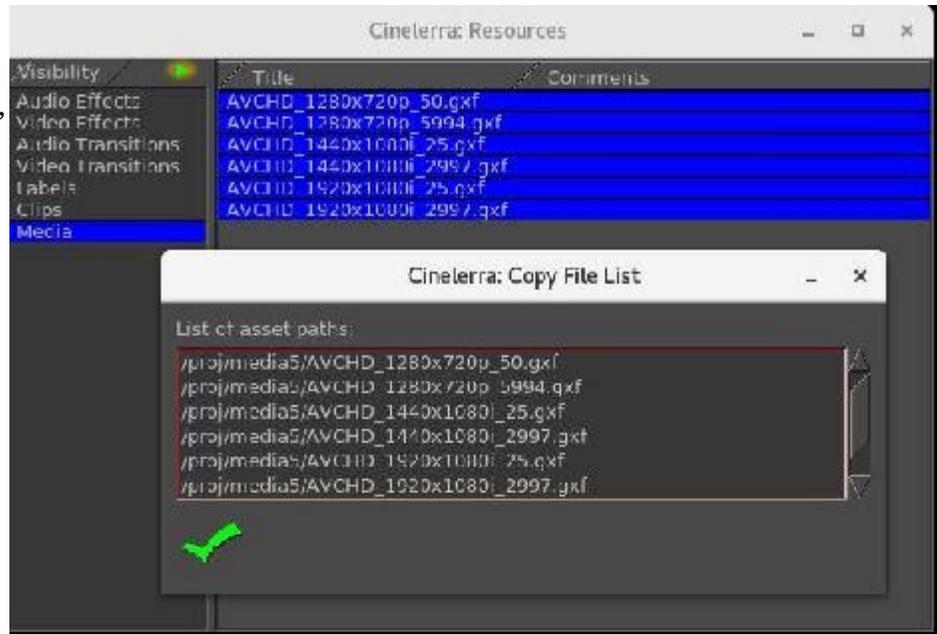
- 1) create a highlighted selection of the desired media files in the media Resources window
- 2) right click on an unused portion of that window to bring up the popup menu
- 3) select the “Copy file list” item and a file list box will appear that contains the full path filenames
- 4) wipe the textbox using your standard copy/paste method to put the list of files in the copy buffer
- 5) in another cinelerra instance, choose the “Paste file list” of the media Resources window
- 6) paste the list of files, again using your standard paste method, into the new file list box; press OK
- 7) the status bar of the main window will be updated as the file list is loaded to the media folder (the purpose of displaying the status is simply to show that the load is progressing normally.

Obviously this “Paste file list” feature means you can create a list of files outside of cinelerra using an editor, wipe the names, and then use “Paste file list” to load them into the media Resources window.

It is important to note that in the steps above, the Operating System cut and paste capabilities are in use for steps 4 and 6 as opposed to Cinelerra’s c/v shortcuts. Since the procedure varies among the distros, you will have to adapt to your specific one. For example, a usage for ubuntu consists of:

- 4) Ctrl-c to copy the list of files; open gedit; Ctrl-v to paste the list of files into gedit
- 5) Ctrl-c or the standard way using the right click to copy this list from gedit
Ctrl-v paste the list of files into the new file list box, and press OK

In the screenshot to the right, one instance of cinelerra has 6 items in the Media area highlighted that were copied to the file list. Note how it includes the full pathname.



In this screenshot on another instance of cinelerra, there are only 2 items in the media but the “Paste file list” box is ready to have the items inserted via the standard text box paste method. When that is done, the additional 6 media files will be available on this other instance too.



Another possible usage of this capability:

- 1) Right Click on the Clips Resources window and use the “Paste Clip” option to paste the Copy selection as a clip.
- 2) Similarly, by highlighting a clip in the Resources window and selecting its copy popup menu item using the right mouse button, that copy buffer can now be loaded onto the timeline.

3.5.6 Snapshot / Grabshot

To take a snapshot, perform the following steps:

- 1) set your timeline insert marker where you want the snapshot - this frame shows in the compositor
- 2) right click in an empty spot in the media folder and the popup shows snapshot as the 5th item down
- 3) highlight that and the submenu comes up allowing you to choose png, jpg, ppm or tiff

The snapshot shows up in the Media folder. It is saved by default in /tmp as snap_date-time.ext BUT you can change the default directory path in Settings->Preferences-> Interface tab in the right hand side of the Editing section.

Grabshot is the 6th menu item. A red circle reticle can be moved to the area to grab; use left mouse drag to surround an area; and right click to grab.



Figure 12: Snapshot menu and choices

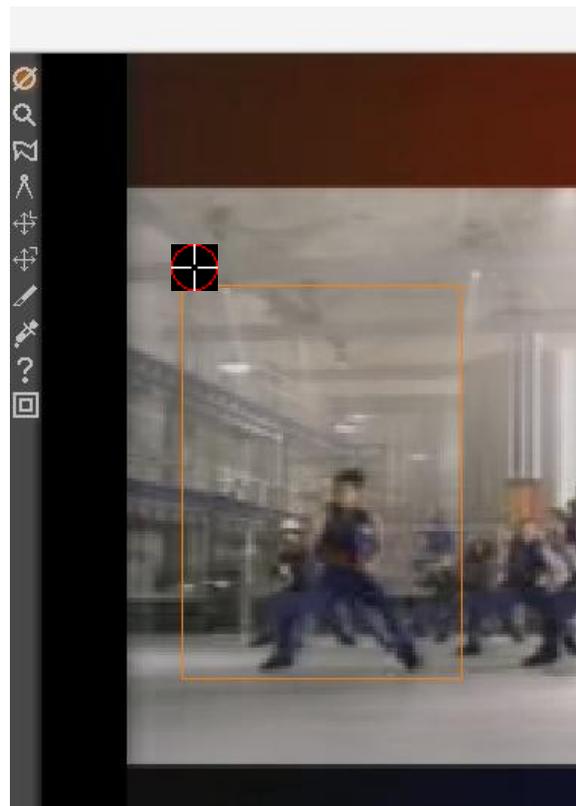


Figure 11: Grabshot reticle & orange box

3.6 Other Options and Other Windows

3.6.1 Transport Controls

Transport controls are useful for navigation and for playing media. Each of the Viewer, Compositor, and Program windows has its own transport panel. The controls generally all contain a yellow colored tooltip when you mouse over the control, providing a hint of their function and shortcuts for usage.

The transport panel is controlled by the keyboard as well as the graphical interface. For each of the operations it performs, the starting position is the position of the insertion point in the Program window and the slider in the Compositor and Viewer windows. The ending position is either the end or start of the timeline or the end or start of the selected region if there is one.

The orientation of the end or start depends on the direction of playback. If it is forward the end position is the end of the selected region. If it is backward the end position is the start of the selected region. The insertion point moves to track playback. When playback stops, the insertion point stays where playback stopped. Thus, by playing back you change the position of the insertion point. The keyboard interface of either the numeric pad or alternative keys has more speeds with the addition of *Forward Slow* (2) and *Reverse Slow* (5). Hitting any key on the keyboard twice pauses it. The shortcuts section of this manual as well as a Shell Command available from the Cinelerra main window has a listing of each of the keys.

When using frame advance functions the behavior may seem odd. If you frame advance forward and then frame advance backward, the displayed frame does not change. This is because the playback position is not the frame but the time between two frames. The rendered frame is the area that the playback position crosses. When you increment the time between two frames by one and decrement it by one, you cross the same frame both times and so the same frame is displayed. There is an option in Settings → Preferences, Appearance tab to “Always show next frame” that may help make this clearer for some users.

The transport behavior changes if you hold down Ctrl when issuing any of the transport commands. This causes the starting point to be the In point if playing forward and the Out point if playing backward. If playing forward, the Out point becomes the ending point and if playing backward, the In point becomes the ending point. If no In/Out points are specified, the behavior falls back to using the insertion point and track boundaries as the starting and ending points.

The transport behavior also changes if you hold down the Shift key along with KeyPad 1-6. If normally audio is included in the play, it will be removed and if normally audio is not included in the play, it will be added.

3.6.2 Zoombar

The compositor has zoom capability. The pull-down menu on the bottom of the compositor window has a number of zoom options. When set to Auto the video is zoomed to match the compositor window size as closely as possible. When set to any other percentage, the video is zoomed a power of 2 and scrollbars can be used to scroll around the output. When the video is zoomed bigger than the window size, you can use scrollbars to scan around or if the zoom icon is enabled, the middle mouse button can be used to zoom in or out the video.

The zoom toggle also causes the Compositor window to enter zoom mode. In zoom mode, clicking in the video output zooms in while a Ctrl-click in the video output zooms out. If you have a wheel mouse, rotating the wheel zooms in or out too. Zooming in or out with the zoom tool does not change the rendered output. It is merely for scrutinizing video or fitting it in the desktop. Playing video on the compositor when zoomed to any size other than 100%, the original size, requires Cinelerra to do extra processing steps. This could affect performance on slower systems

3.6.3 Show Overlays

Color Coded Keyframe Curves are a big feature in the “Show Overlays” window because by changing the colors to suit the user, it helps to remove confusion from multiple curves on the track canvas. They can be viewed from the pulldown menu of Window → Show overlays but they will operate the same as when used from the View pulldown menu. The “Color Coded Keyframe Curves” have distinct colors associated with each type for ease of identification. By clicking button 1 on the “Color Ball” to the right of any keyframe type in the “Show overlays” menu you have the ability to change the colors to whatever works best for your video. The color ball changes made will be retained across sessions.

There is a line separating the first 4 items, which are just non-automation type settable values as opposed to “auto” keyframe types. The color is not changeable for the 3 items of Mode, Pan, and Mask which simply display their symbol icon.

Screenshot below displays the Show overlays popup with all of its options and color coded types such as yellow for Speed and blue for Camera Z. Upon clicking on the associated “color ball” to the right of any keyframe type, for example “Fade” in this screenshot, the color wheel palette window pops up so that you can manipulate the color as desired.

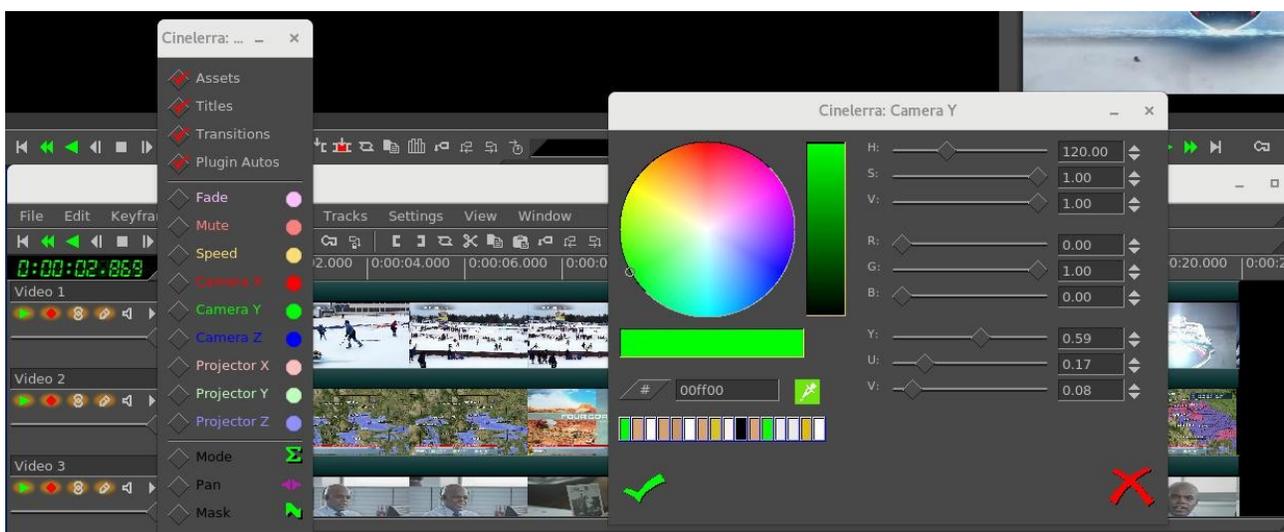


Figure 13: Show Overlays window on the left with the Color ball window to the right to set color

Screenshot below shows several color coded lines for different keyframes along with the Fade slider for manipulation. The slider is in the same color as the color coded keyframe type line which is the same color as in the “Show overlays” window.



Figure 14: Lines are colored here on the timeline as designated in Show Overlays

Overlays Window Nuances:

The Overlays window is an alternative to the main track canvas View pulldown, and thus the order is mostly maintained to match each other. To make it easier to get a quick temporary look at a specific option, there is a shortcut of Shift-LMB (left mouse button) that can be used as opposed to having to uncheck everything that is currently checked and then having to recheck them on when done. Here is a list of how they work. Keep in mind that if the Expander on the patchbay is enabled, you still see the track.

- Shift+LMB (left mouse button) in the Overlays Window on a checkbox will turn off all other checkboxes except for the one you are on. Then this named box will have outline for a "hot" spot.
- Shift+LMB on this "hot" spot will return to "cool" of the previous settings with all of the previous checkboxes checked again.
- Shift+LMB on a non-"hot" spot will simply check or uncheck a box and there is no previous state.
- This all works in conjunction with the View pulldown menu which, of course, has no hot spots.
- Caveat #1 - Shift+LMB on the top 4 choices of Assets, Titles, Transitions, Plugin Keyframes will turn off all of the checkboxes below because it makes sense to do so.
- Caveat #2 - Shift+LMB on the Autos will not turn off Assets, Titles, Transitions, or Plugin Keyframes because you need to be able to see what is going on.

- Caveat #3 - XYZ toggle on/off of Camera and Projector are not affected.

Original Settings – cool spot



Note Titles box hot spot



Cam/Proj XYZ toggle to fine tune



3.6.4 Sound Level Meters Window

An additional window, the levels window, can be brought up from the Window pulldown. The levels window displays the output audio levels after all mixing is done. The visible range of the sound level meters is configurable in Settings → Preferences, Interface tab under the Operations section.

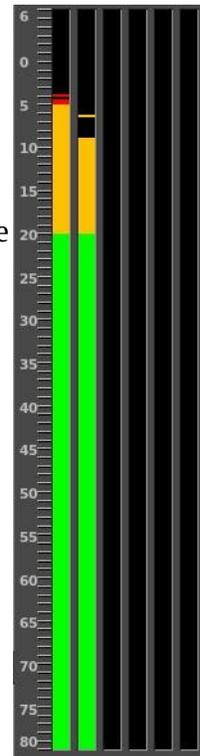
Sound level meters can be toggled in the viewer and compositor windows with the *show meters* button. They also appear in the patchbay when the track is expanded and in the recording monitor when audio is being recorded.

The sound levels in the levels window, compositor, and viewer correspond to the final output levels before they are clipped to the sound card range. In the record monitor they are the input values from the sound card. In the patchbay they are the sound levels for each track after all effects are processed

and before down-mixing for the output. Most of the time, audio levels have numerical markings in dB but in the patchbay there is not enough room.

The sound level is color coded as an extra means of determining the sound level. Even without numerical markings, the sound level color can distinguish between several ranges and overload. Look at the color codings in a meter with numerical markings to see what colors correspond to what sound level. Then for meters in the patchbay in expanded audio tracks, use the color codings to see if it is overloading.

Be aware that sound levels in Cinelerra can go above 0 dB. This allows for not only seeing if a track is overloading but how much information is being lost by the overloading. Overloading by less than 3 dB is usually acceptable. While overloading is treated as positive numbers in Cinelerra, it is clipped to 0 when sent to a sound card or file.



This page intentionally left blank so that chapters begin on a right facing page.

4 - Project and Media Attributes

When you play media files in Cinelerra, the media files have a certain number of tracks, frame size, sample size, and so on. No matter what attributes the media file has, it is played back according to the project attributes. So, if an audio file's sample rate is different than the project attributes, it is resampled. Similarly, if a video file's frame size is different than the project attributes, the video is composited on a black frame, either cropped or bordered with black.

The project attributes are adjusted in Settings->Set Format or can be created in File->New. When you adjust project settings in File->New, a new empty timeline is created. Every timeline created from this point on uses the same settings. When you adjust settings in Settings->Format, media on the timeline is left unchanged. But every timeline created from this point uses the same settings.

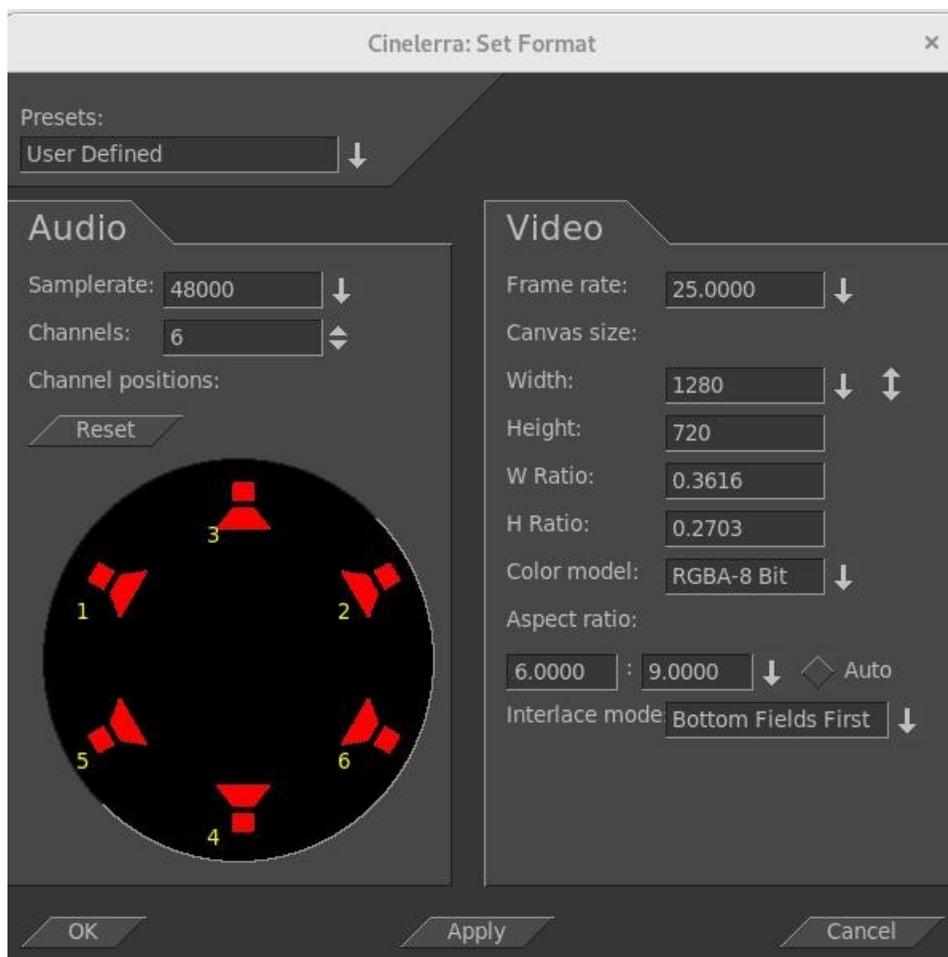


Figure 15: Set Format window - note the Audio Channel positions

In addition to the standard settings for sample rate, frame rate, and frame size, Cinelerra uses some less traditional settings like channel positions, color model, and aspect ratio. The aspect ratio refers to the screen aspect ratio.

Edit decision lists, the EDL stored in XML, save the project settings. Formats which contain media but no edit decisions just add data to the tracks. Keep in mind details such as if your project sample rate

is 48 kHz and you load a sound file with 96kHz, you will still be playing it at 48 kHz. Or if you load an EDL file at 96kHz and the current project sample rate is 48 kHz, you will change it to 96 kHz.

The New Project window has some options that are different than the Set Format window as you can see by comparing Figure 1 above with this Figure 2. Mostly notably is the field for a directory path and a Project Name.

Explanation of the various fields is described next.

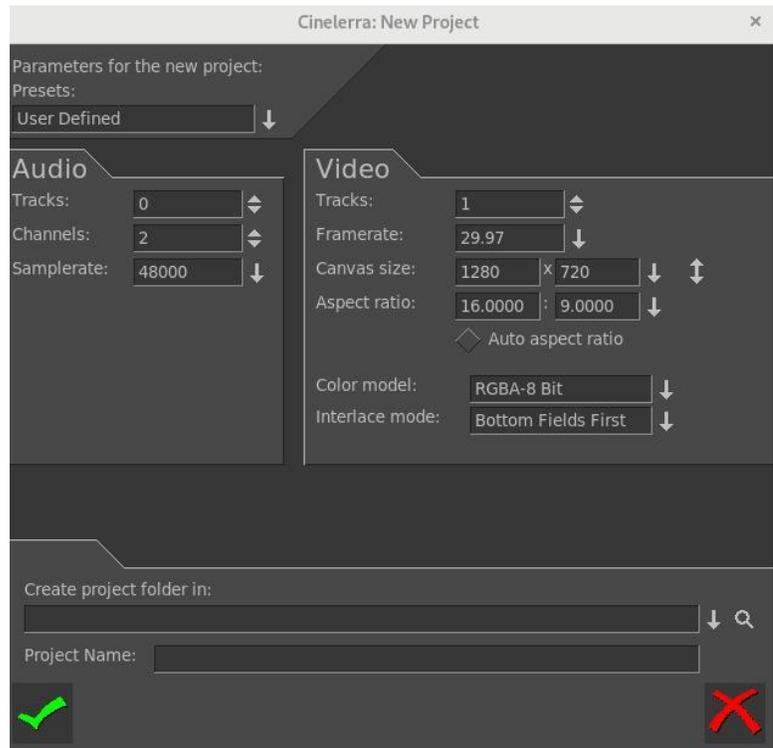


Figure 16: New Project dialog window

Audio attributes

Presets: select an option from this menu to have all the project settings set to one of the known standards. Some of the options are 1080P/24, 1080I, 720P/60, PAL, NTSC, YouTube, and CD audio.

Tracks: (in New Project menu only) sets the number of audio tracks for the new project. Tracks can be added or deleted later, but this option is on the New Project menu for convenience.

Samplerate: sets the samplerate of the audio. The project samplerate does not have to be the same as the media sample rate that you load. Media is resampled to match the project sample rate.

Channels: sets the number of audio channels for the new project. The number of audio channels does not have to be the same as the number of tracks.

Channel positions: the currently enabled audio channels and their positions in the audio panning boxes in the track patchbay are displayed in the channel position widget in the Set Format window. You can see this display on the left side in Figure 1 above. Channel positions are not in New Project window.

The channels are numbered. When rendered, the output from channel 1 is rendered to the first output track in the file or the first sound card channel of the sound card. Later channels are rendered to output tracks numbered consecutively. The audio channel positions correspond to where in the panning widgets each of the audio outputs is located. The closer the panning position is to one of the audio outputs, the more signal that speaker gets. Click on a speaker icon and drag to change the audio channel location. The speakers can be in any orientation. A different speaker arrangement is stored for every

number of audio channels since normally you do not want the same speaker arrangement for different numbers of channels.

Channel positions is the only setting that does not affect the output necessarily. It is merely a convenience, so that when more than two channels are used, the pan controls on the timeline can distinguish between them. It has nothing to do with the actual arrangement of speakers. Different channels can be positioned very close together to make them have the same output.

Video attributes

Tracks: (in New Project menu only) sets the number of video tracks the new project is assigned. Tracks can be added or deleted later, but options are provided here for convenience.

Framerate: sets the framerate of the video. The project framerate does not have to be the same as an individual media file frame rate that you load. Media is reframed to match the project framerate.

Canvas size: sets the size of the video output. In addition, each track also has its own frame size. Initially, the New Project dialog creates video tracks whose size match the video output. The video track sizes can be changed later without changing the video output.

Aspect ratio: sets the aspect ratio; this aspect ratio refers to the screen aspect ratio. The aspect ratio is applied to the video output. The aspect ratio can be different than the ratio that results from the formula: h / v (the number of horizontal pixels divided into the number of vertical pixels). If the aspect ratio differs from the results of the formula above, your output will be in non-square pixels.

Auto aspect ratio: if this option is checked, the Set Format dialog always recalculates the Aspect ratio setting based upon the given Canvas size. This ensures pixels are always square.

Color model: the project will be stored in the color model video that is selected in the dropdown. Color model is important for video playback because video has the disadvantage of being slow compared to audio. Video is stored on disk in one colormodel, usually a YUV derivative. When played back, Cinelerra decompresses it from the file format directly into the format of the output device. If effects are processed, the program decompresses the video into an intermediate colormodel first and then converts it to the format of the output device. The selection of an intermediate colormodel determines how fast and accurate the effects are. A list of the current colormodel choices follows.

RGB-8 bit

Allocates 8 bits for the R, G, and B channels and no alpha. This is normally used for uncompressed media with low dynamic range.

RGBA-8 bit

Allocates an alpha channel to the 8 bit RGB colormodel. It can be used for overlaying multiple tracks.

RGB-Float

Allocates a 32 bit float for the R, G, and B channels and no alpha. This is used for high dynamic range processing with no transparency.

RGBA-Float

This adds a 32 bit float for alpha to RGB-Float. It is used for high dynamic range processing with transparency.

YUV-8 bit

Allocates 8 bits for Y, U, and V. This is used for low dynamic range operations in which the media is compressed in the YUV color space. Most compressed media is in YUV and this derivative allows video to be processed fast with the least color degradation.

YUVA-8 bit

Allocates an alpha channel to the 8 bit YUV colormodel for transparency.

In order to do effects which involve alpha channels, a colormodel with an alpha channel must be selected. These are RGBA-8 bit, YUVA-8 bit, and RGBA-Float. The 4 channel colormodels are slower than 3 channel colormodels, with the slowest being RGBA-Float. Some effects, like fade, work around the need for alpha channels while other effects, like chromakey, require an alpha channel in order to be functional. So in order to get faster results, it is always a good idea to try the effect without alpha channels to see if it works before settling on an alpha channel and slowing it down.

When using compressed footage, YUV colormodels are usually faster than RGB colormodels. They also destroy fewer colors than RGB colormodels. If footage stored as JPEG or MPEG is processed many times in RGB, the colors will fade whereas they will not fade if processed in YUV. Years of working with high dynamic range footage has shown floating point RGB to be the best format for high dynamic range. 16 bit integers were used in the past and were too lossy and slow for the amount of improvement. RGB float does not destroy information when used with YUV source footage and also supports brightness above 100%. Be aware that some effects, like Histogram, still clip above 100% when in floating point.

Interlace mode: this is mostly obsolete in the modern digital age, but may be needed for older media such as that from broadcast TV. Interlacing uses two fields to create a frame. One field contains all odd-numbered lines in the image; the other contains all even-numbered lines. Interlaced fields are stored in alternating lines of interlaced source footage. The alternating lines missing on each output frame are interpolated.

5 - Load, Save, and the EDL

There are many supported file formats that can be loaded and rendered to, that is to say, imported and exported. The format of the file affects what Cinelerra does with it. Some file formats are very slow to display on the timeline, especially video which is highly compressed. Drawing video thumbnails, picons, on the timeline can be very slow. (You can disable picon drawing for these files with the *draw media* toggle in the patchbay to speed up operations).

5.1 EDL - Edit Decision List

When Cinelerra saves a file, it saves the EDL, Edit Decision List, of the current project but does not save any media. Edit decision lists, more commonly referred to as the EDL, are generated by Cinelerra for storing projects. The EDL contains all the project settings and locations of every edit. Instead of media, the file contains pointers to the original media files on disk. EDL files are specific to Cinelerra.

The EDL files generally have an extension of .xml. The purpose of the EDL is to separate the media from all of the editing operations so that the original media remains intact. When the .xml file is loaded, changes to the attributes of the current project are made based on the EDL. Edit decision lists are text files which means they can be edited in a text editor. EDL and XML are used interchangeably.

5.2 Supported File Formats

There are basically 2 kinds of supported file formats, native and ffmpeg. With the addition of ffmpeg, the majority of the supported file formats you will be using comes via this thirdparty package. There are hundreds of ffmpeg file format and codec combinations. This set of possibilities includes qt (quicktime), avi (audio-video interleave), mp4, mp3, mov, mpeg, m2ts, ts, wmv, mts, mpg, flv, mkv, webm, webp, and many more.

The other supported formats, referred to as “native”, include the following:

AC3
Apple/SGI AIFF
Sun/NeXT AU
FLAC
Microsoft WAV
Raw DV and PCM
MPEG Audio and Video - more information follows
+
Still Images
JPEG/EXR/PNG/PPM/TGA/TIFF
+ this list as a Sequence - more information follows
Raw Camera Mode - more information follows

MPEG Files

What is an MPEG file? A very common file format is MPEG because it works with many cameras and televisions. Mpeg2 video, an elementary codec stream for mpeg files, is the most common format. To read this format you need to decode the mpeg stream. You can read and write mpeg natively. Mpeg video encoding is done separately from mpeg audio encoding when using the native file format, meaning that 2 passes are required and then they have to be muxed together. However, if using ffmpeg it is rendered in only 1 pass. DVD uses MPEG as does NTSC and Pal.

5.2.1 Working with Still Images

Still images are played from 1 to any number of times, over and over; they have no duration. You can load still images on video tracks just like you do for any video file. When loaded on the track, use the down arrow on the timeline so you can see the single frame. To extend the length of the image, drag its boundaries just as you would do with regular video media. You can drag the boundaries of a still image as much as you want. Images in Cinelerra have the ability to be dragged to an infinite length.

Alternatively, you can define the initial length of the loaded images. The parameter is set in the Images section of the Settings->Preferences->Recording window.

Unless your original material comes from a digital source using its best resolution (like a digital camera), the first thing you might have to do before you can use it is to somehow capture the assets into a usable digital medium. For old photos, paper maps, drawings or diagrams, you can scan them into a file format like PNG, TIF, TGA or JPG files by using a digital scanner.

Rendering a video to a single image causes the final image file to be overwritten for every timeline position. The rendered file is a single still image of the last frame of the video.

5.2.2 Timelapse Sequence of Images, and Animation

The next areas covered in this section are the following:

- 1) file lists formats such as jpeglist
- 2) and image2 ffmpeg

In order to be reasonably fast to use, you will most likely want to prepare them by creating a script and then load by utilizing this file generated script.

An image sequence is a series of ordered still pictures; for example a bunch of camera shots, frames of an animated scene, or series of frame shots. These can be loaded as multiple files. For timelapse sequences, as the size of camera images increases to 70 megabytes and beyond, and more images can be stored on a memory stick, more cache, memory, and system resources (such as file descriptors) are used by cinelerra to load the images when you use the “concatenate tracks” or “paste at insertion point” strategies. It is very time consuming and resource consuming when each of the image files is loaded and concatenated as edits, and it also plays super poorly. Here is an alternative to the usual “load”. This technique may also be useful for just a bunch of pictures.

File lists formats can be utilized in some way for the following list of types of “Sequence files” The first line of the sequence list file identifies the list codec.

PNGLIST = *.png
EXALIST = *.exa

PPMLIST = *.ppm
CR2LIST = *.cr2

TGALIST = *.tga
JPEGLIST = *.jpg

TIFFLIST = *.tiff
GIFLIST = *.gif

Using the example of jpeg's, the jpeg list sequence file type is the easiest and fastest way to access a sequence of jpg images as a single asset. First build a jpeglist sequence file and name it something like jpeglist.sh. There is an example script of how to do this in the Auxiliary Programs section of the Appendix. Once the jpeglist.sh file is built you can then run it similar to this line:

```
jpeglist.sh /<path>/file.jpg /<path>/DSC*.jpg
```

If <path> is the same on both outfile and infiles, then file.jpg is created in the same directory as infiles, the directory contains the entire asset, and the file list uses relative paths; otherwise the file list contains absolute paths. Since this creates outfile list as a single asset, the memory demand and access time is much lower. When you load the outfile in cinelerra, you will need to set "Try ffmpeg last" since ffmpeg does not work with jpeglist sequence files.

An example output file from running this script residing in the directory where DSC*.jpg files exist is shown below. To use this, turn off ffmpeg probes first, and open timelapse.jpg using File->Load files..

Example: timelapse.jpg

```
JPEGLIST
# First line is always JPEGLIST
# Frame rate:
29.970030
# Width:
6016
# Height:
4016
# List of image files follows
./DSC04948.jpg
./DSC04949.jpg
./DSC04950.jpg
./DSC04951.jpg
... (files in between)
./DSC04997.jpg
./DSC04998.jpg
```

Image2 file format is an alternative method to open an image sequence via ffmpeg. To do this, create 2 files in the same directory as the DSC*.jpg files named: DSC0%04d.opts, and DSC0%04d.jpg . DSC0%04d.opts should contain the following lines which have to be modified to fit your exact requirements for duration, start_number, and frame_rate. Example DSC0%04d.opts :

```
loglevel=verbose
threads=auto
format=image2
codec=mjpeg
start_number=4948
frame_rate=29.97
duration=17.36
```

Example of the contents of the file DSC0%04d.jpg would be just a single line as: JPEG
In this case, make sure "Try ffmpeg first" is enabled, and load DSC0%04d.jpg. This will access the media using ffmpeg which is slower so be patient.

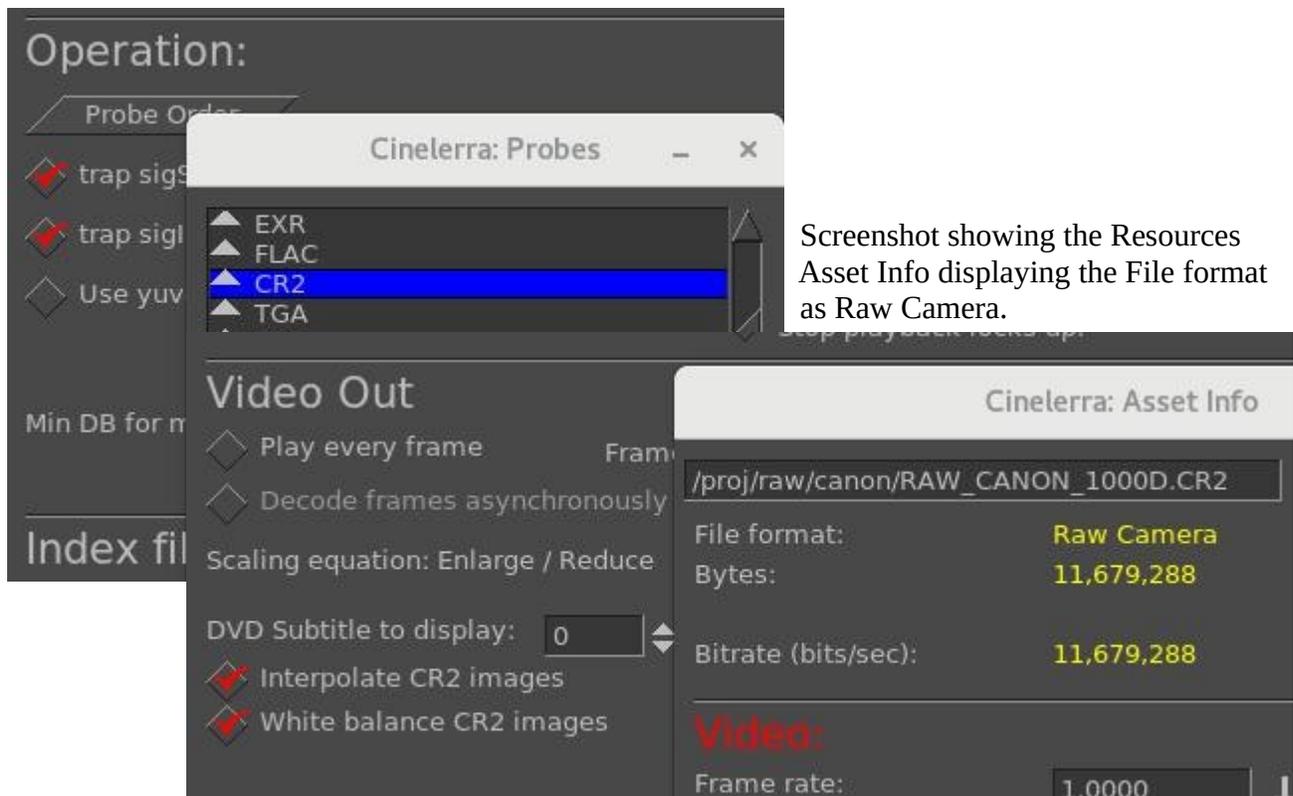
5.2.3 Raw Image Format of Some Digital Cameras & Probe Order

Warning: Expert Usage only. Raw digital camera images are a special kind of image file that Cinelerra can load. Dcraw, as used by Cinelerra, is Dave Coffin's open-source computer program which reads many raw-image formats typically produced by high-end digital cameras. Currently almost 700 of the types of cameras it recognizes are listed at:

<https://www.cybercom.net/~dcoffin/dcraw/>

For example, included is the Canon Powershot SX60 (newly available in August, 2014). Because ffmpeg tries to load "any and every" file if "Try Ffmpeg first" is enabled, it will make an attempt to load Raw Camera files first before any other file driver gets the chance. In addition, there is the possibility that dcraw could conflict with the standard TIFF format, since it might be seen as format type "tiff-pipe". Therefore it is necessary to specifically enable CR2 and either move it to the top or disable "FFMPEG_Early" and enable "FFMPEG_late" in the "Probe Order" as described in another section. These changed settings will be retained across cinelerra sessions in .bcast5. Raw Camera mode is most likely going to be used by expert camera users.

Screenshot below showing CR2 for Raw Camera highlighted/enabled in the Preferences Probes' screen.



Screenshot showing the Resources Asset Info displaying the File format as Raw Camera.

This screenshot [^] as in Settings → Preferences, Playback A Tab, shows the default checked settings of "Interpolate CR2 images" and "White balance CR2 images" which display the raw images in a way that you expect. However, you may want to uncheck them to ensure that no program manipulation has modified your images so that you can add plugins or make your own modifications. Unchecked indicates that the images are as closest as possible to unadulterated raw.

5.3 Loading Files

All data that you work with in Cinelerra is acquired either by loading from disk or recording from a device. This section describes loading. To bring up the Load window go to the File pulldown and choose Load Files. Next *Select files to load*, and click ok (the green checkmark) or *Apply*. When you use the *Apply* button, the Load window remains active for easily loading more files later. Depending on the setting of the Insertion Strategy list box, your file will be either loaded directly on the Program window or in the Resources Media window. If the file is a still image, the project's attributes are not changed and the first frame of the track becomes the image. Cinelerra usually builds an index file if one does not already exist, in order to speed up drawing. You can edit and play the file while the index is being built.

To load a file, you will need to:

- 1) set a directory path
- 2) pick a file selection
- 3) choose a selection strategy for loading the file

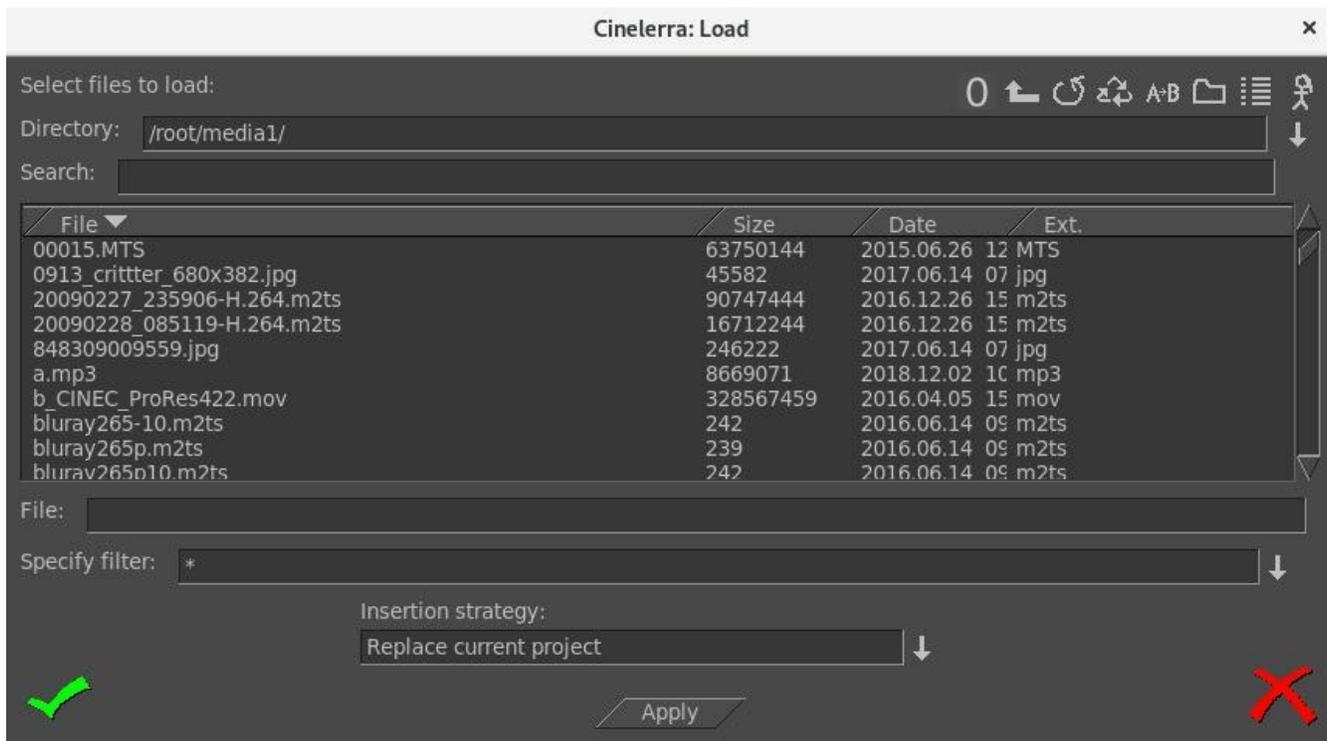


Figure 17: Load file menu. Note the green checkmark for OK and the middle Apply option

Once you have completed making your choices and clicking OK or Apply, by default 3 things occur:

- 1) The existing project is cleared from the screen.
- 2) The project's attributes are changed to match the file's attributes.
- 3) The new file's tracks are created in the timeline.

Insertion Strategy

Cinelerra lets you change what happens when you load a file. In the Load dialog window go to the Insertion strategy box and select one of the options in the drop down menu. Each of these options loads the file a different way.

Replace current project: all tracks in the current project are deleted and a set of new tracks are created to match the source file. Project attributes are only changed when loading XML. If multiple files are selected for loading, Cinelerra adds a set of new tracks for each file. New resources are created in the Resources Window, replacing the current ones.

Replace current project and concatenate tracks: same as replace current project, except that if multiple files are selected, Cinelerra will concatenate the tracks of each file, inserting different source files in the same set of tracks, one after another, in alphanumeric order, starting at 0. New resources are created in the Resources Window, replacing the current ones. Files go across the timeline.

For ffmpeg and mpeg files, when the Insertion strategy methodology in the File → Load files pulldown is chosen to be either *Replace current project* or *Replace current project and concatenate tracks*, the basic session format parameters are reinitialized to match new media. This selects the default asset and determines its width, height, and video length, frame rate, calculates the colormodel, and assumes square pixels to make an intelligent guess about aspect ratio for video. For audio, the sample rate, audio length, and channel count (mono, stereo, or 5.1) are reinitialized. In addition the “Track Size” will be computed and is reinitialized to match the new loaded media. When using “replace” type insertion strategy, the new asset list is the only media in use so that this update saves the user from immediately needing to change the session format to match the only possibility.

Append in new tracks: the current project is not deleted and new tracks are created for the source, one set of tracks for each file. New resources are created in the Resources Window. Files go down tracks.

Concatenate to existing tracks: the current project is not deleted and new files are concatenated to the existing armed tracks, inserted in the same set of tracks of the current project, one after another, in alphanumeric order, starting at the end of the tracks. If the current project has more tracks than the source, the source file will be inserted in the first set of armed tracks. If no tracks are armed, no files will be inserted. New resources are created in the Resources Window.

Paste at insertion point: the file is pasted into the timeline at the insertion point, on the first set of armed tracks. If multiple files are selected for loading, they will be inserted on the same set of tracks, one after the other. New resources are created in the Resources Window.

Create new resources only: the timeline is unchanged and new resources are created in the Resources Window only.

Nest sequence: nested assets are added to the timeline by using the *Nest sequence* insertion strategy. The file will be pasted into the timeline over the current selection or at the insertion point. A nested sequence is media that had already been saved as an EDL earlier. See the section on Nesting.

The insertion strategy is a recurring option in many of Cinelerra's functions. In each place the options do the same thing. If you load files by passing command line arguments to Cinelerra, the files are loaded with *Replace current project* by default.

Loading Multiple Files

In the Load dialog go to the list of files. Selecting files utilizes the motif style selection method.

- 1) Select a single file by highlighting.
- 2) Select multiple files by moving to another file and select it while holding down Ctrl. This selects one additional file.
- 3) Or move to another file and select it while holding down Shift. This selects every intervening file. This behavior is available in most listboxes. It is an especially useful method when used with *Concatenate to existing tracks* insertion strategy to create an images slideshow or a song playlist.

Loading files from the command prompt

Another way to load files is to pass the filenames as arguments on the command line. This starts the program with all the arguments loaded and creates new tracks for every file. For example:

```
{your_cinelerra_program_path} video1.mp4 video2.mp4
```

Finding Files by Extension, Sub-list, or with Search

If there are too many files in your media directory, it can be difficult to find the file you want. For this reason, the Load window allows you to filter which files are displayed in the list box by extension name. Click the dropdown box on the right side of the *Specify filter* list box below the file name text box, and select the file extension of your media (for example: mp4, mov, mp3, avi, jpg, etc). The file list now shows only files with the selected extension. Perhaps even easier is to use the Search box on the top underneath the *Select files to load* listbox. Here you can keyin a character or string to look for.

You can also get a sub-list of potential files to choose from. For example, you know that the file you are looking for begins with the capital letter "C". If you keyin "C" into the selection box immediately below the list of files, and then click the left mouse button, a sub-list of files beginning with the "C" shows up under the selection box. Clicking the right mouse button cancels this sub-list.

Loading the backup

There is one special XML file on disk at all times. After every editing operation, Cinelerra saves the current project to a backup in ``$HOME/.bcast/backup.xml'`. In the event of a crash, the first thing you should do after restarting Cinelerra is select File->Load backup in order to load the backup. This will start Cinelerra at the point in your editing operations directly before the program crashed. It is important after a crash to restart Cinelerra without performing any editing operations as you will overwrite the backup. Note that the backup.xml file is always a single file which means that when you are working with two instances of Cinelerra open at the same time, they use the same backup file. In this case, the last operation made in whatever instance will overwrite the backup.

5.3.1 Sort within Sort in File Load Dialog

When you use the File pulldown to load files, you can do a sort within a sort when you click on the labeled header box. This is useful, for example, when you want to find the smallest file for a specific extension. In the screenshots below, the first illustrates the default “File” sorted alphabetically; the second shows the “Size” is now sorted; the third shows how after sorting on Size, you sort on Ext. The size sort is maintained within the extension sort so that c.d comes before a.d in the File header box because the size is smaller.

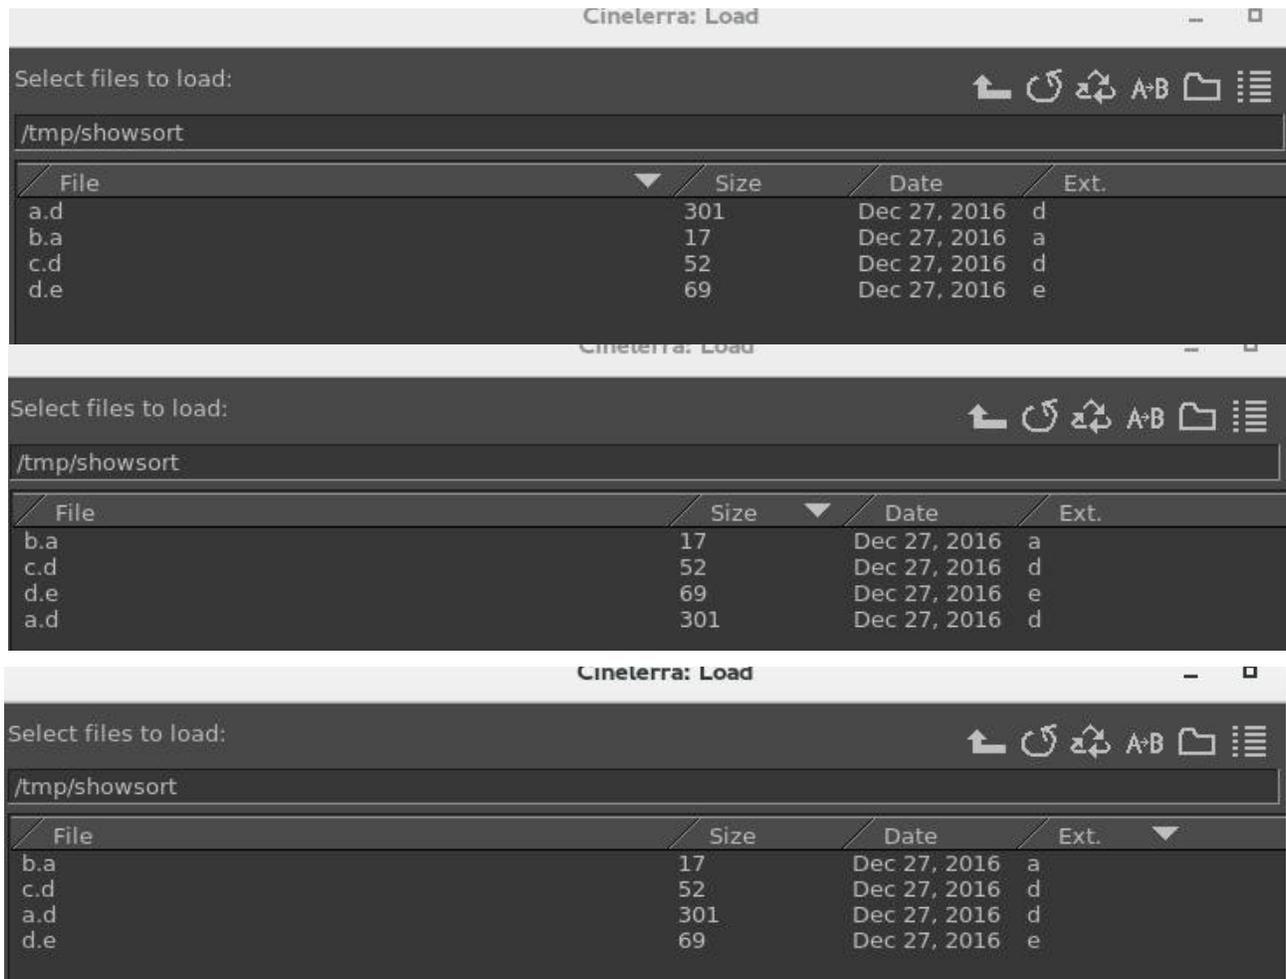


Figure 18: Load - Sort by File name, sort by file Size, and within Extension after a previous Size sort

5.3.2 Size Numeric Format Displayed in File Load

There are several icon buttons at the top on the right hand side of the Load window. Each has a tooltip to explain what it is for. You can see these in the previous figure. One is for File size format. There are 4 numerical representation variations for reporting the file size in the File → Load pulldown. You can see the options in the Load window to the right of the top line that read “Select files to loads:”

- 0** - this is the default and current behavior and shows bytes the same as the `ls -l` command.
- b** - 3 significant digits suffixed by lower case k,m,g,t,b for representing magnitudes in 10^3 (1000)
- B** - 3 significant digits followed by upper case characters for representing magnitudes in 2^{10} (1024)

; - like the exact default byte representation but with comma separators for easy reading. Periods can not be used as separators due to locale conflict with ffmpeg coding.

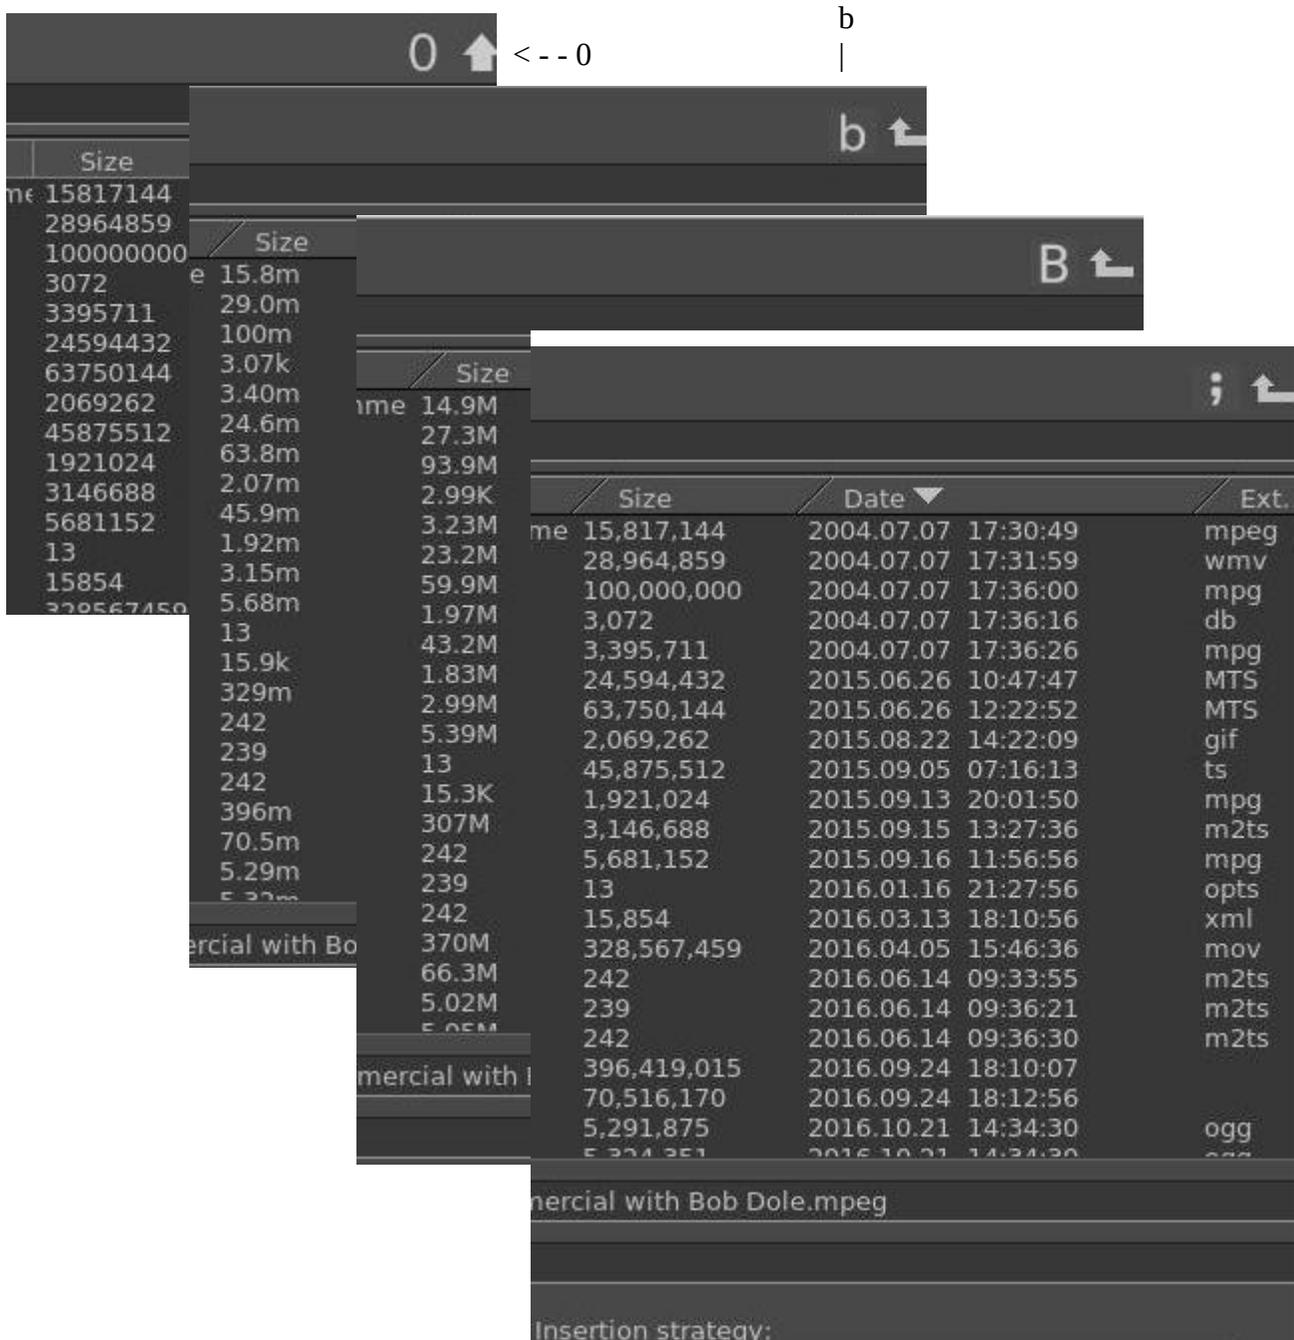


Figure 19: Load windows with various Numeric Sizes

5.3.3 Probe Order when Loading Media

Why is this mentioned here? So many programs have been written whose functionalities overlap and you may want to ensure that the one you wish to use is actually used. Over time which one matches first may vary. Ffmpeg is so generic that if your setting is “Try ffmpeg first” it will almost certainly get

used and it leaves little chance that other methods will even get a chance. Some of the codec file drivers can open a variety of media, and some of the more common methods may have more than one file driver which could be useful to decode your media file, for example Tiff. For expert specialized usage, when you want to guarantee that a certain method is used, you can change the “probe order”. Use the pulldown Settings → Preferences to get to the Interface tab where you will see a box in the Operation section on the left side called “Probe Order”. Click on the box and use the up/down/enabled boxes to change the order of the item you have highlighted.

Up – move the item up 1 (if the item is currently on the top, it will be moved to the bottom)

Down – move the item down 1 (similarly, if on the bottom, it will be moved to the top)

Enabled – there will be a check mark if the item is currently enabled. If you disable it, it will not be used to probe the media to determine what it is. Double left mouse click will toggle the enabled/disabled state of the highlighted item. If both FFMPEG Early/Late are enabled, the FFMPEG code could be run twice if the 1st FFMPEG run failed (but so will the 2nd!).

The default setup is set to duplicate the past expected behavior with the exception that CR2 for Raw Camera mode is disabled. Changes made in the settings will be retained in the .bcast5 file.

Screencast showing the first few probe items. Note that the up arrow on the left, signifies “enabled”.



Scrolling down shows the next 2 pages of possible drivers for a total of 18.

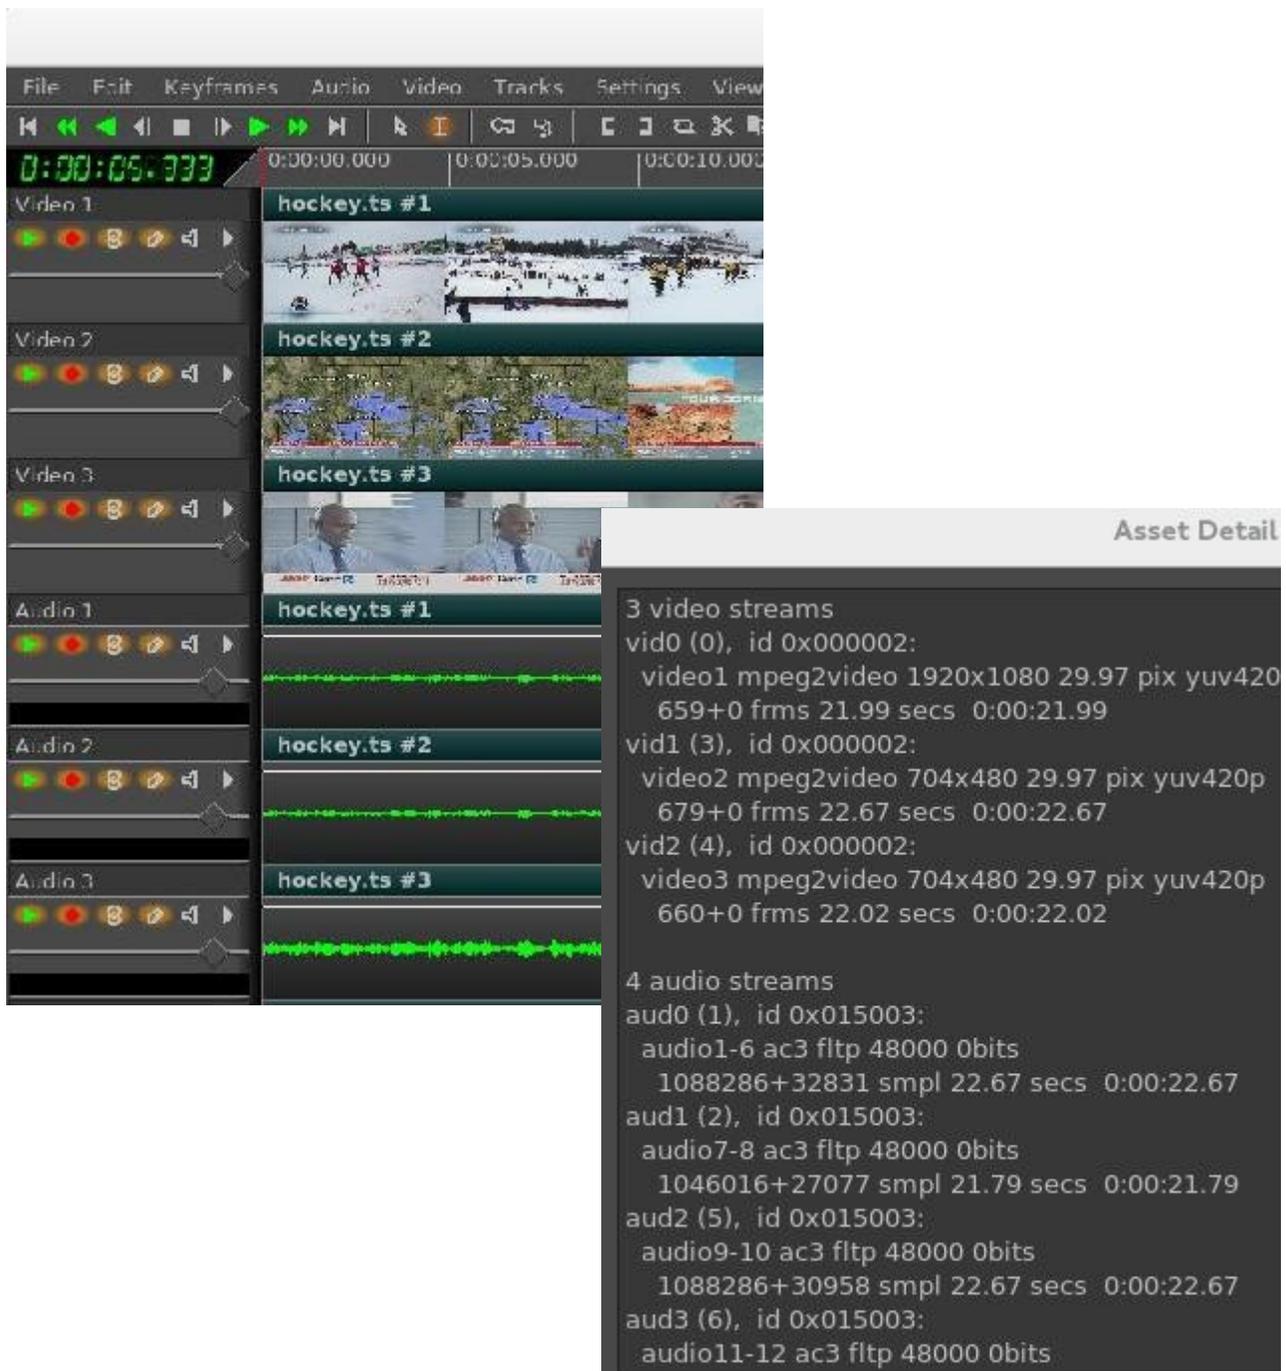
The order change will not take effect until you click on the checkmark in both the Probes window and the Preferences window. When you click on the FF button, which is in the upper right-hand corner of the main page to change “Try FFMpeg first/last”, enabling of FFMPEG_Late or FFMPEG_Early will be toggled automatically in Probes to match that choice but does NOT change its position in the table. Be sure to only click on the FF button without the Preference/Probes window up to avoid unexpected results. It is also recommended to leave FFMPEG_Early/FFMPEG_Late close to the top/bottom

positions. There is one case where you may want to disable all of the probes if you want to force PCM – Pulse Code Modulator. This code is always run when all other probes fail.

5.3.4 Program Selection Support after Load

Some kinds of media have “program” streams, like captured mpeg broadcast stream data. For example, you may be able to “tune” to channel 9, but be able to see 9-1, 9-2, and 9-3 on your TV. If you open a capture of this kind of media, all of the channels are present in the timeline. To select and view just one program, you can use Alt-1 to select program 1, or Alt-2 to select program 2, etc. up to Alt-8. This will remove all of the other unrelated tracks and reset the format. This feature can be used even if there is only one program, by pressing Alt-1, and the effect will be to reset the session format to the parameters from the media probe. Note that there may be several audio “programs” associated to a video stream; for example, there may be dialog in another language or some kind of descriptive dialog. Since the first associated audio is always selected, this may not produce the intended results.

Below are screenshots illustrating multiple program streams. The left screenshot is a partial main Cinelerra window showing a pre-recorded broadcast TV media/audio stream with 3 programs plus several associated audio tracks. The second screenshot of “Asset Detail” provides detailed information on each of the streams obtained through executing the Info Details as explained in the section “Info Asset Details”.



5.4 Saving Your Work

You can save your work as a project, which is what is loaded in Cinelerra now, or as an export, which is all the media it takes to reproduce your project space.

5.4.1 Saving Project Files

Saving XML files is useful to save the current state of Cinelerra before quitting an editing session. Cinelerra saves projects as XML files. There are a few options you can use to save your work via the File pulldown menu: Save, Save as..., Export project, Save backup. You can either overwrite an existing file or enter a new filename. Cinelerra automatically concatenates .xml to the filename if no .xml extension is given.

When Cinelerra saves a file, it saves the EDL of the current project but does not save any media, instead just pointers to the original media files. For each media file, the XML file stores either an absolute path or just the relative path. If the media is in the same directory as the XML file, a relative path is saved. If it is in a different directory, an absolute path is saved.

You have to be careful when moving files around to avoid breaking the media linkages. You can keep the media and the XML file in the same directory forever and freely move the whole directory, since relative paths are saved. Alternatively you can save the XML file in a different directory than the media but then you can't move the media. In this case you can freely move your XML file around, since absolute paths are saved. If you saved your XML file in the same directory as your media but you would like to move location, you can change the paths from relative to absolute by going to File->Save as... and entering the new location. Similarly if you saved your project outside your media directory but you would like to move your media to another location, you can change the paths from absolute to relative by going to File->Save as... and saving your XML file in the same directory as the media.

You can also repair broken media linkage by editing the XML file in a text editor. For every media you moved, search for the old path and replace it with the new one. You should make a backup copy of your XML file before editing. You can also replace the path of every asset whose source file you moved also within the program, by entering the new location in the Asset info window. To open this window, right click on the asset in the Resources window and choose Info... in the popup menu. Directly type the path in the first field of the dialog or click on the magnifier on the right to browse your files. Operating from the GUI is convenient only when a very small number of changes is needed.

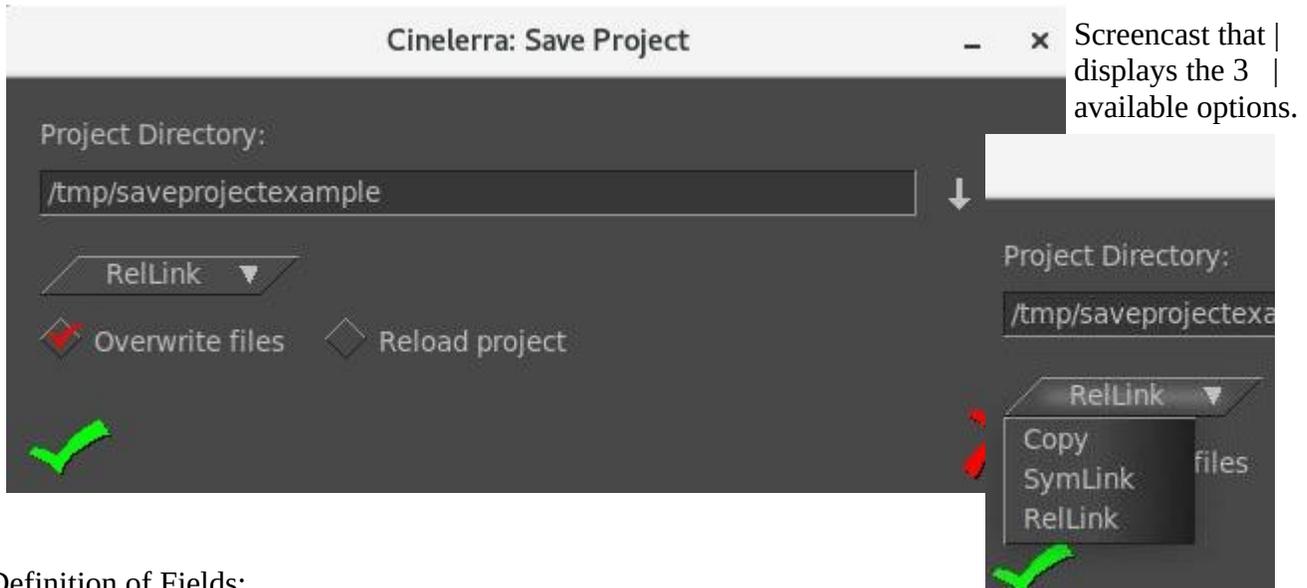
Real-time effects in an XML file have to be re-created every time you play it back. The XML file also contains copies of all the source assets on disk, which takes up space. Render your projects to a final format for more persistent storage of the output.

5.4.2 Export Project – Save or Moving Project to another Computer

A File pulldown called “Export Project...” is also available. Although, it can be used in the same manner as the other “save” options, it is very useful when it is necessary to move a project to another computer that may have a different top level directory structure or if you want to include subdirectories to better organize your files.

Originally, the easiest way to maintain a project for moving to another computer, was to put all of the files in a single directory with no subdirectories along with the EDL saved .xml file. This is commonly called a “flat” file structure. So if the media was in the same directory as the XML file, a relative path was saved. If it was in a different directory, an absolute path was saved.

Screencast of “Export Project...” option popup.



Definition of Fields:

Project Directory – name of the directory where you want the xml file to be saved. It will only create a subdirectory in 1 level of the defined directory.

Available option types for saving a project:

- 1) Copy – all files are copied to the project directory and the xml is saved; same as original “flat”. Option is very useful to ensure all files needed to illustrate a problem are available for analysis.
- 2) SymLink – symbolic links are created for absolute paths of media in their current location.
- 3) RelLink – symbolic links are created for relative paths of media in their current location. This option allows for using relative paths without the requirement to maintain a “flat” file structure and makes it easy to move to another computer.

Overwrite files – when checked, if any files with the same name currently exist in the directory, they will be overwritten. In any case, the named XML file will always be overwritten.

Reload project – when checked, after the save option the new saved project will be loaded.

Default is not to do so.

Keep in mind that to maintain the integrity of your project xml file for easy moving to another computer, do not delete the symbolic links. You will want to use “cp -a” to maintain the links for moving to a USB key or another computer.

5.4.3 Information about Backups and Perpetual Session

In an effort to minimize loss of work due to user, hardware, or software issues, cinelerra has some automatic backup capabilities.

Cinelerra automatically saves every *editing operation* to the current project on disk continuously to a file named \$HOME/.bcast5/backup.xml. In the unlikely event of a crash, when you restart cinelerra, you should select File->Load backup in order to continue with the operations that were recorded before

the crash. If you have more than 1 instance of Cinelerra running, only the last editing operation made in whichever instance it was last made, will overwrite the backup.

There is still 1 more backup that may save you. If for some reason you forgot to use “Load backup” immediately when restarting, you have a second chance to use File → Load and select \$HOME/.bcast5/backup.prev as long as you only loaded a different file and have performed no editing operations. This same file is also used by multiple instances of cinelerra.

Perpetual session is very useful for working on a project over many days so you can just quit before shutting down and the next time you start up Cinelerra you will be right back where you left off. You will retain all of your undo’s and redo’s. The binary file name is \$HOME/.bcast5/perpetual.dat and as long as Settings → Preferences, the Appearance tab has the Flag “Perpetual session” set this capability takes effect. It is very important to understand that this is not the same as the continuously editing-operation-updated backup.xml file. The perpetual.dat file is *only* updated when you Quit cinelerra in the normal manner. Which means if you interrupt the program, or kill it, or there is a segv or system crash, the perpetual.dat file will only reflect the state of your project from when you last started cinelerra and none of the editing/undo’s/redo’s you executed during the current session which was not ended normally.

Some notes to keep in mind about Perpetual session are:

- when you Quit in the normal manner, it does not have to ask whether or not to save a backup
- takes disk space in .bcast5 area and this could get really big
- after you complete a project, it is advisable to turn off the Perpetual session flag before quitting so that when you start a new project, you can start with a fresh perpetual.dat by turning the flag on or after stopping cinelerra, delete the current \$HOME/.bcast5/perpetual.dat file
- only session data is backed up (not program feature setup)
- the files backup.xml and backup.prev will operate the same as before so that if there is a crash, you will want to use File → Load backup in order to continue where you were interrupted.

This page intentionally left blank so that chapters begin on a right facing page.

6 - Editing

Editing comprises both time and track space. The timeline consists of the time certain media appear on the track going left to right and a set of tracks from the top to the bottom. There are 2 methods of timeline editing – drag and drop editing, also called “arrow mode”, and cut and paste editing or “I-beam mode”. Cut and Paste is the default editing mode. An additional, but not often considered editing method is called two-screen editing where the Viewer is used to view media and then the desired media is transferred to the timeline.

The timeline is where all editing decisions are made. This is a stack of tracks in the center of the main window. It can be scrolled up, down, left and right with the scrollbars on the right and bottom. It can also be scrolled up and down with a mouse wheel, or left and right while holding down the Ctrl key and using the mouse wheel.

The active region is the range of time which is affected by editing commands on the timeline. The active region is determined first by the presence of in/out points on the timeline. If those do not exist the highlighted region is used. If no highlighted region exists, the insertion point is used as the start of the active region. Some commands treat all the space to the right of the insertion point as active while others treat the active length as 0 if no end point for the active region is defined.

Most importantly, editing decisions never affect source material meaning that it is non-destructive editing. So not only does your original media stay completely untouched, it is much faster than if you had to copy all the media affected by an edit. Editing only affects pointers to source material, so if you want to have a new modified media file at the end of your editing session which represents the editing decisions, you need to render it. Saving and loading your edit decisions is explained in the Load, Save and the EDL section and rendering is explained in the section on Rendering.

In the following editing sections, references to common operations are scattered within any of the modes where they seem pertinent. However, many of the editing operations work in different modes.

6.1 The Patchbay

On the left of the timeline is a region known as the patchbay. The patchbay enables features specific to each track as described next.

Textbox for naming the track. The default names will usually be Video #, Audio #, or Mixer # if using the multi-camera/mixer operations. A # will be designated for subsequent tracks as in 1, 2, 3 and so on.

Expander which is a down arrow on the right side, is for viewing more options on the patchbay and for viewing the effects represented on the track. You can just click on the expander to expand or collapse the patchbay and the track. If it is pointing sideways, the track is collapsed. If it is pointing down, the track is expanded. Existing effects appear below the media for the track.

Below the textbox name are several toggles referred to as “attributes” for different features (currently there are 5 as shown in this figure). If the toggle button is shadowed by a color, the feature is enabled. If the toggle is the background color of most of the window, it is disabled. Click on the toggle to enable/disable the feature.

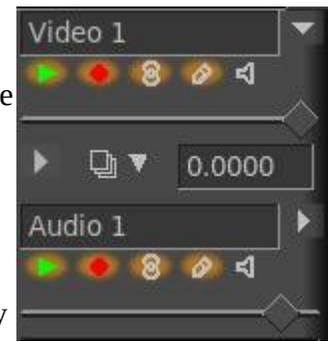


Figure 20: Patchbay

Several mouse operations speed up the configuration of several tracks at a time. Click on an attribute and drag the cursor across adjacent tracks to copy the same attribute to those tracks. Hold down Shift while clicking a track's attribute to enable the attribute in the current track and toggle the attribute in all the other tracks. Or you can (1) hold down Shift while clicking an attribute, (2) click until all the tracks except the selected one are disabled, (3) then drag the cursor over the adjacent track to enable the attribute in the adjacent track.

The “attributes” are described here next.

1) *Play track* determines whether the track is rendered or not. If it is off, the track is not rendered. For example if you turn it off in all the video tracks, the rendered media file will have only audio tracks. If the track is chained to any other tracks by a shared track effect, the other tracks perform all the effects in this shared track, regardless of play status of the shared track that in this particular case affects the media output but not fade and effects.

2) *Arm track* determines whether the track is armed or not. Only the armed tracks are affected by editing operations. Make sure you have enough armed destination tracks when you paste or splice material or some tracks in the material will get left out. In addition to restricting editing operations, the armed tracks in combination with the active region determine where material is inserted when loading files. If the files are loaded with one of the insertion strategies which do not delete the existing project, the armed tracks will be used as destination tracks.

Press the Tab key while the cursor is anywhere over a track to toggle the track arming status. Press Shift-Tab while the cursor is over a track to toggle the arming status of every other track.

3) *Gang faders* cause the fader to track the movement of whatever other fader you are adjusting by dragging either the fader or the curve on the track. It doesn't affect the editing made with menu controls. A fader is only ganged if the arm track is also on. This is often used to adjust audio levels on

all the tracks simultaneously. Gang also causes Nudge parameters to synchronize across all the ganged tracks.

4) *Draw media* determines if picons or waveforms are drawn on the asset in the track. You may want to disable this if you know that the media/format takes a long time to draw on the timeline. By default it is set to on in order to see picons on the timeline.

5) *Don't send to output*, more commonly called *mute*, causes the output to be thrown away once the track is completely rendered. This happens whether or not *Play track* is on. For example if you mute all the video tracks, the rendered media file will have a blank video track. Mute track is represented on the timeline with a line that has the default color of blue. Use the pulldown View -> Mute to have the line displayed. It is a keyframable attribute, but Mute track keyframing is a toggle and it has only the two values of on or off. If a track is part of a shared track effect, the output of the track with the shared track effect is overlaid on the final output even though it is routed back to another track (the shared track). Mute track is used to keep the track with the shared track effect from overlapping the output of the source track (the shared track) where the shared track effect is not present.

Fader slider fade values are represented on the timeline with a white curve that is keyframable. All tracks have a fader, but the units of each fader depend on whether it is audio or video. Audio fade values are in dB. They represent relative levels, where 0 is the unaltered original sound level, -40 is silence, -80 the minimum value set by default. You can move fader and keyframes down to -80 but the parameter's curve won't go below -40. For your convenience you can set a different fade range with the curve zoom. Audio fader's main purpose is to "fade out" sound or to lower the sound level smoothly to silence, or "fade in" to make sounds appear gradually instead of suddenly. Video fade values are the percentage of opacity of the image in normal overlay mode, the percentage of the layer that is mixed into the render pipeline in the other overlay modes.

Click and drag the fader to fade the track in and out. If it is ganged to other tracks of the same media type, with the arm option enabled, the other faders should follow. Hold down the Shift key and drag a fader to center it on the original source value (0 for audio, 100 for video).

Mixer in the expanded patchbay for that track designate the multi-camera mixer mode.

Overlay mode in the expanded patchbay is used for porter-duff operations and is full explained in the section on Overlay Modes.



Figure 21: Video Overlay and Nudge



Figure 22: Audio Pan and Nudge

Nudge is in the expanded patchbay. The nudge value is the amount the track is shifted left or right during playback. The track is not displayed shifted on the timeline, but it is shifted when it is played back. This is useful for synchronizing audio with video, creating fake stereo, or compensating for an effect which shifts time, all without altering any edits.

Enter the amount of time to shift to instantly shift the track. Negative numbers make the track play later. Positive numbers make the track play sooner. The nudge units are either seconds or the native units for the track (frames or samples). Select the units by right clicking on the nudge textbox and using

the context sensitive menu. Nudge settings are ganged with the Gang faders toggle and the Arm track toggle. Use the mouse wheel over the nudge textbox to increment and decrement the value.

Pan is available in the expanded patchbay for audio tracks via a panning box. Position the pointer in the panning box and click/drag to reposition the audio output among the speaker arrangement. The loudness of each speaker is printed on the relative icon during the dragging operation. The panning box uses a special algorithm to try to allow audio to be focused through one speaker or branched between the nearest speakers when more than 2 speakers are used.

Automatic audio mappings

Several convenience functions are provided for automatically setting the panning to several common standards. They are listed in the Audio menu. These functions only affect armed audio tracks. They are:

Audio->Map 1:1

This maps every track to its own channel and wraps around when all the channels are allocated. It is most useful for making 2 tracks with 2 channels map to stereo and for making 6 tracks with 6 channels map to a 6 channel sound card.

Audio->Map 5.1:2

This maps 6 tracks to 2 channels. The project should have 2 channels when using this function. Go to Settings->Format to set the output channels to 2. This is most useful for down-mixing 5.1 audio to stereo.

Standard audio mappings

Although Cinelerra lets you map any audio track to any speaker, there are standard mappings you should use to ensure the media can be played back elsewhere. Also, most audio encoders require the audio tracks to be mapped to standard speaker numbers or they will not work.

In the channel position widget, the channels are numbered to correspond to the output tracks they are rendered to. For stereo, the source of channel 1 needs to be the left track and the source of channel 2 needs to be the right track. For 5.1 surround sound, the sources of the 6 channels need to be in the order of center, front left, front right, back left, back right, low frequency effects. If the right tracks are not mapped to the right speakers, most audio encoders will not encode the right information if they encode anything at all. The low frequency effects track specifically can not store high frequencies in most cases.

6.2 Manipulating Tracks

Tracks in Cinelerra either contain audio or video. There is no special designation for tracks other than the type of media they contain. When you create a new project, it contains three default tracks: one video track and two audio tracks. You can still add and delete tracks from the menus. The Tracks menu contains a number of options for dealing with multiple tracks simultaneously. Each track itself has a popup menu which affects one track.

Operations in the **Tracks pulldown** affect only tracks which are armed. *Move tracks up* and *Move tracks down* shift all the armed tracks up or down the stack. *Delete tracks* deletes the armed tracks. *Delete last track* deletes the last track, whether it is armed or not.

Concatenate tracks operation copies all the assets of every disarmed but playable track and concatenates it by pasting those assets at the end of the first set of armed tracks. They are pasted one

after the other, keeping the same order they have on the stack. If there are two armed tracks followed by two disarmed tracks, the concatenate operation copies the assets of the two disarmed tracks and pastes them after the assets of the two armed tracks. If there are three disarmed tracks instead, the assets of two tracks are pasted after the assets of the armed tracks and the assets of the third track are pasted at the end of the first armed track. The destination track wraps around until all the disarmed tracks are concatenated. Disarmed tracks that are not playable are not concatenated.

Append to project allows for creating new tracks after any existing tracks.
Add subttl will add a track for subtitles at the top of the other tracks.

The **Audio and Video pulldowns** each contain an option to add a track of their specific type. In the case of audio, the new track is put on the bottom of the timeline and the output channel of the audio track is incremented by one. In the case of video, the new track is put on the top of the timeline. This way, video has a natural compositing order. New video tracks are overlaid on top of old tracks.

6.3 Two Screen Editing

This is a fast way to construct a program out of movie files. The idea consists of viewing a movie file in one window and viewing the program in another window. Subsections of the movie file are defined in the viewer window and transferred to the end of the program in the program window. Two screen editing can be done simply by using keyboard shortcuts. To get familiar with which keys to use, move the mouse pointer over the transport panel and a tooltip appears, showing what key is bound to that button.

To begin a two screen editing session, load your media resources by using the main menu **File pulldown** and choose *Load files*; make sure the insertion mode is set to “Create new resources only”. This insertion strategy is to ensure that the timeline stays unchanged while new resources are brought in. Go to the Resources window and select the Media folder. The newly loaded resources will appear. Double click on a resource or drag it from the media side of the window over to the Viewer window.

Check to make sure there are enough armed tracks on the timeline to put the subsections of source material that you want. Usually this would be one video track and two audio tracks, but if there are not enough, just create new tracks or arm more tracks.

Now to start your 2 screen editing, in the viewer window, define a clip from the media file:

- Set the starting point with the In pointer button. You will see a left hand bracket on the timebar.
- Move your cursor to the ending point of the clip you want to use.
- Set the ending point with the Out pointer right hand bracket.

These In/Out points define a clip. You can now use this in a couple of different ways.

Splice – The splice icon, or shortcut letter “v”, inserts the selected area in the timeline after the insertion point. After the splice has taken effect, the insertion point moves to the end of the edit ready to be used as the next splice location. This way you can continuously build up the program by splicing. If an In point or an Out point exists on the timeline the clip is inserted after the In point or after the Out point. If both In and Out points are set on the timeline, the clip is inserted after the In point. If there are edits after your chosen splice location on the timeline, they will be moved to the right.

Overwrite – The overwrite icon, or shortcut letter “b”, overwrites the region of the timeline after the insertion point with the clip. If an In point or an Out point exists on the timeline the clip is overwritten after the In point or after the Out point. If both In and Out points are set on the timeline, the clip is inserted after the In point. If a region is highlighted or both In and Out points exist they limit the region of the overwriting and the clip may therefore be shortened. Here is a detailed explanation to take advantage of this method.

To overwrite exactly on a precise region of the timeline:

- Arm only tracks to change.
- Define the destination region on the timeline with [and], the In and Out points.
- Define the clip you want to use in the viewer with [and], the In and Out points..
- Overwrite from Viewer to the timeline.

If the destination region is shorter than the clip defined in the viewer, the portion of the clip longer than the destination region won't be inserted and on the timeline the following edits won't move.

If the destination region is longer than the clip defined in the viewer, the destination region will shrink and on the timeline the following edits will move to the left.

Clip – The clip icon, or shortcut letter “i”, generates a new clip for the resource window containing the affected region but does not change the timeline. Every clip has an optional/default title and description.

Copy – The copy icon, or shortcut letter “c”, copies the selection into the copy buffer.

6.4 Cut and Paste Editing

This is the more traditional method of editing and therefore is the default. To enable the cut and paste editing mode on the timeline, select the I-beam toggle on the control bar at the top of the main program window. You can copy edits in the same track, copy from different tracks in the same instance, start a second instance of Cinelerra and copy from one instance to the other or load a media file into the Viewer and copy from there.

To start editing, load some files onto the timeline. Select a region of the timeline by click dragging on it and select the cut button to cut it. Move the insertion point to another point in the timeline and select the paste button. Assuming no In/Out points are defined on the timeline this performs a cut and paste operation.

Most editing operations are listed in the **Edit pulldown**. Some of them have a button on the program control toolbar as well as a keyboard shortcut. The keyboard shortcut is in parenthesis here.

- *Split | Cut* (x) - Delete the selected area and put it in the cut buffer for future pasting.
- *Copy* (c) - Copy the selected area and put it in the cut buffer for future pasting.
- *Paste* (v) - Paste the material that is in the cut buffer.
- *Clear* (Del) - Clear the selected area. If the insertion point is over an edit boundary and the edits on each side of the edit boundary are the same resource, the edits are combined into one edit comprised by the resource. The start of this one edit is the start of the first edit and the end of this one edit is the end of the second edit. This either results in the edit expanding or shrinking.
- *Paste silence* (Shift+Space) - Paste blank audio/video for the length of the selected area. Following

edits will be pushed to the right.

- *Mute Region* (m) - Overwrite blank audio/video on the selected area. Following edits don't move.
- *Trim Selection* - Delete everything but the selected region.
- *Select All* (a) - Select the whole timeline.

In Cut and Paste editing mode you can edit labels as well. By enabling *Edit labels* in the **Settings pulldown**, or by disabling the Lock labels from moving button on the Program Control Tool Bar, labels will be cut, copied or pasted along with the selected regions of the armed tracks.

Using labels and In/Out points are useful in editing audio. You can set In/Out points for the source region of the source waveform and set labels for the destination region of the destination waveform. Perform a cut, clear the In/Out points, select the region between the labels, and perform a paste.

In and Out Points

The In and Out bracket placement is explained here to illustrate their usage. Because of the shape of the markers "[" and "]" you may assume that they are inclusive - that everything placed in between would be included in the clip, such as in the case of being transferred to the timeline from the Viewer. In reality, one of the two markers will not include the frame that was visible at the time the marker was affixed. Depending on whether the "Always show next frame" option is used or not, it is the In or Out marker that will not be inclusive.

To obtain a clip on the timeline exactly as you saw in the Viewer, you must necessarily move the In mark back from the beginning before the first desired frame or move the Out mark forward after the last desired frame, depending on the "Always show next frame" setting.

Some of the confusion can be attributed to the fact that the Viewer shows frames, while the markers determine spaces, i.e. times, that are not visible between frames. You have to think of each frame as being delimited by two spaces - one preceding and one following. The In mark is always placed before the displayed frame and the Out mark is always placed after the displayed frame, while taking into account in its calculations whether the "Always show next frame" option is used or not. If you just remember that the reference of the markers is in the middle of the icon, you will avoid confusion.

Other editing operations:

- *Overwrite* - To perform overwriting within the timeline paste on a selected region (highlighted or between In/Out points). The selected region will be overwritten. If the clip pasted from the clipboard is shorter than the selected region, the selected region will be shrunk. Following edits will move. If the clip pasted from the clipboard is longer than the selected region, the selected region will be overwritten with the first part of the clip and the remaining part of the clip will be written after the overwriting. Following edits will move.
- Tracks -> *Concatenate tracks*. This operation copies all the assets of every disarmed but playable track and concatenates it by pasting those assets at the end of the first set of armed tracks. They are pasted one after the other, keeping the same order they have on the stack.

Split – blade cut and hard edges:

You can cut the tracks into 2 pieces on the timeline by putting the hairline cursor on the place you want to do a cut and then using the character “x” or the scissors tool. A “cut” uses a non-empty selection region, where the “blade cut” or “split” has no duration in the selection, just a hairline. As usual the use of cut when a selection is set, deletes/cuts the highlighted area. In the case where an In point or an Out point exists on the timeline, the clip is split at the location of the In/Out point since it has priority over the cursor location. A blade cut simply splits the edit into two edits. In order to have the video and audio aligned, it works best to have Settings → Align cursor on frames. When a blade cut occurs, the edges are created as “hard edges”. These are edges that cannot be deleted by track optimizations.

When you do a blade cut/split, all armed tracks will be included in the cut and green-colored triangles will show on the bottom of the track on both the left and the right side of the cut. This is a “hard edge” marker toggle, as opposed to the soft edge designation for an ordinary edit. The “hard edge” marker can be toggled off/on if so desired. In order to not interfere with the usual drag handles, only a few pixels are used for the toggle so you have to be sure you have the cursor right over the hard edge triangle – when in position, it will be obvious because you can see an arrow pointing to the corner. Use Shift-left mouse button 1 to toggle off/on the hard edge marker on all tracks simultaneously.



Figure 23: blade cut

6.5 Drag and Drop Editing

To enable the drag and drop editing mode on the timeline, select the arrow toggle on the control bar at the top of the main program window. Drag and drop editing is a quick and simple way of working in Cinelerra, using mostly only the mouse. The basic idea is to create a bunch of clips, then drag them in order into the timeline, thus building prototype media that you can watch in the compositor. If after watching it, you wish to re-arrange your clips, set effects, add transitions or insert/delete material, just drag and drop them on the timeline.

To simply get started, perform the following operations which are useful for working in a drag and drop editing session. First load your media by using the main menu File pulldown and choose *Load files*; make sure the insertion mode is set to “Create new resources only”. This loads the files into the Resources window.

- Create some video and audio tracks on the timeline using the Video and Audio pulldowns.
- Open the Media folder in the Resources window. Make sure the necessary tracks are armed and drag a media file from the Resources window to the timeline. If the media has video, drag it onto a video track or if just audio, drag it onto an audio track. For a still image, drag it onto a video track.

You can also drag multiple files from the Resources window. When dropped in the timeline they are concatenated. If you have *Display Icons* selected in the Resources window, drawing a box around the files selects contiguous files. If you have *Display Text* selected, Ctrl-clicking on media files selects additional files one at a time; Shift-clicking on media files extends the number of highlighted selections. In addition to dragging media files, if you create clips and open the clip folder you can drag clips onto the timeline.

Cinelerra fills out the audio and video tracks below the dragging cursor with data from the file. This affects what tracks you should create initially and which track to drag the media onto. To drag and drop a file on the Program window, you need to create on the timeline the same set of tracks as your media file.

When you drag your chosen media from the media folder to the timeline, your mouse pointer will drag a thumbnail and, once over the timeline, the outline of a white rectangle, as big as the edit you are going to have appears. Drag the media to the desired position of an empty track of the timeline and drop it. If there are other edits on that track, when you move the white outline over an edit, you will see a bow tie symbol >< appearing at edit boundaries. If you drop the media there, the new edit will start from the edit boundary indicated by the center of the bow tie ><.

Since the mouse pointer is in the middle of the white outline, when this rectangle is bigger than the visible part of the timeline, it is quite cumbersome to precisely insert it for long media. Lengthening the duration visible in the timeline by changing the sample zoom in the zoom panel will reduce the size of the white rectangle, making a precise insertion possible.

When you drag and drop edits within the timeline:

- If you drop an edit when bow ties >< are shown, that edit will be cut and pasted starting at the edit boundary indicated by the center of the bow tie ><. Following edits on the same track will move.
- If you drop an edit when there are no bow ties >< shown, the original edit will be muted and pasted where you dropped it. No edits will move. A silence will appear in place of the original edit.
- If you have more armed tracks on the timeline than in the asset you are dragging, only the following edits of the tracks affected by the drag and drop operation will move to the right. This will cause loss of synchronization. To restore it, disarm the tracks affected by the drag and drop operation, highlight the just dropped edit and paste silence over it using the Edit pulldown, *Paste Silence*.

Labels sometimes work differently in Drag and Drop editing mode in that you can't drag and drop them. They might be locked to the timebar, even with the Edit labels option enabled. Although with the Edit labels option enabled, if a selected area of a resource is spliced from the Viewer to the timeline in a position before labels, these labels will be pushed to the right for the length of the selected area.

In/Out points can be used to perform Cut and Paste operations in Drag and Drop mode as well as in Cut and Paste mode. Use the Edit pulldown to view the list and their keyboard shortcuts.

6.5.1 Copy/Paste Behavior

There are many options for moving, copying, pasting, inserting, and deleting selected “edits”, more commonly referred to by the user as “clips”, when in the Drag and Drop (arrow) editing mode. This makes it easier to avoid constantly having to disarm/arm tracks. To create a selection move the cursor over the clip and just click the left mouse button; remove a selection by left mouse button click again. This will mark your selection with a colored border which contains some red. The easiest way to initially use the various modes is to click on the middle mouse button when your cursor is over a track and a popup displays the modes and shortcuts. However, for those users who prefer the addition of the Ctrl key to add multiple selections as is commonly done for listbox operations, there is a preference in Settings → Preferences, Appearance tab, called “Clears before toggle” that changes the behavior.

When an edit is marked as selected, it can be cut/copied into the paste clip buffer. The constructed clip buffer will begin with the leftmost edit and end with the rightmost edit. The edits may contain media, or be silence, or skipped if they are not selected. The clip tracks are copied from the first track with an active edit selection to the last track with an active edit selection. A clip track can be completely empty if no selection was made on the track. The word “packed” means that the silent edits and empty tracks are not included in the copy to the clip buffer, and all of the elements are packed together, no gaps. Packing a clip buffer makes it easier to move “blobs” of data around. Once the edits have moved and have a relative relationship applied, an unpacked clip buffer allows the media to be copied with the relative positions of the edits preserved.

The “edits” popup is activated on a track and a red and yellow colored reticle appears to temporarily mark the location when you click on the middle mouse button. An expanded explanation is provided below.

<u>Key</u>	<u>Operation</u>
Drag	Hold down the Ctrl key while dragging to move a single edit elsewhere.
Left mouse button	Selects and highlights the edit under the cursor with a red selection box. Left mouse button also will toggle that clip selection off if it is on.

<u>Popup Label</u>	<u>Key</u>	<u>Operation</u>
Clear Select	Ctrl-Shift-A	Deselect all selected edits – ones that have the red lines around them.
Copy	Ctrl-c	Copy the selected edits into the copy buffer.
Cut	Ctrl-x	Delete the selected edits after copying them into the buffer. The edits after the selection will move left to occupy the vacated space.
Mute	Ctrl-m	Delete the selected edits after copying them into the buffer. The space previously occupied by the edits will be replaced with silence.
Copy Pack	Ctrl-Shift-C	Copy the selected edits into the buffer and remove any silent sections.
Cut Pack	Ctrl-z	Delete the selected edits after copying them into the buffer. The edits after the selection will move left to occupy the vacated space. The edits in the copy buffer will be packed together within each track.
Mute Pack	Ctrl-Shift-M	Delete the selected edits after copying them into the buffer. The deleted edits will be replaced with silence. The edits in the copy buffer will be packed together within each track.
Paste	Ctrl-v	Paste contents of the copy buffer at the insertion point marked by the red & yellow reticle of the popup menu or the position of the hairline cursor. This is a splice operation which creates space for the edits.

Overwrite	Ctrl-b	Paste contents of the copy buffer at the insertion point marked by the red & yellow reticle of the popup menu or the position of the hairline cursor. This destroys the current edits in that space.
Overwrite Plugins	Ctrl-Shift-P	Pastes plugins that are in the Copy buffer to current location but no clip.

The copy/paste behavior respects the armed/disarmed tracks state. A paste of audio on a video track will fail and vice versa. In addition if you attempt to paste edits consisting of more tracks than what is available at that location it will not allowed.

Attaching transitions to multiple selected edits via the Video or Audio pull-downs is also available. The new transitions are attached at the start of the edits, and will replace any existing attached transitions.

6.5.2 Groups

Cinelerra recognizes as a group, the edits of different armed tracks that have aligned beginnings, regardless of whether they have the same source or aligned ends. You can drag these edits around on the timeline to construct your movie by rearranging scenes. If more than one track is armed, Cinelerra will drag any edits which start on the same positions the edit the mouse pointer is currently over.

Another method of Grouping of edits is performed as follows:

- Select each of the clips you would like to be part of a group.
- Use the desired Copy mode as described above to get into the buffer.
- Go to the Resources window Clip folder and right click on an empty spot.
- Choose the Paste Clip option. Now you have a named clip of the current selection.

A more traditional Grouping of edits which make a Permanent Group works as described next. A temporary group is just a set of selected clips and works the same as a single selection.

- Create a group:
 - Select multiple clips/edits with left mouse button over the clip;
 - Hold down shift and left mouse click over any of the selected clips to create a group;
 - A Group Id and color are assigned to this group making it easy to distinguish.
- Move a group:
 - Move your cursor over any of the clips within the group; click to select the clip group.
 - Click and Hold down the left mouse button to drag.
- Dissolve a permanent group:
 - To ungroup select any of the group edits/clips and shift left mouse click to ungroup (same as creating).
 - The edits will be ungrouped, and the current selection will be replaced with the ungrouped clips.

The color of the created groups are not muted and are assigned by Group Id going sequentially through the number of 1 to 64. Therefore if you dissolve a group and make it again, it will have a different Group Id and a different color.

6.5.3 Dragging Groups

Dragging while in *Drop and Drag editing mode (arrow mode)* is really easy. Just select the clip or clips you want to drag using the left mouse button, then put your cursor over one of them and drag

while holding down the left mouse button. Keyframes, autos, labels, and plugins will also be dragged. Dragging honors armed/disarmed tracks. When you drag there will be 3 possible colors as defined:

- Red color means can not drop here because it will not fit in the space.
- Green color means OK to drop in that position.
- Yellow color means you can drop here and when you do it will be exactly next to that existing edit.

The original (older) method of dragging while in Arrow mode, lets you just left mouse click on a single clip or aligned clips and just drag. This older method of dragging does not move any of its effects with it at this time. There will only be a white outline while dragging and it will let you drop only if it fits. You can also perform some dragging and grouping while in the *Cut and Paste editing mode (ibeam mode)* by taking advantage of the Ctrl button in conjunction with the left mouse button.

- Double click selects a column so you can move, for example, the audio and video together by holding down the Ctrl key and dragging.
- A single clip can be dragged without any of its aligned clips, by holding down and Ctrl key and drag.

This last section on Dragging, outlines the difference between “column selection” and “marking selection”. Column selection is available to make it easy to still be able to do some dragging in I-beam mode whereas Marking selection makes it easy to drag clips together that are not columnated.

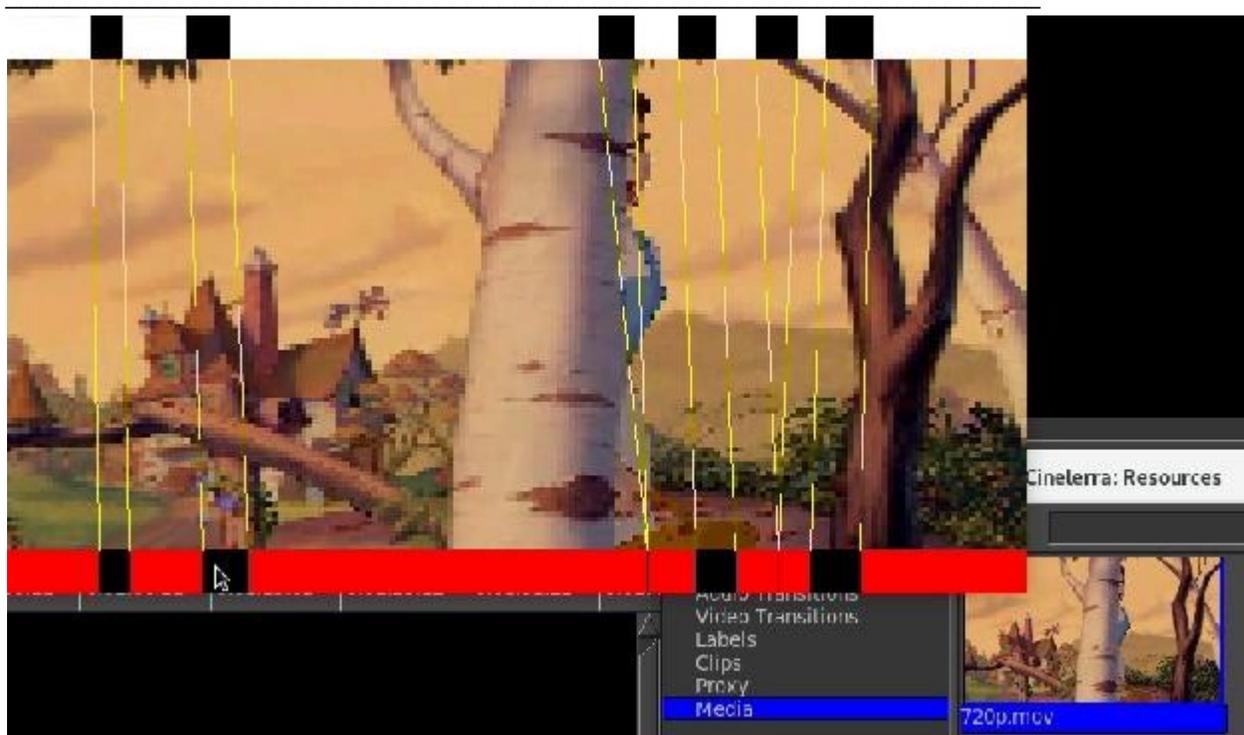
6.6 Inter-View Mode / Identifying Source Targets Editing

Inter-View mode provides a mapping of a particular media file to its timeline usages. It is somewhat similar to Two Screen Editing in that you make use of the Viewer. It makes it possible to precisely trace and indicate in the media the origin of a particular segment of the timeline and visually indicate the use and distribution that the timeline makes of a particular media. A good example usage would be in the case of a 30 minute interview where you use a few short pieces to make a shorter 10 minute section, find out that you have made the timeline 12 minutes instead and need to cut out another 2 minutes.

This feature provides the following capabilities.

- You can see on the timeline all of the places where a particular piece of media was used.
- And you can see which parts of that particular media are already used so you do not reuse that same piece again.

Screencast below shows an example of the Inter-View mode mapping preview mini-window.



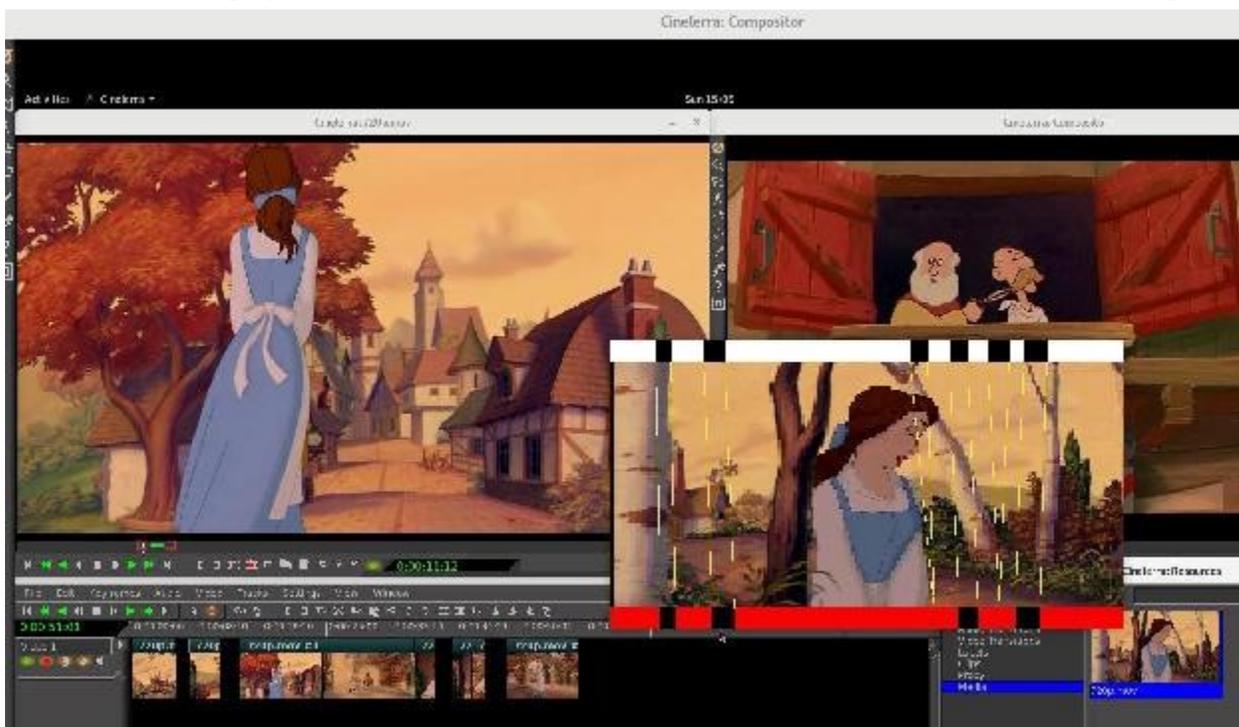
Explanation of how to use Inter-View mode will be described here next.

- Do your editing as usual on the timeline until you are ready to see what is used or unused.
- Make sure you are in Preview mode in the Resources window; you enable this mode using the right facing green colored arrow to the left of the word Search. When enabled the arrow is surrounded by a gold color that looks like this --> .
- Middle mouse click on a thumbnail in the Resources window and a popup occurs of that media with a white colored bar at the top and a red colored bar at the bottom with black sections.

The red/white bars represent the presence and the black sections represent the absence of where that media is used on the timeline. To get to a bigger representation, use the “f” key for a full screen. Now you can operate the following buttons to display what you need to see and to move around. It is important to note that “locked tracks” will not be represented. This makes it easy to ignore the audio track segments if you want so there is less confusion in the display.

- Clicking on the top white or black spaces in the top time bar loads the Viewer with the source media, and sets the In/Out [/] pointers to be the selection of that edit.
- Click on a location in the bottom red or black bar, and the main cursor and composer will re-position to the corresponding location on the session timeline.
- Dragging on the red/black bar will correspondingly update the position in the timeline and composer.
- Ctrl-click on the bottom bar and the timeline and composer are re-positioned to the beginning of that edit.
- Shift-click on the bottom bar and a "selection" is made of that section in the timeline and the composer is updated with that start position.

Screencast below displays Inter-View window and its relation to the timeline, viewer, and compositor.



The Inter-View mode works for Media, Proxy, and User Bins. When the preview window has only black bars on the top and bottom, it means that this particular media is not loaded in the timeline. So when you are in Proxy, meaning that the Proxy files are loaded on the timeline, there will be only black bars for the corresponding Media file UNLESS there is an audio track associated with the video. Because audio tracks are not proxied, they will show for Media but not for Proxy.

6.7 Some Specific Editing Tools

This section covers some more detailed editing tools and scenarios for edit management.

6.7.1 Edit Length

To set the length of an edit in the timeline, select the region which contains the edit to be modified. Now select the menu bar Edit → Edit Length... menu item to activate the “edit length” popup. The duration of the edit can be reset by entering the desired edit length in seconds. Pressing OK will change all of the selected edits (in armed tracks) to the specified length.

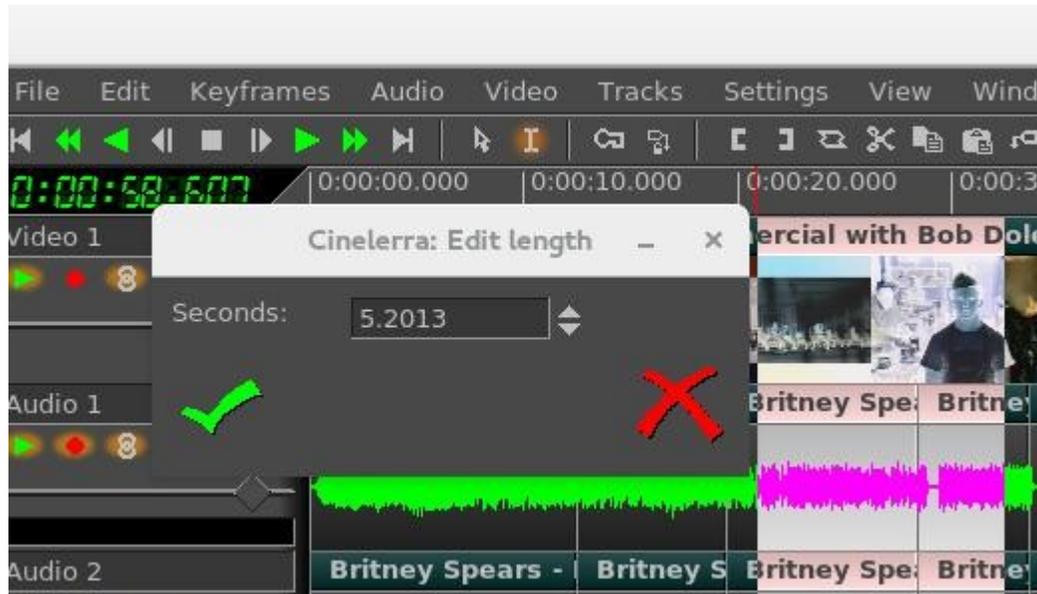


Figure 24: Edit Length window

6.7.2 Align Edits

When loading media, a common problem is that the various audio/video tracks do not always have exactly the same lengths. For example, you might load audio/video recordings from your camera and be dismayed to see that the audio for each segment is a half second longer than the video. If you load a large set of media clips by concatenation, the audio and video will be more skewed as more media is loaded. Align Edits makes it possible to adjust the edits so the audio and/or video align by adjusting the edits so that the track lengths are consistent. To use this feature, load all of the desired media and select a region which contains all of the edits to be aligned in the timeline. Now select the menu bar Edit → Align Edits menu item to operate the change. The topmost armed track is used as a template reference, and the rest of the tracks are either cut or padded to align the edit boundaries. Besides aligning audio with the video, you can also align video with the audio if the first armed track is audio.

The code performs the following algorithm:

- Use the first armed track as the master track (it must contain data).

- Collect the “edit project start times” on the selected master track.

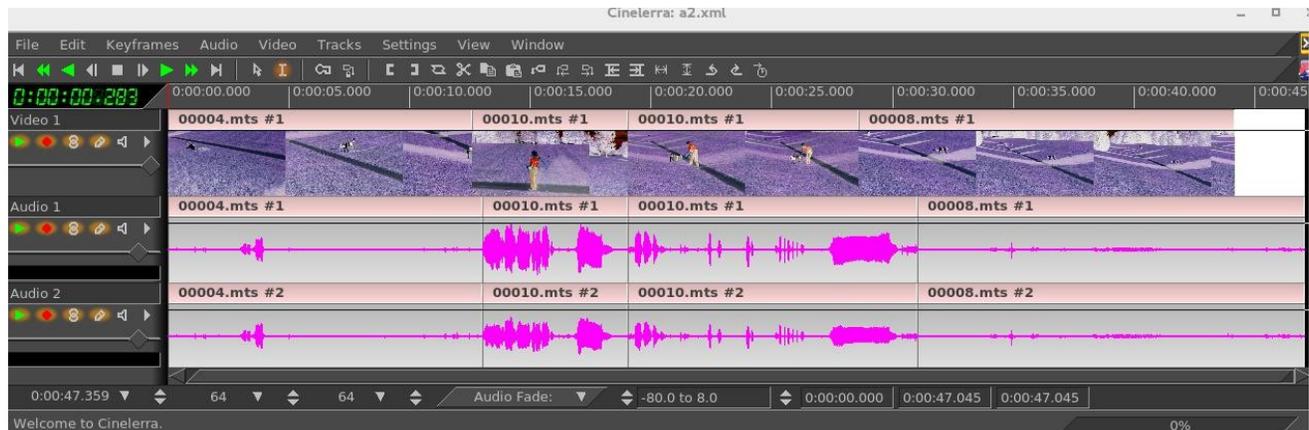
 - Only edits that are 100% inside the selected area will be used.

- Set all other tracks to match the “edit times” of the template track, either by putting in silence or cutting the region to align the edits on the “edit times” of the master track.

The start time sequence of media and silence edits along the master track are collected as the target alignment boundaries. All armed tracks after the master track are modified so that if the next edit edge

is too soon, it adds silence; if it is too late, edits are shortened or deleted past the point of the next target alignment boundary time. Align Edits works best if there are an equal number of Video and Audio sections. Also, it is better to use cuts instead of adding silence – if there are silence edits together, the algorithm will combine the silence edits into a single edit and results may not be as desired.

Screenshots below show the Before, the Highlighted Edits to be manipulated, and the After results for the Align Edits.

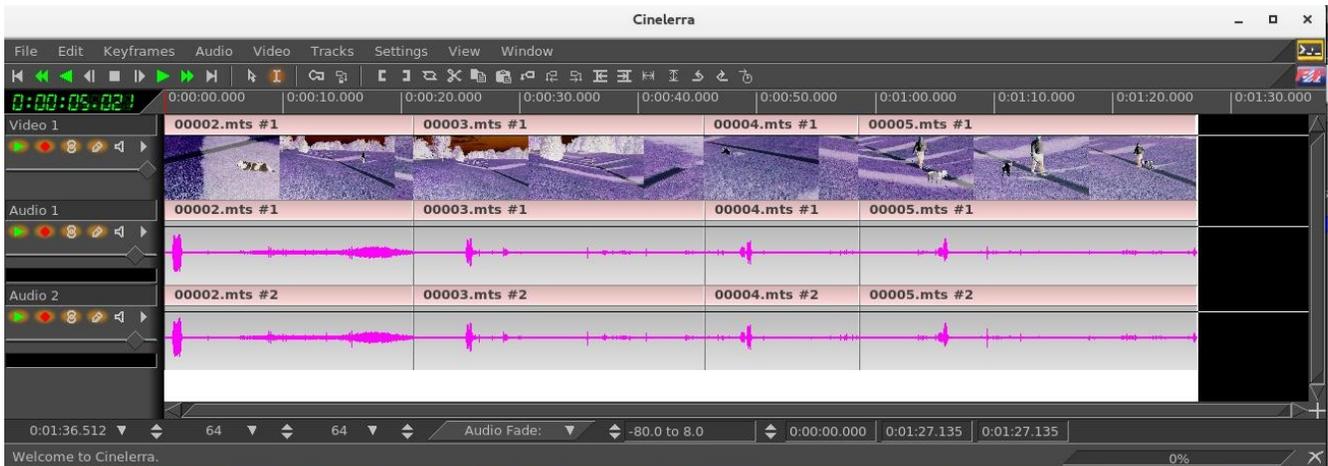


This After screenshot “adds silence” in the second section as noted in red letters.

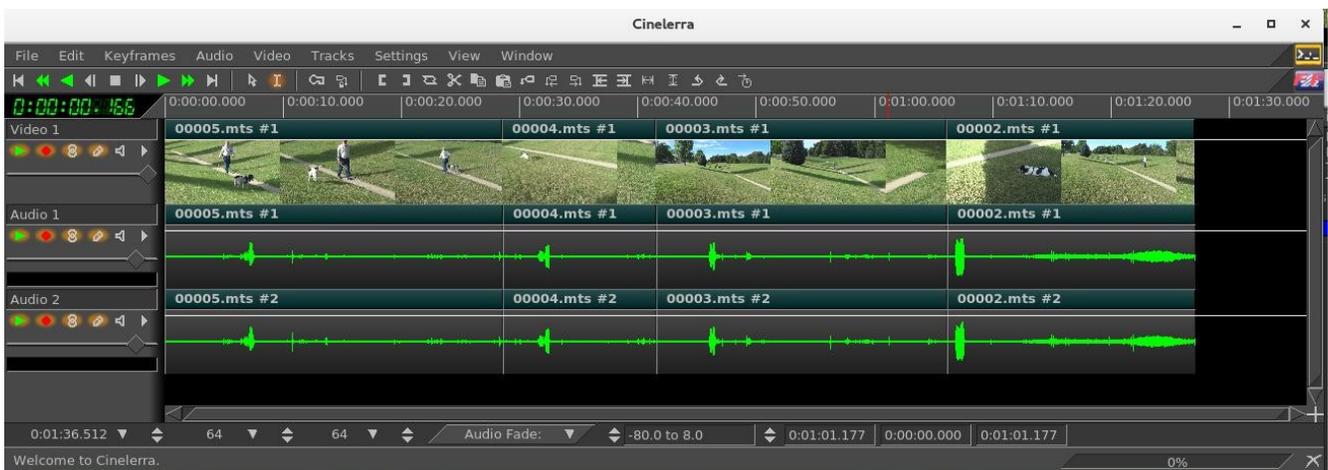
6.7.3 Reverse Edits

The Reverse Edits can be useful to change the order of 2 edits in the case where you would like to put a “teaser” section that occurred in the middle of a movie at the beginning instead, that is, reversed positions. To operate, highlight completely the edit areas you would like reversed and then use the pulldown Edit → Reverse Edits.

Screenshot immediately below shows the selected/ highlighted area to which Edits will be applied. Note the first edit is 00002, followed by 00003, 00004, and 00005 in that order.



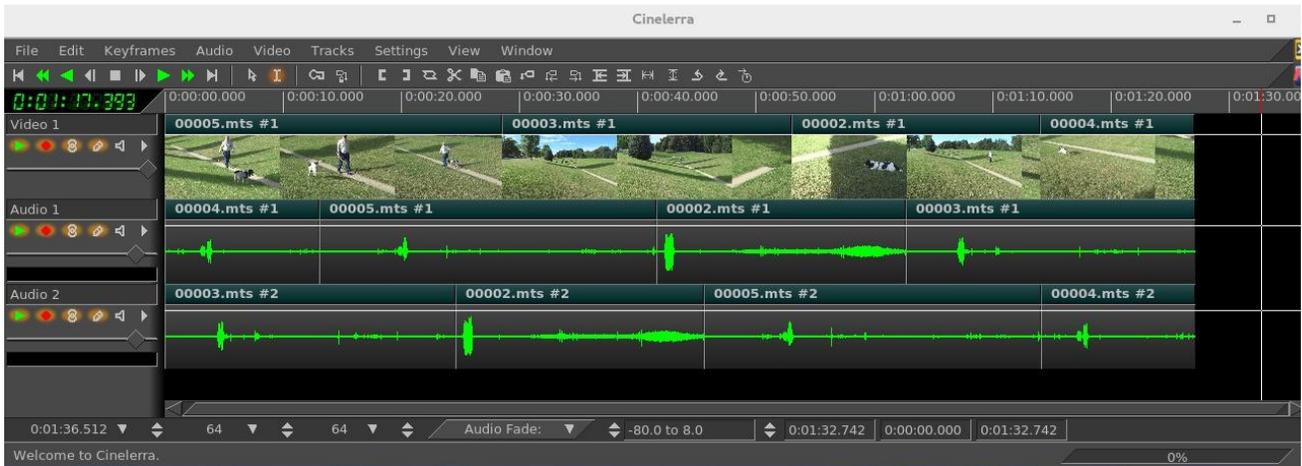
Screenshot illustrating the results of executing “Reverse Edits”. Now you will see the reversed order of 00005, 00004, 00003, and last 00002.



6.7.4 Shuffle Edits

The file pulldown Edit → Shuffle Edits will randomly exchange the location of the edits. This feature can be used to change the order of the music like you would do from your MP4 player where you have a playlist of your favorite music. Or perhaps you are creating an advertisement background, you can randomly change it, thus the viewer sees a different order of scenes each time shown.

Screenshot illustrating Shuffle Edits of the highlighted area of the first screenshot on the page. Note the permutation of the fragments resulting in 00003 now being first, then 00005, 00002, and 00004 last.



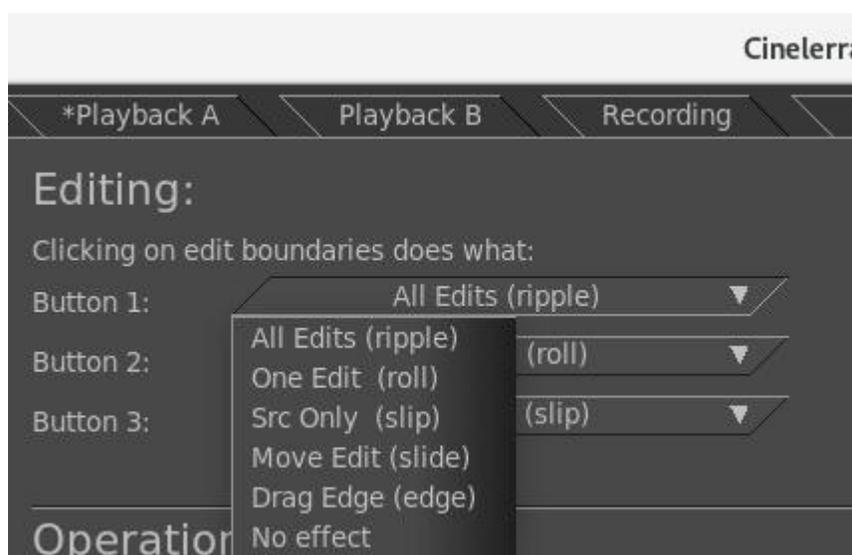
6.7.5 Drag Handle Management / Trimming

With some edits on the timeline it is possible to do trimming. By trimming you shrink or grow the edit boundaries by dragging them. In drag and drop mode or cut and paste mode, move the cursor over an edit boundary until it changes shape. The drag handle shows as a left or right facing fat arrow when you cursor near the clip start or end. If the cursor faces left, the dragging operation affects the beginning of the edit. If the cursor faces right, the dragging operation affects the end of the edit.

The effect of each drag operation not only depends on the behavior button but whether the beginning or end of the edit is being dragged. When you release the mouse button, the trimming operation is performed.

For all file formats, other than still images, the extent of the trimming operation is limited to the source file length. Attempting to drag the start of the edit beyond the start of the source, limits it to the source start. In all trimming operations, all edits which start on the same position as the cursor when the drag operation begins are affected. You have to disarm tracks in order to prevent edits from being affected.

You have 6 different choices of which mouse button to use for specific types of editing while using the drag handle. You change the drag handle mouse effects by using the Settings → Preferences, Interface tab and modifying the Editing section as shown in the next figure. The drag handle affects not only the clip you are working on but also frequently the entire duration of all clips on the timeline. A description of the fundamental/common terminology for choices follows.



Default choices for mouse:

Ripple for button 1

Roll for button 2

Slip for button 3

Screencast above for Settings → Preferences, Interface Button choices

All Edits (ripple) – shorten or lengthen the start or end of a single piece of media while moving all media to the right of that clip up or down on the timeline correspondingly. Timeline duration is modified. In a drag *All Edits* operation, the beginning of the edit either cuts data from the edit if you move it forward or pastes new data from before the edit if you move it backward. The end of the edit pastes data into the edit if you move it forward or cuts data from the end of the edit if you move it backward. All the following edits shift. If you drag the end of the edit past the start of the edit, the edit is deleted.

One Edit (roll) – shorten or lengthen the start or end of a clip which changes the clip to the left and the clip to the right to take up or down the change depending on movement. Timeline duration is not modified. In a drag *One Edit* operation, the behavior is the same when you drag the beginning or end of an edit. The only difference is none of the other edits in the track shift. Instead, anything adjacent to the current edit expands or shrinks to fill gaps left by the drag operation.

Src Only (slip) – move the in and out point of a single clip without changing the timeline duration. In a drag *Src Only* operation, nothing is cut or pasted. If you move the beginning or end of the edit forward, the source reference in the edit shifts forward. If you move the beginning or end of the edit backward, the source reference shifts backward. The edit remains in the same spot in the timeline but the source shifts.

Slide – a single clip is moved but retains its current in and out point; however the out point of the clip to the left changes and the in point of the clip to the right also changes. Timeline duration remains the same.

Edge Left / Right – moves the edge of the clips.

No effect – no changes are made. You might want to use this choice to prevent accidental movements.

The next table displays the options and results with the Key Table here first.

Key Table:

s = src media start
 p = proj position
 l = length
 c = cut distance
 rest == p+=c: for rest of clips
 01 = flags edits_moved, rest_moved

Cinelerra-GG:

	drag left	drag right
curr s+=c,l-=c; +rest	<	> ... rest
+-----+ all left	+-----+	+-----+
+abc 12345 xyz+ ----->	+abc 012345 xyz+	+abc 2345 xyz+
+-----+ 11	+-----+	+-----+
curr l+=c; +rest	<	> ... rest
+-----+ ripple right	+-----+	+-----+
+abc 12345 xyz+ ----->	+abc 1234 xyz+	+abc 123456 xyz+
+-----+ 01	+-----+	+-----+
prev l+=c; curr ps+=c,l-=c	<	>
+-----+ roll left	+-----+	+-----+
+abc 12345 xyz+ ----->	+ab 012345 xyz+	+abcd 2345 xyz+
+-----+ 00	+-----+	+-----+
curr l+=c; next ps+=c,l-=c	<	>
+-----+ roll right	+-----+	+-----+
+abc 12345 xyz+ ----->	+abc 1234 wxyz+	+abc 123456 yz+
+-----+ 00	+-----+	+-----+

```

s-=c
+-----+ slip left +-----+ +-----+
+abc|12345|xyz+ -----> +abc|23456|xyz+ +abc|01234|xyz+
+-----+ 10 +-----+ +-----+
s-=c
+-----+ slip right +-----+ +-----+
+abc|12345|xyz+ -----> +abc|23456|xyz+ +abc|01234|xyz+
+-----+ 10 +-----+ +-----+

```

```

prev l+=c; curr p+=c; next ps+=c, l-=c < >
+-----+ slide left +-----+ +-----+
+abc|12345|xyz+ -----> +ab|12345|wxyz+ +abcd|12345|yz+
+-----+ 10 +-----+ +-----+
prev l+=c; curr p+=c; next ps+=c, l-=c < >
+-----+ slide right +-----+ +-----+
+abc|12345|xyz+ -----> +ab|12345|wxyz+ +abcd|12345|yz+
+-----+ 10 +-----+ +-----+

```

```

curr s-=c, l+=c; +rest < > ... rest
+-----+ edge left +-----+ +-----+
+abc|12345|xyz+ -----> +abc|2345|xyz+ +abc|012345|xyz+
+-----+ 11 +-----+ +-----+
curr l+=c; +rest < > ... rest
+-----+ edge right +-----+ +-----+
+abc|12345|xyz+ -----> +abc|1234|xyz+ +abc|123456|xyz+
+-----+ 01 +-----+ +-----+

```

ver 2

Cinelerra=HV/CV:

```

s = src media start
p = proj position
l = length
c = cut distance
rest == p+=c: for rest of clips
01 = flags edits_moved, rest_moved

```

```

curr s+=c, l-=c; next p-=c; +rest drag left drag right
+-----+ all left +-----+ +-----+ < > ... rest
+abc|12345|xyz+ -----> +abc|012345|xyz++abc|2345|xyz+
+-----+ 11 +-----+ +-----+
curr l+=c; +rest < > ... rest
+-----+ all right +-----+ +-----+
+abc|12345|xyz+ -----> +abc|1234|xyz+ +abc|123456|xyz+
+-----+ 01 +-----+ +-----+

```

```

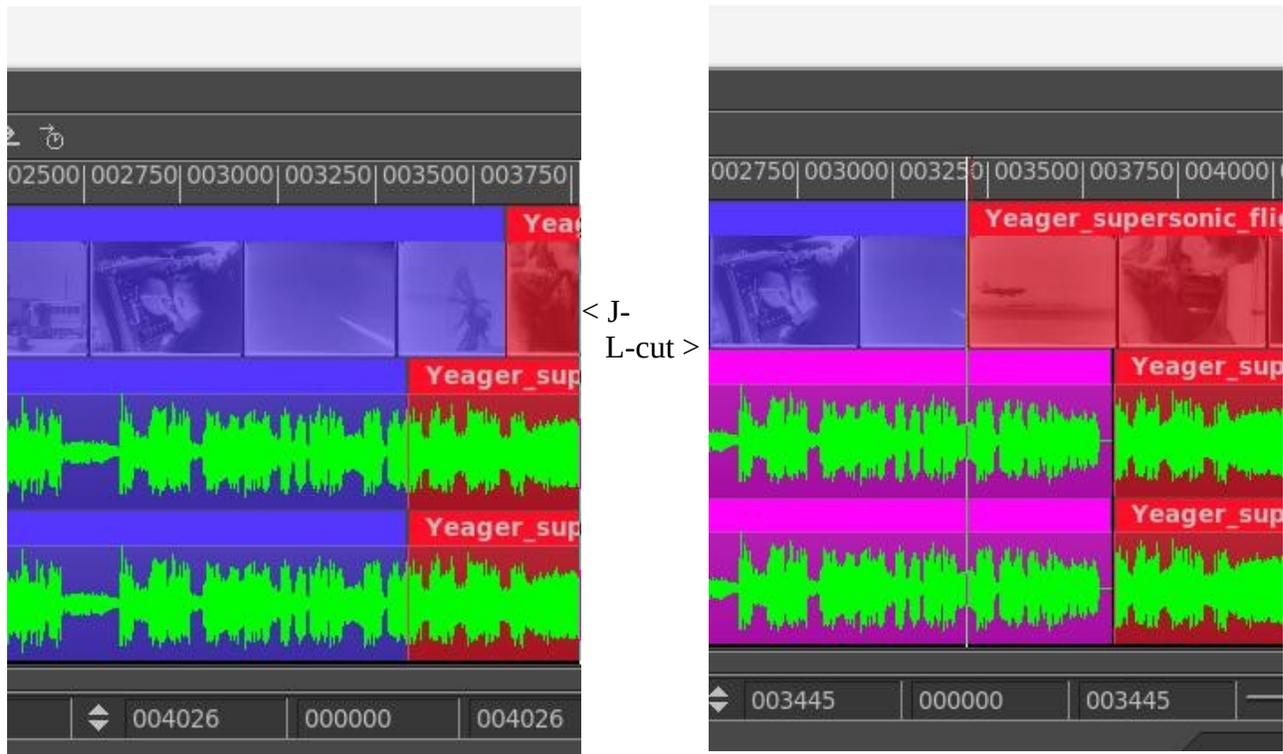
prev l+=c; curr ps+=c,l-=c          <          >
+-----+      one left  +-----+ +-----+
+abc|12345|xyz+ -----> +ab|012345|xyz+ +abcd|2345|xyz+
+-----+          00    +-----+ +-----+
curr l+=c; next ps+=c,l-=c          <          >
+-----+      one right +-----+ +-----+
+abc|12345|xyz+ -----> +abc|1234|wxyz+ +abc|123456|yz+
+-----+          00    +-----+ +-----+

s+=c          <          >
+-----+      src left  +-----+ +-----+
+abc|12345|xyz+ -----> +abc|01234|xyz+ +abc|23456|xyz+
+-----+          10    +-----+ +-----+
s+=c          <          >
+-----+      src right +-----+ +-----+
+abc|12345|xyz+ -----> +abc|01234|xyz+ +abc|23456|xyz+
+-----+          10    +-----+ +-----+

```

How to do a J-cut or L-cut

A J-cut is a split edit film editing technique in which the audio from a following scene overlaps the picture from the preceding scene, so that the audio portion of the later scene starts playing before its picture as a lead-in to the visual cut. An L-cut is a different split edit film editing technique in which the audio from preceding scene overlaps the picture from the following scene, so that the audio cuts after the picture, and continues playing over the beginning of the next scene. To do either a J-cut or an L-cut, you first shorten the first or second video a little. Then you block the audio tracks from changing by disarming the appropriate tracks. Finally use *One Edit (roll)* the cutting edge off the videos. Moving to the right creates a J-cut and moving to the left creates an L-cut.



6.7.6 Split View in Compositor Using the Drag Handle with Trim

The Trim Feature using the drag handle provides some good ways to view your video while editing. The playback position in the compositor is updated live and the view in the compositor can be split so that in the left half of the compositor you can see the last frame of the left clip and in the right half the first frame of the right clip. Dragging edits can not be extended past the beginning or the end.

First familiarize yourself with button operation; check your setup by executing the following step. In the Settings-> Preferences, Interface tab, Editing section, clicking on the edit boundaries can be set for Button 1, 2, 3 as one of the following:

All Edits / Ripple, One Edit / Roll, Src Only / Slip, Move Edit / Slide, Drage Edge / Edit, No effect.

Now to use this feature, create a track with edits that have trims on the left and/or the right. The edit boundary can be modified using "drag handles" at the boundary between the edits.

Left Mouse Button (LMB) usage:

- If you grab the edit handle from the right side, you will see a left arrow and dragging the boundary will modify the right edit playback starting time
- If you grab the edit handle from the left side, you will see a right arrow and dragging the boundary will modify the left edit playback ending time.
- In both cases, the compositor will show the edit endpoint of the changed edit.

Shift LMB usage:

- The effect on the edits is the same as described above, but the compositor will show a split screen of the left and right edits as they appear at the drag handle boundary. Dragging will only change one of the two images, since only one edit is being changed.

Middle Mouse Button (MMB) usage:

- Both the left and the right edit ending/starting times are updated. The image shown in the compositor will be drawn from the side of the drag grab, that is the left if it is grabbed from the left, and the right if it is grabbed from the right.

Shift MMB usage:

- The effect on the edits is the same as described above, but the compositor will show a split screen of the left and right edits as they appear at the drag handle boundary. Dragging will change both of the two images, since both edits are being changed.

Right Mouse Button (RMB) usage:

- the start/end point of the current edit is moved, but the edit length is unchanged only one image changes since only one edit endpoint is view is updated.

Shift RMB usage:

- The effect on the edits is the same as described above, but the compositor will show a split screen of the left and right edits as they appear at the drag handle boundary. Dragging will only change one of the two images, since only one edit is being changed.

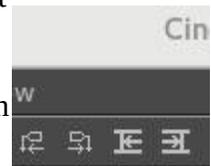


Figure 25: Split compositor screen showing the result of the Trim feature

6.7.7 Snapping while Cutting and Dragging

Cutting/Snapping edits - cuts from an edit handle to the insert point.

There are Edit Panel buttons which normally are used to move to the previous or next edit handle/label. They look like tags and the letter E on the menu bar and are oriented forward/backward. These same buttons can be used to “cut” from the insert pointer to the previous or next edit/label when the ctrl+alt keys are both pressed when the buttons are used. They “snap” off the media instead of doing the standard re-positioning. This is useful to minimize the number of operations necessary to cut between edits/labels.



Instead of using the edit panel buttons, you can more easily use the following keyboard shortcuts to perform the same functions:

- | | |
|------------------|--|
| snap_right_edit | ctrl+alt+ '.' |
| snap_left_edit | ctrl+alt+ ',' |
| snap_right_label | ctrl+alt +shift '.' (shift+period is the > sign on US keyboards) |
| snap_left_label | ctrl+alt +shift ',' (shift+comma is the < sign on US keyboards) |

Drag Snapping - if you hold down the Ctrl + Alt keys while dragging using the mouse, once the clip gets near to an edit, a label, an in/out pointer or the start/end of the timeline, the dragged clip will snap next to that marker. The 2 will now be exactly aligned with no gap and no overlap. As you drag the clip close to one of the markers, when they are within a short distance they start to stick and stay that way until you move further away from that distance. Also, the line will turn color from green to yellow while in the sticky phase.

6.7.8 Nesting

Nesting is a good way to see what media would look like without having to render it first since that is done automatically. You can nest clips within nested clips at any level. It is **IMPORTANT** to note that nesting will use certain additional resources and is meant to be used while working to test the results.

Nested Assets

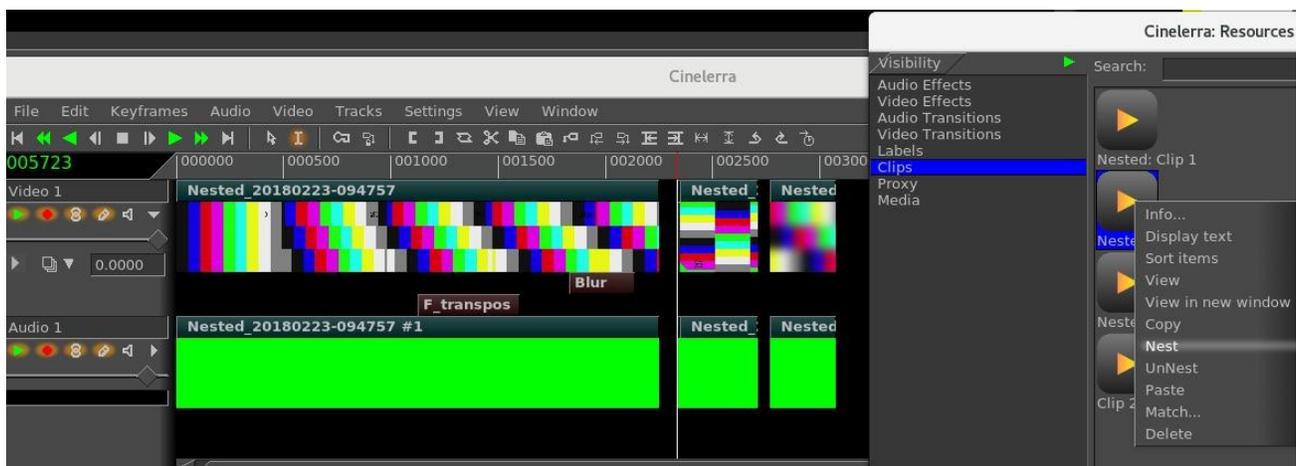
A nested asset is an EDL session that embeds an existing EDL session, all tracks, all plugins, editing, and effects into a media object that appears as one audio/video media object, no plugins, editing, or effects. It is as if the existing EDL was rendered, and loaded in its place. This has several interesting side effects. First, you don't have to render the entire media file to see any portion. Second, it requires no rendering compute time or storage. Third, it changes the precedence of the composer so that you get more control over the projection and automation, so that the results can be sent into another rendering step, not simply part of the current stack. It groups the plugin stack in much the same way that an arithmetic expression is grouped by parenthesis.

The EDL session and the rendered output are visually equivalent. Nested assets allow for complex grouping and stacking of effects, and makes media access much more flexible. This feature can be used recursively, that is, any number of sessions may be stacked and referenced as an asset, as long as all of the rendering resources are available. Nested assets are added to the timeline by using the pulldown File → Load files... on the main menu and selecting the "Insertion strategy" of "Nest asset". The file will be pasted into the timeline over the current selection or at the insertion point.

It is somewhat important to note that nested assets and nested clips will have index files automatically created. These index files can start to clutter up your \$HOME/.bcast5 directory with files named Nested_###.idx and you may want to periodically delete any index files which are no longer in use.

Nested Clips

It is also possible to create "clips" and convert them to "nested edl". This is done by first creating a clip using the standard cut, clipboard, paste, and/or edit panel buttons. Now, using the resources "clip" folder, select a clip to be nested, and use the right mouse button to select a clip. This activates the clip popup menu. Select the "Nest" menu item, and the clip will be converted to a "Nested: Clip". Conversely, you can select a "Nested: Clip", use the "UnNest" menu item, and the clip will be reverted to a "Clip". This works similarly to the group / un-group editing features of many graphic design editing programs, but in this case the groups are rendered compositions. They can not be proxied.



6.7.9 Copy/Paste clips/medias across Multiple Instances

It is easy to copy/paste clips/media within a single instance of cinelerra or across multiple instances. The reason this works is because there are hidden X cut buffers and these are used to transmit EDL from 1 instance to another.

Steps to copy from a source timeline and paste to a target timeline:

- 1) highlight a selection on the timeline in 1 instance of Cinelerra
- 2) use the Copy icon (shortcut c) on the main menu bar to copy into a buffer
- 3) move the pointer to another instance of Cinelerra and set an insertion point in its timeline
- 4) use the Paste icon (shortcut v) to paste the clip to that other instance selection target

6.8 ShuttlePROv2 and ShuttleXpress Jog Wheels for Editing

(credit Eric Messick - FixedImagePhoto.com/contact)

The ShuttlePROv2 and ShuttleXpress are affordable jog wheels which can be useful for working with Cin, especially if you do a lot of playing forward/backward, fast/slow/normal, and single frames. Directions for using the ShuttlePROv2 and the ShuttleXpress with Cinelerra are described next. These devices work by sending keystrokes used in Cin, corresponding to the shuttle action, to the keyboard buffer. The shuttle has been fully integrated into the Cinelerra code so that after the one initial setup, no further intervention is required. The multi-speed outer wheel works the same and has the same number of S positions on both shuttles but the shuttle Xpress has only 5 keys. Since the majority of user operations will most likely be with the use of the 2 wheels, the slightly smaller Xpress could be a better choice with its 5 easy to reach keys. The Pro is approximately 4x7 inches whereas the Xpress is about 4x4 inches.



Figure 26: ShuttlePROv2



Figure 27: ShuttleXpress

The vendor supplied “string” device names for the shuttles are currently:

```
/dev/input/by-id/usb-Contour_Design_ShuttlePRO_v2-event-if00  
/dev/input/by-id/usb-Contour_Design_ShuttleXpress-event-if00  
/dev/input/by-id/usb-Contour_Design_ShuttlePro-event-if00
```

Only 1 necessary initial setup is required due to permission settings for non-root usage. As root, just copy a file that provides the necessary permissions to use the shuttle, then reboot, Example copy:

```
sudo cp {cindat_path}/doc/99-ShuttlePRO.rules /etc/udev/rules.d/.
```

then the next time after you reboot, the permissions should be correct. This file only needs to contain one of the following lines depending on which shuttle version you have/use, but all will be in the file.

```
ATTRS{name}=="Contour Design ShuttlePro" MODE="0644" # for newer PRO model  
ATTRS{name}=="Contour Design ShuttlePRO v2" MODE="0644" # for older PRO model  
ATTRS{name}=="Contour Design ShuttleXpress" MODE="0644" # for the Xpress model  
SUBSYSTEMS=="usb", ATTRS{idVendor}=="0b33", ATTRS{idProduct}=="0020", MODE="0666"  
SUBSYSTEMS=="usb", ATTRS{idVendor}=="0b33", ATTRS{idProduct}=="0030", MODE="0666"
```

If you swap your shuttle, for example upgrade from an Xpress to a PROv2, just stop Cin, unplug the original shuttle, plug in the replacement shuttle, and restart Cin. If you start the cinelerra program and

the shuttle does not function as before, stop cinelerra and then simply unplug it and plug it in again. There are a couple of reasons why it may stop functioning. One is because cinelerra was not stopped with the usual Quit command and the shuttle was improperly shut down when there was a crash. The other possibility is that a static discharge occurred in the area.

A default shuttlerc file is automatically used when a shuttle device is plugged in when Cin is started. This file sets up the key bindings for Cinelerra to use. You can override any default settings by having a local file in your \$HOME directory, named .shuttlerc to reflect your personal preferences.

How to Modify the Default Key Settings

Detailed information on how to modify your local .shuttlerc file is described next, but if you need help you can request more information in the forum at <https://cinelerra-gg.org> . In the shuttlerc file, a # always represents a comment and blank lines are ignored. The first thing you must do is copy the system supplied shuttlerc file to your \$HOME directory and rename it as .shuttlerc (with a period).

The shuttlerc file has sections that in the case of Cinelerra, represent different windows allowing you to set the keys, K1-K15 for the Pro and K5-K9 for the Xpress, the shuttle wheel positions of S0/S1/S-1 for stop, S2 through S7 for wheeling to the right, and S-7 through S-2 for wheeling to the left for reverse. Then there is JR to jog right (clockwise) and JL to jog left (counter-clockwise) for the inner smaller wheel for single frame movement. See the key arrangement on a later page for location of the keys for each of the two different shuttles.

The sections are surrounded by brackets for windows such as Cinelerra (the main window), Viewer, Composer, Resources, Load, and Default. If you want the keys to be defined the same in every window, you can bracket each window on lines one right after the other and then just define one set of keys. The other lines will have the key name/shuttle position followed by its assigned value. The values you use for the keys are usually shortcuts and have to be operationally defined within Cinelerra. For example, the shortcut “f” to go fullscreen is defined so can be used; however the shortcut “h” is not defined so will not do anything. You can check the file, shortcuts.html, for some options to use.

Next are a few actual examples from the default {cindat_path}/shuttlerc file.

The next brackets represent sections. Default, Resources, Load windows all use the same key values.

[Default]

[Resources]

[Load]

K5 XK_Home

K6 XK_Button_1 # same as mouse button 1

K7 XK_Button_2 # same operation as mouse button 2

K8 XK_Button_3

K9 XK_End

JL XK_Scroll_Up # for example, in the Load menu, use scroll up to get to the next file name

JR XK_Scroll_Down

Cinelerra with brackets around it next, is the section with some key definitions for the main window.

[Cinelerra]

```

# Most useful functions have to be on K5-K9 because Xpress only has 5 keys
K5 XK_Home      # Beginning
K6 XK_KP_6     # Reverse, or if playing Stop
K7 XK_KP_0     # Stop
K8 XK_KP_3     # Play, or if playing Stop
K9 XK_End      # End
...
S-7 REV_16     # Next 6 are reverse keys
S-6 REV_8      # the number on the end represents speed
S-5 REV_4      # number can be decimal up to 64
S-4 REV_2      # 2 means 2x or double speed
S-3 REV_1
S-2 REV_0.5    # 0.5 represents 1/2 speed
S-1 XK_KP_0    # Because the Shuttle does not generate S0, have to use S-1
S0 XK_KP_0     # Hardware does not generate S0
S1 XK_KP_0     # Because the Shuttle does not generate S0, have to use S1
S2 FWD_0.5
S3 FWD_1
...

```

An explanation for the above REV and FWD key symbol values is necessary to facilitate user preferences. Obviously REV stands for reverse and FWD for forward. You can set any speed up to and including 64x (that is, 64 times the normal speed) on any of the S keys. First in the line is the key name such as S-3 and then the key direction of FWD or REV followed by the symbol for underscore (_) and then the numerical value to use. For example, if you want the 5th forward position, S5, to play 10 ½ times faster, you would use the statement “S5 FWD_10.5”. Integer or decimal numbers are legal.

For the **Viewer**, you may want keys defined to do a Splice or an Overwrite so define differently. Note that assignments that contain single character letters must be enclosed in quotes.

```

[Viewer]
K2 "v"      # Splice - Viewer only; may be defined differently than Composer or Cinelerra
K4 "b"      # Overwrite

```

To change any key value to an alternative value, just edit the file and make the changes. Besides just keys and alphabetic letters of numbers, you can also use any cinelerra value that contains the combination with Shift, Alt, and Ctrl. For keys that are not printable characters, you can look up the symbol name to use for a specific operation in the file called: /usr/include/X11/keysymdef.h .

Some examples:

```

K10 Alt-XK_Left      # Go to previous edit
K13 Ctrl-XK_Right    # Go to next label

```

For sequences of one or more *printable* characters, you can just enclose them in double quotes. For example in the [Composer] section, to go into or out of fullscreen mode, automatically start playing and put a label there, you could define a key like this: K7 “f l” - that is printable character f, a space, and printable character l.

After modifying .shuttle.rc, the next time you use the shuttle, your changes will automatically take affect without even having to stop and restart Cin. However, the first thing to try if problems is to stop cinelerra, unplug the shuttle, wait a few seconds, plug it in again, and then restart cin. If for some

reason, the shuttle keys still do not work after that, you may have an incorrect setup and you will have to correct that first. For example, if you define S5 twice within the Cinelerra setup, it will fail. It is suggested that if you make changes, you should initially uncomment DEBUG in the .shuttlerc file and start up cinelerra from a terminal window so that you can make sure it is working and has no output errors. An error might look like:

```
dupl key name: [Cinelerra]K1
shuttle config err file: /root/.shuttlerc, line:37
```

Keep in mind when changing the values, that the ShuttleXpress has fewer buttons so if you define K1 it will only work for the ShuttlePro.

Any time you are having trouble with your shuttle, you can copy the default shuttlerc file from {cindat_path}/shuttlerc to your local .shuttlerc file, and edit that to switch to DEBUG mode by removing the # comment from the DEBUG line. But you will have to have started Cin from a terminal window to see the key values. The first time you use the shuttle or after you change the file, the current assignments will show in the terminal window so will look something like:

```
[Cinelerra] # 1
K5[D]: XK_KP_0/U
K5[U]: XK_KP_0/U
```

When you are in DEBUG mode and are just working away, what you will see is something like this:

```
key: 0058 1
key: 0055 0
```

Or:

```
shuttle: 00 00 00 00 00
key: XK_Home 0
```

When you change the focus from one window to another, you will see something like this:

```
new focus: 04c00137
new translation: Viewer
key: 0059 1
```

You can also set an environment variable to temporarily use an alternative shuttle configuration file for testing as in: `export SHUTTLE_CONFIG_FILE=/tmp/shuttlerc_test`

The shuttle wheel occasionally will not “stop” after you have wheeled it to play forward. This is a documented known problem from the original code so you just have to joggle it a little in the other direction and then it will stop. S0 does not always generate a signal to do a stop and that is why S1 and S-1 have to be used to relay the stop instead. Also, if you have a fullscreen Composer or Viewer up and the regular one also, the fullscreen takes precedence.

Troubleshooting auxilliary information:

In order to see if you hardware was recognized by the operating system, keyin:

```
lsusb -v -d 0b33:0030      # for the Shuttle Pro or PROv2
lsusb -v -d 0b33:0020      # for the Shuttle Xpress
```

Note 1: Currently, the keys K14 and K15 do not function on the “Contour Design ShuttlePro” but do on the “ Contour Design ShuttlePRO v2” due to a Report Descriptor error. You can workaround this by uncommenting USB_DIRECT in your local .shuttle file. This directly uses libusb rather than the generic Linux hid driver. USB_DIRECT works for any of the currently tested shuttles.

Note 2: If you are not sure if your shuttle is fully functional, you can verify that the hardware device has been seen by your operating system with this procedure.

1) From a terminal window as an ordinary user keyin: lsusb (the first character is a lower case L – for list). You will see something like the following depending on which usb device you have the ShuttlePro plugged into:

```
Bus 003 Device 002: ID 0b33:0030 Contour Design, Inc. ShuttlePro v2
```

2) To make sure you have usbmon installed keyin: sudo modprobe usbmon

3) Next keyin the following: sudo od -tx1 /dev/usbmon3
where the last 3 is the same # as the Bus in above. If it lists Bus 002, then use /dev/usbmon2 instead.

4) Now with focus in that same terminal window, press any shuttle key just to see what happens and should see about 12 lines similar to these below – a new set every time you press a single key or the wheel. The lines are usually not important, just the fact that you get a response is. However if you have multiple devices on the same bus, you will get responses from any and all of them. Attempt to isolate your shuttle by temporarily unplugging unnecessary devices on the same bus or plug the shuttle into a different usb port that has fewer devices.

```
00000000 80 70 99 75 53 8c ff ff 43 01 81 02 03 00 2d 00
00000020 4e 61 5c 5c 00 00 00 00 8d 2c 06 00 00 00 00
00000040 05 00 00 00 05 00 00 00 00 00 00 00 00 00 00
00000060 01 ff 00 00 00 80 70 99 75 53 8c ff ff 53 01 81
00000100 02 03 00 2d 3c 4e 61 5c 5c 00 00 00 00 b1 2c 06
00000120 00 8d ff ff ff 05 00 00 00 00 00 00 00 00 00
00000140 00 00 00 00 00 80 70 99 75 53 8c ff ff 43 01 81
00000160 02 03 00 2d 00 4e 61 5c 5c 00 00 00 00 3d d7 09
00000200 00 00 00 00 00 05 00 00 00 05 00 00 00 00 00
00000220 00 00 00 00 00 00 ff 00 00 00 80 70 99 75 53 8c
00000240 ff ff 53 01 81 02 03 00 2d 3c 4e 61 5c 5c 00 00
00000260 00 00 64 d7 09 00 8d ff ff ff 05 00 00 00 00 00
```

5) Next press the key that you want to verify is functioning -- if no new lines show up, then the key is non-functional so there is a hardware problem. If you get output, then perhaps there is a problem with your software setup.

6) Use Ctrl-C on the terminal window when done to get back to the prompt.

Note 3: Another method for testing to make sure your model of the Shuttle does not have different key definitions than the one that Cinelerra was coded for is to do the following.

- 1) Locate the shudmp.C program in your Cinelerra directory.
- 2) Compile that with the command: `c++ shdmp.C -o shudmp`
- 3) Make the file executable with the command: `chmod +x shudmp`
- 4) `sudo ./shdmp /dev/input/by-id/usb-Contour_Design_ShuttlePro-event-if00` # substitute your shuttle

Then press your shuttle key that is having problems and check the results. They should look like:

Example for K7:

event: (4, 4, 0x90007)

The last number, 7, is the expected Key number.

event: (1, 262, 0x1)

event: (0, 0, 0x0)

event: (4, 4, 0x90007)

event: (1, 262, 0x0)

event: (0, 0, 0x0)

Example for K15:

event: (4, 4, 0x9000f)

The last number f is 15 in octal and is the expected Key #.

event: (1, 270, 0x1)

event: (0, 0, 0x0)

event: (4, 4, 0x9000f)

event: (1, 270, 0x0)

event: (0, 0, 0x0)

When done, you will have to Ctrl-C to get out of the program.

Note 4: For developers, it is also possible to do the following for further in depth testing:

- 1) `ls /sys/kernel/debug/hid` # to locate numerical value of the shuttle, e.g. 0003:0B33.0030.0006
- 2) `cat "/sys/kernel/debug/hid/0003:0B33.0030.0006/rdesc"` # substitute your own numerical value
- 3) `cat "/sys/kernel/debug/hid/0003:0B33.0030.0006/events"` # press keys to see the results

This page intentionally left blank so that chapters begin on a right facing page.

7 – Rendering

Rendering takes a section of the timeline, performs all the editing, effects and compositing, and creates a new media file. You can then delete all the source assets, play the rendered file, or bring it back into Cinelerra for more editing. All rendering operations are based on a region of the timeline to be rendered. You need to define this region on the timeline. The rendering functions define the region based on a set of rules. When a region is highlighted or in/out points are set, the affected region is rendered. When no region is highlighted, everything after the insertion point is rendered. By positioning the insertion point at the beginning of a track and unsetting all in/out points, the entire track is rendered. But you also have the choice to render “one frame”.

7.1 Single File Rendering

Use the File pulldown and select Render to start the render dialog. Then choose the desired parameters.

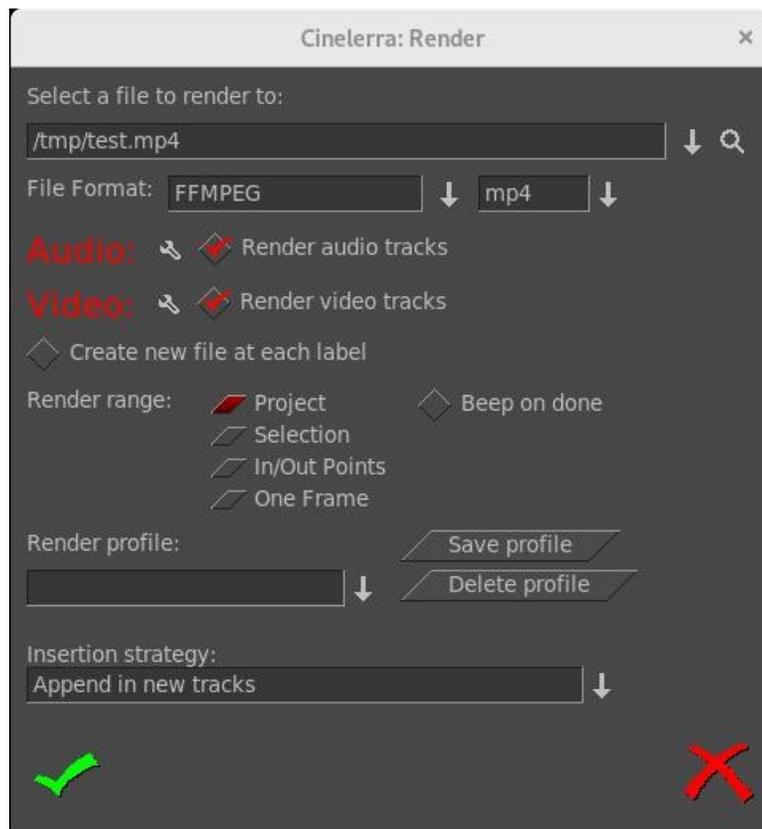


Figure 28: Example of the Render menu

Select a file to render to: enter the path and filename to write the rendered file to in the textbox below.

File Format: use the down arrow to see file format options. For ffmpeg, which has its own set of options, you will then have to select an ffmpeg file type from the down arrow choices.

The format of the file determines whether you can render audio or video or both. Check the *Render audio tracks* toggle to generate audio tracks and check the *Render video tracks* to generate video tracks. The Render window will sometimes automatically update the Render Audio Tracks or Render Video Tracks checkbox as allowed by the chosen file format, but you should always check. For example, if the PNG file format is selected, only the "Render Video Tracks" will be checked. Or if an ffmpeg format is chosen and the file format does not render audio, the "Render Audio Tracks" will be unchecked. The invalid choices will be ghosted out.

Select the *wrench* next to each toggle to set compression parameters. If the file format can not store audio or video the compression parameters will be blank. If *Render audio tracks* or *Render video tracks* is selected and the file format does not support it, trying to render will result in an error message.

The *Create new file at each label* option causes a new file to be created when every label in the timeline is encountered – a separate file for each. This is useful for dividing long audio recordings into individual tracks. When using the Render Farm (described later), *Create new file at each label* causes one render farm job to be created at every label instead of using the internal load balancing algorithm to space jobs. If the filename given in the render dialog has a 2 digit number in it, the 2 digit number is overwritten with a different incremental number for every output file. If no 2 digit number is given, Cinelerra automatically concatenates a number to the end of the given filename for every output file. For example, in the filename “/movies/track01.wav” the “01” would be overwritten for every output file. The filename “/movies/track.wav”; however, would become “/movies/track.wav001” and so on. Filename regeneration is only used when either render farm mode is active or creating new files for every label is active.

Render range: choices are Project, Selection, In/Out points, and One Frame for single images like Tiff. For these images, Render range will have "One Frame" automatically checked and all of the others ghosted since nothing else makes sense. This makes it easy to set the insertion point where you want the 1 frame to be rendered rather than having to precisely zoom in to set the in/out pointers. Note that whichever Render range is checked, remains checked so that if “One Frame” gets automatically checked, the next time you render it will still be checked and you will have to select a different one if desired. That is why you should always check the settings.

Beep on done: as a convenience when a render is complete, check this box. It gives you the chance to work on something else while waiting and still be immediately notified when the render is complete.

Render Profile: another convenience feature to take advantage of if you use specific render formats frequently, is to save that profile for future usage without having to set it up again.

Save Profile: after setting up your render preference formats, use the save profile button to save it.

Delete Profile: if you want to delete a saved profile, highlight the one you no longer want and delete.

Insertion strategy: select an insertion mode from the available choices as seen when you click on the down arrow on the right hand side of the option. The insertion modes are the same as with loading files. In the case if you select “insert nothing” the file will be written out to disk without changing the current project. For other insertion strategies be sure to prepare the timeline to have the output inserted at the right position before the rendering operation is finished.

Even if you only have audio or only have video rendered, a paste insertion strategy will behave like a normal paste operation, erasing any selected region of the timeline and pasting just the data that was

rendered. If you render only audio and have some video tracks armed, the video tracks will get truncated while the audio output is pasted into the audio tracks.

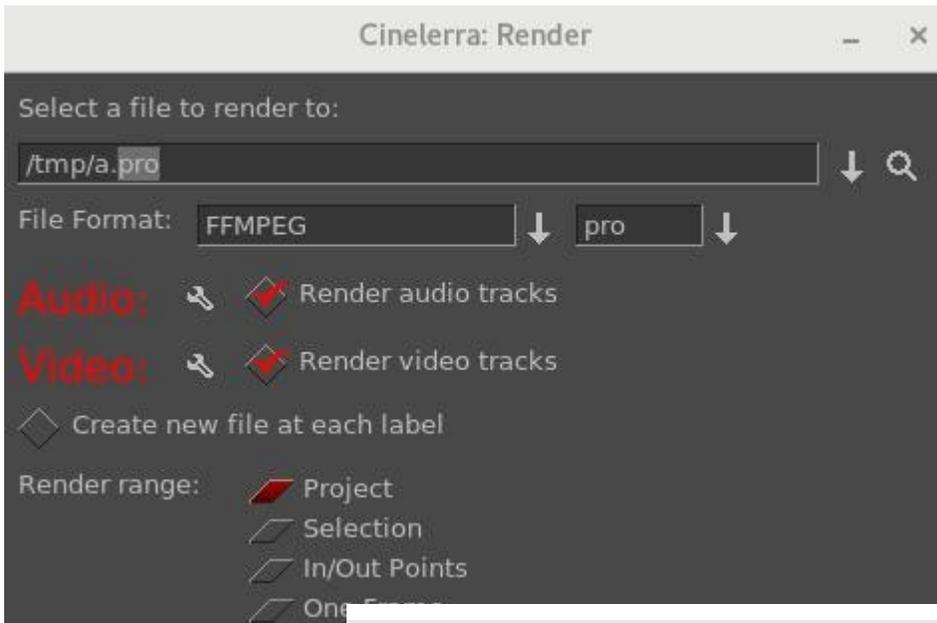


Figure 29: Audio and Video type

Screenshot showing the File Format of PNG which can only do 1 frame and is only for video. Note that Render audio tracks is ghosted out as are all of the Render range options since the only choice is One Frame which was automatically chosen.

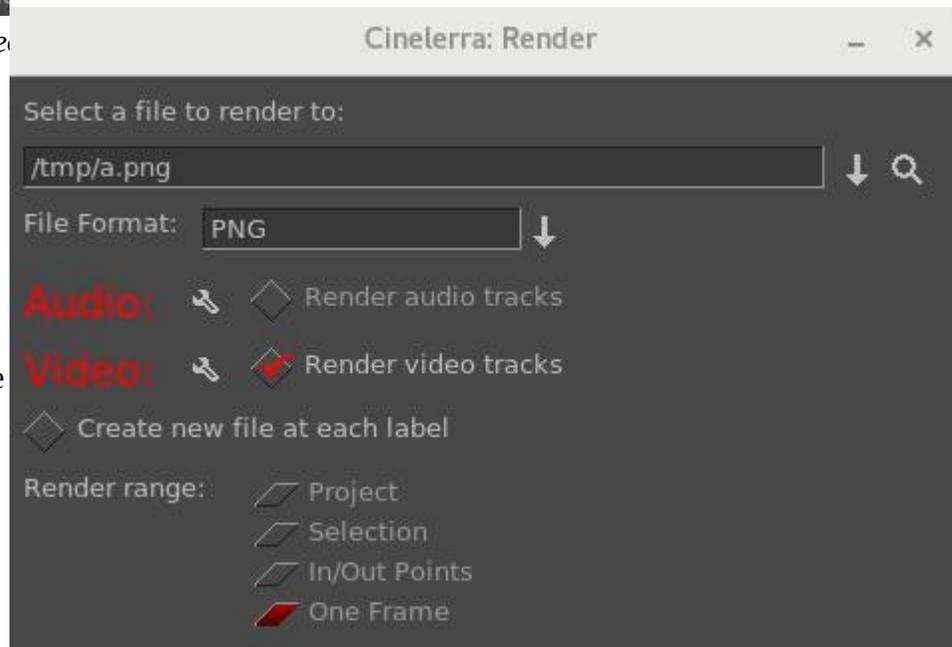


Figure 30: Render menu displaying a PNG "one frame" option

7.2 Batch Rendering

Batch Rendering automates the rendering of audio/video files in that you can establish a set of job parameters, save them, and use them repeatedly. It also allows for Cinelerra to be run by external programs, with no need for the user to manually interact with the user interface.

If you want to render many projects to media files without having to constantly set up the render dialog for each one, batch rendering is a more efficient method of rendering. In the Batch Render menu, you specify one or more Cinelerra project XML files, the EDL, to render and unique output files for each. (The EDL is the Edit Decision List or the set of changes to be applied to the project and media files.) Then Cinelerra loads each project file and renders it automatically. The project XML files, combined with the settings for rendering an output file, are called a batch. This allows a large amount of media to be processed without user intervention.

The first thing to do when preparing to do batch rendering is to create one or more Cinelerra projects to be rendered and save them as a normal project, such as ProjectA.xml. The batch renderer requires a separate project file for every batch to be rendered. You can use the same Cinelerra project file if you are rendering to different output files, as in an example where you might be creating the same output video in different file formats.

To create a project file which can be used in batch render, set up your project and define the region to be rendered either by highlighting it, setting in/out points around it, or positioning the insertion point before it. Then save the project as usual to your project.xml file. Define as many projects as needed this way. The batch renderer takes the active region from the EDL file for rendering.

With all the Cinelerra xml project files prepared with active regions, go to File → Batch Render. This brings up the batch render dialog. The interface for batch rendering is more complex than for single file rendering. A list of batches must be defined before starting a batch rendering operation. The table of batches appears on the bottom of the batch render dialog and is called *Batches to render*. Above this are the configuration parameters for a single batch; a batch is simply a pairing of a project file with a choice of output file and render settings.

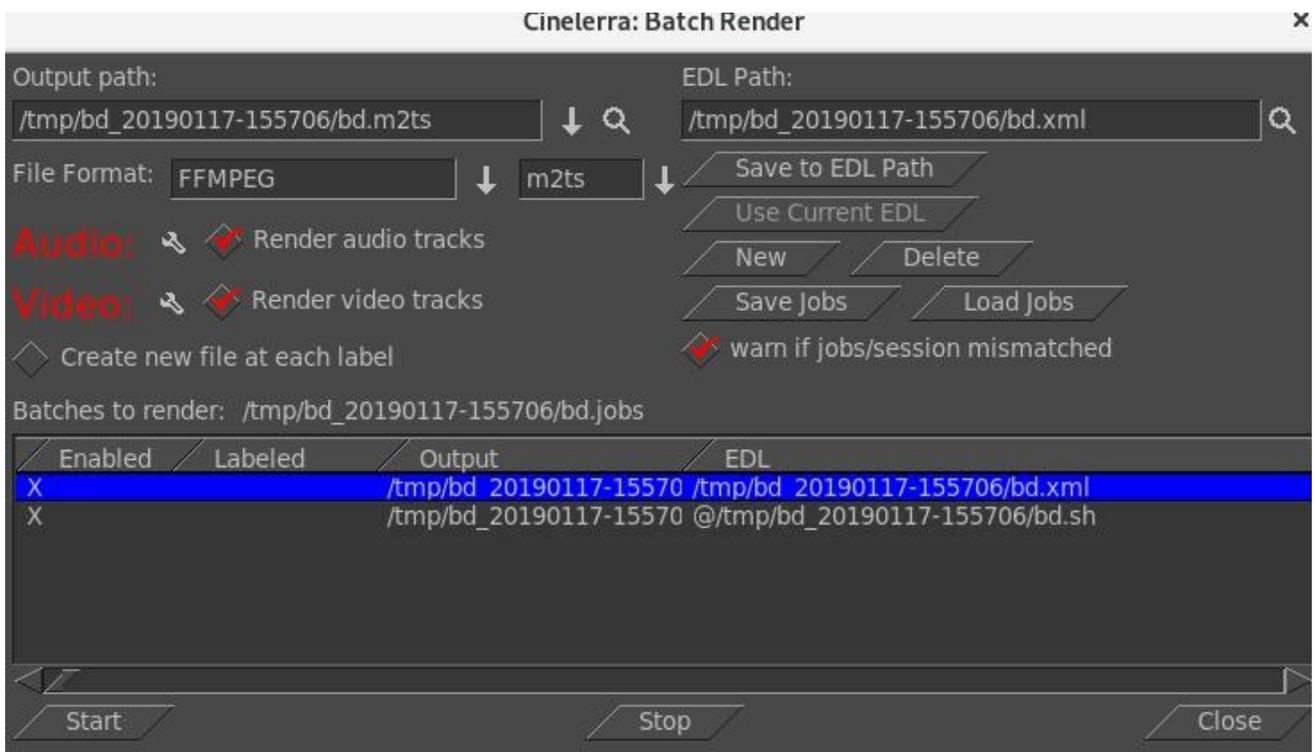


Figure 31: Example of the Batch Render menu

Set the *Output path*, *File format*, *Audio*, *Video*, and *Create new file at each label* parameters as if you were rendering a single file. These parameters apply to only one batch. In addition to the standard rendering parameters, you must select the *EDL Path* to be the project file (such as ProjectA.xml) that will be used in the batch job. In this case, *EDL Path* is not related in anyway with the EDL files as created by File/Export EDL. In batch render mode the program will not overwrite an existing output file and will simply fail, so make sure that no files with the same name as the output files exist before starting.

If the batches to render list is empty or nothing is highlighted, click *New* to create a new batch. The new batch will contain all the parameters you just set. Repeatedly press the *New* button to create more batches with the same parameters. When you highlight any batch, you can edit the configuration on the top of the batch render window. The highlighted batch is always synchronized to the information displayed. You can easily change the order in which the batch jobs are rendered, by clicking and dragging a batch to a different position. Hit *Delete* to permanently remove a highlighted batch. In the list box is a column which enables or disables the batch with an X meaning the batch job is enabled and will be run. This way batches can be skipped without being deleted. Click on the Enabled column in the list box to enable or disable a batch.

The description of each of the columns in the batch list are as follows:

Enabled: an X in this column means the batch job will be run.

Labeled: an X in this column goes hand in hand with create new file at each label.

Output: path and filename for the generated output.

EDL: the path and filename of the source EDL for the batch job.

Elapsed: the amount of time taken to render the batch if finished. If field is empty, it did not run.

To start rendering from the first enabled batch, hit *Start*. Once rendering, the main window shows the progress of the batch. After each batch finishes, the elapsed column in the batch list is updated and the next batch is rendered until all the enabled batches are finished. The currently rendering batch is always highlighted red. To stop rendering before the batches are finished without closing the batch render dialog, hit *Stop*. To stop rendering before the batches are finished and close the batch render dialog, hit *Close*. Or you can exit the batch render dialog whether or not anything is being rendered, by hitting *Close*.

You can automate Cinelerra batch renders from other programs. In the batch render dialog, once you have created your list of batch render jobs, you can click the button *Save Jobs* and choose a file to save your batch render list to. Once you have created this file, you can start up a batch render without needing to interact with the Cinelerra user interface. From a shell prompt, from a script, or other program, execute:

```
{path_to_cinelerra} -r batchjob.xml
```

substituting your actual filename for batchjob.xml. When invoked with these parameters, Cinelerra will start up and perform the rendering jobs in that list, without creating its usual windows.

Command Line Rendering

The command line rendering method consists of a way to load the current set of batch rendering jobs and process them without a gui. This is useful if you want to do rendering on the other side of a low bandwidth network and you have access to a high powered computer located elsewhere. Setting up all the parameters for this operation is somewhat difficult. That is why the command line aborts if any output files already exist.

To perform rendering from the command line, first run Cinelerra in graphical mode. Go to File->Batch Render. Create the batches you intend to render in the batch window and close the window. This saves the batches in a file. Set up the desired render farm attributes in Settings->Preferences and quit out of Cinelerra if you want to use the Render Farm capability. These settings are used the next time command line rendering is used to process the current set of batch jobs without a GUI.

On the command line run: cinelerra -r

More about Save/Use EDL and Save/Load Jobs

The File->Batch Render pulldown brings up the Batch Render window to be used for batch rendering as well as DVD/BD creation. There are some additional buttons that can save time and mistakes. These are described next.

The *Save to EDL Path* and *Use Current EDL* buttons can be valuable tools for advanced usage or for developers doing testing. Description of how you can expect them to work will help to illustrate how to take advantage of their capabilities.

- 1) *Save to EDL Path* – if you have made a change to the EDL, use this button to save the changes so that they will be used in the render operation. Although you can get the same results by using File → Save..., this capability was initially added to assist developers in testing the batch jobs needed to create dvd/bluray media as it keeps the work focused in a single window and retains the original job name. An example -- you have everything all set up with a new job in the Batch Render window using generic.xml for the EDL path and with a job name of original_name.xml. Then you realize that you forgot to cut out a section in the media that is not wanted in the final product. You can cut that out and then *Save to EDL Path* so your change will be in effect for the rendering. Without this button, you would be using the EDL you started with and the cut would be ignored. Alternatively, if the cut changes are saved via File → Save as... with a filename of new.xml and then you use *Save to EDL Path*, the current highlighted job displayed in the window as “original_name.xml” will be replaced with new.xml. However, it is important to note that the result will be saved with the name original_name – that is, the new content from new.xml but with the old name of original_name.xml.
- 2) *Use Current EDL* – if you are working on media and still testing out the results, you can take advantage of this click-box to quickly get results. Basically, you change the media, save that change with another name (in order to preserve the original name in case you don't like the changes), and press *Use Current EDL*. As an example, a user creates a new job in the Batch Render window using the current media, previously defined in generic.xml, with the EDL path of generic.xml. The user then changes the media on the timeline, saves the changes via File → Save as... with a new name, such as new_name.xml, and then clicks on *Use Current EDL*. In this case, the EDL path listbox will be automatically updated to the new_name.xml and the current existing highlighted job will be replaced with the new_name.xml in the EDL column.
- 3) *Save Jobs* – when you have set up the batch jobs the way you want and you think you may have to

run them more than once, it is beneficial to save the jobs for later use so you easily run them again.

- 4) *Load Jobs* – reload a previous set of saved jobs. This can come in handy if you did not have the time to render them when you originally set them up, if you need to rerun, or if you got interrupted.
- 5) Checkbox *warn if jobs/session mismatched*. After you set up your render and press Start, the program checks to see if the current EDL session matches your Batch Render job. If the EDL has been changed since the batch job was created, it warns you so that you have the opportunity to *Save to EDL path* to record those changes. Otherwise, you can dismiss that warning box, disable the warning message by unchecking the box and use the original values. If you never want to be warned about the mismatches, leave the box unchecked.

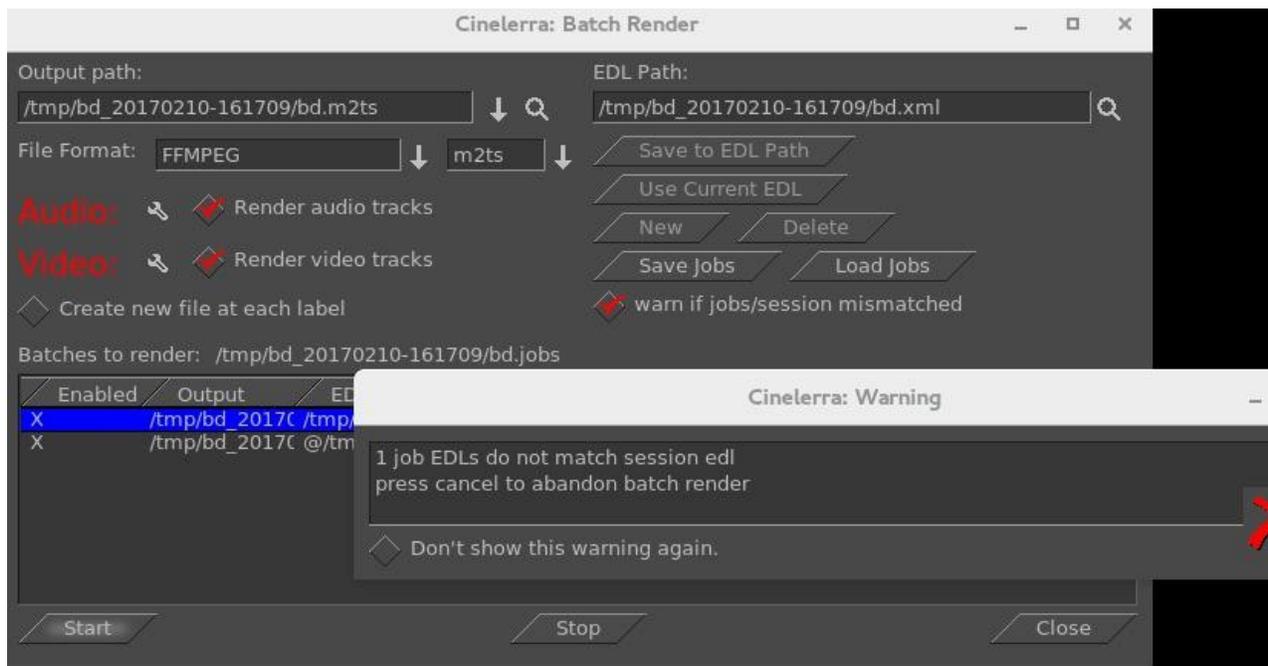


Figure 32: Batch render with the 4 ghosted buttons on the right side + the Warning message below

7.3 Background Rendering

Background rendering causes temporary output to be rendered constantly while the timeline is being modified. The temporary output is displayed during playback whenever possible. This is useful for transitions and previewing effects that are too slow to display in real time. If a Render Farm is enabled, the render farm is used for background rendering. This gives you the potential for real-time effects if enough network bandwidth and CPU nodes exist.

Background rendering is enabled in the Performance tab of the Preferences window. It has one interactive function Settings menu -> Set background render. This sets the point where background rendering starts up to the position of the insertion point. If any video exists, a red bar appears in the time ruler showing what has been background rendered.

It is often useful to insert an effect or a transition and then select Settings menu -> Set background render right before the effect to preview it in real time and full frame rates.

Frames per background rendering job

This only works if a Render Farm is being used; otherwise, background rendering creates a single job for the entire timeline. The number of frames specified here is scaled to the relative CPU speed of rendering nodes and used in a single render farm job. The optimum number is 10 - 30 since network bandwidth is used to initialize each job.

Frames to preroll background

This is the number of frames to render ahead of each background rendering job. Background rendering is degraded when preroll is used since the jobs are small. When using background rendering, this number is ideally 0. Some effects may require 3 frames of preroll.

Output for background rendering

Background rendering generates a sequence of image files in a certain directory. This parameter determines the filename prefix of the image files. It should be accessible to every node in the render farm by the same path. Since hundreds of thousands of image files are usually created, ls commands will not work in the background rendering directory. The browse button for this option normally will not work either, but the configuration button for this option works.

File format

The file format for background rendering has to be a sequence of images. The format of the image sequences determines the quality and speed of playback. JPEG generally works well.

7.4 Render Farm Usage

Render Farm uses background rendering, a feature of Cinelerra where the video is rendered in the background, to speed up rendering significantly. Because rendering is memory and cpu intensive, using multiple computers on a network via a render farm is a significant gain. With Cinelerra installed on all nodes, the master node and the clients communicate via a network port that you specify.

Cinelerra can distribute the rendering tasks over the network to the other computers of the Render Farm. The render farm software tries to process all of the rendering in parallel so that several computers can be used to render the results. The *Total jobs to create* in the setup or labels on the timeline are used to divide a render job into that specified number of tasks. Each background job is assigned a timeline segment to process and the jobs are sent to the various computer nodes depending upon the load balance. The jobs are processed by the nodes separately and written to individual files. You will have to put the files back together via a load with concatenation, or typically by using a command line tool from a script.

7.4.1 Basic Steps to Start a Render Farm

The following steps are just a guideline to start your render farm. It is assumed that you already have the master and client nodes communication, shared filesystem, permissions and usernames synched.

- 1) On the master computer, use Settings → Preferences, the Performance tab to set up a Render Farm:
 - a) check the *Use render farm* box
 - b) in the Hostname box, keyin your hostname or ip address such as 192.168.1.12 or localhost;
 - c) enter in a port number such as 401-405 or 1025 and click on *Add Nodes*;
 - d) you will see something like the following in the Nodes listbox to the right:

On	Hostname	Port	Framerate
X	192.168.1.12	401	0.0
X	192.168.1.12	402	0.0
X	192.168.1.12	403	0.0
X	192.168.1.12	404	0.0
X	192.168.1.12	405	0.0
X	localhost	406	0.0
X	localhost	407	0.0

- e) set the Total number of jobs to create;
 - f) click OK on the bottom of the Preferences window.

- 2) On the client computers (192.168.1.12), start 5 background cinelerra tasks via:

```
cd /{path_to_cinelerra}
cin -d 401
cin -d 402
...
cin -d 405
```

- 3) On the master node (localhost), start the 2 background cinelerra tasks via:

```
cd /{path_to_cinelerra}
```

cin -d 406
cin -d 407

- 4) When your video is ready, setup a render job via File “Render” or “Batch Render” and check OK.
- 5) The results will be in the shared file path / filename that you selected in the render menu with the additional numbered job section on the end as 001, 002, 003, ... 099 (example, video.webm001).
- 6) When finished, load your new files on new tracks via File → Load “concatenate to existing tracks” or if you used ffmpeg, run RenderMux from the Shell Scripts icon.
- 7) If you plan on doing more rendering, you can just leave the master/client jobs running to use again and avoid having to restart them. Or you can kill them when you no longer are using them.

7.4.2 Render Farm Menu and Parameter Description

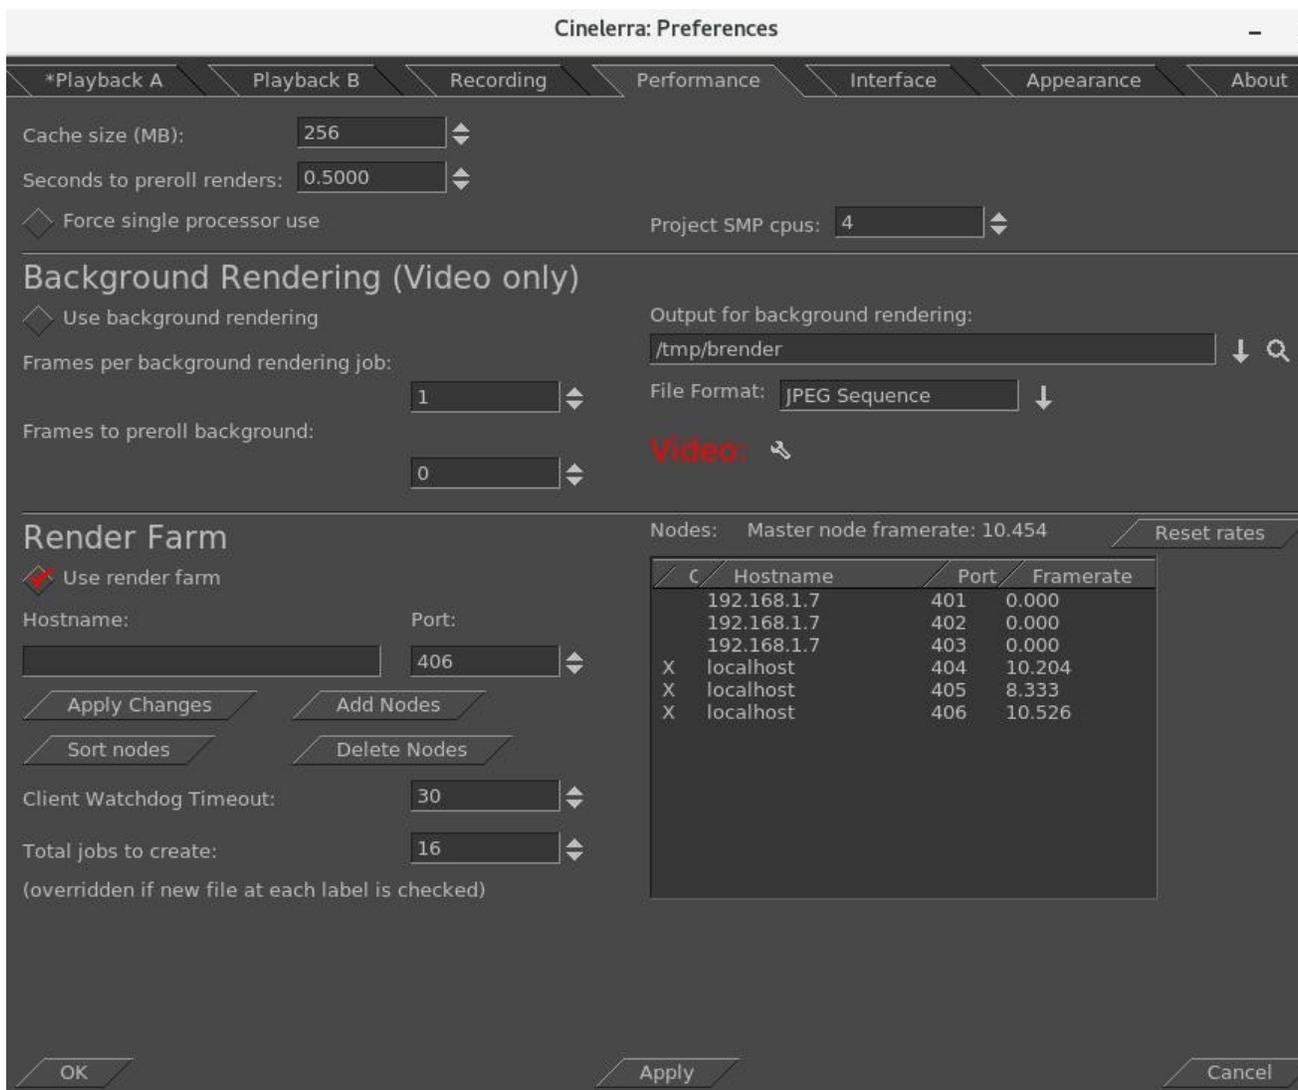


Figure 33: Settings → Preferences, Performance tab, menu to set up your Render Farm

Project SMP cpus - although this field is not Render Farm specific, it is useful for Cinelerra to have the cpu count and for using multiple threads.

Use render farm - check this to turn on the render farm option. Once checked ALL rendering will be done via the farm including the usual Render (Shift-R). You may want to turn it off for small jobs.

Nodes listbox - displays all the nodes on the render farm and shows which ones are currently enabled. The Nodes listbox has 4 columns – On, Hostname, Port, Framerate – which show the current values. An X in the “On” designates that that host is currently enabled; Hostname shows the name of the host; Port shows the port number that host uses; and Framerate will either be zero initially or the current framerate value.

Hostname - this field is used to edit the hostname of an existing node or enter a new node.

Port - key in the port number of an existing or new node here. You can also type in a range of port numbers using a hyphen, for example 1501-1505 when you need to add many.

Apply Changes - this will allow you to edit an existing node and to then commit the changes to hostname and port. The changes will not be committed if you do not click the OK button.

Add Nodes - Create a new node with the hostname and port settings.

Sort nodes - sorts the nodes list based on the hostname.

Delete Nodes - deletes whatever node is highlighted in the nodes list. You can highlight several at once to have them all deleted.

Client Watchdog Timeout - a default value of 15 seconds is used here and the tumbler increments by 15 seconds. A value of 0 (zero) disables the watchdog so that if you have a slow client, it will not kill the render job while waiting for that client to respond.

Total jobs to create - determines the number of jobs to dispatch to the render farm. Total jobs is used to divide a render job into that specified number of tasks. Each background job is assigned a timeline segment to process. The render farm software tries to process all of the rendering in parallel so that several computers can be used to render the results.

To start, if you have computers of similar speed, a good number for *Total jobs to create* is the number of computers multiplied by 3. You will want to adjust this according to the capabilities of your computers and after viewing the framerates. Multiply them by 1 to have one job dispatched for every node. If you have 10 client nodes and one master node, specify 33 to have a well balanced render farm.

(overridden if new file at each label is checked) – instead of the number of jobs being set to *Total jobs to create*, there will be a job created for each labeled section. If in the render menu, the option “Create new file at each label” is selected when no labels exist, only one job will be created. It may be quite advantageous to set labels at certain points in the video to ensure that a key portion of the video will not be split into 2 different jobs.

Reset rates - sets the framerate for all the nodes to 0. Frame rates are used to scale job sizes based on CPU speed of the node. Frame rates are calculated only when render farm is enabled.

Framerates can really affect how the Render Farm works. The first time you use the render farm all of the rates are displayed as 0 in the Settings → Preferences, Performance tab in the Nodes box. As rendering occurs, all of the nodes send back framerate values to the master node and the preferences page is updated with these values. A rate accumulates based on speed. Once all nodes have a rate of non-zero, the program gives out less work to lower rated nodes in an effort to make the total time for the render to be almost constant.

Initially, when the framerate scaling values are zero, the program just uses package length – render size divided by the number of packages to portion out the work (if not labels). If something goes wrong or the rates become suspect, then all of the rest of the work will be dumped into the last job. When this happens, you really should “reset rates” for the next render farm session to restart with a good balance.

`{cinelerra pathname} -h` displays some of the options.

7.4.3 Detailed Setup Description

CAUTION, any exact command lines worked as of 1/2018 on a Fedora system. These can change over time and on different operating systems/levels. Always check/verify any command line before using.

1) Set up Cinelerra

A Cinelerra render farm is organized into a master node and any number of client nodes. The master node is the computer which is running the gui. The client nodes are anywhere else on the network with Cinelerra installed and are run from the command line. Before you start the master node for Cinelerra, you need to set up a shared filesystem on the disk storage node as this is the node that will have the common volume where all the data will be stored.

The location of the project and its files should be the same in the client computers as in the master computer and to avoid problems of permissions, it is better to use the same user in master and clients. For example, if you have the project in `/home /<user>/project-video` you must create the same directory path on the clients, but empty. Sharing the directory of the location of your project on the master computer can be done with NFS as described next. Alternatively, you can look up on the internet how to use Samba to share a directory.

2) Create a shared filesystem and mount using NFS

All nodes in the render farm should use the same filesystem with the same paths to the project files on all of the master and client nodes. This is easiest to do by setting up an NFS shared disk system.

1) On each of the computers, install the nfs software if not already installed. For example, on Debian 9 you will need to run: (be sure to check/verify before using any command line):

```
apt-get install nfs-kernel-server
```

2) On the computer that contains the disk storage to be shared, define the network filesystem. For example to export `/tmp`, edit the `/etc/exports` file to add the following line:

```
/tmp 192.168.1.0/24(rw,fsid=1,no_root_squash,sync,no_subtree_check)
```

3) Next reset the exported nfs directories using: `exportfs -ra`
and you may have to start or restart nfs: `systemctl restart nfs`

4) Each of the render farm computers must mount the exported nfs target path. To see the exports which are visible from a client, login as root to the client machine and keyin:

```
showmount -e <ip-addr> using the ip address of the storage host
```

5) to access the host disk storage from the other computers in the render farm, mount the nfs export on the corresponding target path: (be sure to check/verify before using any command line):

```
mount -t nfs <ip-addr>:/<path> <path>
```

where <path> is the storage host directory, and <ip-addr> is the network address of the storage host.

Because all of the computers must have the same directory path, create that same directory path with the same uid/gid/permissions on each storage client computer ahead of time.

6) To make this permanent across reboots on the client nodes, add the following line to /etc/fstab:

```
{masternode}:/nfsshare /mnt nfs defaults 0 0
```

You can make this permanent on the disk storage host BUT the command lines shown, which were correct in January 2018 on Fedora, may be different for your operating system or in the future. In addition if your network is not up, there may be numerous problems. If you make a mistake, your system may not boot. To make permanent, add the following line to /etc/fstab:

```
192.168.1.12:/tmp /tmp nfs rw,async,hard,intr,noexec,noauto 0 0
```

You will still have to mount the above manually because of the “noauto” parameter but you won’t have to remember all of the other necessary parameters. Depending on your expertise level, you can change that.

Later, to remove access to the storage host filesystem: `umount <path>`

Be aware that you may have to adjust any security or firewalls you have in place. **Most firewalls will require extra rules to allow nfs access.** Many have built-in configurations for this.

3) Configure Rendering on Master Node

There is 1 master node which is running the Cinelerra gui and where the video will be edited and the command given to start up the rendering. Any number of client computers can be run from the command line only, so they can be headless since no X or any graphical libraries are needed. Of course, the cinelerra software must be installed on each of the client computers.

- 1) Assuming you already have Cinelerra installed on the master node, start Cinelerra by clicking on the icon or by typing the following command on the terminal screen: `/{cinelerra_path}/cin .`
- 2) Use the file pulldown Settings → Preferences, the Performance tab, to set up your Render Farm options in the Render Farm pane.
- 3) Check the “Use render farm” option. By default, once you enable the option of Render Farm, rendering is usually done using the render farm. Batch rendering can be done locally, or farmed.
- 4) Add the hostname or the IP address of each of the client nodes in the Hostname textbox and the port number that you want to use in the Port textbox. You can make sure a port number is not already in use by keying in on the command line: `netstat -n -l -4 --protocol inet`. Next, click on the Add Nodes button and then you will see that host appear in the Nodes list box to the right. The ‘X’ in the first column of the nodes box denotes that the node is active. To review the “standard” port allocations,

check the /etc/services file.

- 5) Enter the total jobs that you would like to be used in the “Total job” textbox.
- 6) The default watchdog timer initial state is usually just fine but can be adjusted later if needed.
- 7) Click OK on the Preferences window when done.

4) Create Workflow

While working on the master computer, it is recommended that you keep all the resources being used on the same shared disk. Load your video/audio piece and do your editing and preparation. Add any desired plugins, such as a Title, to fine-tune your work. You want to make sure your video is ready to be rendered into the final product.

5) Start the Client Nodes

To start up the client nodes run Cinelerra from the command line on each of the client computers using the following command:

```
{cinelerra_pathname}/cin -d [port #] ; for example /mnt1/bin/cinelerra -d 401
```

This starts cinelerra in command prompt mode so that it listens to the specified port number for commands from the master node for rendering. When you start each of the clients up, you will see some messages scroll by as each client is created on that computer, such as:

```
RenderFarmClient::main_loop: client started  
RenderFarmClient::main_loop: Session started from 127.0.0.1
```

As it completes its jobs, you will should see:

```
RenderFarmClientThread::run: Session finished.
```

A quick way to start a sequence of clients is to use:
for n in `seq 1501 1505`; do cin -d \$n; done

6) Render Using Render Farm

After you have followed the preceding steps, you are ready to use the render farm. Click on File>Render... which opens the render dialog. The most important point here is to use for the “Output path / Select a file to render to” a path/file name that is on the shared volume that is also mounted on the clients. Click on OK to render. The cinelerra program divides the timeline into the number of jobs specified by the user. These jobs are then dispatched to the various nodes depending upon the load balance. The first segment will always render on the master node and the other segments will be farmed out to the render nodes. Batch Rendering, as well as BD/DVD rendering, may use the render farm. Each line in the batchbay can enable/disable the render farm. Typically, video can be rendered into many file segments and concatenated, but normally audio is rendered as one monolithic file (not farmed).

Another performance feature which can use the Render Farm is “Background Rendering”. This is also enabled on the Performance preferences tab. The background render function generates a set of image files by pre-rendering the timeline data on the fly. As the timeline is update by editing, the image data

is re-rendered to a “background render” storage path. The Render Farm will be used for this operation if it is enabled at the same time as the “background render” feature.

7) Assemble the Output Files

Once all of the computer jobs are complete, you can put the output files together by using the shell script, RenderMux (from the menubar “scripts” button just above FF), if the files were rendered using ffmpeg, or you can load these by creating a new track and specifying concatenate to existing tracks in the load dialog in the correct numerical order. File types which support direct copy can be concatenated into a single file by rendering to the same file format with render farm disabled as long as the track dimensions, output dimensions, and asset dimensions are equal.

7.4.4 Quick and Easy Render Farm Setup – The Buddy System Way

These steps are for quickly setting up render farm with the least amount of additional system work, but it is non-optimal. It is useful in situations where a few people all show up with their laptops to work together on the same video/audio file and you don't want to bother setting up NFS for a shared disk.

- 1) Make sure the Cinelerra program is installed on all of the computers and the network between the main computer and the client computers is working. Use the same version if possible.
- 2) Load your video file on the master node and use “File → Save as...” to save it to /tmp .
- 3) Move that same file with the same name to /tmp on all of the client computers via rsh or sneaker net – the ONLY reason you are doing this is to avoid having to set up NFS or Samba on the buddy client laptops that show up!
- 4) Edit your video/audio file to get it the way you want it and add the plugins, such as a Title, etc.
- 5) Check for a set of unused ports in /etc/services file, if username is root usually 401-425 are available; if non-root, then 1024-1079.
- 6) On the master computer, in Settings → Preferences, the Performance tab
 - a) check the box “Use render farm”
 - b) keyin localhost for the hostname or an ip address of the buddy client node
 - c) keyin the desired port number for each client; and use “Add Node” for each host
 - d) set total jobs to the number of client computers+1 multiplied by 3 (or proportion to client speeds)
 - e) check OK
- 7) On each buddy client, create a job for each port: `{cinelerra_pathname}/cin -d port# .`
- 8) On the master, bring up the render menu and name the output files, for example /tmp/myoutput.mp4.
- 9) The client nodes output results will be on their local /tmp filesystems so you will have to again use rsh/ftp or usb sneaker net to move them over to the main computer. File names will be the render job output file name with port number tacked on (e.g. /tmp/hb.mp4001...mp4005).
- 10) Load the files by concatenate to existing track on the master node or use RenderMux shell script.

7.4.5 Multi-core Computers Render Farm Setup (the “Epyc” method!)

If you are lucky enough to have a computer with a large cpu core count, setting up a render farm can really take advantage of using all of the cpus. This is much faster than the default automatic threading capability. Since you don’t need to communicate with other computers, you will not have to be concerned about TCP communication or shared disks/files. Since you may be doing other work simultaneously while rendering a large job, you will want to leave some of the cpus available for that. Be sure to set “Project SMP cpus” in the Settings → Preferences, Performance tab to your CPU count.

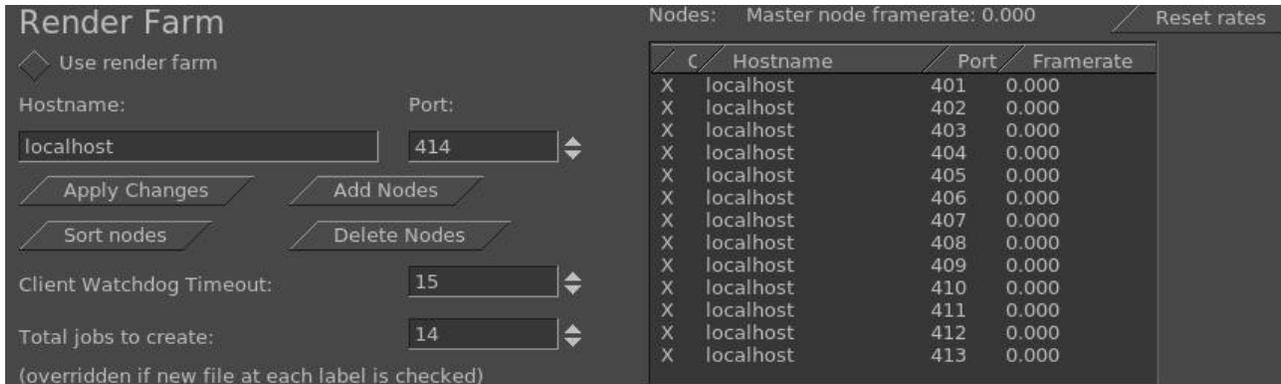


Figure 34: Setup of a Render Farm on a 16-core processor computer

Sharing Index Files

Because index files speed up displaying the video you may want to share these files with the clients on a shared filesystem. There is a configuration option available in Settings → Preferences, the Interface tab, that sets up in your preferences the location of the index files which you can put on a shared disk. Screenshot below shows part of the Preferences menu where you can change the index files setup.

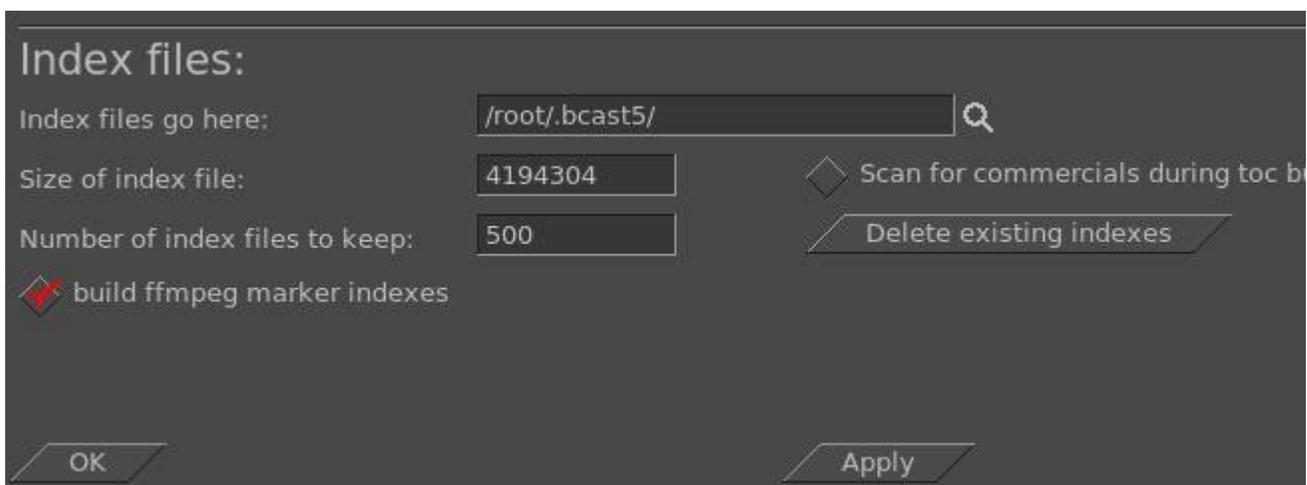


Figure 35: Index file setup for your preferred configuration for Render Farm sharing or anything

Or, one of the convenient features of cin5 is the redirection of the path via CIN_CONFIG as in:
CIN_CONFIG=/<shared_file_pathname>/<filename_such_as_.bcast5> /<cinelerra_pathname>/cin .
This means that you can make project related configurations that do not impact the default \$HOME

config. You can either export your default \$HOME config or the CIN_CONFIG config to use on the render farm.

7.4.6 Troubleshooting Tips and Warnings

If you have problems running the Render Farm. Here is a list of items to check.

- Cinelerra must be installed on the master node and all client machines.
- It is best to have the same username available on all nodes to avoid problems with access rights.
- Check file permissions and ownership to ensure that the clients all have access.
- If a node does not have access to an input asset it will not die, but just display error messages.
- If a node can not access an output asset, the rendering will abort.
- A port in use when stopped may take up to 30 seconds to time out before you can restart the jobs.
- Each of the port combinations have to be unique across clients, and not already in use in the network.
- Cinelerra load balances on a first come, first serve basis. If the last section of the video is sent to the slowest node, the render job will have to wait for the slowest node to finish. It would be better to start on the slowest node with the earlier section of the video so keep that in mind when designating port numbers.
- If not running as root, a port number in the higher range of 1024 and above must be used instead of the 400+ range.
- The master and client jobs on the ports do not go away so if you want to stop them, you will have to kill them via: kill PID# .
- Check to see if there are services listening on the ports to use: netstat -n -l -4 --protocol inet .
- There is a watchdog timer in Cinelerra and if there is no response from a client in the designated number of seconds, it will kill the render job.
- The "localhost" should exist as 127.0.0.1 in /etc/hosts and as the "lo" network device in ifconfig.
- If the job loads become unbalanced, you may want to "reset rates" to start over for new framerates.
- If jobs are split in a key section on the timeline, you may wish to "use labels" to prevent this.
- For testing purposes, you may want to start a client in the foreground using "-f" instead of "-d".
- If one of the client computers is unavailable, check to see if there is an X to the left of the nodename in the Nodes listbox. Check the X to disable it which sets On to Off.
- A red message in the lower left hand corner of the main timeline that reads "Failed to start render farm" often means that the client cinelerra programs were not started up.
- A message of "RenderFarmWatchdog::run 1 killing server thread #address#" means that the client did not respond in time. You can adjust the timer in Settings → Preferences, Performance tab.
- When you get the message "RenderFarmClient::main_loop: bind port 400: Address already in use", use a different port.
- A message of "RenderFarmServerThread::open_client: unknown host abcompany" means that the hostname of abcompany is not in /etc/hosts so you will have to add it or use the ip address instead.
- There are numerous error messages associated with file open/close/status or problems with the file that should be dealt with according to what is printed out.
- Other illustrative messages may be shown such as: "RenderFarmClientThread::run: Session finished.

Warnings

If one of the render farm computers is connected to the internet, you should use a firewall to maintain the safety of all of the computers. The ports have to be reachable for the intranet but you do not want the ports to be open to the outside.

7.5 Some Specific Rendering

The next few pages relate to rendering for specific common cases.

7.5.1 FFmpeg Common H.264 Rendering

Because H.264 is so widely used, the method in Cinelerra-GG Infinity is outlined below. These setup steps make it easy to just get started.

File->Render

File Format->FFMPEG + mp4

Video Wrench->Preset->h264.mp4 + bitrate: 6000000 (or whatever) + OK

Audio Wrench->Preset->h265.mp4 + bitrate: 224000 (or whatever) + OK

Set your target path in: Render=>Select a file to render to

Set your timeline in: Render=>Render range + click Project

Set your insertion strategy: Replace project (or whatever)

Press OK to start rendering.

7.5.2 Lossless Rendering

Lossless means that in the compression of a file, all of the original data, every single bit, can be recovered when the file is uncompressed. This is different than “lossy compression” where some data is permanently deleted so that when uncompressed, all of the original data can not be exactly recovered. Lossy is generally used for video and sound, where a certain amount of information loss will not be detected by most users or the playback hardware does not reproduce it anyway – it is a trade-off between file size and image/sound quality. The files created will be more than 10 times larger than usual. Most players will not be able to decode lossless as the bitrate will overwhelm the device.

For x264 lossless compression to work, the only color model allowed here is yuv420p. Any other specification will be converted to yuv420p and the data will be modified. Also, keep in mind that the YUV color model has to be converted to RGB, which also modifies the data.

To use x264 lossless rendering - choose File format of “ffmpeg”, “m2ts” in the Render window. Click on the Video wrench, which brings up the Video Preset window and scroll down in the Compression filebox and choose “lossless.m2ts”. Preset=medium is the default, but can be varied from ultrafast (least amount of compression, but biggest file size) to veryslow (most amount of compression, but still HUGE) in the parameter box where you see “qp=0”. This option is also available for bluray creation.

7.5.3 Extra “cin_” Options for Render with FFmpeg

There are several special parameters that can be used in the ffmpeg options file to pass values to the codecs that are not normally available. These are explained below.

1) cin_pix_fmt

The Render menus allows you to choose the codec input pixel format. The Pixels selection provides the available pixel format options for the chosen codec type; valid choices vary for the different file types. This list represents the formats that the codec advertises. It is not always complete, and it may include options that are not legal with all parameter configurations.

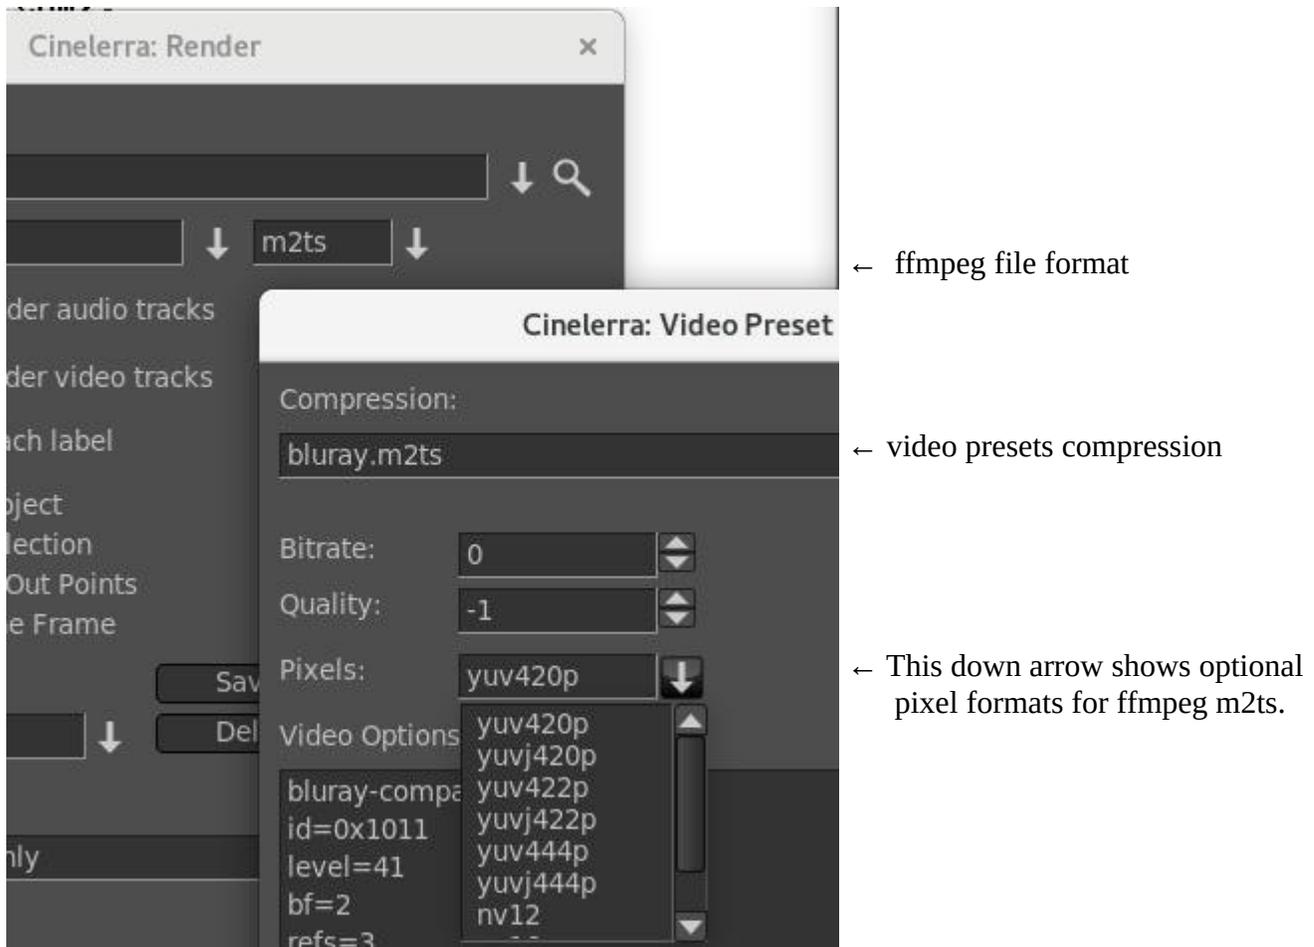


Figure 36: Render & Video Preset menus displaying Pixel choices

- The Bitrate, Quality, and Pixels fields are only updated when the Video Options are reloaded. This occurs when you either change the ffmpeg file format, or video presets compression fields.
- If the video options preset has cin_pix_fmt defined, its value will be loaded as the default. If you override the default, the value you specify will be used.
- If the video options preset does not have cin_pix_fmt, the default pixel format will be computed by ffmpeg (avcodec_find_best_pix_fmt_of_list), using the session format as the source choice. The “best” is usually the format which is most similar in color and depth.
- If no choices are available, yuv420p for video will be used.
- You can also specify ffmpeg pixel formats which are not in the list. The list is provided by ffmpeg as input selection, but is more like suggestions than fact. For example, the raw formats can take almost any format, but the rawvideo codec actually specifies no legal formats.

Some option files provide cin_pix_fmt to suggest a choice for good quality output or to prevent parameter errors when the other provided parameters conflict with the “best” pixel format. This is the case in faststart_h264.mp4 where the profile=high parameter dictates pixel format must be yuv420p.

2) cin_bitrate

If you specify the bitrate, you can not specify the quality.

Example: cin_bitrate=2000000

3) cin_quality

If you specify the quality, you can not specify the bitrate.

Example: cin_quality=7

4) cin_stats_filename

This parameter is useful for 2 pass operations.

Example: cin_stats_filename /tmp/cin_video_vp9_webm

5) cin_sample_fmt

For audio the preset sample format default is computed in a similar way as stated above for video or can be set with the cin_sample_fmt parameter. If no choices are provided, s16 will be used.

Example: cin_sample_fmt=s16

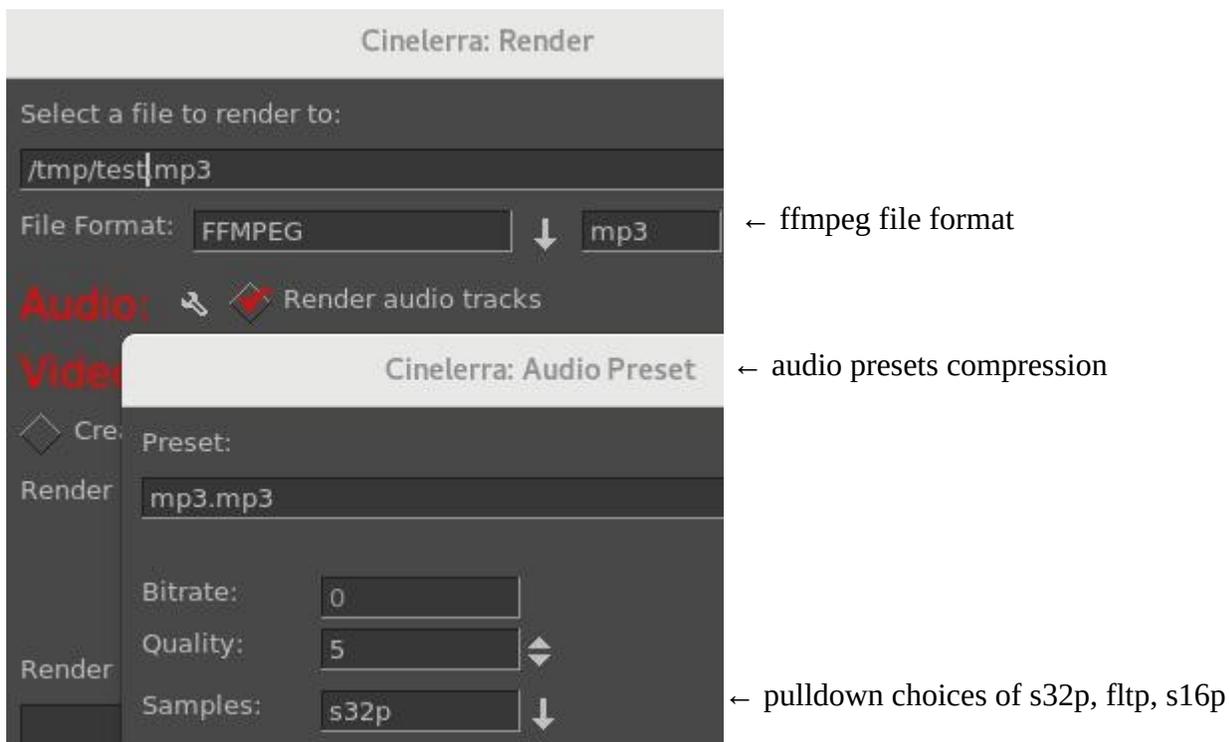


Figure 37: Render menu showing where Samples is

7.5.4 Two-pass Encoding with FFmpeg

In Cinelerra for two-pass, you need to run ffmpeg twice, with the same settings, except for designating the options of pass 1 for the first pass and then pass 2. In pass 1, a logfile that ffmpeg needs for the second pass is created. In pass 1 the audio codec should be specified that will be used in pass 2. For more information on ffmpeg 2-pass, check this url: <https://trac.ffmpeg.org/wiki/Encode/H.264> . Different libraries may have different requirements and you will probably have to determine this by looking online at ffmpeg or looking directly at that code.

This 2 line ffmpeg 2-pass operation can be functionally duplicated in Cinelerra in the steps below them:

```
ffmpeg -y -i input -c:v libx264 -b:v 2600k -pass 1 -c:a aac -b:a 128k -f mp4 /dev/null && \  
ffmpeg -i input -c:v libx264 -b:v 2600k -pass 2 -c:a aac -b:a 128k output.mp4
```

1) After you have completed your editing, do a Save Session with File → Save as...

Before starting, be sure your session is ready for batch render.

That is, positioned at the beginning and nothing selected.

2) Bring up File → Batch Render... where you will do the setup.

3) Click on the “Delete” box to remove old jobs highlighted in the bottom listbox.

For the “File Format” choose ffmpeg and mp4 for the type.

Set “Output path” to the path and filename for the render output file.

Click on “Use Current EDL” to use the designated EDL Path file.

Click on “New” and you will see a new highlighted job show up in the listbox at the bottom.

Use the Audio wrench to set bitrate to 128000 (128k as in ffmpeg example above).

Click checkmark OK. Open the video tools with the video wrench.

Set the Video Compression to h264.mp4 (as seen in the example).

Set the bitrate to 2600000 (2600k as in ffmpeg example above).

Add the following 2 lines after the first line:

```
flags +pass1
```

```
passlogfile /tmp/{temporary log file name}.log
```

Click checkmark OK.

4) Click on “New” to create the second pass job. You will see this second job in the listbox below.

Use the Video wrench and change pass1 to pass2 as follows.

```
flags +pass2
```

5) Click checkmark OK.

6) Click on the “Start” box and watch it go!

7) You can now check the output file for results. At the time this was documented, “rc=2pass” will be in the output.

If you need to re-render this, the Batch Render will still be set up but you have to click on the “Enabled” column in the listbox to re-enable the jobs to run which puts an X there. Click Start again. You can reuse batch job using the “save jobs” and “load jobs” buttons in the batch render dialog.

Render shortcuts for webm, h264, h265 are available by using the option files that are already set up for this purpose. Use the render menu as usual, with ffmpeg/mp4, choose h264 or h265 pass1of2_h26x for the video and passes1and2_h26x for the audio; with ffmpeg/webm, choose pass1of2_vp9. When that is finished, you will have to use the render menu again and this time for video, choose pass2of2_h26x or pass2of2_vp9. The logfile is hard coded in the options file so will write over any currently existing logfile if you do not change it before you start the render.

Requirements for some other libraries are below (used instead of “flags +pass1 & passlogfile”):

x265:

add this line: x265-params=pass=1:stats=/tmp/{temporary log file name}.log
at the time this document was written, you should see in the output: stats-read=2

libvpx-vp9, xvid, and huffyuv: NOTE for vp9, the best Pixels is gbrp.

cin_stats_filename /tmp/{temporary log file name}.log
flags +pass1 (or flags +pass2 for the second pass)

7.5.5 Piping Video to a Command Line

You can pipe a video to any command line on the computer, such as ffmpeg. This can be especially useful with raw video files. Next is an example usage.

1) on a terminal window create a named pipe file, for example:

```
mkknod /tmp/piper.yuv p
```

2) load your video and do your editing

3) set up your Render (Shift-R), you can choose a raw format such as “yuv” or “rgb”

4) for the filename “Select a file to render to”, use the named pipe as created in step 1 (**/tmp/piper.yuv**)

5) for “Insertion Strategy”, you will want to make sure to select “insert nothing”

6) click for OK on the green checkmark

The cinelerra gui will look like it is hanging while waiting for a command line to use the pipe.

7) on the terminal window, keyin your command, for example:

```
/mnt0/build5/cinelerra-5.1/thirdparty/ffmpeg-3.4.1/ffmpeg -f rawvideo -pixel_format yuv420p \ -  
video_size 1280x720 -framerate 30000/1001 -i /tmp/piper.yuv /tmp/pys.mov
```

A slightly different option can be used instead that may be more familiar to some. In the render menu after choosing the File Format of “ffmpeg”, use the pulldown to choose “y4m” as the file type. This choice results in putting a header on the rendered output with some pertinent information that can be used for ffmpeg processing thus alleviating the requirement for pixel_format, video_size, and framerate on the ffmpeg command line. In this case the format is yuv4mpegpipe instead of rawvideo. An example command line would look as follows (assuming the created pipe is called piper.y4m):

```
ffmpeg -f yuv4mpegpipe -i /tmp/piper.y4m -vcodec libx264 /tmp/test.mp4
```

7.5.6 Faststart Option for MOV type files

If you have mov video and want to be able to start playing without having to first load the entire video, -movflags+=faststart is needed for ffmpeg to put the meta-data, known as the “moov atom”, at the

beginning of the file. Otherwise, ffmpeg puts this atom at the end of the video file which means you have to wait to play until the whole video is loaded. Or worse yet, if the file becomes damaged in the middle and you can not get to the end, you won't be able to play anything.

Now you can have the "moov atom" put on the front of the file (automatically via a second pass). To do this, when rendering using ffmpeg & either the mp4 or qt format/container, click on the video/audio wrenches and choose faststart_h264. With the "qt" format, settings will just be the default whereas the "mp4" format uses the highest quality and lowest file size as possible, but you can easily modify these options in the associated Video Preset textbox.

This page intentionally left blank so that chapters begin on a right facing page.

8 - Keyframes

The word "keyframe" has at least 3 contextual meanings in the NLE environment. First, the oldest meaning, is the "I-Frame" definition used in codecs algorithms. These are "key" frames that begin a new sequence of pictures, and are anchor points for repositioning (seeks). Next are the automation parameter data points. These are usually input to primitive math forms, like translation and zoom. And last are blobs of data that are chunks of parameters to plugins that can do almost anything. The data can be a simple value, like a fader value, or more complex like a group of points and colors in a sketcher plugin keyframe. The word keyframe has changed a lot in meaning. In the context of Cinelerra, keyframes are data values that have been associated to the timeline which affect the media presentation. So a keyframe no longer refers to a frame, but to a position on the timeline.

In Cinelerra, there are two general types of keyframe data, *automation keyframes* (autos) which are drawn as colored lines and box icons overlaid at a point on a media track, and *plugin keyframes* which are drawn as gold key symbols on a plugin bar of a track. Auto ← Keyframe → Plugin 

8.1 Automation Keyframes / Autos

The "autos" are created by clicking on an "automation curve" to establish the time position for the new keyframe anchor point. The basic nature of these simple auto values make them primitive operations that are easy to apply when needed.

There are many automation curve types, and most are not normally visible or clickable. To make them visible, use the "View" pulldown, or open the "Show Overlays" window in the "Window" menu pulldown. This window allows toggling of the parameters in the View pulldown but is more convenient because you can leave the window up to change values quickly. If all of the automation curves are turned on, the timeline will be quite cluttered, and so usually only the parameters of interest are enabled during use. When keyframes are selected, they are drawn on the timeline over the tracks to which they apply. The keyframe is represented on the timeline as a little square on the curve, for example as in fade, or as a symbol as in a mask. This square, timeline attachment point, can be used for positioning by clicking on a keyframe anchor and using drag and drop to set the new position.

The automation keyframes include:

mute/play audio, camera translation x,y and zoom, projector translation
x,y and zoom, fade blending, audio panning, overlay mode, mask point sets, and sampling speed

Except for the mask auto, the values are all simple numbers. Mute is different from the other autos in that it is simply a toggle of either on or off. Mute keyframes determine where the track is processed but not rendered to the output. An example usage would be to use auto keyframes to fade in a clip by setting the transparency to 100% at the first keyframe and adding another keyframe 5 seconds later in the timeline with a transparency of 0%.

The Keyframes pulldown on the main timeline is used for Cut, Copy, Paste, Clear, Change to linear, Change to smooth, Create curve type of Smooth, Linear, Tangent, or Disjoint, Copy default keyframe or Paste default keyframe. If you right click on a curve keyframe on the timeline, a set of options popup including the choices "keyframe type" (such as Fade, Speed, etc.), Hide keyframe type, Delete keyframe, Copy keyframe, smooth curve, linear segments, tangent edit, or disjoint edit.

Usually, the use of the keyframe values are more pleasing when the data varies smoothly between keyframe anchors on the timeline. This is useful in many cases that are familiar, like a video fade in/out, or audio pan between channels. To make the auto value change smoothly as the media is played, the keyframes auto values are points on curves that are created according to the design of the effect. Most of the primitive types can create anchor points on curves that are piecewise linear, smooth, sloped, or broken at the keyframe anchor points.

Curve smoothing is called "interpolation" and it uses keyframe point values and control values that determine how the curve will react at the time the media is played or rendered. Interpolation uses 2 keyframes to create a set of intermediates which are used as active values between the "previous" and "next" keyframe anchors on the timeline. The way the intermediate data is generated depends on the type of curve used to invent these values. Cinelerra interpolates the intermediate values making the change happen smoothly and gradually over time. The simple linear mathematical formula for interpolation is: $a*(1-t) + b*t$ where $0 \leq t \leq 1$ uniformly.

8.1.1 Using Autos

The first click on the curve, creates a keyframe which you can click drag on to reposition. The second click at a later position, generates the smoothing by creating a smooth ramp. Ctrl-dragging on a keyframe round control point handle changes the value of either the input control or the output control. This affects the sharpness of the curve. While the input control and the output control can be moved horizontally as well as vertically, the horizontal movement is only for legibility and is not used in the curve value. When you Shift-drag on a timeline curve, the keyframe snaps to the value of either the next or previous keyframe, depending on which exists. It will snap up or down depending on direction of movement. This lets you set a constant curve value without having to copy the next or previous keyframe.

To make it easier to navigate curve keyframes, since there is not much room on the timeline for a wide range of curve values, you need to zoom the curves in and out vertically to have any variability. This is done by 2 tools: the automation fit button, Alt-f, and automation zoom menu which is seen at the bottom of the main window. The automation fit button scales and offsets the vertical range so the selected curve area appears in the timeline. If a region of the timeline is highlighted by the cursor, only that region is scaled. In/out points do not affect the zoomed region. The automation zoom menu manually changes the vertical scaling of the curves in multiples of 2. Click on its tumbler to change the zoom. Alt-up arrow and Alt-down arrow change the automation zoom from the keyboard.



Figure 38: Automation zoom menu with tumbler

Other mouse actions have the following effects:

- Double left mouse click on a curve Fade or Speed line will create ganged keyframes so that there is a keyframe on each of the tracks in the exact same position.
- Left mouse click on a keyframe position will show the numerical value in a yellow tooltip-like box.
- Right mouse click on the curve type line will bring up the option of "Hide keyframe type". This provides the same functionality as disabling the keyframe type in the View pulldown menu. Often it helps to use this in order to be able to see other things on the timeline once it gets cluttered.

- Fade and speed allow for setting a specific value for the keyframe using the following:
 - 1) click using the right mouse button on one of the auto speed keyframes on the timeline;
 - 2) a popup menu comes up with speed or fade auto type as the first menu item;
 - 3) click on speed or fade and a colored slider bar will appear (default speed=yellow; fade=purple);
 - 4) click the slider, press and hold the left mouse button and move the slider to update the value or simply type in a value in the textbox followed by pressing the Enter key or click the checkmark;
 - 5) a tooltip shows the keyframe value;
 - 6) release the button, the slider will be deleted from the canvas, and the value will be updated.

You can click mouse button 3 on a keyframe box and a menu pops up with the first menu item showing the keyframe type. The top menu item can be activated for immediate access to update the automation keyframe value. Some keyframe types, which have values that can be manipulated in another way than by dragging the color coded line, now show up with a different colored background to make them more visible. **Keep in mind** that Zoombar ranges/values must be set to appropriate values when working with specific keyframe types, such as Fade or Speed. If you do not see the auto line in the visible area of the video track, try the key combination Alt-f or select the speed in the "Automation Type" drop-down menu at the bottom of the main window. To the right of this field is "Automation Range" where you can set the display ratio of these lines. Simply change the values until the lines are visible again.

Screenshot below shows several color coded lines for different keyframes and specifically the slider bar for the Fade keyframe. It is in the same color as the color coded keyframe type line which is the same color which would be shown in the "Show overlays" window.



Figure 39: Fade curve with pink colored box Keyframe and Slider bar

In the "Editing" section of Settings → Preferences, the Interface tab there is "Keyframe reticle" with options of Never, Dragging, or Always. This is used to help in checking edit alignment across tracks. (A reticle is a sighting line used to line up visual items, like cross hairs in a eyepiece.) The appearance and function of sighting lines can be changed when dragging auto keyframes. To see the effect, create some fader autos and drag a few to see the reticles drawn – you will see something similar to the next screencast. "Always" renders a line over all plugins, and "dragging" only over the drag icon. "Never" draws nothing. The default is "dragging".

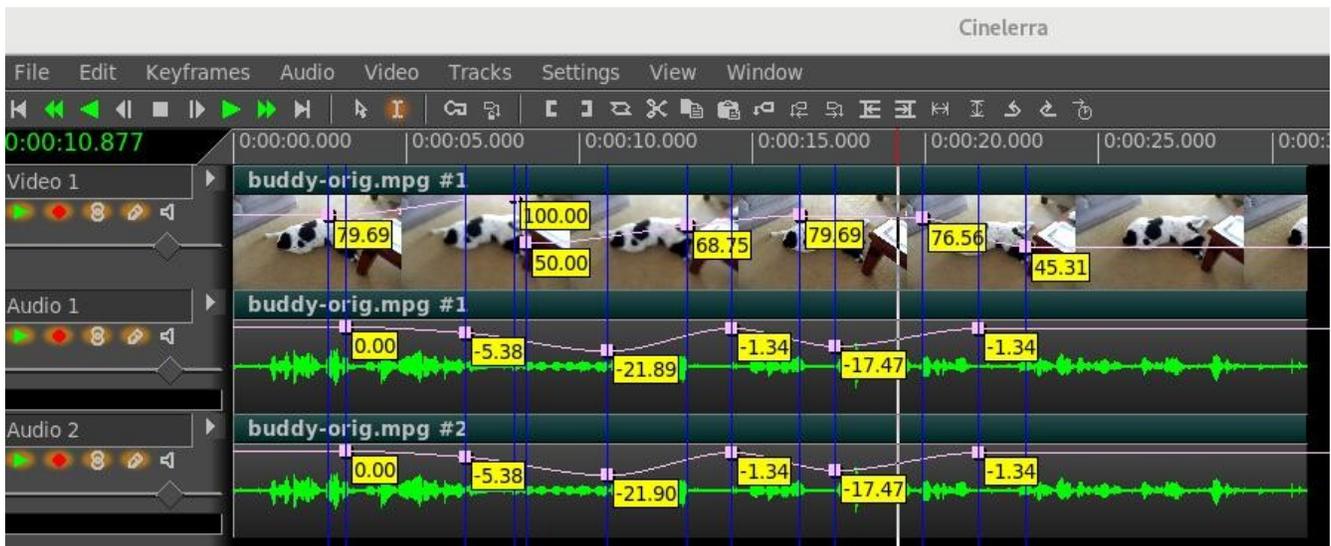


Figure 40: Fade Auto with "Always" set for easy to see numeric values

Control points allow for setting the slope of auto curves and then subsequently adjusting that slope. To modify a current keyframe you just right mouse it and change to either Tangent or Disjoint edit. In the screencast to the right, the Fade Auto has pink colored curves and control points are seen as dashed lines next to the keyframe box with black filled circles on each end of the line. Use the Ctrl key with the left mouse button to modify the control point lines.

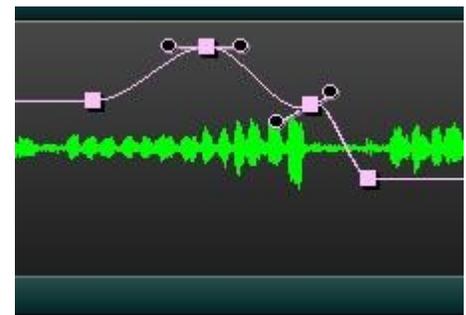


Figure 41: Fade Auto with Controls

8.1.2 Speed / Fade Automation Usage and Auto Gang

Speed automation resamples the data at a higher or lower playback rate. Speed automation can operate on all tracks of the same type, either video or audio, with a single click; or all tracks, both video and audio, with a double click. The curves to be affected must be both armed and ganged in the track patchbays to be included in the operation. You can see the advantage of having different media types ganged the same so that they will remain synchronized. However, if as recommended you have Settings of "Align cursor on frames" set, video will always snap to a frame, whereas audio does not. Because the smallest pieces are audio and not video, when video is present in order to maintain

synchronization be sure to do the double click on the video track. The single/double click feature also works in the same manner for Fade Automation.

The conversion algorithms that calculate the duration are approximate, not exact. Speed is not checked when inserting media on a track that already has speed adjustments so duration will not be corrected in that case. You can not do negative time and you can not play backwards.

Steps to demonstrate Auto Gang of Speed/Fade on all of the audio and video tracks are listed below:

- 1) Use the pulldown menu "View" on the main canvas and ensure that there is a checkmark left of the Speed or Fade selection.
- 2) Double click and hold button 1 on any point of the Speed line on any track of the main track canvas and drag the handle where you want it, then release button 1.
- 3) Note how all video and audio boxes move together simultaneously for synchronization.
- 4) You can adjust the sampling rate (speed/fade) at any time by doing the same single/double click on a keyframe.
- 5) To move the location left or right, hold down button 1 while moving it back and forth.
- 6) Use the status bar in the lower left hand corner to see position and playback rate.

Releasing button 1 ends the current dragging operation, but single clicking on a speed/fade keyframe handle again will restart the drag operation on tracks of the same type. Double clicking will select all armed tracks which have handles in the exact same cursor position. It is sometimes difficult to get the desired value for the speed and it may take a few tries to get to the desired endpoint, but you can use an auto slider bar to get better control for more precision. This is not ganged so you will have to do this on each track to achieve the same value.

An easy way to get an exact position is to set the "Automation range" in the bottom bar of the main window to the low and high ends of the desired range. For example, 1.000 to 50.000, which makes it so you can drag the speed/fade handle all the way to the left for 1.000 or all the way to the right for 50.000. Screencast to the right shows yellow keyframes and lines for the ganged speed automation on all video/audio tracks. Similarly the purple keyframes and lines show the ganged fade automation. Note the status bar numbers in the lower left-hand corner displaying 22.319 seconds as the location and 18.00 as the playback rate.



Figure 42: Ganged Speed Auto on all tracks



Figure 43: Speed auto with yellow Slider bar

Screencast to the left shows the yellow speed slider bar triangle with the current value displayed via the tooltip.

One other helpful item is a small colored indicator in the zoombar "Automation Type" box when you have clicked on an auto line. It will be the same color as that line so you can easily see out of the corner of your eye what that line refers to. This is especially useful for the Camera and Projector X,Y,Z lines which are often right on top of each other. In addition the status bar will have the colored indicator to the left of the last auto value when that status is visible.

8.2 Plugin Keyframes

The Plugin keyframes are structured. The individual data values are named parameters to the keyframe function. For example, the hue plugin has keyframe parameters of hue, saturation, and value. Each plugin has its own parameters, and what they do depends on the plugin. Most of the time, it is pretty obvious what the value controls, like the audio gain plugin with the level parameter. Some of the

plugins have a wide variety of controls, like the titler which can setup a wide number of controls, like formats, fonts, styles, placement, and so on.

Plugins also may use interpolation to smooth the data between the keyframe parameters on the timeline. Just exactly how this is done varies with the plugin and is not always easy to predict, but this is usually due to the nature of the keyframe parameter data. For example, it does not make sense to interpolate which font is in use.

8.2.1 Default Keyframe

For plugins, there is a special hidden keyframe, called the "default keyframe", that is used when no previous keyframe exists. It is like keyframe zero on the timeline, and it is persistent and shared on all sessions. The intent is to make a parameter set that is likely to be reused on all initial instances of the plugin. An example may be to color correct a set of media that was taken in low light, or needs resampling to be played correctly. The default keyframe is "off the bar on the left side" of the plugin title bar and can not be seen. It is used when there is no "previous" keyframe for its default values.

It may be useful to create a default keyframe which has specific desirable values for later use. To do this, set the timeline to position 0 and be sure to disable "generate keyframes while tweaking". This will create a default keyframe at the beginning of the timeline which contains global parameters for the entire duration. Or if you have copied a non-default keyframe via Keyframes pulldown "copy default keyframe", it can be stored as the default keyframe by calling keyframes->paste default keyframe. After using paste default keyframe to convert a non-default keyframe into a default keyframe, you will not see the value of the default keyframe reflected until all the non-default keyframes are removed.

The keyframes->copy default keyframe and keyframes->paste default keyframe allow conversion of the default keyframe to a non-default keyframe. Keyframes->copy default keyframe copies the default keyframe to the clipboard, no matter what region of the timeline is selected. The keyframes->paste keyframes function may then be used to paste the clipboard as a non-default keyframe. However, typeless keyframes are not described

Typeless keyframes enabled under the Settings pulldown allow keyframes from any track to be pasted on either audio or video tracks. Ordinarily audio keyframes can only be pasted to another audio track and video keyframes can only be pasted to another video track.

8.2.2 Keyframe “Edit Params” for Plugins

Keyframe values can be set using the various plugin “show controls” (magnifying glass) icon on the plugin track. It is possible to see all of the keyframe data in a raw format using the *Edit Params* popup menu item which you will see when you right mouse click the keyframe icon on the timeline. The keyframe data is stored in “xml” format, and the “edit params” feature allows you to view and modify the xml data directly. This normally should not be necessary since the plugin's control gui displays the intended parameters, but this will let you view and specify just about anything that can be specified in xml. There is no validation checking of the modified data.

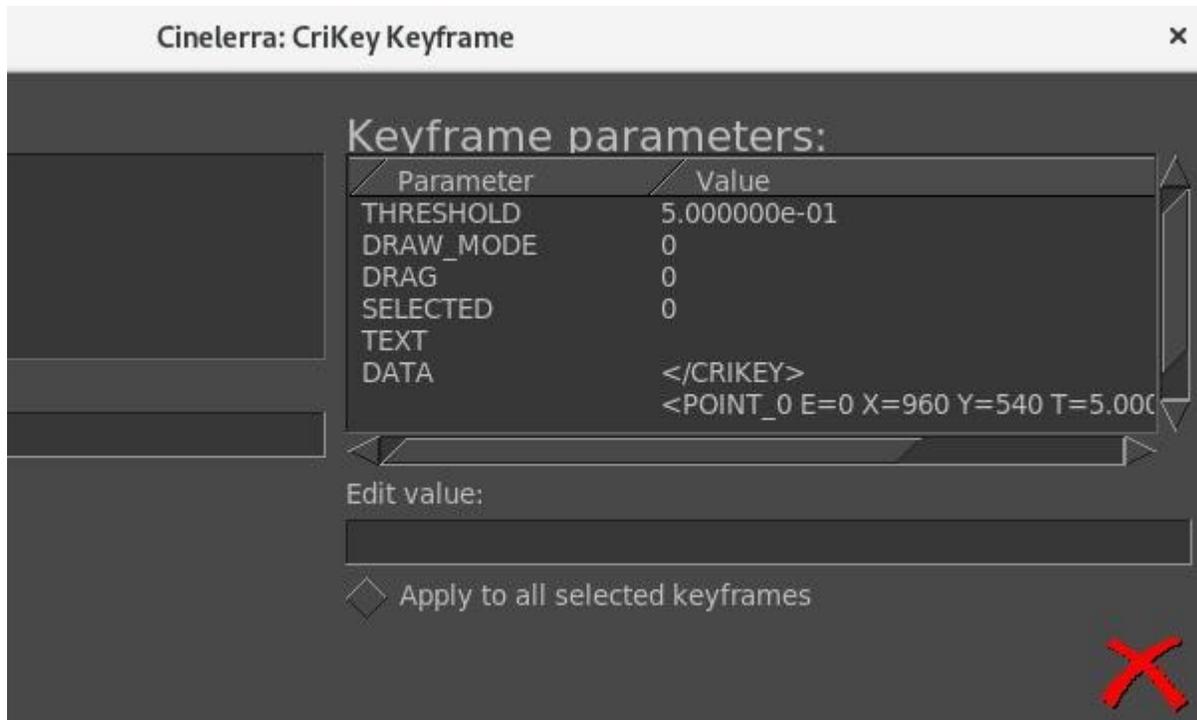


Figure 44: "Keyframe Parameters" window for a CriKey plugin added to the video track

8.3 Generate Keyframes while Tweaking / Automatic Keyframe Mode

Tweaking is defined as changing a parameter while playing – it can be small changes or small motion increments performed in a series. These changes are recorded as a series of new keyframes on the timeline. Enable automatic keyframe mode by enabling the automatic keyframe toggle, that is “Generate keyframes while tweaking”. In automatic keyframe mode, every time you tweak a keyframable parameter it creates a keyframe on the timeline. Since keyframes affect playback, you should enable generate keyframes just before you need a keyframe and disable when your parameter changes are complete. So turn on when ready to make changes and turn off when done!

Also, before making a change, be sure to check the View pulldown and make the desired parameter visible with the enable checkmark. The location where the automatic keyframe is generated is under the insertion point. If the timeline is playing back during a tweak, several automatic keyframes may be generated as you change the parameter. When automatic keyframe mode is disabled, adjusting a

parameter adjusts the keyframe immediately preceding the insertion point. For example, if two fade keyframes exist and the insertion point is between them, changing the fader changes the first keyframe.

8.4 Compositor Keyframes

Camera and projector translation is represented by two parameters: x and y, making it difficult to adjust with curves. Cinelerra solves this problem by relying on automatic keyframes. With a video track loaded, move the insertion point to the beginning of the track and enable automatic keyframe mode. Move the projector slightly in the compositor window to create a keyframe. Then go forward several seconds. Move the projector a long distance to create another keyframe and emphasize motion. This creates a second projector box in the compositor, with a line joining the two boxes. The joining line is the motion path. If you create more keyframes, more boxes are created. Once all the desired keyframes are created, disable automatic keyframe mode.

Dragging the auto curve with tweaking off, adjusts the previous keyframe. With tweaking on, if no keyframe is at the timeline position, then a new one is created at that position and the modification value is applied. If you are halfway between two keyframes, the first projector box is adjusted while the second one stays the same. The video does not appear to move in step with the first keyframe. This is because halfway between two keyframes, the projector translation is interpolated. In order to set the second keyframe you will need to move after the second keyframe.

Image translation, motion direction, and speed determine the results. Motion varies based on the interpolation type that is set, for example, if set to Linear, the result will be uniform straight line motion. Smooth is the default so that interpolation will generate curved lines with control points. Ctrl-drag the endpoint of the control handle to adjust the curved slope.

8.5 More about Editing Keyframes

Keyframes can be shifted around and moved between tracks on the timeline using similar cut and paste operations to editing media. Only the keyframes selected in the View menu are affected by keyframe editing operations.

An often used, keyframe editing operation is replication of some curve from one track to the other to make a stereo pair. The first step is to solo the source track's record patch by Shift-clicking on the "arm track" icon in the patchbay. Then either set In/Out points or highlight the desired region of keyframes. Go to keyframes->copy keyframes to copy them to the clipboard. Solo the destination track's record patch by Shift-clicking on it and go to keyframes->paste keyframes to paste the clipboard. Another common application for keyframe modification is to highlight a region on the timeline which contains multiple keyframes that you want to modify. Then when you adjust a parameter or set of parameters, the change will be applied to all keyframes within the selection instead of a new keyframe being created. This only works when the keyframe stores multiple parameters and only for mask and effect keyframes. Other types of keyframes are generated as usual.

And there is an easy way to delete keyframes besides selecting a region and using keyframes->clear keyframes. Click-drag a keyframe before its preceding keyframe or after its following keyframe on the track. This is the only way you can simultaneously delete keyframes on ganged tracks.

9 - FFmpeg Interactions

Cinelerra-GG uses ffmpeg for decoding and encoding media, thus there are many opportunities available to manipulate options.

9.1 FFmpeg Early Probe Explanation

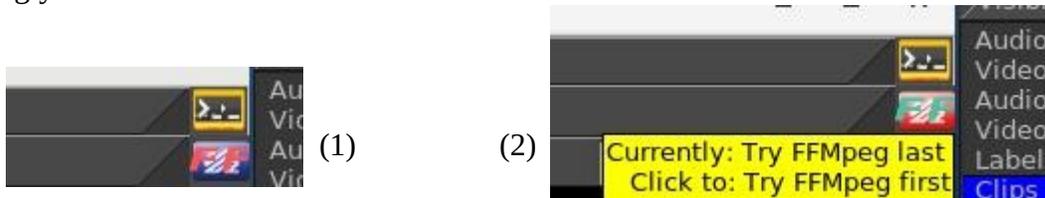
When you open media, a series of libraries and codec functions are used to "probe" the data, to see if it can determine the type of file format and codec parameters needed to correctly decode the file. If ffmpeg probes early - "Try FFMpeg first" is in effect for the button - it will usually find some way to try to decode just about any contemporary media file. But there are some times that the built in codecs are actually a better choice. A lot of this may fall into the category of personal preference. For example, some may prefer the mpeg library in the cinelerra code over the ffmpeg code because it has more decoding capability and seems to be more robust when the media is damaged. In that case you will want the *FF* button to read "Try FFMpeg last" in the upper right hand corner of the main window.

So, if ffmpeg probes early, you will never get to use the built in libraries, and if you want to skip over buggy old libraries, use ffmpeg early probe enabled so that the newest code will be tried first.

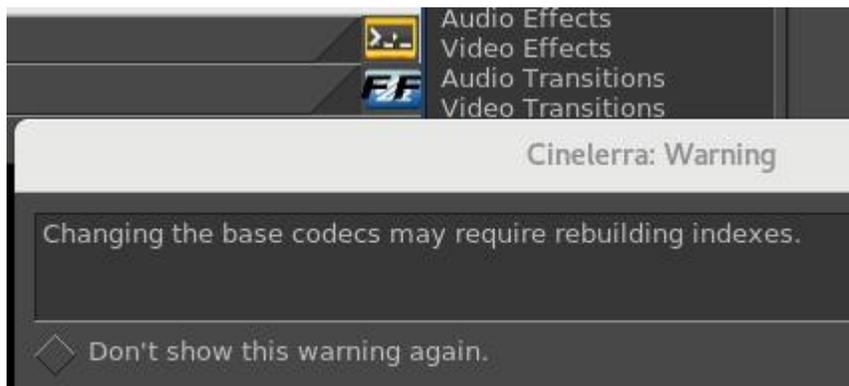
When the icon is red, ffmpeg probes early is enabled and you will see it reads "Currently: Try FFMpeg first" when moving over the *FF* button in the upper right hand corner of the screen. When the icon is black, ffmpeg probes early is disabled so that ffmpeg probes late and it reads "Currently: Try FFMpeg last". The initial default state of the icon is on, that is, ffmpeg probes first. This is the original code behavior before ffmpeg code was added. Suggestion is to leave it on except in a few special cases where it may be better to have early probes disabled. When you mouse over the main menu ff toggle button, the text displays ffmpeg's "Currently" set position and "Click to" change it.

The ffmpeg early probe state is saved between sessions and is also affected by choices made in Probe Order (in another section). It is important to note that the various file indexes may need to be rebuilt if you change which codec is being used to decode the file. There is a warning popup to remind you when you change the default ffmpeg early probe state (unless you have checked the box to no longer show you the warning). You can easily rebuild the index for a specific media file by going to the Resources window, right mouse click on that media, and choose "Rebuild Index" from the popup choices.

Screenshots below show (1) reddish colored *FF* in upper right hand corner of main window indicating that ffmpeg early probes is enabled; (2) “Try FFMpeg last” indicator message for ffmpeg early probes enabled (note that the color is different because you highlighted the icon); and (3) black colored *FF* indicates ffmpeg will be used last and you are changing the behavior so that Cinelerra warns you accordingly.



(3)



9.2 How to Create FFmpeg Options Files

This section describes how the FFMpeg options files work for decoding and encoding and goes into great detail. It will make more sense if you look at cinelerra's ffmpeg config directory and the cinelerra menus at the same time. It is meant to include everything necessary for complete understanding. You will be able to personalize your own options files without knowing all of the information included below if you know the basics. The word encoding is used interchangeably with the word rendering.

The possible combinations for ffmpeg options files are literally combinatorial – that is a lot (factorial!). The allowed media file format / codec choices are much more flexible than you might realize. When the ffmpeg design was initially added, some parameter files which describe the choices which the program uses had to be created. There are way too many to enumerate in the deliverable Cinelerra package. Some quite detailed information for how ffmpeg options work is given here and hopefully, enough basics for simple understanding. It may all seem complicated at first, but will become obvious.

9.2.1 File naming convention

In Cinelerra's ffmpeg configuration directory you will see files as listed and described below. File type and extension names are the key for Cinelerra's use of ffmpeg. Basically the .opts file extension represents options; .dfl represents defaults; and all the rest are media types. For example one media type is quicktime so that *.qt file names would be the “quicktime” choices. In the file names below, “ext” refers to a set of files with file names matching the *.ext file extension. And “typ” refers to a type of format / codec combination used, that is, the media type.

In the ffmpeg configuration directory there are a series of options files used when encoding or decoding audio or video. They are read in the order from top to bottom and only the files needed for the current operation are added to the active configuration.

```
ffmpeg/ffmpeg.opts = global ffmpeg options, always used
ffmpeg/decode.opts = global decoder options, used when opening existing files for decoding
ffmpeg/encode.opts = global encoder options, used when creating new files for encoding
ffmpeg/audio/audio.opts = audio encoder options, used when creating audio streams
ffmpeg/video/video.opts = video encoder options, used when creating video streams
ffmpeg/plugin.opts = parameters for ffmpeg filters as audio/video plugins
```

Decoder options:

Normally, only ffmpeg.opts and decode.opts are used when reading/decoding files, but may be specialized if a <path>/media.opts exists for a given <path>/media.ext file. For example, if you want to only fail on fatal errors and to always use the video filter, edgedetect, when working with your media file “dreaming.y4m”, then create a file “dreaming.opts” in the same directory with the contents of “loglevel=fatal” on the first line and “video_filter=edgedetect” on the next. These specialized settings will override the defaults. The fatal loglevel is especially handy for lesser quality media.

Encoder Options:

Within the audio /video subdirectories of the first level ffmpeg directory, the typ.ext files are for encoder (rendering) setups.

```
ffmpeg/audio          = directory of audio encoder settings
  audio.opts          = options used by all audio encoders
  typ1.ext, typ2.ext, ... = are all *.ext type choices for encoding audio
  ext.dfl             = contains the default selection used when ext is first selected

ffmpeg/video          = directory of video encoder settings
  video.opts          = options used by all video encoders
  typ1.ext, typ2.ext, ... = are all *.ext type choices for encoding video
  ext.dfl             = contains the default selection used when ext is first selected

ffmpeg/format         = presets needed to initialize audio / video formats
```

9.2.2 Option File Format / Content

For the option files a specific format must be followed in creating the file content.

In typ.ext encoder parameter files, the first line is defined as:

```
muxer codec  
or muxer codec | bitstream filter [ bitstream filter options ]
```

where the | represents piping the codec data through the bitstream filter.

The rest of the lines in the file should look as follows:

```
# in column one is a comment  
id1 value1  
or id2 = value2
```

Only one equals sign is allowed and it is just for readability. There may be any number of id / value pair lines in a media definition, including zero. A typical line might be:

```
bitrate 4000000  
or bitrate = 5000000
```

There are 4 special id's recognized by cinelerra which cause special processing. They are:

```
duration      (overrides the probe duration when opening media for decoding)  
video_filter  (adds an video stream filter, eg. edgedetect, ... at the stream level)  
audio_filter  (adds an audio stream filter, eg. echo, ... at the stream level)  
loglevel      (sets the library logging level, as quiet, panic, ... verbose, debug)
```

All other id's should be in the ffmpeg documentation, and correspond to the global, muxer, and codec option names and values used by ffmpeg. For example to set the aspect ratio to 4:3, use:

```
aspect 4:3
```

Below shows an example: "decode.opts" which is used when the ffmpeg decoder is initialized.

```
# apply at init decode  
loglevel=fatal  
formatprobesize=5000000  
scan_all_pmts=1  
threads=auto
```

The encoder options you see in the cinelerra menus depend on the files in these directories, NOT THE CODE. If you add files, you will get to use more variety.

In the "cinelerra" directory, which contains the ffmpeg configuration folder, there are the choices the program uses. When you open an ffmpeg format popup dialog, the listbox contains all of the codec types which are identified by the file.ext extensions. Decoding has only a few options, since the

ffmpeg file probes determine most of the options by looking at the media being opened, but encoding media requires a lot of setup. Below are some of the folders and files used to determine the configurations used by ffmpeg to decode and encode files.

These extensions create audio / video media classes:

```
dvd m2ts mkv mp3 mp4 mpeg qt pro
```

which become the choices in the render pulldown menu

So if you want to create a "mov" codec class, add two new files to the ffmpeg configuration directory:

```
audio/aud.mov and video/vid.mov
```

Now you will see this as what you can choose in the rendering choices for ffmpeg.

Inside the file you will see that the first line is special. It is the muxer and codec. For example:

```
h264 libx265
```

The contents may be something like: <path>/video/vid.mov =

```
mp4 libx265  
bitrate 4000000
```

This will code an mp4 formatted file using the lib264 codec encoder.

For audio and video together, **the mux format must agree** between the aud.mov and vid.mov files when they are to be used together. The stream muxer must be the same for the all the streams in the file being written.

for example: <path>/audio/aud.mov =

```
mp4 pcm_mulaw
```

This will create mp4 media using audio format pcm_mulaw coding.

Both the audio and the video are using mp4 mux format, and so there will be 2 streams:

- 1) x265 video
- 2) pcm_mulaw audio

When the menu popup is created, there may be many choices for that class type, so you may want defaults. That can be specified as:

```
audio/<class>.dfl and video/<class>.dfl  
so: audio/mov.dft =  
    aud.mov  
and: video/mov.dft =  
    vid.mov
```

The above will be the default choice when the menu opens.

When you see problems in using the new options files you have created and put into place, add the following line to ffmpeg/encoder.opts:

```
loglevel=verbose
```

sometimes that will be enough to see what has caused a failure, or even catch unexpected results.

Now there is an EXCEPTION to all of the above because of a conflict between ffmpeg and the x264 person making the detection of default ffmpeg settings terminate with an error. If you get this error, you must work around this termination by including parameters that don't match 5 or more of the normal expected values. So you just have to change a few parameters to avoid the probe detection. Here is an example where you will notice the x264opts line tweaking values to throw off the detection/error termination code.

```
<path>/ffmpegvideo/test.mp4 =  
  mp4 libx264  
  preset=slow  
  x264opts keyint=25:min-keyint=4:qpmin=3:qpmax=33:qp_step=4:merange=8  
  crf 20
```

Another EXCEPTION is for some forms of aac, which is the reason bitstream filters were added. If you think that you are in this situation, use:

```
<path>/ffmpeg/audio/test.mp4 =  
  
  mp4 libfdk_aac
```

For more examples, look around the ffmpeg directory for examples which may be close to what you are trying to use, and see if the parameters look usable.

This is quite complicated, but that is because ffmpeg has a lot of parameters and history. Good results are not that hard to create. Initially you should mostly use the defaults. If you send any new options files to cin@lists.cinelerra-gg.org, it will be given consideration to being added to the baseline for future deliverables.

To get a listing of the current ffmpeg supported formats and codecs that can be made to work with cinelerra, provided there are option files added, run the following commands. This should be done from the <build> directory substituting the location of <build> where you have installed cinelerra on your system and the ffmpeg may be a different version than 4.1.1 as used below. Then look at the output created in /tmp/ff-formats.txt and codecs.txt.

```
/<build>/cinelerra-5.1/thirdparty/ffmpeg-4.1.1/ffmpeg -formats > /tmp/ff-formats.txt  
/<build>/cinelerra-5.1/thirdparty/ffmpeg-4.1.1/ffmpeg -codecs > /tmp/ff-codecs.txt
```

9.2.3 Complete Options File Example

For illustrative purposes, here is an example of the options files that need to be added for using the ffmpeg ProRes 422 format. This makes it possible to transcode to h264.mov with FFmpeg retaining 10-bit yuv422p from the source to the target output video.

Add the file named “./ffmpeg/audio/acc256k.pro” which contains the following lines:

```
mov aac
strict -2
b 256000
```

(Note that in the example above, even though the bitrate is set here to 256000, it can be overridden by the render menu settings).

Add the file named “./ffmpeg/audio/pro.dfl” which contains the following lines:

```
acc256k.pro
```

Add the file named “./ffmpeg/video/med422p10.pro” which contains the following lines:

```
mov prores
preset medium
pix_fmt yuv422p10le
```

Add the file named “./ffmpeg/video/pro.dfl” which contains the following lines:

```
med422p10.pro
```

Then to use and to get 10 bit depth and preserve depth from decode to encode:

- 1) load media
- 2) use settings->format to set the frame rate, sample rate/channels, aspect ratio, color model = rgb_float or rgba_float if blending
- 3) press Shift-R and select FFmpeg format type “pro”
- 4) select target path
- 5) check ok, and watch for messages in the terminal window

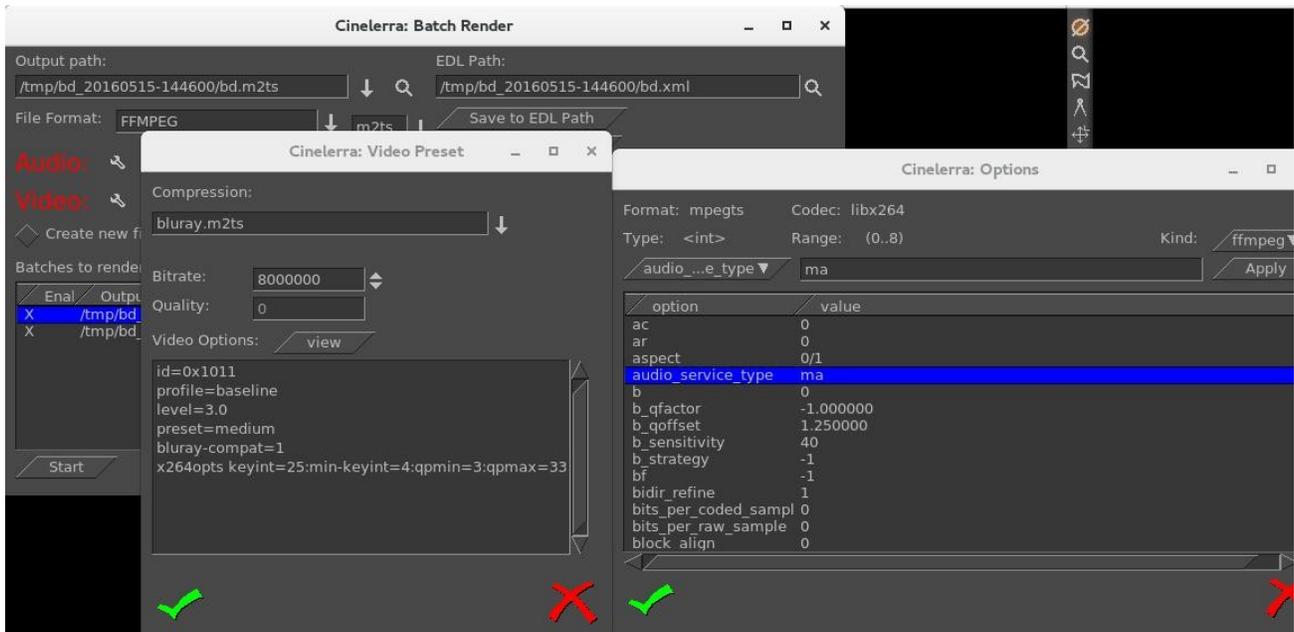
9.2.4 Viewing & Modifying FFmpeg Format Options inside Cinelerra

There are thousands of options for using ffmpeg. Now it is possible to “view” the available options for a particular video and audio choice by using the “wrench icon” and then clicking on the “view” box. Ffmpeg has to be the selected file format for this feature to be visible. It makes it a lot easier since only the applicable options show up as opposed to everything that ffmpeg can do. These options are just “Hints” and some may be missing due to the way that ffmpeg options are coded – Cinelerra shows the option data ffmpeg has exposed.

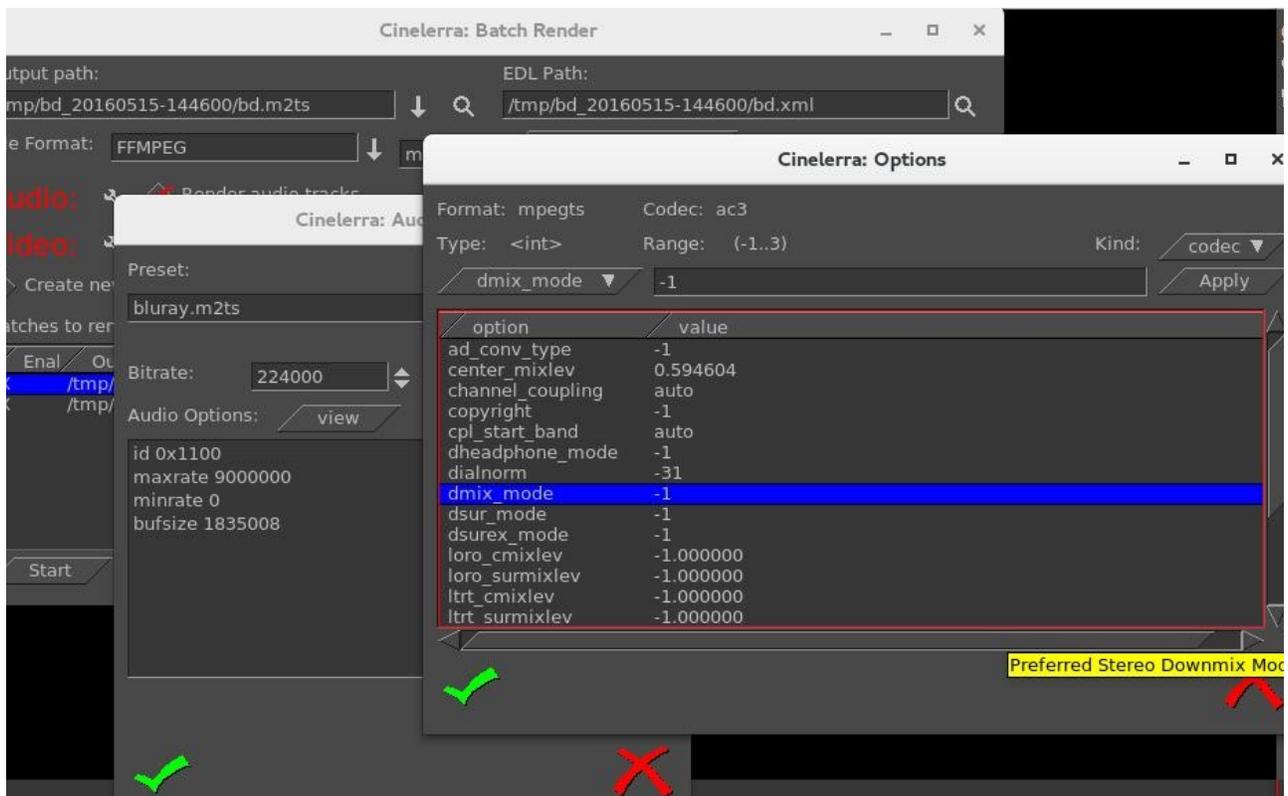
As an example, instead of reading the entire 264 library information, you only have to look at the shown available options. Both the video and the audio are browsable. The options visible in the Audio/Video Preset textbox are the final values which are used when rendering once you have checked OK. For assistance in choosing the options you want, use the view popup to see the objects that go with the selected format tool, highlight the option, modify the parameter value in the box at the top of the “Options” window based on what you want, and then click apply. Updates modify parameter values or new parameters will be appended at the bottom. Note that when you highlight an option, a tooltip will show up when available in the lower right hand corner which describes the option. Also note that the Format and Codec types are shown on the top line of the Options window.

Parameters exist in 3 layers: ffmpeg, codec, and an interface layer. You can apply parameters to each layer. The top 2 layers are accessed with the Kind popup menu. The ffmpeg layer is topmost, and is selected as Kind: ffmpeg. It can specify many of the more common parameters, such as the bitrate, quality, and so on. The middle layer is selected as Kind: codec. These options can specialize your choices, and frequently includes presets and profiles useful for coding well known parameter sets, like profile=high422, preset=medium, or tune=film, etc. The interface layer may or may not be available. It is usually accessible only by an “opts” parameter, like x264opts=key=value:key=value:... These options are passed directly to the low level codec library.

The 2 screenshots next show “ffmpeg” for as the Kind: for Video first and then “codec” Audio second. Note the x264opts in the Video Preset window immediately below.



Note the yellow tooltip in the lower right hand corner describing the highlighted option. Also note the allowed “Range” values above the box provided for keyins.



9.3 The FFmpeg Image2 Streams

Another feature gained from using ffmpeg in cinelerra takes advantage of what is being referred to as the “%d trick”. This trick uses the ffmpeg muxer image2 and a filename template to create a series of image files of a given kind. A specific example is described below.

To encode a series of 48 bit tiff output image files, add a file to the cinelerra data ffmpeg/video subdirectory as in:

```
filename: ../ffmpeg/video/tiff.dfl containing the line:  
tiff48.tif
```

Then create an ffmpeg video encoder parameters file in the same directory:

```
filename: ../ffmpeg/video/tiff48.tif containing the next 2 lines:  
image2 tiff  
pixel_format=rgb48
```

This will define a new ffmpeg encoder format which is a video image file format that uses the tiff codec for encoding, and a pixel_format of rgb48 (or a similar equivalent such as rgb48le). Next load up your project and set up for a Render using File → Render in the usual way. Now the tricky part; the output file name should contain a %d which will be the frame number used in the image output filename as in: Select a file to render to = /tmp/tiff_images/img%03d.tif. You will get multiple files as output – one for each frame!

The resulting directory of images can be opened for reading by simply opening the template path. As in: File → Load Files = /tmp/tiff_images/img%03d.tif. You will notice a file named the same as the template, which has been automatically created, is empty, is needed, and has to remain with the set.

9.4 Raw Input Opts File for Video/Audio

Raw video is not affected by decoding on read in. This makes it very attractive to provide raw image data for editing and rendering media. A wide variety of raw formats are available via the ffmpeg file interface. To load media in a raw format, select “try ffmpeg first” and create an accompanying “opts” file. The opts files must be in the same directory as your media, with the same base name, and the .opts extension. The opts file contents should reflect your video setup. An example follows:

Video file name = /tmp/buddy.mov

Opts file name = /tmp/buddy.opts

Contents of opts file =

```
format=rawvideo  
codec=rawvideo  
video_size=352x240  
pixel_format=yuv420p  
duration=90.25
```

9.5 FFmpeg Items of Note

Quality Option when rendering: Ffmpeg responds variably to the quality option in the render option but seems to respond well to bitrate. The subranges used by quality even seem to vary somewhat depending on how old the codec is. Some use 0 to 35, some use 0 to 500 or so. The quality is supposed to cause the codec to output data until the noise level is below a limit determined by the quality setting. Your specific results may vary.

Previous Changes when rendering: With ffmpeg there are 2 cases that the defaults will be used. The first time when you have nothing set up and any other time when you reset the render File Format in the Render Menu. Otherwise with ffmpeg if you change a video compression type for the render (for example h264.mp4 to h265.mp4), the settings will be from the previous session settings.

Outstanding Issues with ffmpeg: There are some problems that need to be addressed by the ffmpeg developer group that adversely affect Cinelerra. These are stated below with the hopes that that group will fix them as time permits.

- 1) Make all the default parameters operational. When they are not, the Cinelerra plugins can't be initialized since the initial state of the filter is not operational. If that is not possible, then provide a set of nominal parameters for each plugin, so that they can be used as the plugins initial default state.
- 2) Make the filter config function project the new parameter data into the filter function at any point during filter operation. This is so that continuous updates can be done as the plugin operates.
- 3) Improve seek codec restarts. The past predictor must be reset or reconstructed after a seek. The only documented way to seek is open/seek/play. Reopening the format layer is very expensive.

This page intentionally left blank so that chapters begin on a right facing page.

10 - Effect Plugins

There are realtime effects – these are the most useful and probably all you will ever need - and rendered effects. The rendered effects are discussed separately in the “Rendered Effects” section. Effect plugins modify the track when played, according to how they are set, with no permanent storage of the output except when the project is rendered. There are many Plugins in Cinelerra-GG Infinity which are actually quite easy to use just by experimenting with them. The plugins are shown and selected from the *Resources window*. They are described in more detail later.



Fig. 45: Screenshot of the native Video plugins in the default Cinfinity icon set.

There is a choice of plugin icons which can be displayed. In **Settings** → **Preferences** → **Appearance** tab, there is a pulldown for *Plugin icons* where the user can choose between the *original* icons, *regular* or *smoother*, *cinfinity* - the default modernized set, or *cinfinity2*.

(Cinfinity icon set is credited to Sam - Creative Commons By <https://creativecommons.org/licenses/by/3.0/>)

Note that when you change the plugin icons, your session will automatically save the backup, stop, restart, and reload.



Fig. 46: Cinfinity2 audio plugins

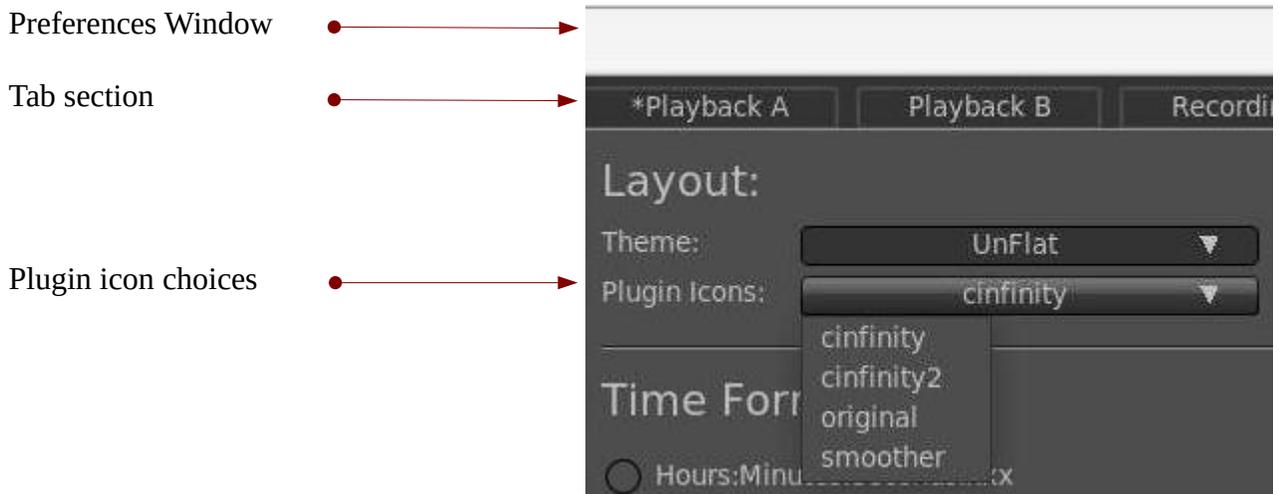


Fig. 47: Screencast showing the screen to change your plugin icons set

10.1 How to Use Plugins

Realtime effect plugins are listed in the *Resources* window as *Audio Effects* and *Video Effects*. Effect plugins are used by dragging them from the *Resources* window onto an audio track if it is an audio effect or video track if it is a video effect. You will see a colored bar appear beneath the track with the plugin name on it. If there is data on the destination track, the effect is applied to the entire track, unless a region of the track is selected in which case the effect is pasted into that region only. If there is no data on the track the effect is not added.

Plugins are layered under the track they apply to. When dragging more than one effect onto a track, you will see the effects layering from *top to bottom*, on the bottom of that track. When the track is played back, effects are processed from top to bottom. The output of the top effect becomes the input of the bottom effect and so on.

Instead of dragging from the *Resources* window, effects may be applied to a track via a popup menu. Right click on a track and select *Attach effect* from the popup. The attach effect dialog gives you more capability than just dragging and dropping. For example, the attach effect dialog lets you attach two more types of effects: *shared effects* and *shared tracks* which are explained in a later section. Select a plugin from the *Plugins* column and hit the green colored checkmark under the plugins column to attach it. The result is the same as if the effect was dragged from the *Resources* window.

After attaching an effect to a track, it often needs to be configured. There are two ways to get to the configuration controls. Click on the *magnifying glass* symbol on the right side of the effect bar - this is the middle symbol on the bar as you can see in the picture below. Alternatively, you can right click on the *effect* bar to bring up the effect popup which has a *Show* option. Either method causes the *GUI* for the effect to appear in a separate window. There will not be a popup if the plugin has no *GUI*.



Fig. 48: magnifying glass

Besides the *magnifying glass*, for *Show Controls*, on the effect colored bar beneath the track, there are two more symbols. The rightmost knob is used to **Turn Off/Turn On** the effect where the default is **On**. This is useful to easily see that the plugin is doing what you expect. The leftmost symbol that looks like a gear is for *Preset Edit* and its usage is described in the section *Saved Plugin Presets*.

10.2 Editing Effects

Many operations exist for manipulating effects once they are on the *timeline*. Because mixing effects and media is quite complex, the methods used in editing effects are not as concise as *cutting and pasting*. Some of the editing happens by dragging *in/out points*, some of the editing happens through popup menus, and some of it happens by dragging effects.

When enabled, which is the default, and you edit tracks, the effects follow the editing decisions. If you cut from a track, the effect shrinks. If you drag edit *in/out points*, the effect changes length. This behavior can be disabled by selecting **Settings** → **Preferences** → **Interface tab** → **Editing** section.



Fig. 49: Screenshot which shows the choices of Editing via the down arrow pulldown.

To edit effects, you can move the timeline cursor over the effect borders until it changes to a resize left or resize right icon. In this state, if you drag the end of the effect, it performs an edit just like dragging an edit edge. The three editing behaviors of track trimming apply to effect trimming and they are bound to the mouse buttons that you set in interface preferences as shown in the previous screenshot.

Trimming simply means changes the duration.

When you perform a trim edit on an effect, the effect boundary is moved by dragging it. Unlike track editing, the effect has no *source length*. You can extend the end of an effect as much as you want. Also unlike track editing, the starting position of the drag operation does not bind the edit decision to media. The media the effect is bound to does not follow effect edits. Other effects, however, do follow editing decisions made on an effect. You can disable effects from being subject to the edit decisions by using the pulldown *Settings* and toggling off *Edit effects*. If you drag the end of an effect which is lined up to effects on other tracks, the effects on the other tracks will be edited while the media stays the same. When you drag an effect in from the *Resources* window you can insert the effect in the portion of the row unoccupied by the trimming operation. In some cases you will want a trimming operation to change only one row of effects. This can be achieved by first positioning the *insertion point* on the start

or end of the effect. Then press the *shift* key while beginning the trimming operation. This causes the operation to change only one row of effects.

You can move effects up or down. Every track can have a stack of effects under it. By moving an effect up or down you change the order in which effects are processed in the stack. Go to an effect and right click to bring up the *effect menu*. The *Move up* and *Move down* options move the effect up or down. When you are moving effects up or down, be aware that if they are shared as *shared effects*, any references will be pointing to a different effect after the move operation.

Finally, there is dragging of effects. Dragging effects works just like dragging edits. You must select the *arrow* in the *main window transport buttons* line to enter *drag and drop* mode before dragging effects. Dragging a plugin causes a highlight outline to be drawn over a targetable *timeline* region, and the plugin can be re-positioned into any plugin track. The effects snap to media boundaries, effect boundaries, and tracks. If you drag a reference to a shared effect, the reference may point to the wrong effect afterwards. It is recommended that you re-construct shared effect track references.

Screenshot showing 2 plugins, one have already been dragged, and the other *Deinterlace* in the process of being dragged. Note the gold-colored arrow which enables *allow drag and drop editing mode*.

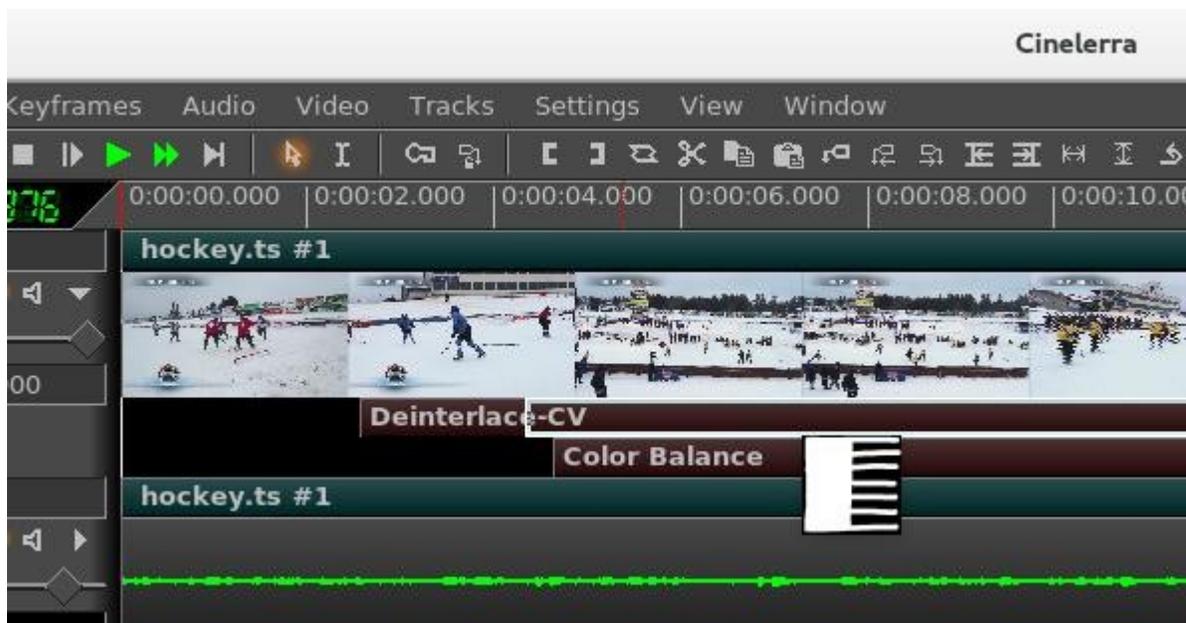


Fig. 50: Dragging the *Deinterlace-Cv* effect

10.3 Shared Effects and Shared Tracks

Two other effect types available in the *Attach Effect* dialog are *Shared effects* and *Shared tracks*. In the case of a *shared effect*, the following conditions must be true:

- ✓ There must be other effects in the timeline.
- ✓ The other effects must be of the same type as the track you are attaching an effect to. That is for *audio tracks*, effect must be audio and for *video tracks*, effect must be a video effect.
- ✓ The *insertion point* or selected region must start inside the other effects.

In the case of a *shared track*, there must be another track on the timeline of the same type as the track you are applying an effect to. If you right clicked on a video track to attach an effect, there will not be anything in the *shared tracks column* if no other video track exists. The same applies equally to audio tracks in that another audio track must exist. *Shared tracks* are often used as *adjustment layers*.

If *shared effects* or *shared tracks* are available, they appear in the shared effects and shared tracks columns when you used the *Attach effect* option (**RMB** on a track). When the green colored checkmark is clicked **OK**, anything highlighted in the column is attached under the current track.

Shared effects and *shared tracks* allow very unique things to be done. In the case of a *shared effect*, the *shared effect* is treated like a copy of the *original effect*, except that in the shared effect the *GUI* can not be brought up. All configuration of the *shared effect* is determined by the *GUI* of the *original effect* and only the *GUI* of the original effect can be brought up.

When a *shared effect* is played back, it is processed just like a normal effect except the configuration is copied from the *original effect*. Some effects detect when they are being shared. These effects determine what tracks are sharing them and either mix the two tracks together or use one track to stage some value.

When an *original track* has a *shared track* as one of its effects, the *shared track* itself is used as a *realtime* effect. This is more commonly known as *bouncing tracks* but Cinelerra achieves the same operation by attaching *shared tracks*. The *fade* and any effects in the *shared track* are applied to the *original track*. Once the *shared track* has processed the data, the *original track* performs any effects which come below the *shared track* and then composites it on the output.

In order to prevent the *shared track* from mixing the same data as the *original track* on the output, enable the output *mute* toggle in the *patchbay* next to each track for which you do not want to mix on the output. If you are making a video and you do want the *shared track* to composite the *original track's* data on the output a second time, the video from the *shared track* would always appear under the video from the *original track*, regardless of whether it was on top of the *original track*. This is because *shared tracks* are composited in order of their attachment. Since it is part of the *original track* it has to be composited before the *original track* is.

10.4 Saved Plugin Presets

Presets and *Factory Presets* for *Plugin* settings are now combined with the *Preset Keyframe* Parameters allowing you to choose, apply, delete, and edit your own *Presets* which can then be easily saved in the file **\$HOME/.bcast5/Cinelerra_presets**. In addition to your own saved *presets*, there are automatically available *Factory presets* for some plugins, for example the *Lens* video plugin. The *Factory presets* are preceded by an asterisk (*) and can not be modified permanently.

Note that using this is directly changing a *keyframe* object so you will only want to modify parameters you are familiar with. Most of the data is obvious and safe to change.

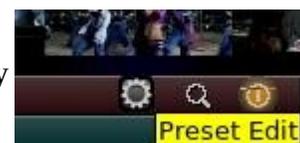


Fig. 51: Preset edit

A *Presets* button on the plugin bar to the left of the *Controls* and *On/Off* button allows for quick access to this feature. The symbol resembles a gear.

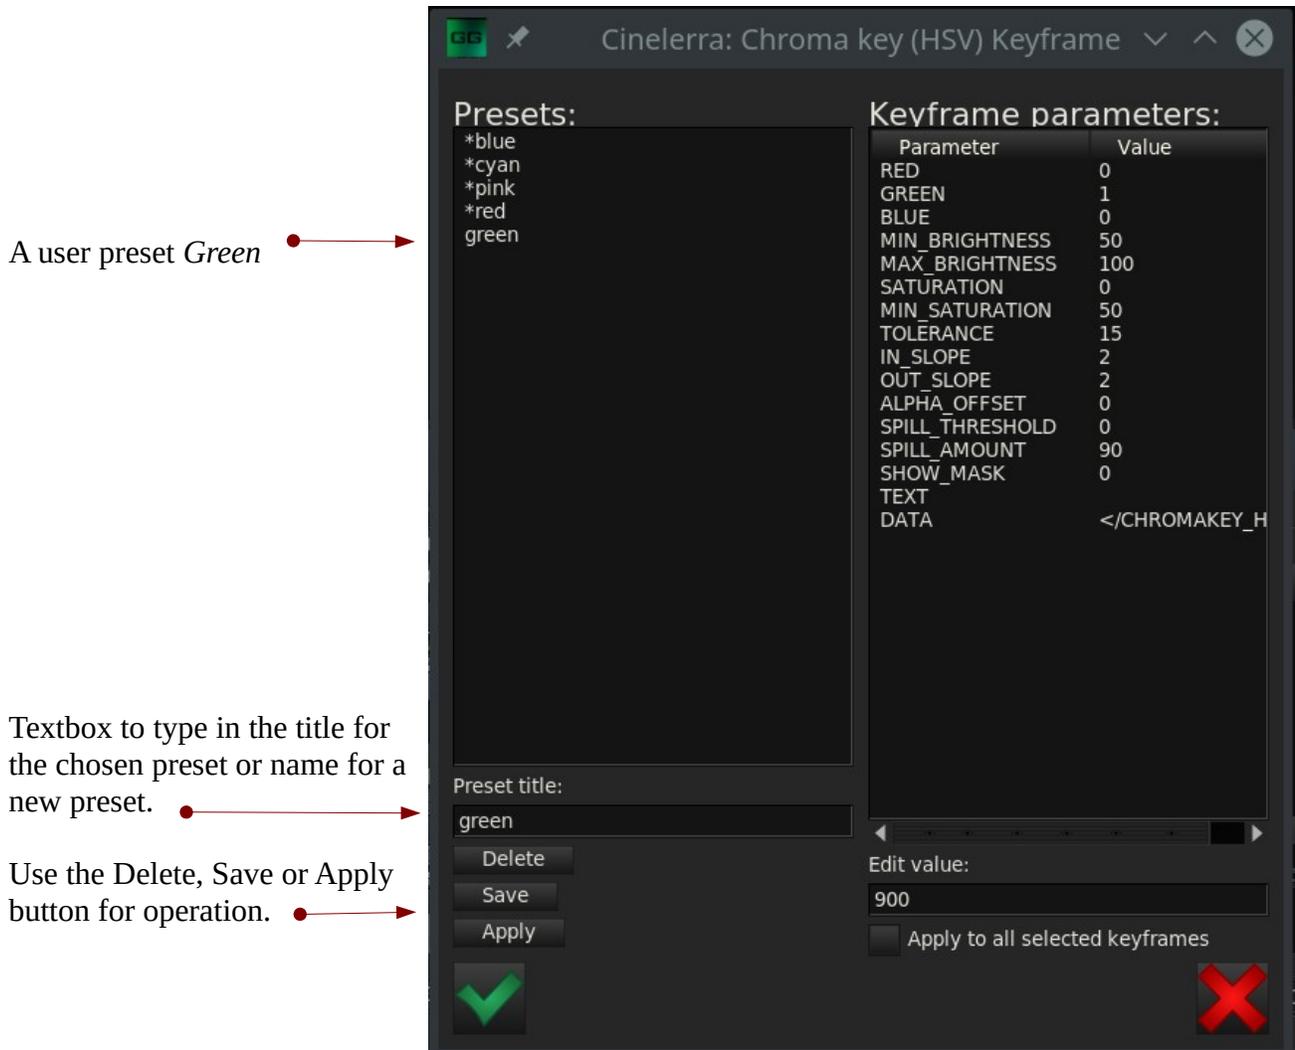


Fig. 52: Screencast shows 4 Factory presets as preceded by an *.

10.5 Some specific details concerning Plugins

These next few sections explain some details about the plugins that are not directly related to actually using them but help to work with them.

10.5.1 How to see short Description of a Plugin

To get a short one or a few lines description of a plugin, right click on that plugin in the *Resources* window and when the popup menu appears, select *Info*. Some of the plugins may not have any description included. An example screenshot is next.

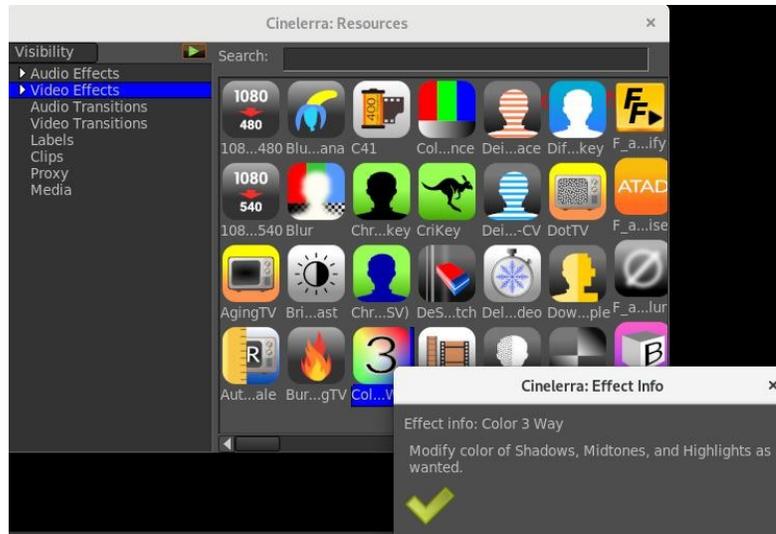


Fig. 53: Effect Info for Color 3 Way

10.5.2 Delete Plugins to save Resources Space or make them Unavailable

Maybe you just don't ever use certain plugins or would prefer to only find the ones that are useful to you. To save space in the *Resources Window* so you don't have to scroll to find the plugins you want as much, a feature to delete others is available. If you have a *System install*, you will have to be root for this function to be usable. The plugins will be *permanently* deleted, but only until you rebuild or download a new set of Cinelerra binaries. To delete a plugin, highlight the plugin you no longer want in the *Resources* window then press **Ctrl-Shift-delete**. A small window will come up allowing you to change your mind and *red-X* out or *check-OK* to remove plugin. This feature may come in handy if you have personnel working on media for you and you only want them to exercise certain functions. Or maybe you can't remember which is the *good deinterlace* plugin out of the available five or so and want to delete the extras so as not to be confused. The *ffmpeg*, *ladspa*, and *lv2* plugins can not be deleted in this manner but, of course, you can always turn them off from view by clicking on *Visibility* and unchecking them.

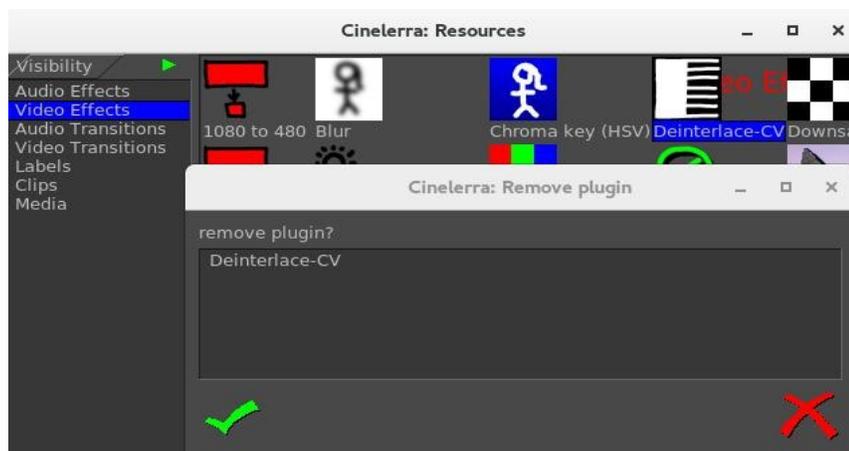


Fig. 54: Remove Deinterlace-CV plugin

10.5.3 Updatable Icon Image Support

When running Cinelerra-GG Infinity builtin icons are loaded before the program starts. *Png* files in the path `<target_directory>picon/picon_set_name` are searched before the images loaded into memory. Override *icon.png* files must be put into the `<target_directory>/picon/picon_set_name` directory. There are currently 4 sets of icons and their directory names are *cinfinity* (the default) and *cinfinity2*, *original* (the long-time original set), and *smoother* (generally was in use by some of the themes). An example, to replace the *cinfinity* icon of *Blue Banana* with a *red apple* instead, create your *.png* file as desired, and replace the file in `<target_directory>/bin/plugins/picon/cinfinity/bluebanana.png`.

For most User installs, the `<plugin_name>.png` file will be located at:

`<cinlib_path>/bin/plugins/picon/cinfinity` (or *cinfinity2*, *original* or *smoother*)

For some System installs, the files might be located at:

`/usr/lib/cin/plugins/picon/cinfinity` (or *cinfinity2*, *original* or *smoother*) (*ubuntu* distros)

`/usr/lib64/cin/plugins/picon/cinfinity` (or *cinfinity2*, *original* or *smoother*) (*Leap* distro)

10.5.4 Details on where to put your own Plugin Icons

In order to make the icons available to all themes, which would thus be the default when no theme-specific icon is available, put the *png* file in the `<cinlib_path>/bin/plugins/picon/cinfinity` directory (or *cinfinity2*, *original* or *smoother*).

The Cinelerra program looks for a plugin icon in two places:

1. First, it tries to find a *png* file in `<cinlib>/plugins/picon/cinfinity(2)` or *original*, *smoother* directory.
2. If there is no corresponding *.png* file for a plugin, the program uses a built-in default:
 - *ordinary* video plugins use 3 vertical color bars as a default;
 - *ffmpeg* plugins use the words *FF* on yellow colored background as a default icon;
 - *audio* and *ladspa* plugins use a green-colored audio wave for a default.



Keep in mind these points for newly created plugin icons:

- All included icon images become part of *open source*, in the *public domain*, and *not proprietary*.
- The preferred format is **52 x 52**, **8-bit/color RGB** or **RGBA**, non-interlaced.
- Since plugin icons are used by different themes, it is recommended that a *transparent background* be used. Otherwise some color background that looks good for one theme may not for another.
- In order to test a new icon, you have to have write permission in the `<cinlib_path>/plugins` directory so you may have to become the root user to copy the *.png* file to the correct location.
- If there is currently no theme-specific *.png* files present, it may be necessary to first create the theme directory in `<cinlib_path>plugins` as `<theme_name>` in order to put the *.png* files in that subdirectory.
- Make sure that the *ownership* and *file permissions* match the existing directory and files.

- All *ffmpeg* icons must begin with **ff_<plugin_name>.png** (*Resources* window title will still be F_...)
- For *ladspa*, check in the **<cin_config>** directory (**\$HOME/.bcast5** normally) and look for the text file **\$HOME/.bcast5/ladspa_plugins...** for the names of the *ladspa* libraries which correspond to plugin names where the needed name is the basename of the **.so** file. For example *phasers_1217.so* would need to have a *phasers_1217.png* file. There may be multiple plugins in a single “so” file which means that you can only have 1 icon to represent all of the plugins in that file; again as in *phasers*.
- Once you have placed the .png file in the correct spot, you will have to **restart** Cinelerra to test it.
- To submit your .png file for inclusion into *Cinelerra-GG Infinity* for all to enjoy, it is best to upload it to any datafilehost and notify the community via email with any informative documentation.

10.5.5 Example of new Plugin Icon Testing

For a simple test just copy an existing *<plugin_name>.png* file into the *cinfinity* directory with the name *bluebanana.png* to write over the existing file. This icon will now show up in cinelerra and still execute the *Blue Banana* function.

For an *ffmpeg* plugin, create *ff_loop.png* and copy it to **<cinlib_path>/plugins/picon/original**. This icon will show up in cinelerra if *original* is selected and execute the *F_loop* function.

For a *ladspa* plugin, the text line in **\$HOME/.bcast5/ladspa_plugins...** as seen below:

2 "am_pitchshift_1433.so" "AM pitchshifter" 1504922321 0 1 0 0 1 0 1 0 1 0 0 indicates that you would create the icon **<cinlib_path>/plugins/picon/cinifinity/am_pitchshift_1433.png**

For your own *personal* plugins, you can create a directory on your system and put any plugin png files you like into that directory. For example, if you want a specialized picon for *F_aeval*, create a picon named *ff_aeval.png* in **<cinlib_path>/plugins/picon/yournamehere**.

```
cd <cinlib>/plugins           ! go to the correct directory
mkdir -p picon/yournamehere  ! create subdirectory if does not exist
ls -l picon/*                ! list the picon directories
                              ! check for existence (and permissions)
cp yourpicon.png ff_aeval.png ! Copy your example .png file
```

Restart cin by changing Settings → Preferences → Appearance & in *Plugins* icons choose a directory.

10.5.6 Plugins/Effects Visibility

Cinelerra contains many plugins, especially with the addition of *ffmpeg*, and it is somewhat difficult to find the one you are looking for in the *Resources* window. In *Cinelerra-GG Infinity*, the plugins have been categorized into the following subsets in the *Visibility* section of the *Resources* window to make it easier to locate a particular one:

Audio Effects, Video Effects, Audio Transitions, Video Transitions



Fig. 55: Screenshot showing on the left hand side the Visibility box with Audio Effects highlighted.

The *Visibility* tool in the *Resources* window gives you the ability to **turn off** or **on** any of several sets of plugins. If you left-click the *Visibility* box, you will see the various categories of plugins, such as *ladspa*, *ffmpeg*, *audio*, *lv2*, and *video*. Highlight the set you want to **turn on** and a check mark appears to show it is active. Highlight again to toggle it **off**. See the next screenshot which illustrates that all of the plugins are turned **off** (not visible) except for *audio*. There is also the ability to add your own personal directory of plugins which will show up here. All you have to do to have these plugins become visible is to create a directory, with some name that is meaningful to you, and put your *.png* files in your **cinelerra_path bin/plugins/<your_directory_name>**.



Fig. 56: Screenshot showing the Visibility categories of plugins with all toggled off except audio.

10.5.7 Expanders for Plugin Subtrees in the Resources Window

To accentuate a set of common plugins, there are *expander arrows* on the left side of the *Resources* window. You will see these expanders only when in *Display text* mode, not *icon* mode. Cinelerra's default setup is in the file `$CIN_DAT/expanders.txt` but if the user wants their own specific setup and if the file in `$HOME/.bcast5/expanders.txt` exists, it will take precedence. If there are recommendations for other relevant categories, they can be added. The *subtree structure* is applicable to any of the *Video Effects/Transitions* or *Audio Effects/Transitions*. You can not sort once an expansion is in effect.



The *expanders.txt* file has very specific requirements. The most specific is that there are no blanks – you must use *tabs* only. A # (pound sign) can be used in column 1 to indicate a *comment*. Here is a short example:

Video Effects

- Color_Correction

- Blue Banana

#

- Color 3 Way

- Color Balance

Audio Effects

- Calf

- Instruments / Generators

- L2_Calf Organ

- L2_Calf Monosynth

- L2_Calf Fluidsynth

10.5.8 Speed-up of Ffmpeg plugin usage with OPTS files

You can speed up some *ffmpeg* plugins that are quite time-consuming and use a lot of CPU. For a specific color-based example, Cinelerra uses 6 primary rendering *color models*. All of them have 3 components at full scale. Direct usage of a particular *ffmpeg* plugin from the *ffmpeg command line* might handle the planar at less than full scale chroma (*yuv420*), which means there is less data to manipulate. But when cinelerra loads a video it uses *full scale color models*. In other words:

```
cinelerra uses yuv444
ffmpeg uses yuv420
```

if using an *ffmpeg* plugin that uses filters and many passes over the data, the amount of data is a big factor. If you load a file in cinelerra with a **same_directory_path_and_filename.opts** file containing the following line, the full scale color modeling upgrade will not be performed until after any plugin, and then the render is faster:

```
video_filter=xxxxxx=threads=8 {where xxxxxx is the desired filter)
```

When the file loads, however, it will initially take longer because it is running through the video filter. The format *rgb* in *ffmpeg* uses more cpu time. For comparison, *ffmpeg* line that might be used:

```
ffmpeg -i /tmp/filename.mpeg -threads 15 -vf format=rgb24,xxxxxxs=threads=8 \  
-acodec ac3 -vcodec libx265 - y /tmp/x.mp4
```

This converts the input to *rgb* before *xxxxxx* runs, and so it too is slower (because there is more color data). You would ordinarily avoid this conversion by omitting the *format=rgb24* parameter. An example *ffmpeg* plugin that could easily take advantage of an auxilliary opts file is *nlmeans*.

10.6 Audio Effects - Native

AudioScope

Convert *input audio* to *video output* representing the *audio power spectrum*. Shows you the *sound wave*.

Compressor

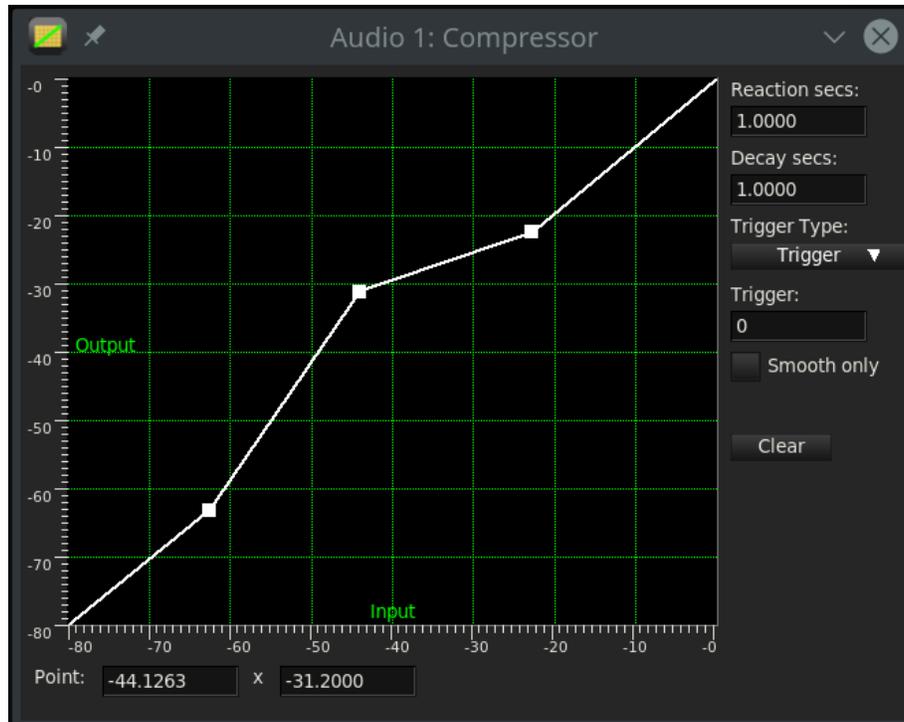


Fig. 57: GUI of configuration for Compressor plugin

The audio *compressor* reduces the dynamic range of the audio, not the amount of data required to store the audio. In Cinelerra the compressor actually performs the function of an *expander* and *compressor*.

The *compressor* works by calculating the *maximum sound level* within a certain time period of the current position. The maximum sound level is taken as the *input sound level*. For every input sound level there is an *output sound level* specified by the user. The *gain* at the current position is adjusted so the maximum sound level in the time range is the user specified value.

The compressor has a *graph* which correlates every input sound level to an output level. The *horizontal* direction is the input sound level in *dB*. The vertical direction is the output sound level in *dB*. The user specifies output sound levels by creating points on the graph. **Click** in the graph to create a point. If two points exist, drag one point across another point to delete it. The most recent point selected has its values displayed in *textboxes* (X) for more precise adjustment.

To have the *compressor* reduce the *dynamic range* of the audio, make all the output values greater than the input values except 0 dB. To make the compressor expand the *dynamic range* of the audio, make all the output values except 0 dB less than the input values. The algorithm currently limits all sound levels above 0 dB to 0 dB, so to get an overloaded effect put a *gain* effect before the compressor to reduce all the levels and follow it with another *gain* effect to amplify all the levels back over 0 dB.

Reaction secs - this determines where in relation to the current position the *maximum sound level* is taken and how fast the *gain* is adjusted to reach that peak. It is in seconds. If the *reaction time* is *negative* the *compressor* reads ahead of the current position to get the future peak. The *gain* is ramped to that peak over one reaction time. This allows it to hit the desired output level exactly when the input peak occurs at the current position. If the *reaction time* is *positive* the *compressor* scans only the current position for the *gain* and ramps gain over one reaction time to hit the desired output level. It hits the output level exactly one reaction time after detecting the input peak.

Decay secs - if the peak is higher than the current level, the *compressor* ramps the gain up to the peak value. Then if a future peak is less than the current peak it ramps the gain down. The time taken to ramp the gain down can be greater than the time taken to ramp the gain up. This ramping down time is the decay seconds.

Trigger type - the compressor is a *multi-channel effect*. Several tracks can share one compressor. How the signal from many tracks is interpreted is determined by the *trigger type*. The *Trigger* type uses the value supplied in the *Trigger textbox* as the number of the track to use as input for the compressor. This allows a track which is not even heard to determine the *loudness* of the other tracks. The *maximum trigger* takes the loudest track and uses it as the input for the compressor. The *Total trigger* type adds the signals from all the tracks and uses the total as the input for the compressor. This is the most natural sounding compression and is ideal when multiple tracks are averaged into single speakers.

Trigger - This parameter is used in conjunction with *trigger type* as described previously. Normally only one track is scanned for the input peak. This track is specified by the *Trigger*. By sharing several tracks and playing with the trigger value, you can make a sine wave on one track follow the amplitude of a *drum* on another track, for example.

Smooth only - for visualizing what the compressor is doing to the *sound-level*, this option causes it to replace the *sound wave* with just the current *peak value*. It makes it very easy to see how *reaction secs* affects the detected peak values.

DC Offset

Use this to remove *DC Offset*, which is usually an undesirable characteristic of a recording normally caused by defective equipment. This effect works like a *high pass* filter and has no controls. *DC* stands for *Direct Current* which is the average amplitude of the *waveform*. It sounds best when it is absent, represented by *zero*, so that there is no imbalance in the audio.

Delay Audio

In the *Delay Audio* effect you can specify the number of seconds you want to *delay* the video track.

Denoise

Reduce audio *background noise*. There is only 1 parameter which is used to regulate the *level dial* with a range of 0 to 1.

DenoiseFFT

Noise removal from audio using FFT editing. Set the *Denoise Power* dial in *dB* and choose the number of *reference samples*.

Despike

Detect and eliminate *out of range impulse values*.

Maximum level: slider to set the maximum value in *dB* above which the frequency cutting takes place.

Maximum rate of change: to adjust peak delete in *dB*.

EQ Graphic

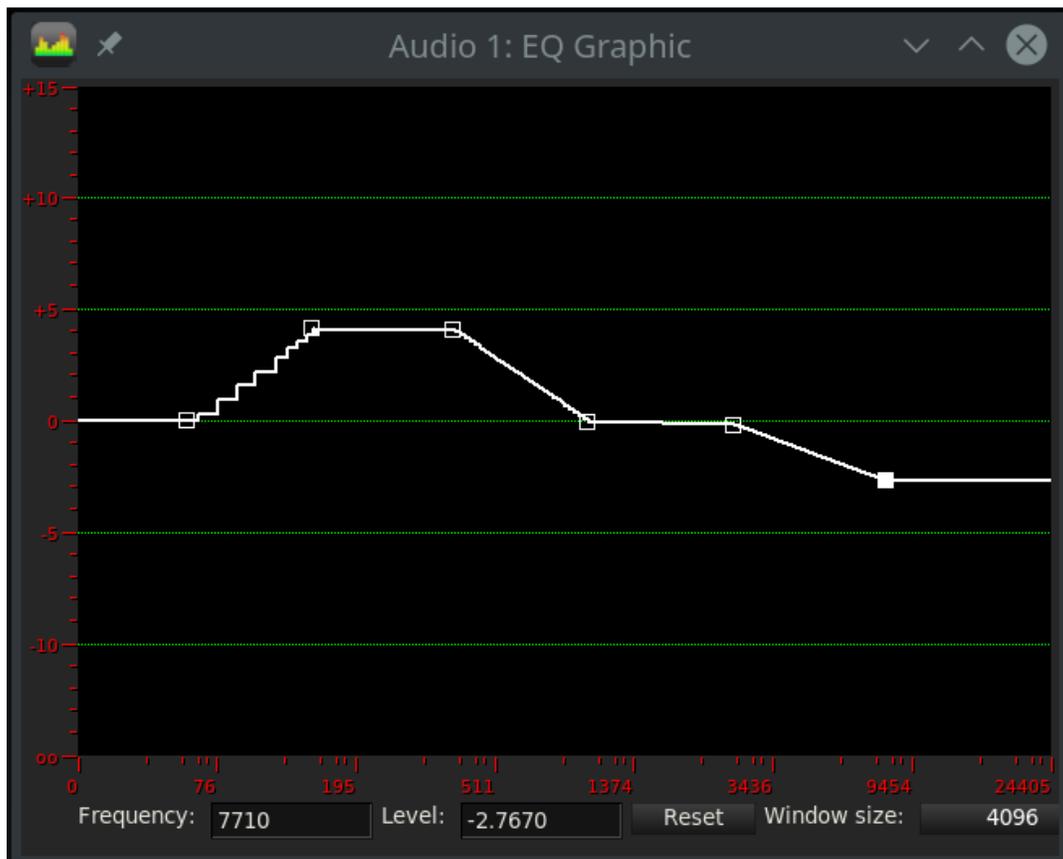


Fig. 58: Graphic Equalizer audio plugin

Graphic equalizer sets the *output levels* for specified frequency bands. This effect works by setting *control points* when you click the left mouse button and drag to the desired value. In the *textboxes* at the bottom can be seen the frequency of the *active control point*, the level of the signal to be set by entering the numerical value or by dragging the control point, and the *number of samples* to act on.

EQ Parametric

Parametric equalizer shows and outputs levels for *frequency, quality, level, mode, and wetness.*

Echo

Echo is reflection of sound. This plugin could be used to add *echoing* to video of your canyon hike.



Fig. 59: The 3 dials of Echo plugin

The 3 dials shown here are used to vary these attributes.

Level: represents the *volume* adjustment.

Atten: is attenuation which is a general term that refers to any reduction in the *echo* reflection.

Sometimes called *loss*, attenuation is a natural consequence of signal transmission over long distances.

Offset: is the *lag* in the attenuated echo signal. *Offset* means adding a *DC level* to a signal. It offsets the signal up or down in a DC sense without changing the size of the AC part of the signal. When you add an audio clip to the Timeline, the clip plays back from the beginning of the source audio file. The point in the audio file where the clip starts playing is called the offset. By default, a clip's offset is zero, the beginning of the source audio file. You can change the offset so that the clip starts playing from a later point in the source audio file.

EchoCancel

EchoCancel is the process of removing echos from audio in order to improve the quality. *Echo cancel* may be needed because an audio recording was done in a room that led to echo generation or there was some kind of unwanted feedback. There are many controls for the *EchoCancel* plugin which are defined here. However, the first thing you will see when you bring up the plugin, is the top portion that is black which will show a + in the middle when you mouse over it. Once you start playing audio, you will see the *cepstrum spectral data* inside the window. A *cepstrum* results from taking the *inverse Fourier transform* (IFT) of the logarithm of the estimated spectrum of a signal. It is used to identify the period of the echo in the audio. It is recommended to just set the *Mode* to **On** but the below defined parameters can be utilized by professionals.

On/Off Boxes

Normalize: audio normalization adds variable amounts of gain to an audio recording to bring the average or peak amplitude to a target level (the *normal* amount), on an ongoing buffer by buffer basis. This is to make the *cepstrum* graphical data appear between 0 and 1. Checkmark appears if **On**.

Dial Parameters

Level: scale factor used to draw the cepstrum output when *normalize* is not in effect.

History: number of previous cepstrum outputs redrawn as fading graphical data.

X Zoom: X axis scale factor to magnify low frequency cepstrum graphical output.

Damp: echo envelope decay factor used to smooth the cepstrum/correlation data.
Peaks: number of maximal envelope values used in the *echo gain* calculation.
Cutoff Hz: low frequency cutoff value to prevent beat frequency (*heterodyne*) echo canceling.

Down Arrow Boxes

Mode: *MAN*, *Off*, or *On*. When *Off* is selected, the plugin is not active. When *MAN* is used, the only one peak is used for the echo gain envelope. It is set by pressing mouse button 1 in the cepstrum graphical output. The *Gain* and *Offset* are updated as the pointer drag operation resets the indicated gain and offset values. When *On* is selected, the echo gain envelope is automatically calculated by cepstrum and auto-correlation of the input audio in the last “window” size audio samples.

Window size: parameter can be set to Default, 1024, 2048, ... doubled values... up to 262144.

Values:

Amplitude: the cepstrum value at the drag point during manual envelope selection.

Gain: echo gain setting determined by manual selection.

Offset: echo period setting determined by manual selection. The *Hz* (frequency), *ms* (millisecond duration), and sample *offset* (audio samples) as determined by manual selection.

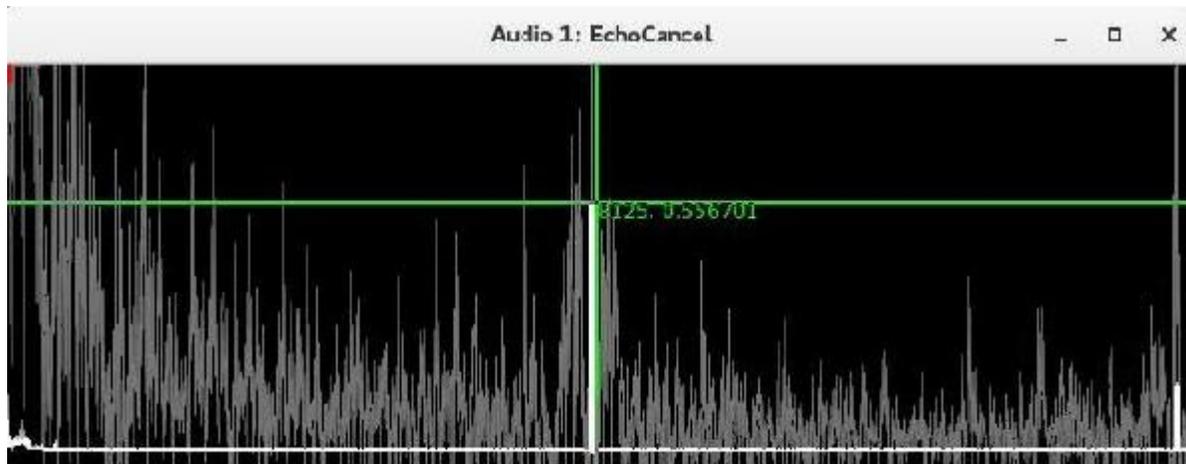


Fig. 60: EchoCancel screenshot showing the crosshairs (+) with the offset/gain values in the middle

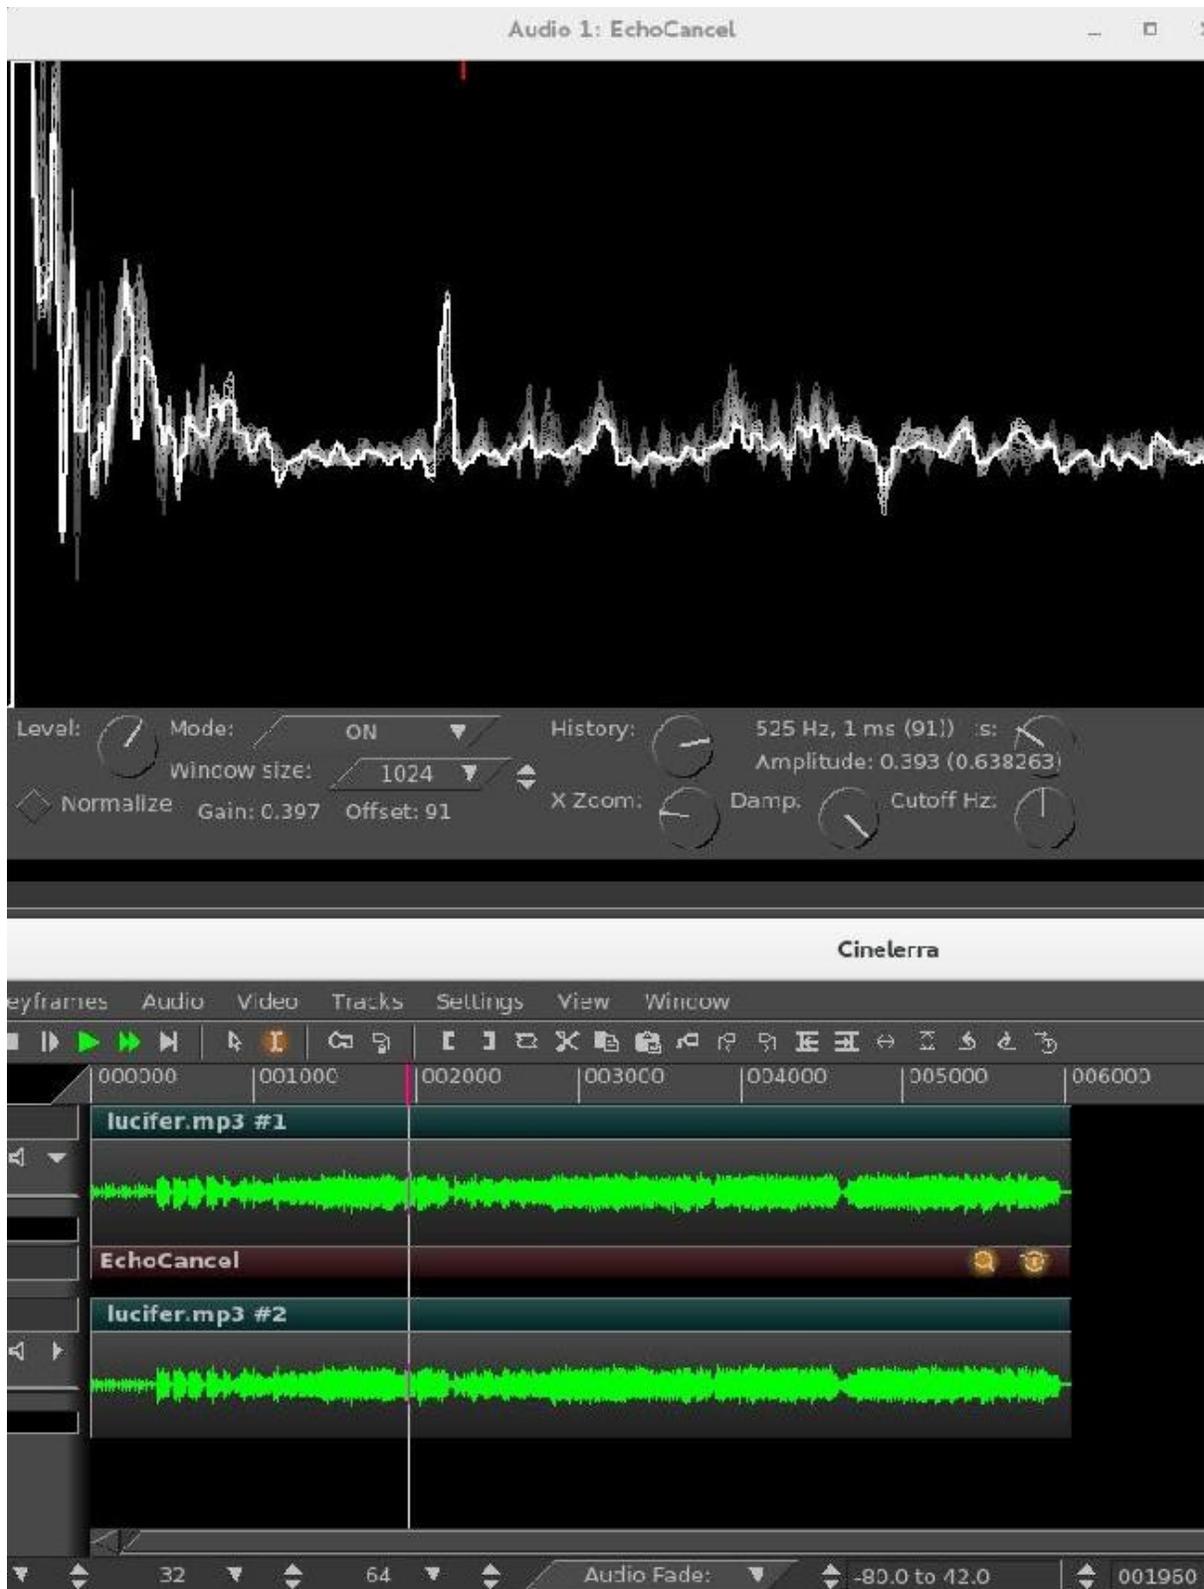


Fig. 61: Screenshot with Mode set to “On” for the usual situation for most users. This is highly recommended

Freeverb

Adds effect of *multiple decaying echoes* to audio signals based on a specific algorithm. Common use of *reverb* is to simulate music played in a closed room.

Gain

Add *gain, input level*, to increase/decrease *loudness*.

Interpolate

Generate a smooth curve based on sound creating a certain softness. There are no controls.

Invert Audio

Reverses the numerical sign of the digital audio. There are no controls.

Live Audio

The *Live Audio* effect reads audio directly from the *sound card* input. It replaces any audio on the track so it is normally applied to an empty track. To use *Live Audio*, highlight a horizontal region of an audio track or define in and *out points*. Then drop the Live Audio effect into it. Create extra tracks and attach *shared copies* of the first *Live Audio* effect to the other tracks to have extra channels recorded. Live Audio uses the sound driver selected in **Settings->Preferences->Playback->Audio Out for recording**, but unlike recording it uses the playback buffer size as the recording buffer size and it uses the project sample rate as the sampling rate. These settings are critical since some sound drivers can not record in the same sized buffer they play back in.

Live audio has been most reliable when ALSA is the recording driver and the playback fragment size is 2048. Drop other effects after Live Audio to process sound card input in realtime. With live audio there is no read-ahead, so effects like *compressor* will either delay if they have read-ahead enabled or playback will under-run. A potential problem is that sometimes the recording clock on the sound card is slightly slower than the playback clock. The recording eventually falls behind and playback sounds choppy. Live Audio does not work in reverse.

Loop audio

Loop some number of samples of audio over and over.

Overlay

Overlay has parameter settings of *top* or *bottom* for the track and *add* or *multiply* for the operation.

Pitch Shift

Like the *time stretching* methods, there are three *pitch shifting* methods: *Pitch shift*, *Resample*, and *Asset info dialog*. *Pitch shift* is a realtime effect which can be dragged and dropped onto recordable audio tracks. Pitch shift uses a *fast Fourier transform* (FFT) to try to change the pitch without changing the duration, but this introduces windowing artifacts. Because the windowing artifacts are less obtrusive in audio which is obviously pitch shifted, Pitch Shift is mainly useful for extreme pitch changes. For mild pitch changes, use *Resample* instead. Another way to change pitch slightly is to go to

the *Resources* window, highlight the *media* folder, right click on an audio file, click on *Info*, then adjust the *sample rate* in the *Info* dialog to adjust the pitch. This method also requires left clicking on the right boundary of the audio tracks and dragging left or right to correspond to the length changes.

Remove Gaps

Remove silent gap (below *DB threshold*) which persist for more than the time limit.

ResampleRT

Allows you to convert an audio file from one *sample rate* to another. This effect works similarly to *ReframeRT* in videos.

Input / output > 1 => fast rate

Input / output < 1 => Slow rate.

Reverb

Reflections of sound to add *depth* and *fullness*. Simulates creation of a large number of reflections (lots of walls) which build up and then decay. You can use the *reverb* plugin to mix tracks together to simulate *ambiance*.

Reverse Audio

Apply *reverse* audio to an audio track and play it backwards. The sound plays forward. Be aware when reversing audio that the *waveform* on the *timeline* does not reflect the actual reversed output.

SoundLevel

Displays the *Max/RMS sound level* in decibels.

Spectrogram

Visual representation of the *sound levels* at specified frequencies as they vary with time.

Synthesizer

Generate *synthesizer sounds*; to set key data, turn on *Generate keyframes* while tweaking.

Time Stretch RT

Change the *speed* of an audio signal without affecting its *pitch*.

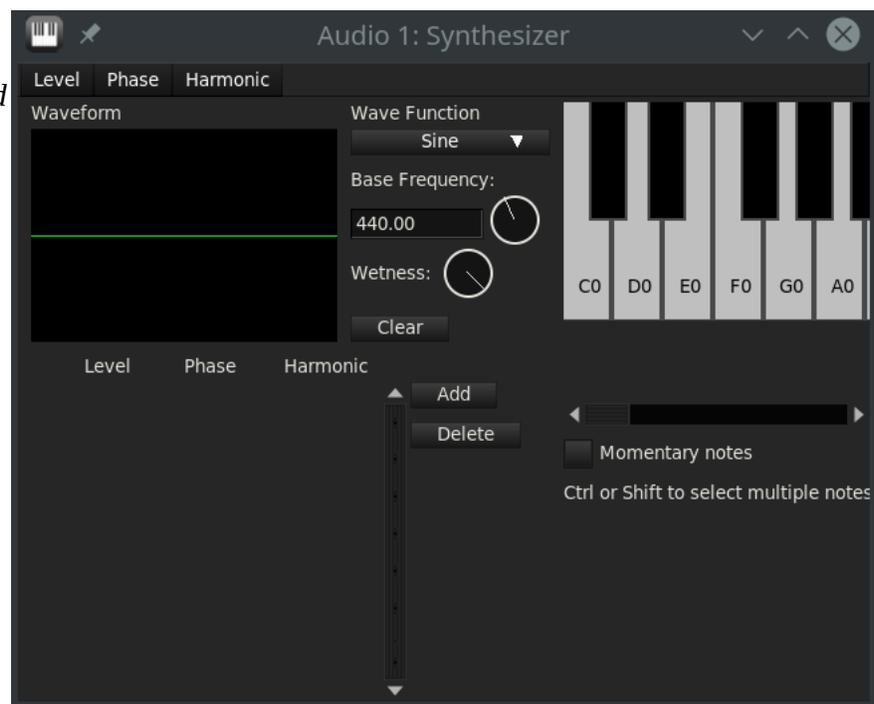


Fig. 62: Synthesizer plugin

10.7 Audio Ladspa Effects

Ladspa effects are supported in realtime and rendered mode for audio. These audio effects are supported since Cinelerra implements the LADSPA interface as accurately as possible. Besides the supplied LADSPA effects (credit Steve Harris), additional LADSPA effects can be enabled by setting the LADSPA_PATH environment variable to the location of your LADSPA plugins:

```
export LADSPA_PATH=/usr/lib/ladspa .
```

10.8 Audio LV2 / Calf Plugins (Optional Feature - OS dependent)

LV2 is an open standard for audio plugins using a simple interface with extensions which add functionality to support audio software. These plugins were written by external developers and provide additional audio effects to Cinelerra audio without having to change Cinelerra every time. Because the LV2 plugins are separate from Cinelerra-GG Infinity, if one fails or does not perform as expected, cinelerra should stay running and you will have to contact the programmers responsible for that plugin for a fix.

Typically, a user OS has specialized package groups installed. It is difficult to create one build of cinelerra to accommodate all potential LV2 plugins. Specifically for the *Calf-Studio LV2* plugins, you should install the *Calf Plugins* package. The user's computer must have *gtk-2-runtime* installed, which seems to be automatically done already for most distros. For users doing their own builds, you can build cinelerra without LV2 support by including **-without-lv2** in the configure step. The default build is **--with-lv2=yes** and requires that *GTK-2-devel* must be installed or the build will fail and notify you.

LV2 plugins have their own category in the *Audio Plugins Visibility* as lv2. There is a simple *text interface* which is available via the usual *Show controls* button when the plugin is attached to the audio track. This window has a *Reset* button to get back to the default settings. To change a value of one of the parameters, highlight that parameter and type in the new value in the topmost *text box* and then hit *apply* to take effect - the reason for requiring hitting apply is so that the audio is not moving all over the place while you are still typing a value. More easily, you can just move the *pot dial* or the *slider bar* which take effect automatically.

Cinelerra's *buffer* size setting may cause a delay in activation of the changes you make taking effect, so you can lessen the time by using a small buffer. Notice that 1024 samples at 48000 samples per sec is only 1/50th a second. This is not a lot of time to shuffle a bunch of stuff. Short buffers produce low latency, but no time for complex programs or lots of stacked effects. Bigger buffers allow for more complex setups.

To set the buffer size:

Settings → Preferences, tab Playback A, section Audio Out, variable Playback buffer samples

However, be forewarned that due to variability in the lv2 plugin programming code, some of the plugins only work with the minimum buffer size of 1024. In these cases, what you will see is the main track canvas cursor just bounces back and forth over a very small area in the timeline. This does not crash cinelerra but you will have to remove the plugin to continue working.

You can specify a certain set of LV2 plugins to use by setting **LV2_PATH** as shown below before starting cinelerra - include a colon (:) separator for multiple paths. The default path for most operating systems is **/usr/lib64/lv2**. To list the system installed lv2 plugins key in: **lv2ls** .

```
export LV2_PATH=/tmp/j/balance.lv2/usr/local/lib/lv2:/usr/local/lv2
```

If there is no default LV2_PATH set automatically, the value will be **\$CIN_DAT/lv2**, which is a *placeholder* only so that no lv2 plugins will be loaded. When there is no system LV2_PATH set it is important to note, that if you do want lv2 plugins loaded, you must set the correct path in **Settings → Preferences, Interface tab, Default LV2_PATH directory path name**. When you change this field, cin will automatically restart and load the newly specified lv2 plugins. If when switching LV2_PATH or if the lv2 audio plugins are not displayed/usable in the Resources window, you can execute a reload via **Settings → Preferences, Interface tab “Reload plugin index”** or else before you bring up cinelerra, delete **\$HOME/.bcast5/Cinelerra_plugins** so that the plugins get properly reloaded.

There are some lv2 plugins that display a *glitzy UI* (User Interface); for example the *Calf plugins*. For these LV2 plugins, if you want that to automatically come up without having to click on the UI button on the simplified UI interface, there is a flag to enable that. It is at:

Settings → Preferences, Operations tab, then check the “Auto start lv2 gui” Flag

Below is a screencast showing the auto start gui flag and the LV2_PATH default directory path on the bottom line. Note the highlighted “Reload plugin index” which will be executed if OKed.

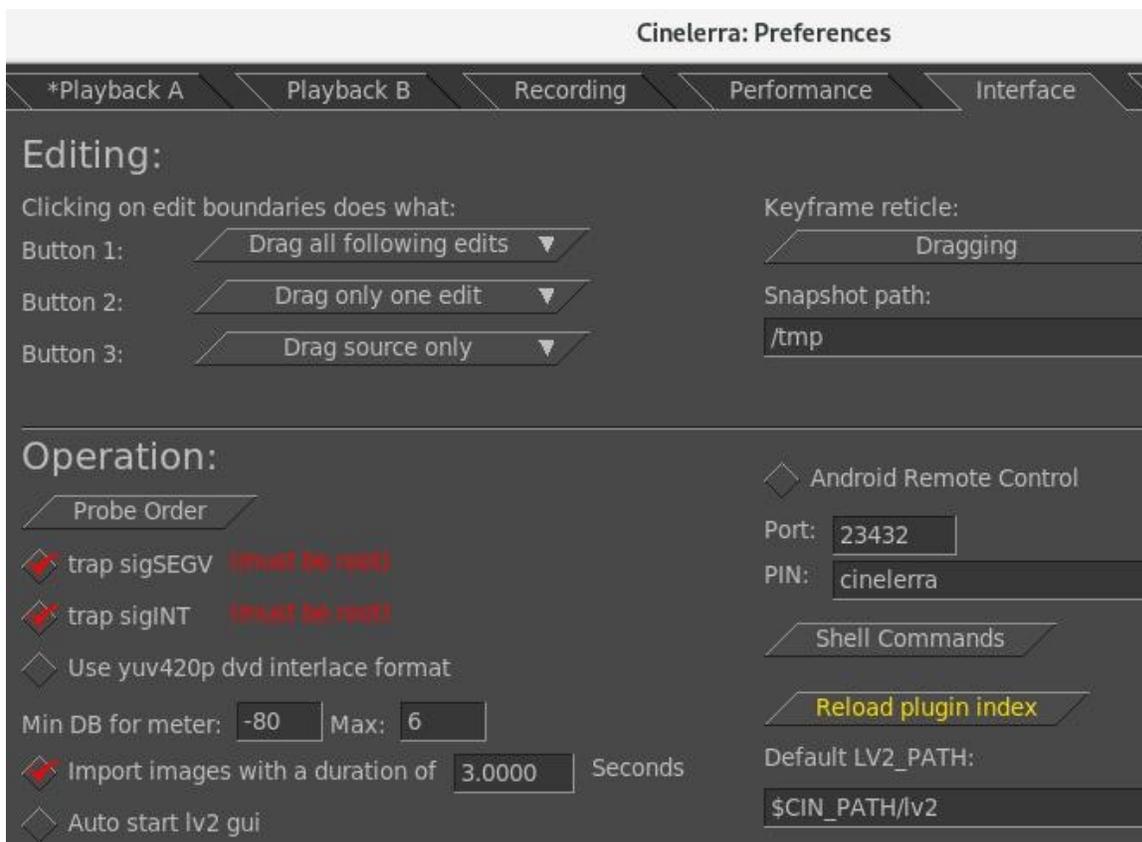


Fig. 63: "Reload plugin index" in yellow and "Auto start lv2 gui" unchecked

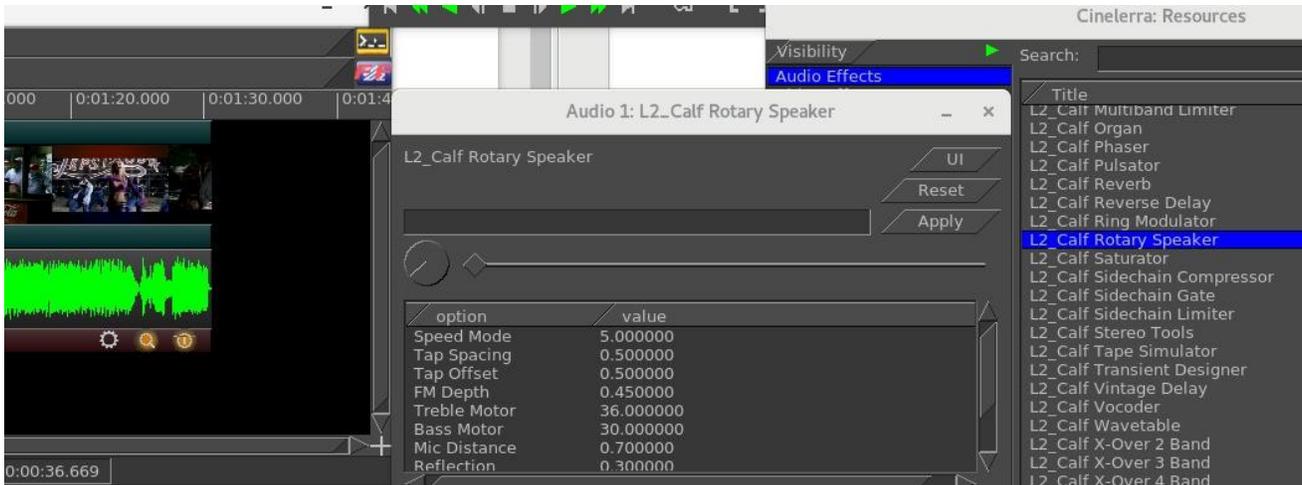


Fig. 64: Screencast of simple text interface in the middle of the screen for a Calf LV2 plugin

There is also a *blacklist* that prevents known problematic-for-cinelerra lv2 plugins from loading to avoid crashes. If others are found to have problems, once informed about them, they will be added to this *blacklist*. In order to determine which lv2 plugin causes a **SEGV** on cinelerra startup, you can start from a terminal window and you will see each plugin that is being loaded and the last one shown before the crash is a *bad plugin*. However, many of the plugins causing a crash are due to not having been compiled on your current system with the current compiler so may actually work correctly on other user systems and so will not be added to the cinelerra-wide blacklist. You can either recompile the problematic plugin, or modify your own *blacklist* which you will have to maintain and save so as not to be written over when loading a new build.

Note the *UI button* on the above screencast in the upper right hand corner above the *Reset* button. If you click this button, a glitzy interface window comes up (if available) for changing variable values. It is possible that a bug in the LV2 plugin causes the glitzy window to appear as blank and then die, but in that case the original simple text window might still work - in either case, if the timeline movement hangs, just detach the plugin to continue your current session. There is an environment variable that you can set, **BC_TRAP_LV2_SEGV**, to get a *dump* of the failure which may be helpful for debugging.

When the glitzy *ui* is up, the simple text window remains up also since it is the cinelerra side and keeps track of the value changes so they remain in effect for further usage of the plugin. Changes to one or the other will occur in both with the exception of certain features in the glitzy window which are not communicated correctly back to cinelerra; for example a *reset* button - the simple interface *Reset* button must be used instead. To change values in the glitzy window you use the mouse and move up or down unlike a knob that turns!



Fig. 65: Screenshot with a Calf plugin glitzy window that appears when clicking the simple interface UI button.

In order to test a particular plugin without bringing up cinelerra, especially for ones that do not operate, it is possible to manually display an lv2ui gui with: `/cin-path/lv2ui <lv2-uri>` . For example:

`/tmp/cinelerra-5.1/bin/lv2ui http://calf.sourceforge.net/plugins/Flanger`

10.9 Video Effects - Native

(credit to Andrea Paz for reviewing, and numerous plugin descriptions and figures)

1080 to 480

Most TV broadcasts are received with a *1920x1080* resolution but originate from a *720x480* source at the studio. It is a waste of space to compress the entire *1920x1080* if the only resolvable details are *720x480*. Unfortunately resizing *1920x1080* video to *720x480* is not as simple as shrinking it.

At the TV station the original *720x480* footage was first converted to fields of *720x240*. Each field was then scaled up to *1920x540*. The two *1920x540* fields were finally combined with *interlacing* to form the *1920x1080* image. This technique allows a consumer TV to display the re-sampled image without extra circuitry to handle *720x480* interlacing in a *1920x1080* image.

If you merely deinterlace the *1920x1080* images, you would end up with resolution of *720x240*. The *1080 to 480* effect properly extracts two *1920x540* size fields from the image, resizes them separately, and combines them again to restore a *1920x480* interlaced image. The scale effect must then be applied to reduce the horizontal size to *960* or *720* depending on the original aspect ratio.

The *tracks* to which *1080 to 480* is applied need to be at *1920x1080* resolution. The project settings in **settings->format** should be at least *720x480* resolution. The effect does not know if the first row in the *1920x1080* image belongs to the first row of the *720x480* original. You have to specify what the first row is in the effect configuration. The output of this effect is a small image in the middle of the original *1920x1080* frame. Use the *projector* to center the output image in the playback.

Finally, once you have *720x480* interlaced video you can either apply *Frames to Fields* or *Inverse Telecine* to further recover original progressive frames.

1080 to 540

Extracts two *1920x540* fields from *1920x1080* image, resizes them separately, and combines them to *1920x540 interlaced* image.

Aging TV

This effect is the one to use if you want to achieve an *old movie* or *TV show* look. It will put moving lines up and down the movie as well as putting *snow* on the video. Use it along with *Brightness/Contrast* and *Color Balance* to make your movie look like a really old black and white movie. This came from <https://effectv.com>.

Auto Scale

Automatically *scale* to a specified size.

Blue Banana

(credit to Monty Montgomery programmer)

Blue Banana is an *HSL* (hue, saturation, lightness) *Qualifier*, one of the basic tools of any grading software that are based on circumscribing a zone of the frame by extracting a *chromatic key* and producing a *matte* in the *alpha channel*. *Blue Banana* differs not by creating a *real matte* on which you

can act independently or use it from other plugins in other stages of grading, but by creating a *selection mask* exclusively for use within the plugin. The preceding sentence implies that this is a limitation in that you can not just apply a *Blur* to the *matte*. The *BlueBanana* plugin has a couple of useful purposes. It can be used for color transformation or remapping – by isolating a specific color and then performing color change/correction on only that color. Another useful purpose is for chroma-key filtering, using multiple *BlueBanana* plugins on the same track. Also, it can be used in conjunction with the *mask* operation of the *Compositor*. Usage of *BlueBanana* may seem complicated at first, but it is necessarily so in order to get enough control to produce the desired effect simply and quickly. Just changing a single color is actually quite easy. *BlueBanana* is keyframable.

The basic strategy for *BlueBanana* is to:

- Select a specific *target* color.
- Create a *selection* region by expanding color ranges around that color.
- Optionally reduce or expand the *alpha* plane as a regional *selection mask*.
- Optionally apply a color remapping or transformation to the selection.
- Optionally reset the output alpha to opaque, or pass the alpha to another *BlueBanana* plugin.

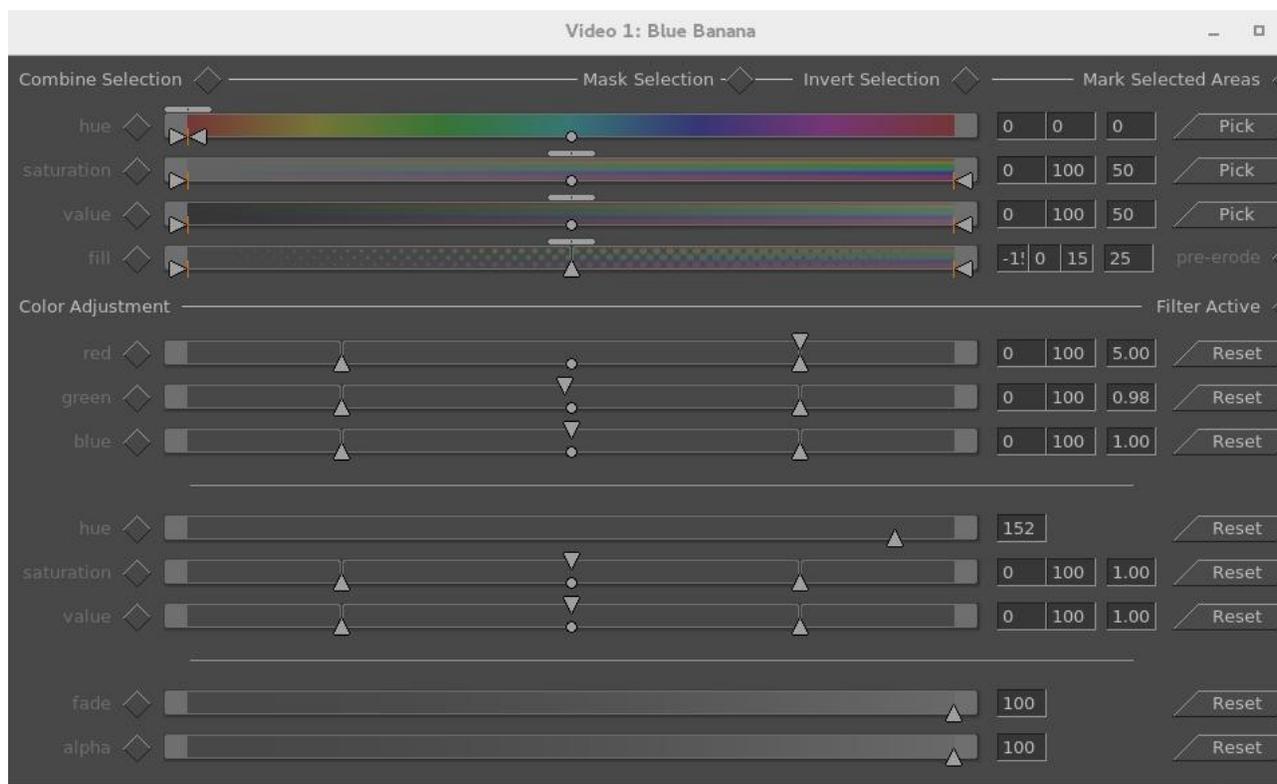


Fig. 66: Screenshot showing the *BlueBanana* plugin control

Just a Warning Note:

May use a lot of *CPU* and *Memory* because it is doing a lot of work. If you turn **off** the plugin on the plugin bar below the video track in the main track canvas it will stop using cpu when not in use. Or once you uncheck *Mark Selected Area*, it will no longer be using the cpu to mark the selected color area in realtime while drawing the diagonal animated pattern in the compositor window.

Example Usage

(from original message by Rebecca)

If you just want to try this, follow these steps.

First – Choose your color.

1. Load your video, add the *BlueBanana* plugin to the track, bring up its *control window*, and uncheck any checked boxes (mostly just to avoid unexpected results).
2. In the *Compositor* window, choose the *eyedropper color picker* tool on the left-hand side and click on the area of the image that shows the color you want to change/correct.
3. In the *BlueBanana* plugin window, to the right of *hue*, click *Pick*. And if you want to modify *saturation* and *value*, also click on the *Pick* button for them. To see what it does, *Pick* them also.
4. Next, check the *Mark Selected Areas* box at the top right of the *BlueBanana* plugin window to see the selected color-matching areas which will become marked in a diagonally striped pattern.
5. You can now manually modify your selection in the *Color Selection* area in the obvious ways for *hue*, *saturation*, *value* and *fill*. The *arrows* to each side of the small circle widen the selected area. Move the *dot* and you move the range. The *slider* on top of the horizontal color strip shifts like the amount of the strip is dedicated to that part of the color spectrum. *Fill* will fill more area or less area in your selected region.

Second – Adjust your color choice.

1. There are color strips under *color Adjustment* which will show color changes as you modify values.
2. Uncheck *Mark Selected Areas* and check the *Filter Active* box to the right of *Color Adjustment*.
3. As needed, you can individually check and uncheck all the various parameters using the boxes to the left of each line. Again, these are intuitive and broadly similar to the above. The *arrows* at the bottom widen the range, the *circle* at the bottom moves the range, and the top *slider*, which is an *arrow* this time, affects distribution. It provides a little *histogram* effect to give you an idea of what you're changing. The *fade* adjusts the level of color blending. The *alpha* is basically the opacity of your changes.

Definition of Wording/Checkboxes/Buttons/Operators are being described next. Some of the commentary was adopted from information provided by *Monty Montgomery* and from questions and answers from email by *Igor Ubuntu*, who did extensive testing.

Operational characteristics for the color-related adjusters:

- *left arrow* slider operates the range minimum; the numerical value shows in the left-most *textbox*.
- *right arrow* slider operates the range maximum; resulting numerical value is in the middle *textbox*.
- *middle circle* slider below, can move the current range up or down and the numerical results will show in the left and middle *textbox*. Move the *dot* and you move the *range*.
- *top pad* slider operates the *edge slopes* (selection attack/decay) and the value will be displayed in the rightmost *textbox*. *Sharp edges* are represented by 0; 100 represents *smooth edges*.
- or *top arrow* affects the *distribution skew*.

Operational characteristics for Fill:

- *left arrow* slider operates *mask erosion filling*. First *textbox* value.
- *center up arrow* slider operates the *fill skew midpoint*. Second *textbox* value.
- *right arrow* slider operates *mask expansion filling*. Third *textbox* value.
- *top pad* slider operates the *edge slopes*. Right *textbox* value.

The *textboxes* are available so that you can directly type in numbers from the *color wheel*. This could be helpful if duplicating previous work as it would be an instantaneous exact numerical match without having to continuously fine-tune the movement of a slider.

There are two panes separated by long horizontal lines (through the middle of the screen) in the *control* window of the BlueBanana plugin, clearly visible in the previous screencast. The top pane is first used to create/modify a selection, and the bottom pane is used to operate a change.

Pane 1

this section is used to *select the target color domain*.

Combine Selection – allows for two *selection modification* modes –

- 1) *intersection*, referring to *inclusion* (unchecked), and
- 2) *union* referring to *combine* (checked).

First, a short explanation about *alpha*. The *alpha channel* used in BlueBanana is not *transparency*; it is used as the *Selection mask*. Alpha plane is the alpha channel of the current image. So that:

RGBA = red/green/blue color planes, alpha data plane.

YUVA = luma/Cb/Cr color values, alpha data plane.

The alpha data normally is *input* to the blending operations in the *patchbay overlay* mode. The alpha data usually creates the appearance of stacking order, and determines which color planes are visible in the rendered result. When BlueBanana is used, the meaning of the alpha data is changed to the *selection*. It is useful to think of the alpha data as more *solid* when it is *transparency* in blending, and more *selected* when it is used in BlueBanana. In both cases, the greater the alpha value, the more the effect is expressed.

Usually, alpha is *normalized* to a range from 0 to 1, *zero = no effect*, *1 = total effect*, *0.5 = partial effect*. In both cases, alpha is what math people call an *auxiliary variable*. It is needed, but is not part of the answer. In this case, the answer is the visible rendered result. Alpha is like meta-data.

Now back to *Combine Selection*. The selection is the *intersection* or *union* of two pixel masks. Mathematically, **A** and **B** are *normalized*, (scaled to between 0 and 1) and used as *selection mask weights*.

Intersection = A*B and **Union = A+B-A*B** where:

A) is the *input* alpha plane as a *mask*, 1=selected, 0.4=partially selected, and 0=not selected.

B) is the *color selection* of trims and feathers made by varying the sliders.

The result is a new alpha plane, which will be *output* (if *End Mask* is not set). The 0..1 selection values are used to weight the color transformation filters if/when they are active and operate a change. The color adjustment filters available in Pane #2 can change "red", "green", "blue", and remap "hue", "saturation", "value" in the pane. There is also *fade* which applies to the color channels and *alpha* which applies to the resulting alpha plane.

The basic plan is to either

reduce a selection area by *intersection* (*Combine selection off*) – **A*B** or

increase a selection area by *union* (*Combine selection on*) – **A+B-A*B**

Mask Selection - applies the current mask to the selection, such that the mask clips/expands the selection. When *mask selection* is enabled, the result of the *and/or* will be stored to the alpha result, but when mask selection is unchecked the mask is ignored and the selection is not modified. The selection is used to weight the effect of the filtering, or to control the output alpha.

End mask – only visible when *Mask Selection* is checked. *End Mask* causes the entire alpha plane to be set to 1. The image becomes opaque. This is usually only set in the last plugin of a stack (the stack may be just one plugin doing only color modification). In the event that a *color selection mask* is used with multiple, layered BlueBanana filters on the same track, the grouped BlueBanana filters may share a single mask by all enabling *Mask Selection*, but with only the last BlueBanana enabling *End Mask*. This usage pattern gives the *End Mask* control its name.

- *End Mask as used in Color Transformation/Remapping:*

In many use cases where you are just remapping color, you are still interested in seeing all of the picture. If this is the case, then checking *End Mask* on the last BlueBanana plugin will show you the entire picture. The alpha plane may be in use as a *selection mask*, but it may not be wanted as part of the result.

- *End Mask as used in Chroma-key Filtering:*

In cases where the selection is for a chroma-key, you are interested in the alpha channel for blending, like *Normal* or *SrcOver*. So for this usage of the BlueBanana, don't check the *End Mask*.

Invert Selection - reverse target color domain, which is 1 minus selection.

Mark Selected Areas – when this box is checked, the chosen colors are presented in an animated diagonally striped pattern.

Hue – select a *hue* domain; click on the *Pick* button to select or check the box to the left of *hue* or uncheck to ignore.

Saturation – select a *saturation* domain; click on the *Pick* button to select or check the box to the left.

Value – select a *value* domain; click on the *Pick* button to select or check the box to the left.

Fill - will fill more area or less area of your selected region. This describes how it works. *Fill* control is an automated way of doing grow and shrink on the selected area, to fill in small holes, or get rid of scattered speckles. If none of the Hue, Saturation, or Value sliders are active - meaning that the whole frame is selected - the *Fill* slider will have no effect even when enabled. The word *fill* will appear ghosted to indicate this.

The three lower handles in the fill slider correspond to *Shrink* (the left hand slider), *Final* (the middle slider), and *Grow* (the right hand slider). These are used in combination to alter the *selection* by first growing it by the amount specified by the right hand *Grow* slider, shrinking it to the amount specified by the left hand *Shrink* slider, and then growing it again to the final size specified by the middle *Final* slider. The *top* slider then *feathers* the resulting selection.

Growing the selection and then shrinking it has the effect of filling small holes in the selected area.

Similarly, shrinking and then growing tends to remove small flecks of unwanted selection. The *Final* slider specifies the overall desired shrinkage or growth of the selection when finished. To specify a pure Grow or Shrink operation, set the Final slider and the Grow/Shrink slider to the same value and leave the other slider at zero.

Pre-erode - this control reverses the order of operation to *Shrink*, then *Grow*, then *Final*. The change is subtle on most images, but overall removes more small features because it first removes flecks before filling in holes.

Pane 2

this section is used to *modify the color* of your selection. Descriptive commentary for this pane.

Filter Active – checkbox to indicate that the modifications will be shown.

Color Adjustment – For Color Adjustment, RGB can be used as *color weights* while the HSV can *transform color*.

For the following items there are three sections on the slider. The *center section* represents the nominal 0%-100% range; the *left section* represents negative values, and the *right section* represents values greater than 100%. Values can be *out-of-range* within BlueBanana without *clipping*, but they will *clip* once they leave the plugin.

The following *RGB filters* affect the color channels individually.

Red – modification color; click the *Reset* button to revert to default. Values are reflected in numerical *textboxes* on the right-hand side.

Green - modification color; click the *Reset* button to revert to default. Values are reflected in numerical *textboxes* on the right-hand side.

Blue – modification color; click the *Reset* button to revert to default. Values are reflected in numerical *textboxes* on the right-hand side.

The following *HSV* reorient the color spectrum, and affect all of the color channels simultaneously.

Hue – a single numerical value will appear in the right-side box. Click the *Reset* button for default.

Saturation – for modifying the saturation; click the *Reset* button to revert to default. Values are reflected in numerical *textboxes* on the right-hand side.

Value – for modifying the value; click the *Reset* button to revert to default. Result is reflected in the numerical *textboxes* on the right-hand side.

Fade – controls the entire color re-mapping, and does nothing if no color adjustment is active.

Alpha - controls the output alpha (this is not available when *End Mask* is set); click the *Reset* button to revert to default. Result is reflected in the numerical *textboxes* on the right-hand side.

BlueBanana Use Case #1

Color Transform/Remapping

- 1) Load a video track, and add the *BlueBanana* plugin to your video. The alpha channel is usually all opaque. This serves as an initial full screen selection mask.
- 2) Open the controls, and start with all boxes unchecked. Now reduce the selection using the top pane in *intersection* mode (that is *Combine Selection* is unchecked) to begin the effect.
- 3) Use the *eyedropper* on the *compositor* window to choose a particular color.
- 4) Click on the 3 plugin *Pick* boxes on the right side of each line of HSV to get the color selection.
- 5) Check *Mark Selected Area*. The affected zones will be identified on the composer.
- 6) Adjust the selection using the *HSV* and *Fill* sliders of the top pane.

The *selection mark* will be updated as you operate the controls.

The *composer mask* striping will be strongest as the mask is nearer *full selection*.

- 6) Now uncheck *Mask selected area* & check *Filter Active* to begin Color Adjustment.
- 7) Enable any needed colorspace modifiers, RGB / HSV sliders, and setup the color changes by moving the sliders. The current output may be the desired output.
- 8) Enable *Mask Section* and the *alpha output* will pass the *selection mask* to the *image alpha channel*. This can be used as a very flexible *chroma-key* filter. It also allows more plugins to be stacked and more selection information to be added, either by *intersections* or *unions* with other selections.
- 9) *End Mask* simply sets the output image alpha to opaque. This is normally used to end a stack of BlueBanana plugins, and render the entire image with a complex selection.

BlueBanana Use Case #2

This case uses stacked BlueBanana plugins working like *chroma-key* filters. It assumes you have already learned how to operate the plugin.

- 1) Bring up 2 tracks of video media – one for *foreground* and one for *background*.
- 2) Add 2 BlueBanana plugins on the *first* track. *Turn off* all checkboxes in both plugins.
- 3) On the top plugin, use the top pane to create a *selection mask*, using *Mark Selected Areas*.
- 4) *Turn off* top plugin *Mark Selected Areas*, and disable the top plugin via the plugin title bar *on/off*.
- 5) Create another selection using the second plugin's mask, using *Mark Selected Areas*.
- 6) *Turn on* the top plugin. Make sure both plugins *Mark Selected Areas* is off.
- 7) Check *Mask Selection* and *Filter Active* in both.
- 8) Check *Combine Selection* on second BlueBanana to see the final results.

You will see that there is *intersection* of the full plane with the first chosen regions, so the alpha = 0 everywhere but the area you picked and you see through. And alpha = 1, where the intersection selection was 1. The *Normal* blend shows you the track on top in these regions (the foreground track where alpha = 1).

If you are building an alpha selection mask by first intersection and then union, the top BlueBanana should not change the colors or the bottom plugin will need to target the remapped colors since that is the input to the lower BlueBanana.

Blur

This is a *Gaussian* type *blur*. Other blur plugins – *Linear*, *Motion*, *Radial*, and *Zoom* – are described later. This plugin is *keyframable*. *Blur* is used to blur a video track via the following parameters:

Horizontal and vertical – values are used to tell which one of the *fields* blurring affects; can be both.

Radius - use this *dial* to define the amount of blur to apply.

Alpha determines radius – use alpha to define the amount of blur to apply. (radius=gray value of alpha)

Blur alpha, red, green, blue - specifies which color channels is to be blurred.

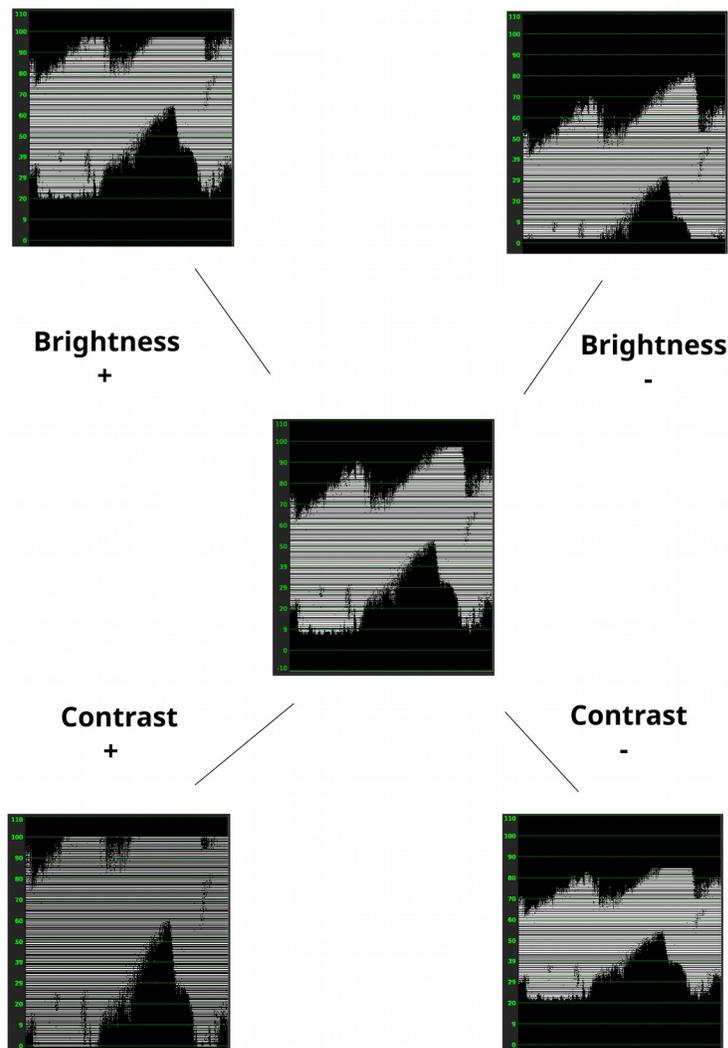
Brightness/Contrast

To brighten a dark shot, or add light, use this plugin. Do not overuse the effect or you risk degrading your video quality.

The *Brightness* slider moves up or down the values of the entire channel and corresponds to the *Master Offset* of the various grading programs.

The *Contrast* slider expands or narrows the brightness values of the entire channel; corresponds to the use of the *cursors* (small triangles) in the *Histogram* plugin.

Use the effect along with keyframing to brighten a long shot that is dark at the beginning but bright at the end. Generally you will want to change the brightness and contrast about the same amount (for example - darkness 28, contrast 26) so that your *original colors* are kept intact. This effect is also keyframable.



BurningTV

Makes your video *burn* where there are small light colored patches of video. This came from <https://effectv.com>.

C41

(credit Florent Delannoy, original program code author, and Edouard Chalaron)

The C41 plugin takes a 16-bit C41 digital intermediate *negative film* as *input* and *outputs* a *positive image*. It became necessary because C-41 negatives can *fade* or *color-shift* over time which was a problem early on. It is still important today because there is a large amount of documentaries, video clips, and other media out there that was shot on super 16 film. This works for RGB-float, RGB, and also YUV variations.

There are two sets of data - the *scanned input values* and your *corrected values*. Simple functionality of the plugin is to compute the data, transform to get corrected values, then apply that.

Basic usage strategy:

1. first time the controls come up, nothing is checked and everything is set to 0
2. check the box *Compute negfix values* to see the current media input values
3. check *Activate processing* and you see a 1-colored screen in the *Compositor* due to zero values
4. check the *Apply values* box to see the input values on the left side propagate to the right side
5. check *Apply default box* if you want to make sure that the borders of the image are not used
6. correct the *output values* as desired on the applied right side

It is important to note as you play or change the frame, the plugin re-computes the data as you move along, but it is not propagated to the applied side.

Checkboxes:

Activate processing – when checked, the c41 operation is used to *render* the image.

Compute negfix values – computes the current negative values of the image (inside the box).

Show active area – draws horizontal and vertical grid lines displaying the boxed area.

Postprocess – when checked, applies contrast/brightness values as defined in coef1/2.

Values:

Compute negfix values (left side) and *negfix values to apply*(right side):

Min/Max R/G/B – minimum and maximum values for Red, Green, and Blue.

Light – value of light; a smaller number is lighter.

Gamma G/B – values for gamma Green and Blue.

Contrast – simple color contrast.

Brightness – white brightness.

Buttons:

Apply values – copies computed RGB/Light/Gamma/Contrast/Bright from negfix to applied values.

Apply default box – copies default computed Box column/row from negfix to applied values.

Shaving box:

The boxing option allows for calculating the inversion of the digital negatives in a given area of the frame as opposed to the entire frame. The program will automatically calculate the columns and rows to shave from the frame when *compute negfix values* is checked. A default box area is initially calculated, called the *shaving box*, based on where the min/max difference in a row/column is less than the program defined tolerance. This row/column minimum and maximum difference must be greater than 0.05. The effect is to cut away the border areas with constant color. If you check the *Show active area*, you can see the box in the *compositor* window. The boundary search is constrained to a range of 0.1 to 0.9 times the frame dimensions, to create a 10 percent shaved margin to avoid over-scan and

negative edge bleeding. Manual adjustment of the shaving box is controlled via the *four sliders* on the bottom right which move each of the left, right, top and bottom shaving margins. The slider bar new values automatically take effect as you move the box and you will see the right-hand side applied values change. When you have either the rows or the columns where the minimum slider is greater than or equal to the maximum slider, the default box will be in effect instead.

Optional postprocessing:

Contrast is the difference in brightness between objects or regions.

Brightness refers to the overall lightness or darkness of the image.

In order to have the values of *Contrast* and *Brightness* take effect, you must check the *Postprocess* checkbox.

Screenshot below shows the C41 controls on the left and part of the *Compositor* window with grid lines showing the default shading box since the Show active area box is checked. Changes have been made to the left-hand side original computed values as seen in the right-hand side such as Gamma G which contains the hairline cursor and has a partial red outline value box.

Negative film, before image

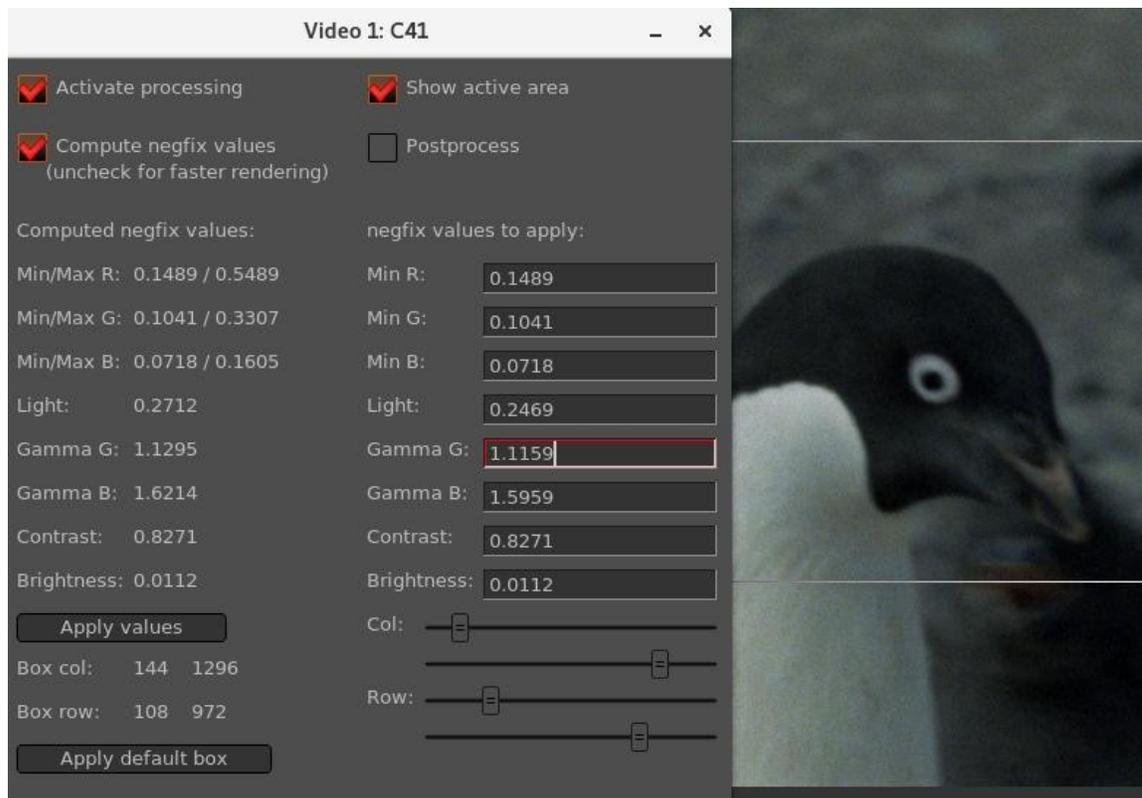


Fig. 67: C41 - Control window and compositor window in action

Chroma key

This effect erases pixels which match the selected color. They are replaced with *black* if there is no *alpha* channel and *transparency* if there is an *alpha* channel. In this case, you create a *matte* in the alpha channel, which is not visible to us. The selection of *color model* is important to determine the behavior.

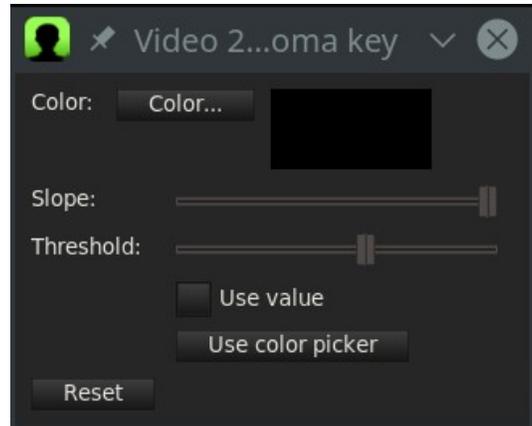


Fig. 68: Chroma Key control window

Chroma key uses either the *lightness* or the *hue* to determine what is erased. *Use value* singles out only the *lightness* to determine transparency.

Select a center color to erase using the *Color button*. Alternatively a color can be picked directly from the output frame by first using the *color picker* in the *compositor* window and then selecting the *Use color picker* button. This sets the chroma key color to the current *color picker* color.

Be aware that the output of the chroma key is fed back to the *compositor*, so selecting a color again from the *compositor* will use the *output* of the *chroma key* effect. The chroma key should be disabled when selecting colors with the color picker.

If the *lightness* or *hue* is within a certain *threshold* it is erased. Increasing the *threshold* determines the *range of colors* to be erased. It is not a simple *on/off switch*. As the color approaches the *edge* of the threshold, it gradually gets erased if the *slope* is *high* or is rapidly erased if the *slope* is *low*. The *slope* as defined here is the number of extra values flanking the *threshold* required to go from opaque to transparent.

Normally threshold is very low when using a high slope. The two parameters tend to be exclusive because *slope* fills in extra *threshold*. The slope tries to soften the *edges* of the *chroma key* but it does not work well for *compressed* sources. A popular softening technique is to use a *maximum slope* and chain a *blur* effect below the *chroma key* effect to blur just the *alpha*.

Chroma key (HSV)

(Credit for Plugin by Jerome Cornet <http://jcornet.free.fr/linux/chromakey.html>)

Chroma Key (HSV) replaces a color with another color or transparency using HSV variables; it is frequently used to remove a color from a video to composite with another image. This process is generally referred to as *green screen* or *blue screen* process (because of the color that is keyed out).

More information: <http://en.wikipedia.org/wiki/Chromakey>

Requirements

The subject in the movie should have a good background. The lighting is crucial and good lighting during production will save you time with much less effort than in post-production.

Here we assume that we have a good video, filmed on green (or blue) screen that we want to use.

Important: Make sure you are using a *color model* that has an *alpha* channel, such as *RGBA8*, *RGBAFloat*, *YUVA8*. To change color model, go to **Settings** → **Format** → **Color Model**.

Usage

As in any other effect, add it to the timeline in the main window. You can tweak each parameter in order to improve the keying.

Start with *Hue Tolerance* at 10%, *Min Brightness* at 0, *Max brightness* at 100%, *Saturation offset* at 0, *Min Saturation* at 0, *In Slope* at 0, *Out Slope* at 0, *Alpha Offset* at 0 (that's mid-way through), *Spill Threshold* at 0, *Spill Compensation* at 100%. At any time, you can check what the *Mask* looks like by clicking on *Show Mask*. This will output a black and white image of the *mask (matte)*.

- *Select the Key color:* Select the key color (green, blue, etc) using the *color wheel* or the *color picker*. Remember, only the *Hue* matters, not *Saturation* or *Value*. To use the color picker, click on the color picker icon in the *Compositor* window, then click on the color you want in the *Compositor* window. Finally in the *Chroma key (HSV)* parameters window, click on *Use Color Picker*.
- *Adjust the Hue Tolerance:* Because there are slight variations in lighting, the background will not be in a uniform key color *hue*. Increase or decrease the Hue tolerance to mask out the background. If there are *dark spots* that are keyed out that shouldn't be, it can be corrected later.
- *Adjust the Brightness:* Increase *Min Brightness* so that only the background is masked out, and not parts of the foreground. You can also reduce *Max Brightness* if some clear areas are keyed out (useful for very dark backgrounds).
- *Adjust the Saturation:* Increase *Min Saturation* so that only the background is masked out, and not parts of the foreground. *Saturation Offset* can be used to change this, but for now leave it set to 0.

Check what it looks like at this stage, your *mask* should be pretty clean. Toggle *Show Mask* to check what it looks like, it should be OK. If not, repeat steps 1 to 4 to get a better key. The rest of the controls are useful to *smear* the *mask* to help compositing later on. They will help you to make your key look much cleaner.

- *Adjust the slope:* For now, the mask is a full on/ full off mask that can be really harsh and not necessarily what you are looking for. *In Slope* and *Out Slope* will help you to smooth that key. *In Slope* leaves *more* colors in the mask, *Out Slope* takes more colors *out* of the mask. The colors that are borderline in the mask will see their alpha channel reduced by half instead of being completely **on** or **off**.

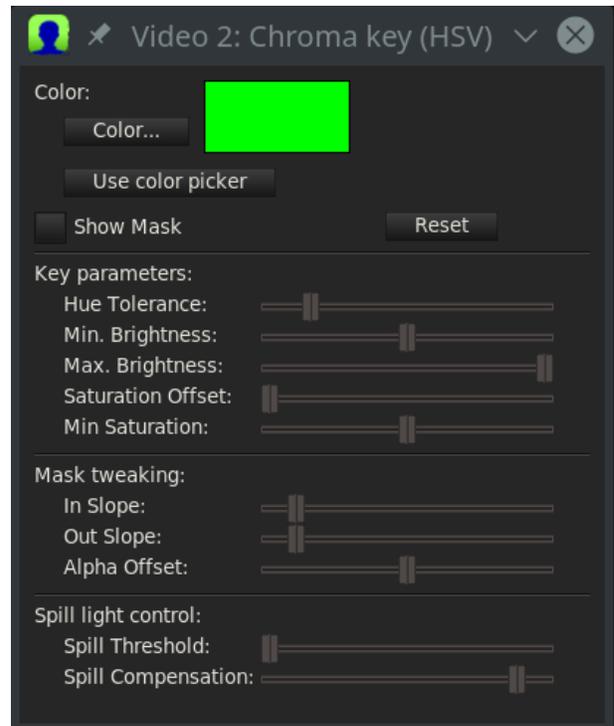


Fig. 69: Keying a green screen with Chroma Key (HSV)

- *Adjust the alpha channel:* This control offsets the whole alpha channel by some amount. Be sure to know what you are doing if you change it from the default value of 0.
- *Remove spill light:* This step helps you remove the green or blue *halo* around the edges of the mask. It does so by removing the *saturation* of pixels that have a similar hue to the key color (turning them into grey instead of green or blue). *Spill Compensation* controls the amount of desaturation. If you start with *Spill Compensation* at 100%, slowly increase the *Spill Threshold* until the remaining green or blue areas turn grey. Then reduce *Spill Compensation* until the image looks good.
- *Blur the alpha channel:* Now the *mask* is probably still very harsh, so just below the *ChromaKey (HSV)* plugin, add a *Blur* effect, and select only the *Alpha* channel, with a radius of 2 or 3 (more if you really want to soften the edges). This will significantly help the keying.

Color 3 Way

Together with *Histogram Bezier / Curves* is the main tool of Color Grading because you can modify the colors of *Shadows*, *Midtones*, and *Highlights* as desired. *Color 3 Way* is keyframable.

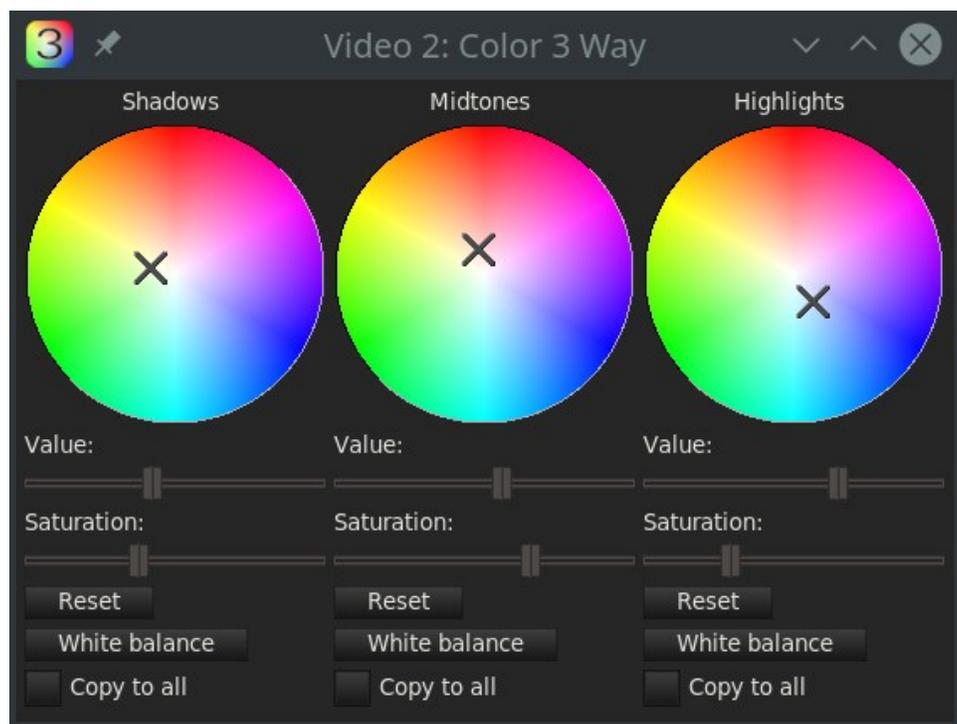


Fig. 70: Color 3 Way control window

- It allows you to vary the *contrast* of the image using the slider *Value*, always acting separately on *shadows*, *midtones*, and *highlights* and thus resulting in very precise application.
- Allows you to automate the *white balance* by simply choosing a *neutral color* in the output of the *Compositing* window using the *Color Picker* and pressing the corresponding button in the plugin.
- Allows you to vary the *Saturation* with sliders in the same manner as *contrast* was varied by the *Value* slider. For instance, to decrease the incidence of color dominants present in the *shadows* or in the *highlights*, vary the *Saturation*.
- With the *color wheels* you can make very sophisticated adjustments to the *shades* of the images, in each of the three main areas of *shadows*, *midtones* and *highlights*.

- Allows you to *copy* exactly the setting of one zone to the other two zones using *Copy to all button*.

This plugin allows maximum control over the result and maximum precision of adjustments when used simultaneously with the *control monitors*, i.e. *Waveform*, *RGB Parade* and *Vectorscope*. It is important to keep in mind that the three zones are not clearly separated, but slightly overlapping. This results in less precision but looks better for more *smooth* shades. By varying the values on the *color wheels* all *RGB channels* are affected simultaneously, which can result in unwanted color dominance. *Saturation* is also affected and must therefore be monitored.

To use more precisely, drag the *crosshair* with the mouse in the desired area and then adjust in steps of 0.001 using the *up/down* and *right/left* arrows on the keyboard.

The most common use cases (but can be adapted to virtually any situation) of the plugin are:

- White balancing.
- Expand/compress contrast.
- Mitigate under and over exposure.
- Balance colors, i.e. eliminate color dominance.
- Color matching Shot to Shot.
- Create a *Stylized look*.

Color Balance

Video *Color Balance* is a great effect to use along with *Brightness/Contrast* and *Hue/saturation* to try to compensate for possible errors in filming (low lighting, for example). It can only do so much without greatly lowering the quality of the video. With it you can change the colors being sent to *output* CMY (Cyan, Magenta, Yellow) or RGB (Red, Green, Blue). Color Balance is also keyframable.

Since *complementary colors* are neutralized, to eliminate a color cast, the pertinent slider is moved in the direction of the complementary color. If you *Lock parameters* you get the same *Color Offset*, that is the fourth color wheel in the grading programs. The parameters of the plugin are:

CMY/RGB sliders - allows you to adjust the colors.

Preserve Luminosity - Adjusts colors while keeping the overall brightness constant.

Lock Parameters - works as a *Color Offset*.

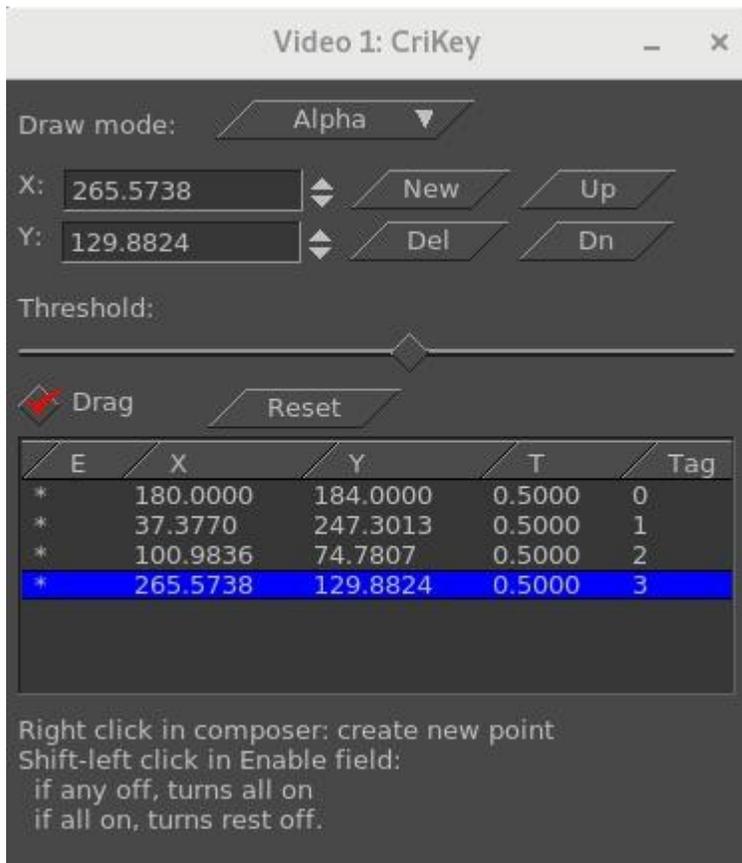
White Balance - used in conjunction with the *color picker* on a *neutral color* of the output, it will automatically balance the white.

CriKey

The *Chroma Interpolation Key* plugin, CriKey, is a *regionally* based *chroma key* with interpolation. This is useful when you only want 1 or some specific zones to be defined by the chroma key as opposed to the entire image. Its most significant feature is that you can select several regions of interests and of different colors as opposed to only 1.

To start, if not already checked, turn on *drag*. In the *composer* window select an area of a certain color by clicking on that point with the *right mouse button* and check to see that it is enabled with an * in the **E** field. The color of the area is used to define the *region of interest* and then you can use the *threshold* slider to designate the tolerance variation. This creates a region that is the *chroma key* selection and a *fill* will be performed in that area, but only within that region. So, say for example, a red colored area was chosen, only the red color inside the region is selected – not that color red in the entire image. The *drag* capability makes it easy to check a point before right clicking it to see the effect. You will want to

turn off *drag* when you are finished with *CriKey* so that it does not interfere with other compositor functions.



← *Alpha, Edge, and Mask* draw modes

← *New* to create a new point, *Up/Dn* to move highlighted point up or down

← *Del* to delete the highlighted point

← *Threshold* slider varies from 0 to 1

← *Drag* for ease of use / *Reset* for defaults

← *E* for Enabled with * marking that

X is the point's x coordinate

Y is the point's y coordinate

T is the threshold value of X, Y point

Tag represents the # of the selected point

← Hints for usage shortcuts

Fig. 71: Four active point created in *CriKey*

Draw mode - options let you use *Alpha* for see-thru, *Edge* to just outline the edges of the region, or *Mask* to block. The pixels which match the selected color are replaced by *black* if *Mask* is chosen or see-thru/transparent if *Alpha*.

Del – delete the highlighted point in the listbox

Threshold - slider goes from 0 to 1. Increasing the *threshold*, increases the area to be filled or masked. You can also use the *mouse wheel* to scroll the slider.

Reset - button to revert to only the default middle point with all others being deleted.

Some notable caveats:

- 1) When choosing an area that has variations of the same color within a region, for less work and for the best results, choose an average color in that region instead of an extreme end of that color.
- 2) If the *threshold* is set appropriately you can see the *edges* which is helpful.
- 3) The mask is computed and shows the *fill* region.
- 4) Use the *Gradient* plugin to substitute a different color for the selected area.

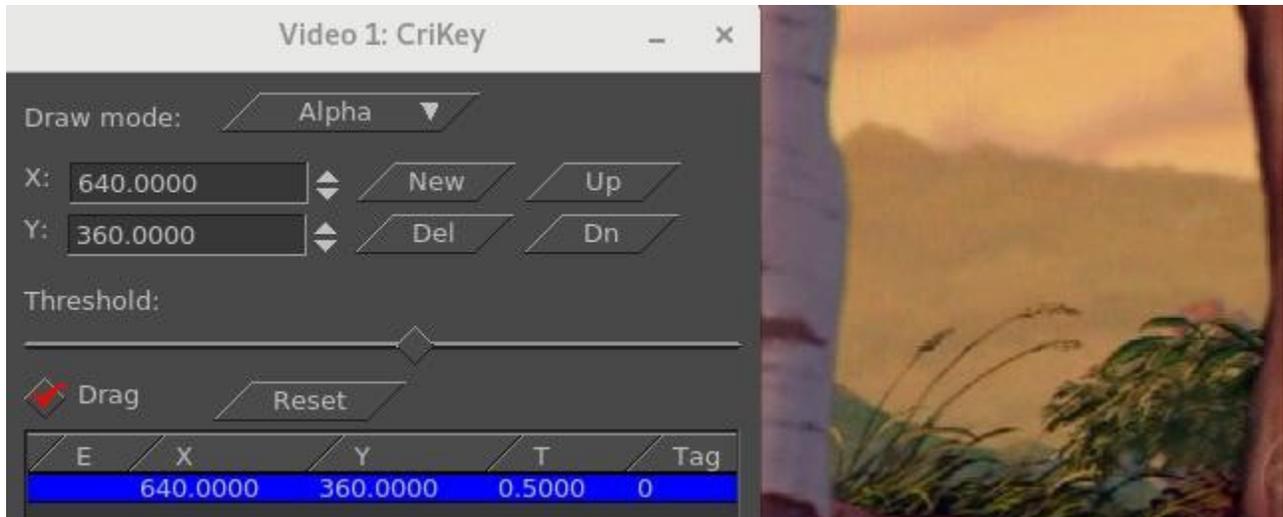


Fig. 72: The screenshot shows the compositor with some default settings in the controls window.

Now this same screen shot shows how moving the *Threshold* slider with the *Point* selected blacks out the single region which has the darker brown hills in it. Because the edge was located, any of the same color in the rest of the video would not be blacked out.



Fig. 73: same screenshot with moving *Threshold*

Usage Steps:

- 1) Click *Reset* (there will be a single X,Y coordinate point that is in the middle and not enabled)
- 2) Check to make sure *Drag* is on.
- 3) In the *Compositor*, right click on area of interest and an X,Y coordinate will appear in the *listbox*.
- 4) Click on the *E Enabled* field next to this latest point and an * asterisk will show.
- 5) Now you will see an area turn *black* so use the *Threshold* slider to only black out the area of interest.
- 6) Repeat steps 3-5 until you have selected all of the desired areas.
- 7) Finally, turn off *drag* so as not to interfere with other compositor functions

DeScratch

The descratch video plugin can be used to remove vertical scratches from film. It can also be used, *after image rotation*, to remove horizontal noise lines that may appear on analog VHS captures. For best results *YUV* should be the video format; however if your format is *RGB*, it will first be converted to *YUV*. There are many tuneable parameters necessary to get good results for your specific film. Screencast of *DeScratch* with a list of the parameter descriptions below.

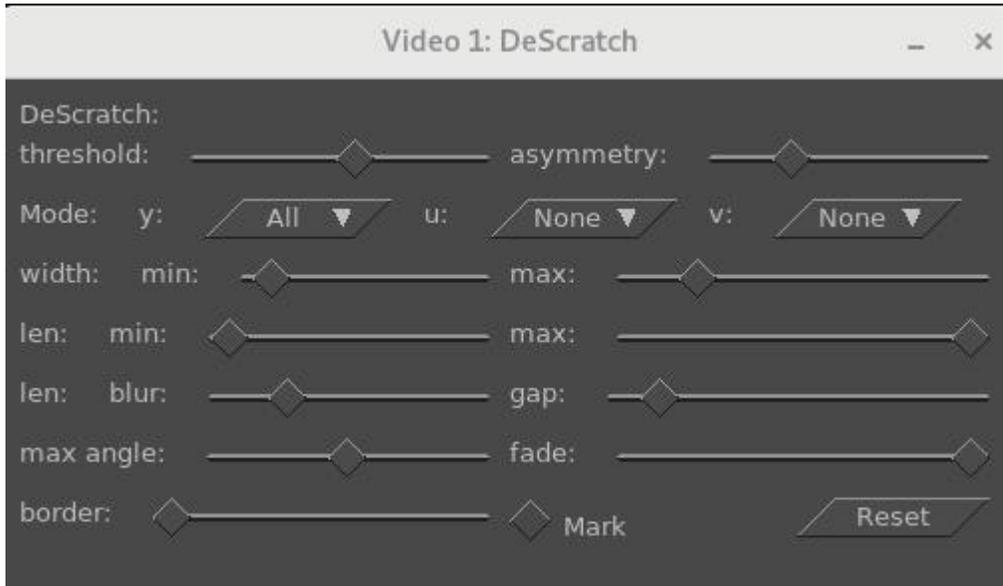


Fig. 74: DeScratch control window

threshold – instantaneous *slope* value; chroma difference in numerical pixels.

asymmetry - maximum asymmetry of surrounding pixels.

Mode - *None*; *Low*=black; *High*=white; *All*=both

y - processing mode for *luma* plane

u - processing mode for chroma *u* plane

v - processing mode for chroma *v* plane

width min -minimal scratch width in pixels.

max - maximum scratch width in pixels.

len min – percent minimal scratch length.

max – percent maximum scratch length.

blur – scaled radius of vertical blur for frame analysis.

gap - number of pixels for maximum vertical gap to be closed.

max angle - maximal angle to vertical in degrees.

fade – percent of how much it fades to and how much it

uses between before image and blurry image.

border - pixel thickness of border near scratch for partial restoration.

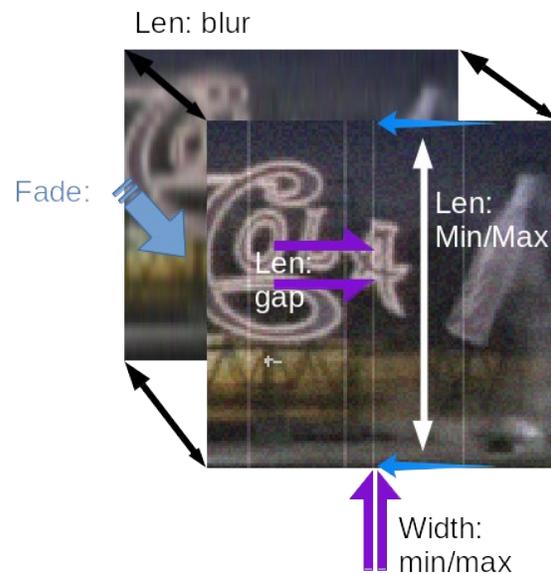


Fig. 75: Various parameters of DeScratch

Mark - shows the potential scratch lines for ease of viewing and for debugging. It shows chosen pixels in the color green, close but still rejected in yellow, and extreme pixels in the color red. This makes it easy to vary some parameters to choose more or fewer scratch lines.

Reset - activating this button returns all of the parameters to their default values.

Screencast below shows a before and after DeScratch scenario.

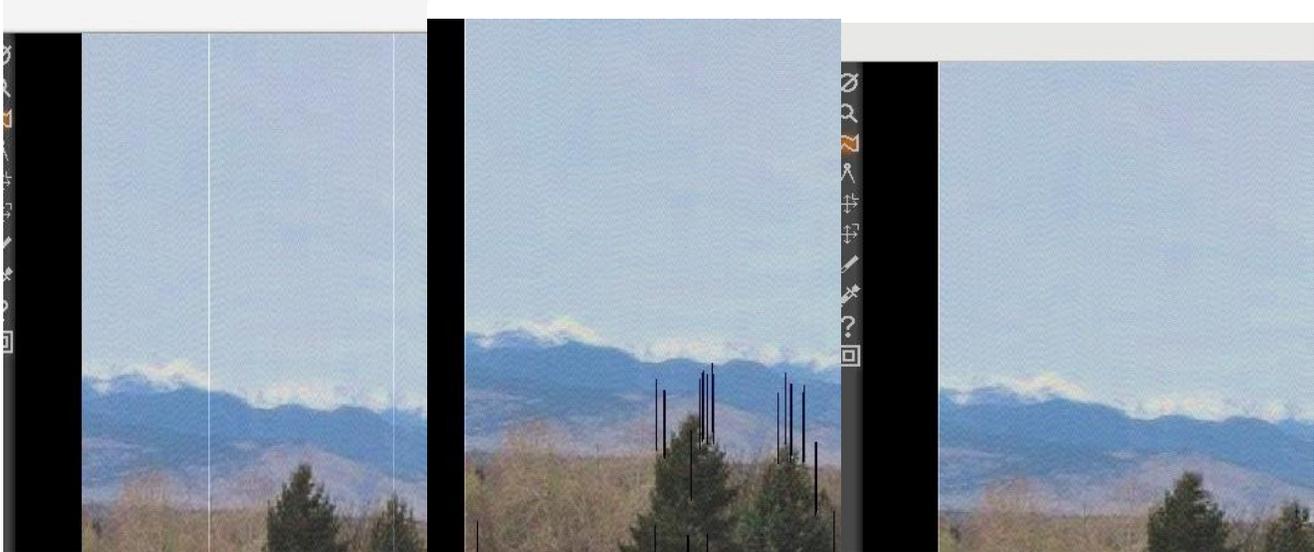


Fig. 76: Original video with scratch

Fig. 77: Option Mark selected

Fig. 78: Final video

with Mark set, you can see the black lines which indicate what the program was looking at to determine the scratches to remove.

Decimate

This is used to decrease the frame rate of a video. Changing the frame rate means eliminating a frame for any given number of frames (*1 in N*); but if frames that are important for visual continuity are deleted, *temporal artifacts* arise: flickering, slowdowns, accelerations, etc. The *Decimate* filter maintains a higher quality because it first eliminates *duplicate* frames or frames that are most similar, thus limiting the appearance of artifacts. It is often used after the *Invert Telecine* plugin to make the video more *smooth*.

One use of the *decimate* effect can be applied to a *DVD* to convert the 29.97 fps video to the 23.97 fps film rate, but the effect can take any *input rate* and convert it to any *lower output rate*. The output rate of *decimate* is the *project* frame rate. The input rate is set in the decimate user interface. To convert 29.97 fps progressive video to 23.97 fps film, apply a decimate effect to the track. Set the decimate input rate to 29.97 and the project rate to 23.97.

Keep in mind that every effect layered before decimate, processes video at the *decimate input rate* and every effect layered after decimate, processes video at the *project frame rate*. Computationally intensive effects should come below decimate.

Deinterlace

The *deinterlace* effect has evolved over the years to deinterlacing and a whole lot more. In fact two of the deinterlacing methods, *Inverse Telecine* and *Frames to Fields*, are separate effects. The *deinterlace* effect offers several variations of line replication to eliminate *comb artifacts* in interlaced video. It also has some *line swapping* tools to fix improperly captured video or make the result of a reverse effect display fields in the right order.

Deinterlace-CV

Selection of deinterlacing mode for your video to eliminate comb artifacts.

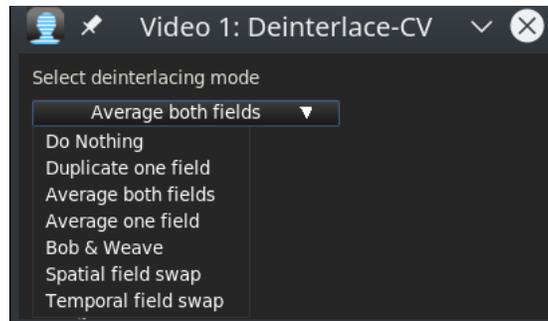


Fig. 79: Pulldown menu

Delay Video

Delay the video by some number of seconds.

Denoise video

Clear the video of noise; specify R,G,B, or Alpha.

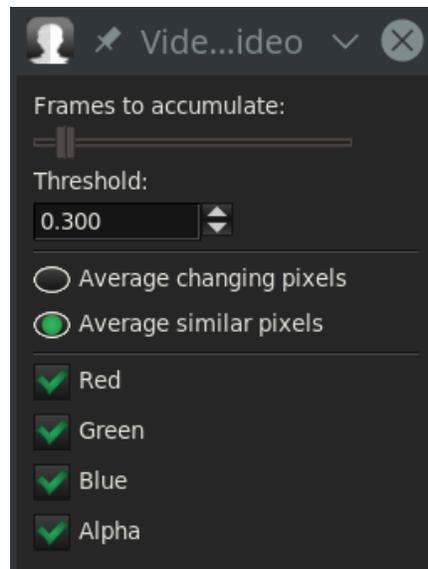


Fig. 80: Control window of the DeNoise plugin

Difference key

The *difference key* creates *transparency* in areas which are similar between two frames. The Difference key effect must be applied to two tracks. One track contains the *action* in front of a constant background and another track contains the background with nothing in front of it. Apply the *difference key* to the track with the action and apply a *shared* effect of it to the track with the background. The track with the background should be **muted** and *underneath* the track with the action and the *color model* should have an *alpha* channel. It's hard to get good results.

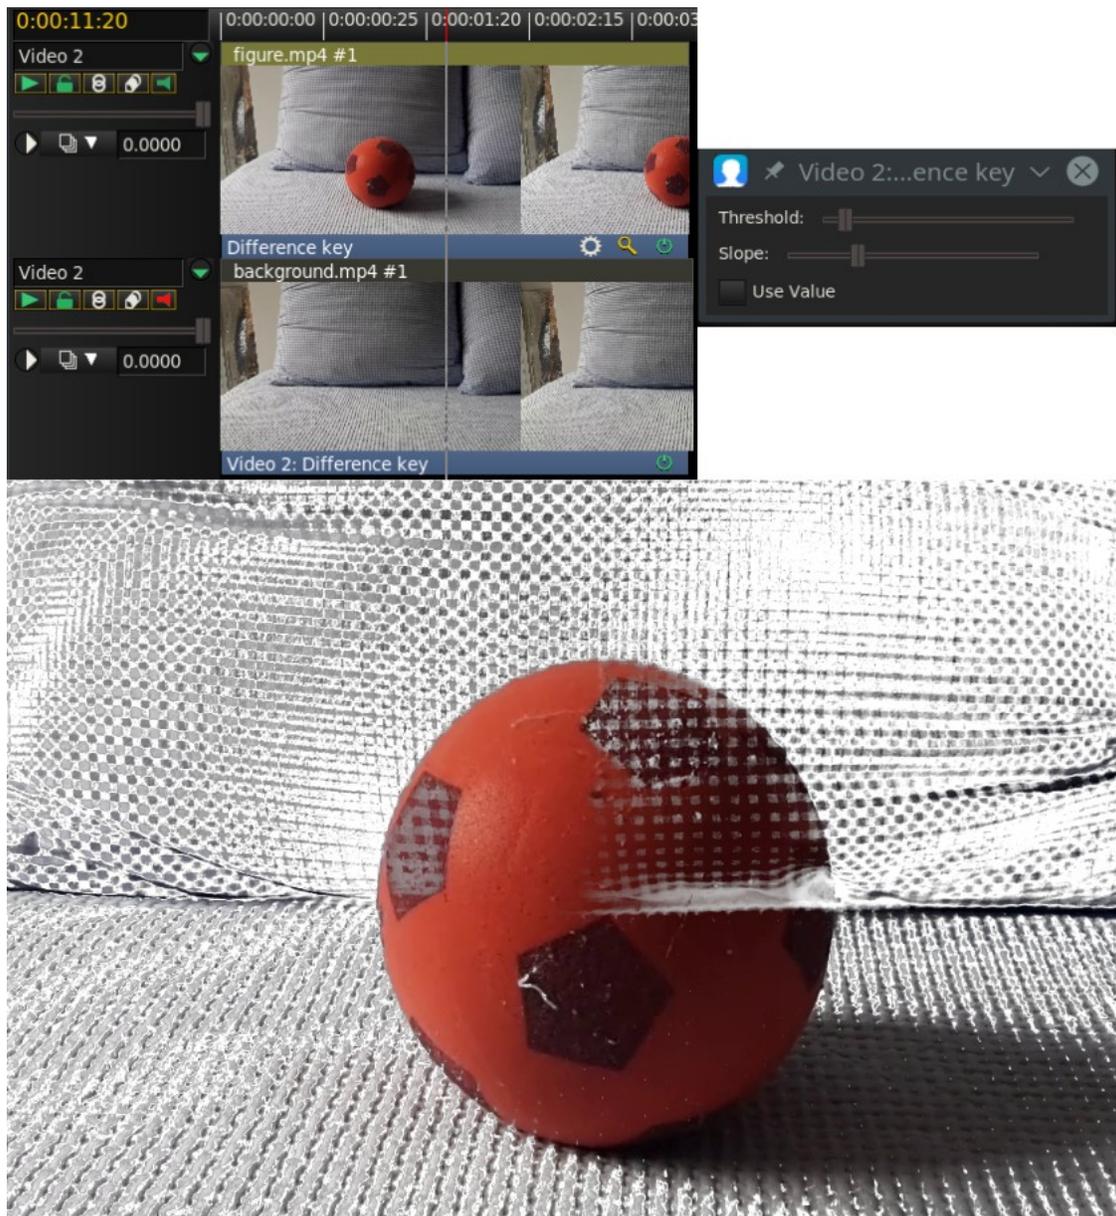


Fig. 81: The two tracks, the control window and the final result

Pixels which are different between the *background* and *action* track are treated as *opaque*. Pixels which are similar are treated as *transparent*. Change *threshold* in the difference key window to make more pixels which are not the same color transparent. Change *slope* to change the rate at which the transparency tapers off as pixels get more different. The slope as defined here is the number of extra

values flanking the threshold required to go from opaque to transparent. A high slope is more useful with a low threshold because slope fills in extra threshold.

Use value causes the intensity of pixels (luma) to be compared instead of the color. Applying a *blur* to the top track with just the *alpha* channel blurred can soften the transparency border.

DotTV

Puts various size *dots* over the picture to simulate TV effect. This came from: <https://effectv.com>.

Downsample

Downsampling is the process of reducing the size of an image by throwing out data, reducing *sampling rate*.

Edge

Display only the edges of the video throughout the image.

Fields to frames

See the theory description in the *Frames to Fields* plugin. This effect reads frames at twice the project framerate, combining two *input* frames into a single *interlaced output* frame. Effects preceding *fields to frames* process frames at twice the project frame rate. Each input frame is called a *field*.

Fields to frames needs to know what *field* corresponds to what *lines* in the output frame. The easiest way to figure it out is to try both options in the window. If the *input fields* are the result of a line doubling process like *frames to fields*, the wrong setting results in *blurrier* output. If the input fields are the result of a standards conversion process like *1080 to 480*, the wrong setting will not make any difference.

Flip

This effect flips a video track either *vertically* or *horizontally*.

Frames to fields

Theory behind the Frames to Fields and Fields to Frames plugins:

Historically, CRT-type TVs used *interlaced* signals to save bandwidth. An *interlaced* video consists of two *fields* that are read and drawn on the screen one after the other. Each *field* must be played at a *framerate* double that of the resulting video. In two steps the complete frame will be reconstructed.

frame 1 => F1-field1 (*Top* or *Odd*), F1-field2 (*Bottom* or *Even*)

frame 2 => F2-field1, F2-field2

Interlaced video reading: **F1-f1** then **F1-f2** then **F2-f1** then **F2-f2** ...

There may be visual problems if the *Top* type interlacing is read according to a *Bottom* scheme. So it's important to know if a video is *Top* or *Bottom*. Generally an *HD* video is *Top*; a *DV* video (both PAL and NTSC) is *Bottom*; *SD* (PAL) is *Top*; *SD* (NTSC) is *Bottom* (but not always). Instead, *high-definition* videos need to be more compressed and this contrasts with the interlacing that is little and badly compressible, so modern videos are mostly *progressive*.

In Cinelerra-GG:

- 1 - upload an interlaced video to the *Timeline* and *Resources* and play it for viewing.
- 2 - The video presents visual *artifacts* because PC monitors are *progressive*.
- 3 - In the *Resources* window, open the *media Info* with the right mouse button. Below you can see that the *Fix interlacing* checkbox is *active*. It has four options that try to deinterlace the video automatically: *Unknown*, *Top Fields first*, *Bottom Fields first*, and *Not interlaced*.
- 4 - If *Fix interlacing* is deactivated, *Interlacing correction* becomes active with three options: *Do nothing*, *Shift up 1 pixel* and *Shift down 1 pixel*. These are *manual* settings and are more effective than *automatic* interlacing fixes.
- 5 - If we want to use the *Frames to Fields* plugin we have to configure it and act manually.

Now for the practical use of this plugin which applies the operation reverse to the *Fields to Frames* plugin. It extracts the two interlaced fields stored in alternating lines of interlaced source footage and outputs them as separate *full frames*. The alternating lines missing on each output frame are *interpolated*.

This plugin is only useful if its output is pulled with doubled framerate with respect to the source footage. One typical usage scenario is to do *masking*, *scaling* and *translating* on interlaced footage without the need to destroy the additional temporal information contained in such source material. This is helpful if your intended *target format* is *interlaced*. If on the other hand, you just want to target a progressive display (for example, you create video for display on a computer monitor solely) then it is much more convenient to *de-interlace* the source material prior to any further processing.

Processing interlaced footage without deinterlacing:

- 1 - Create a new *project* with doubled frame rate. That is, make it 50fps if your source footage is 25i. In **Resources** → **Media** → **Info** uncheck *Fix Interlacing* and set *Interlace Correction* to *Do nothing*.
- 2 - Insert your *source* footage onto a video track in the timeline. Now, Cinelerra will playback each frame of your footage twice. There will be visual artifacts because the video is *interlaced* and the monitor is *progressive*.
- 3 - Apply the *Frames to Fields* effect. Be sure to choose the correct *field* order. If we know or believe that the original video is *Top First* let's try it first, but it doesn't have to be the right solution. The only way is to playback and look for visual artifacts.
- 4 - Then apply any further effects afterwards, including *translations*, *scaling*, *slow motion*, precise frame-wise *masking* or use of the *motion tracker* plugin.
- 5 - **Render** your project to an *intermediate* clip. Be sure to choose a rather *lossless* video codec, for example *Motion-JPEG-A* or even *uncompressed YUV* if you have plenty of storage.
- 6 - Insert the intermediate clip into your original project. Make sure the doubled framerate has been detected correctly by Cinelerra (by looking in the clip's *media info* in the *media resources* folder).
- 7 - Apply the *Fields to frames* effect to the *intermediate* clip. This will combine two adjacent *fields* into one interlaced field with the original frame rate.
- 8 - Do the final **render** on your original project. Now there will be no visual artifacts on the monitor.

Freeze Frame

In its simplest form, *highlight* a region of the track to freeze, drop the *freeze frame* effect on the highlighted region, and the *lowest numbered frame* in the affected area will play throughout the entire region. Freeze Frame has an enabled option which can be *keyframed*. Regions of a freeze frame effect

which are enabled, repeat the lowest numbered frame since the last keyframe. This has unique possibilities.

- If a freeze frame effect has a *keyframe* in the *middle* of it set to enabled, the frame in the middle is repeated in the entire effect.
- If a freeze frame effect has several *keyframes*, each set to enabled, every time a keyframe is encountered the frame under it becomes the frozen one.
- If a freeze frame effect alternates between enabled and disabled, each time an enabled keyframe is encountered the frame under it is replicated until the next disabled keyframe. The disabled regions play through.

Gamma

Raw/Log camera images store colors in a *logarithmic scale*. The *blacks* in these images are nearly 0 and the *whites* are supposed to be infinity. The *graphics card* and most *video codecs* store colors in a *linear scale* but Cinelerra keeps raw/log camera images in their original logarithmic scale when it renders them. This is necessary because the *raw image parser* can not always decode the proper *gamma* (γ) values for the images. It also does its processing in 16 bit integers, which takes away a lot of information.

Mathematically, the *gamma* function is *exponential* (**output = input ^{γ}**) and therefore the inverse of the logarithmic function [**output = log(input)**]. Actually the formula used by the Cinelerra-gg plugin is: **output = input^{1/ γ}** which allows for a range of values 0 – 1.0. The *gamma* effect converts the *logarithmic* colors to *linear* colors through a *gamma* value and a *maximum* value. The *gamma* value determines how *steep* the *output curve* is (i.e. the value of the gamma parameter; for color space *Rec709* is 2.4 ($1/\gamma = 0.41\dots$), for *sRGB* is 2.2 ($1/\gamma = 0.45\dots$), etc.). The *maximum* value is where 1.0 in the *output* corresponds to maximum brightness in the *input*. It serves to avoid *clipped* values because it allows you to set the maximum value of the output, 1.0, whenever range adjustment is done (see the next image). It is important to adjust the two parameters accurately in order to avoid undesired and unexpected effects, such as *excessive values*, *unbalanced image*, *incorrect linearization*, etc.

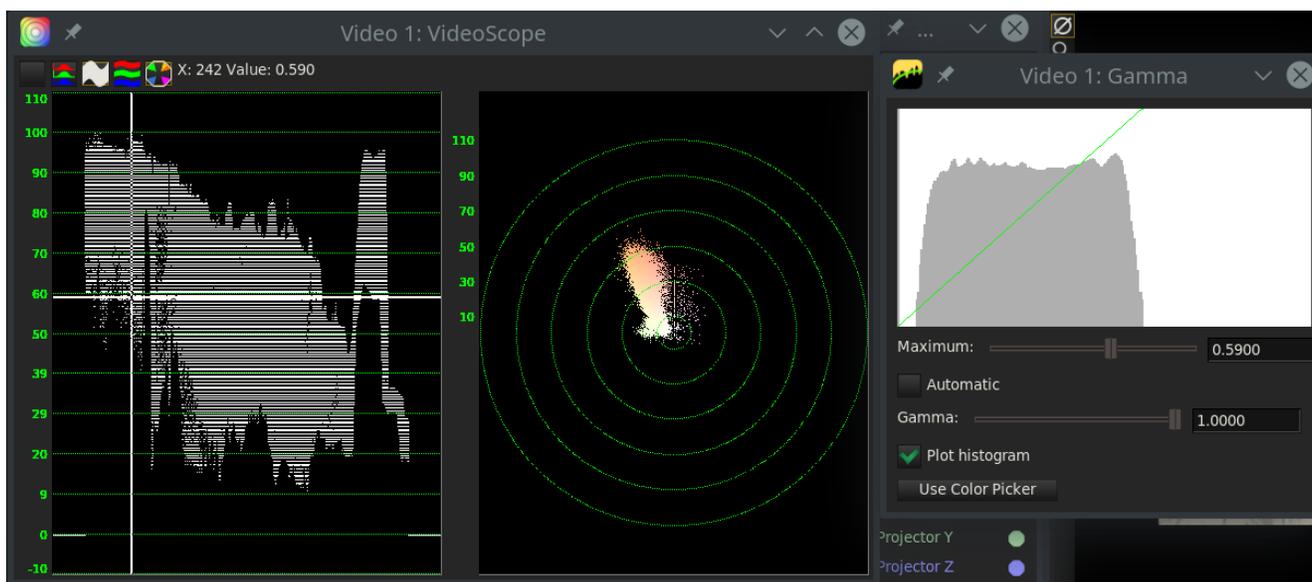


Fig. 82: setting Maximum to 0.5900

The gamma effect has two more parameters to simplify gamma correction. The *automatic* option causes it to calculate max from the *histogram* of the image. Use this when making a preview of a long list of images since it changes for every image. The *use color picker* option uses the value currently in the color picker to set the *maximum* value. Note that every time you pick a color from the compositor window, you need to click on *use color picker* to apply the new value.

The best use of the gamma is *manually* monitoring the waveform as shown in the next image.

- Look at the highest peak on the *waveform* and measure it with the crosshair observing the numerical value at the top left.
- Set this *value* with the *maximum* slider.
- Then adjust the slider of the *gamma* to our liking, always checking the result on the waveform so to be sure never to exceed the values of clipping, 0 – 1.0.

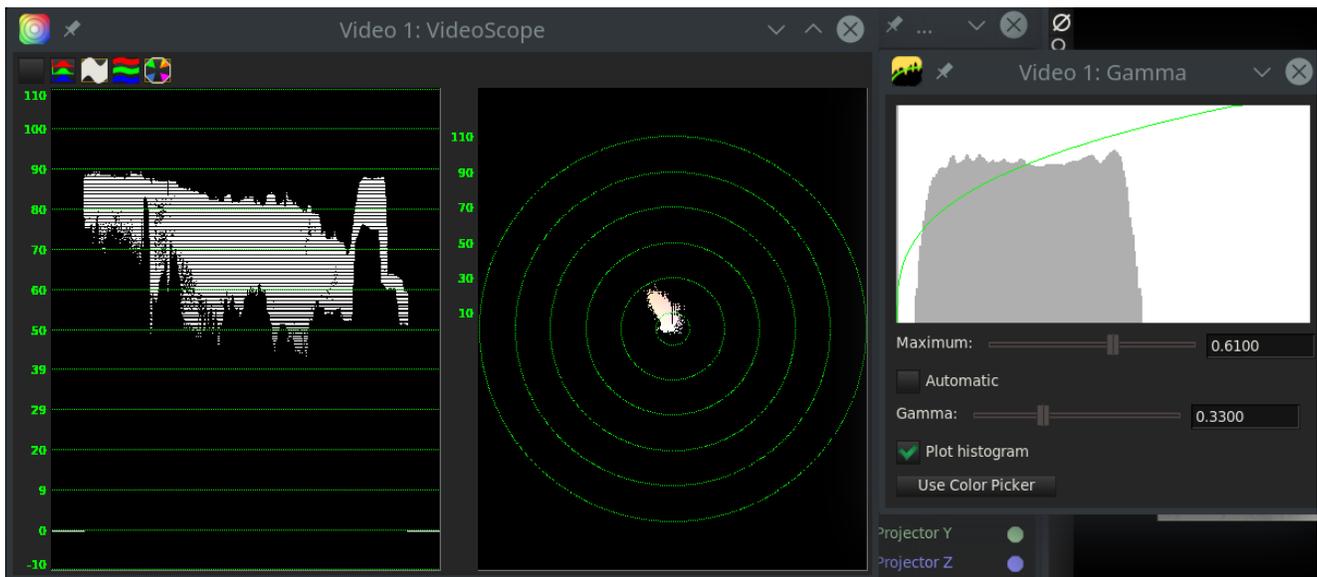


Fig. 83: Setting Maximum to 0.6100 and Gamma to 0.3300

Care must be taken when using *gamma correction*: if the image carries a specific *gamma value*, or if it has already been corrected previously (for example automatically in the *camera*), etc.; then a second application of the *gamma* leads to excessive and artificial results. *Gamma* is keyframable.

Gradient

The *gradient* effect overlays a smooth color gradient on top of every video frame. It is useful for all sorts of *background fills*, for *partially filtering*, adding *depth* to the image, or for adding *moving highlights*. The *Gradient* effect can generate *linear* or *circular color fills / shape*. For *linear fills*, you can choose the *angle*, for *circular fills* the *center (X,Y)* of the created gradient pattern. You can control the *slope* of the color transition by selecting a *transition function (linear, logarithmic, squared)* and by changing the *start (inner)* and *stop (outer)* radius. Note that both colors used in this color transition can contain an arbitrary *Alpha value (transparency)*. All parameters can be keyed and will be interpolated between keyframes.

The first time you use the plugin it may seem complicated, but if you understand that we have to adjust the *gradient* from an *inner spot* we choose to an *outer spot* we also choose, the work will become easy and fast.

Use case (Vignette):

1. Setting the *shape* to radial
2. Setting the *rate* to *Linear* (or *Log* or *Square*)
3. Position *X* and *Y* to center the main figure.
4. Choose inner *color=black*
5. Adjust inner *radius*
6. Bring the inner color *alpha* slider to 0
7. Choose outer *color=black*
8. Check that the outer color *alpha* slider is 1.
9. Adjust outer radius

Note: The inner and outer colors are *visibly mixed* in the gradient area. If you want to make a vignetting of only *black*, you must set the two colors to black and then make the inner one transparent so that it does not cover the figure.

HistEq

Remap colors using *blended histogram weights*. Below shows the GUI and the results in a split screen.



Histeq equalizes the *colorspace* through use of a histogram equalization algorithm - a technique for adjusting image *intensities* to enhance *contrast*. Parameters are:
Gain: when set to 1, the colorspace is *best effort*. If the *gain* is set to 0, the result is the entire regression line of the color map.
Blend - goes between a straight and a twisted line.

Split output – diagonally shows in the *compositor*, the new results on the left and old on the right.
Plot bins/lut – displays a plot of the result.

Histogram

The *histogram* allows an immediate view of the *contrast* amplitude of an image with its distribution of *luma* and *colors* values. If the *columns* of values occupy the whole range 0-100% then we have maximum *contrast*; if the range is smaller the *contrast* is smaller. If most of the values are on the right of the histogram you have an image with *highlights* at the limit with values clamped to 1.0. This is called *overexposure*. However, if most of the values are moved to the left, with the limit of the values clamped to 0, we have a *lowlight* image and we talk about *underexposure*. Histogram is keyframeable.

The *Histogram* is always performed in *floating point RGB* regardless of the *project color space*. The histogram has two sets of *transfer parameters*: the *input* transfer and the *output* transfer. The input transfer has value on the *horizontal* axis of *x*; it is a scale of values ranging from 0 to 255 in the case of an 8-bit image, or it can have *normalized* values in the range (0-1.0) or even be a scale in percentage (0-100%). In the output transfer (the *y-axis*) is represented the number of times (that is, *y*) a given value *x* appears. A higher column (*y* greater) indicates that many pixels have the corresponding value *x*; a lower column indicates that fewer pixels have that value. On the left we have the minimum value 0, which is the *black point*. On the right we have the maximum value 1.0 which is the *white point*. The intermediate values pass smoothly from one extreme to the other. The three important points (including the *midtone*s, i.e. the *Master Offset*) are indicated by *cursors* (small triangles) at the base of the histogram. You can adjust them to change the values of the three points if you want.

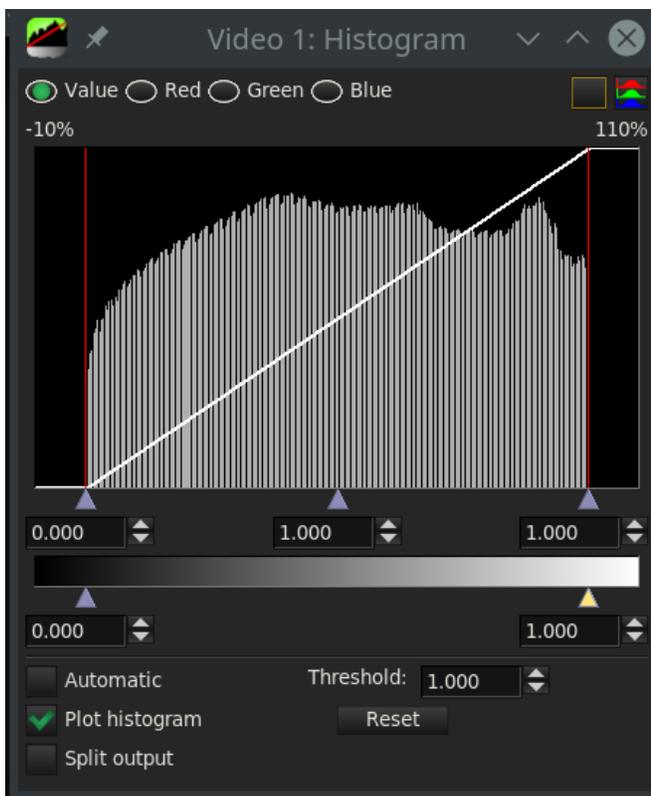


Fig. 85: Master Histogram

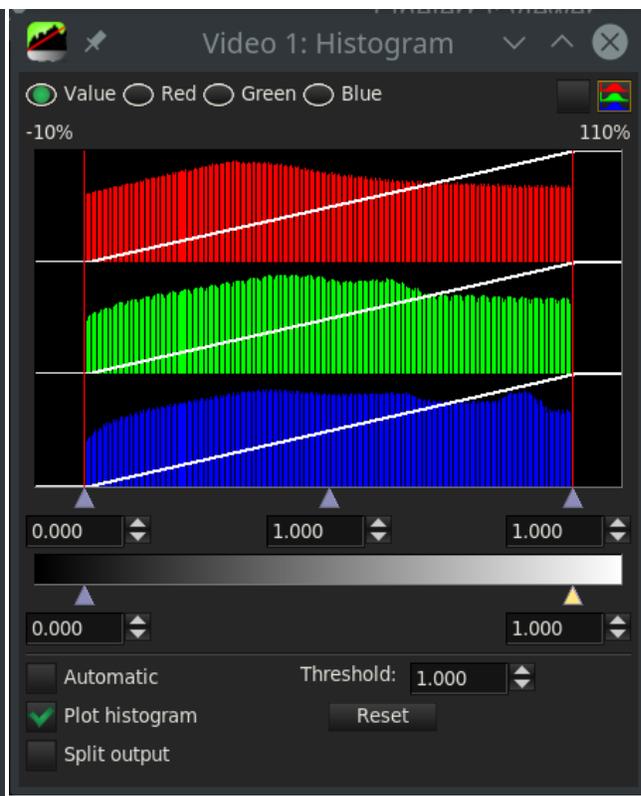


Fig. 84: RGB Histogram

There are 4 possible histograms in the histogram viewer. The red, green, blue histograms show the input histograms for red, green, blue and multiply them by an *input transfer* to get the output red, green, blue. Then the output red, green, blue is scaled by an *output transfer*. The scaled red, green, blue is converted into a value and plotted on the value histogram. The value histogram thus changes depending on the settings for red, green, blue. The value transfers are applied uniformly to R, G, B after their color transfers are applied. Mathematically, it is said that the values of x are linked to the values of y by a *transfer function*. This function is represented by a *line* that leaves the values of x and y unchanged, but we can intervene by modifying this line with the cursors.

You need to select which transfer to view by selecting one of the channels on the top of the histogram. You can also choose whether to display the *master*, i.e. only the values of the *luma*, or show the *Parade*, i.e. the three *RGB* channels. You can switch from one to the other with the two buttons in the upper right corner. The input transfer is defined by a graph overlaid on the histogram; this is a *straight line*. Video entering the histogram is first plotted on the histogram plot, then it is translated so output values now equal the output values for each input value on the input graph.

After the *input transfer*, the image is processed by the *output transfer*. The output transfer is simply a minimum and maximum to *scale* the input colors to. Input values of 100% are scaled down to the output's maximum. Input values of 0% are scaled up to the output minimum. Input values below 0 are always clamped to 0 and input values above 100% are always clamped to 100%. Click and drag on the output gradient's triangles to change it. It also has *textboxes* to enter values into.

Enable the *Automatic* toggle to have the histogram calculate an automatic input transfer for the *red*, *green*, and *blue* but not the *value*. It does this by scaling the middle 99% of the pixels to take 100% of the histogram width. The number of pixels permitted to pass through is set by the *Threshold* textbox. A threshold of 0.99 scales the input so 99% of the pixels pass through. Smaller thresholds permit fewer pixels to pass through and make the output look more contrasty.

Plot histogram is a checkbox that enables plotting the histogram.

Split output is a checkbox that enables a diagonal split showing in the compositor.

Reset returns the four curves to their initial state (neutral) as well as the Value/RGB histogram buttons.

Histogram Bezier / Curves

Histogram Bézier allows an immediate view of the *contrast* amplitude of an image with its distribution of *luma* and *colors* values using a piecewise linear method. In addition it uses a *Bézier curve* (parametric) on the histogram plot. When mapping color spaces, it has a variety of presentations to get smoother transitions and more pleasing output. It uses more general remapping, not just straight lines but more contour lines. *Curves are perhaps the most powerful and sophisticated tool for color correction*. For some repetitive details, see the previous description of the *Histogram* plugin.

The *input graph* is edited by adding and removing any number of *points*. Click and drag anywhere in the input graph to create a *point* and move it. Click on an existing point to make it *active* and move it. The active point is always indicated by being *filled in*. The active point's *input X* and *output Y* values are given in *textboxes* on top of the window. The input and output color of the point can be changed through these *textboxes*. Points can be *deleted* by first selecting a point and then dragging it to the other side of an adjacent point. They can also be deleted by selecting them and hitting *delete*.

Master (value) and *R, G, B* histograms.

Clear button: reset only one color channel.

Textbox for input x (input) and input y (output).

Output min and *output max*: sets *black* or *white* points. If you use both points it works as *Master/Color Offset*. Values can also be less than 0 and greater than 1.0.

Scale for contrast range: sets with cursors shown as little triangles.

Automatic and *Threshold*: enable the *Automatic* toggle to have the histogram calculate an automatic input transfer for the red, green, and blue but not the value. It does this by scaling the middle 99% of the pixels to take 100% of the histogram width. The number of pixels permitted to pass through is set by the *Threshold* textbox. A threshold of 0.99 scales the input so 99% of the pixels pass through. Smaller thresholds permit fewer pixels to pass through and make the output look more contrasty.

Reset: returns the four curves to their initial state (neutral).

Split picture: a checkbox that enables a diagonal split showing in the compositor.

Interpolation of linear, polynomial, or Bézier: type of algorithm for the parametric curves.

Histogram Bézier is keyframable.

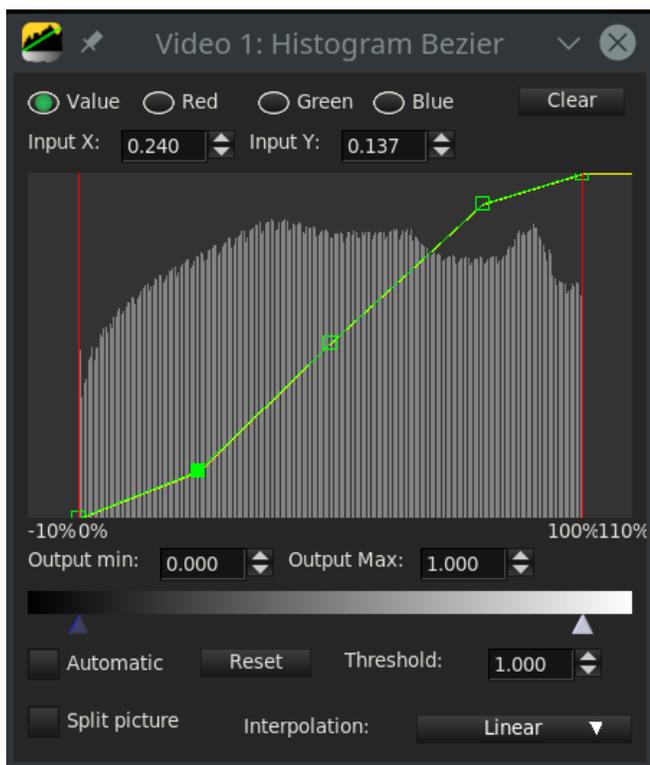


Fig. 87: Interpolation set to Linear

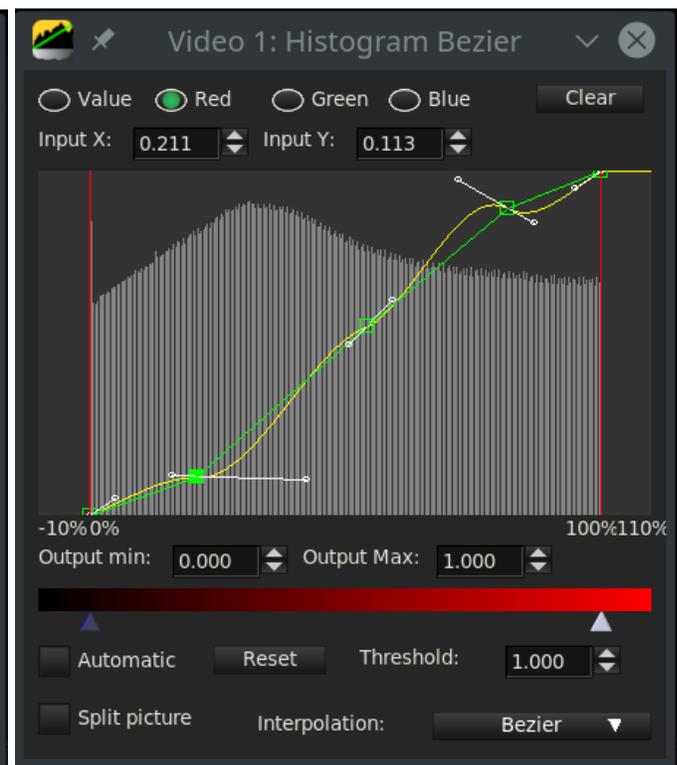


Fig. 86: Interpolation set to Bézier

Curves are used by introducing *control points* simply with the left mouse button and adjusting the value by dragging and dropping. If you drag along the *horizontal* line only, you change the value of x and you can read this value in the *input x* textbox. If you drag along the *vertical* line only, you change the value of y and you can read the value in the *input y* textbox. This is the *output* value. The newly clicked control point becomes *active* and is full green in color. To delete a point we have to make it *active* and then press the **Del** key, or we can drag the point beyond the position of another control point to its right or left or, finally, pressing RMB. The control points corresponding to the *black point* and the *white point* are automatically created from the beginning, to fix their values and prevent clipping.

Curves are generally adjusted by introducing several *control points*, some to be kept fixed (as anchors) to prevent curve modification beyond them, and others to be dragged to make the desired correction. The power of the curves lies in being able to *circumscribe* a small interval at will and intervene only on this without involving the remaining parts of the frame. The precision with which you can work is such that you can almost arrive at a *secondary color correction*.

The most used type of modification is to create a *S curve*. There can be a lot of shapes that use the *S curve*; the simplest is to create a control point in the *shadows*, one in the *midtone* (anchors) and one in the *highlights*. Moving the highlight point upwards and the shadow point downwards increases the *contrast*, making the image sharper and improving the color rendering. With the type of *linear curve* you can make *hard* adjustments, similar to the result of the use of *Color 3 Way*, even if this acts on the *color wheel (Hue)* while the *curves* act on individual *RGB channels*.

The *Polynomial* and *Bézier* types introduce *control handles* that allow for more sophisticated and smoother adjustments. The quality of the result is much better, but they require more experience for their optimal use. Extending the handles away from the control point increases the *radius* of the curve at that point. By varying the *angle* of the handles we change the tangent and thus the curvature of the curve below. The difference between *Polynomial* and *Bézier* lies in the underlying mathematics, but for practical purposes the use is similar.

HolographicTV

Incoming *objects* are projected like *holovision* seen in the movie *Stars Wars* as in R2-D2's video message projector of the Princess Leila. You need a movie or background image and above it a track containing the figure on which to apply the effect. This must have a transparent background. There are no configuration parameters; it only has to be applied to the upper track. This effect originated from <https://effectv.com>.



Hue saturation

With this effect you can change *hue*, *saturation* and *value*. The parameters are modified using 3 simple sliders. The *hue* control shifts the colors circularly in the color plane, normally resulting in *false colors*. The *saturation* control can be used to reduce color footage to *black and white*. The *value* control makes any given colors more bright or more subdued.

Interpolate Bayer

Uses a *Bayer* filter algorithm to *interpolate* (estimate) missing color information. This is needed for some *cameras* where each pixel location only has an R or G or B value instead of all R, G, and B values for each location. The algorithm creates values for each of the three colors at every location by smearing (interpolating) each set of partial R, G and B values to create values at every pixel location.

Interpolate Video

Theory

Each video has its own *framerate*. If we want to change it (for *timelapse* or *slowmotion*) the best thing is to shoot the scene with suitable framerate. But even in *post production* we can do something. The simplest method is to remove some frames to speed up the movie or add some to slow it down (from now on, for simplicity we will consider only the *timelapse*). Needless to say, the result is not smooth and the viewer will notice it immediately. A better method is to use the interpolation, *mediating* the pairs of frames that alternate. For example, if we have a sequence of frames 1, 2, 3, 4, 5, 6, 7, 8... we can make a *timelapse* mixing frames 1 and 2, 3 and 4, 5 and 6, 7 and 8 and so on. So we will have a new sequence of 4 frames instead of the initial 8: 12, 34, 56, 78... We will get 50% acceleration but it will always be of bad quality because of the too rough blending between the pairs of frames. Blending can be improved by weighing it differently by 50% *frame 1* + 50% *frame 2*, but the result is still unsatisfactory. Further improvements can be achieved by using *logarithmic* or *exponential* interpolation instead of *linear* interpolation. But the most sophisticated methods that lead to better results are based on *optical flow analysis*. These analyse the movement of circumscribed areas over a given period of time. With this method the intermediate frames do not derive from an approximate blending, but from the calculation of the *vector of the motion* between two frames that determines the displacement (*warping*) of the moving figure in the new intermediate frame. *Interpolate Video* works this way.

Practice:

The practical use of *Interpolate Video* is a little different than the theory. The interpolate effect tries to create the illusion of a higher frame rate from source footage of very low framerates by averaging frames over time. It averages two *input frames* for each *output frame*. You choose a zone to be evaluated (*macroblock size*) and a radius (*search radius*) where you can search for this macroblock in the following frames. The *Use optic flow* button is activated and playback starts. The plugin will calculate the *motion vector* (which can be made visible by the *draw motion vectors* button) and apply it to *intermediate* frames. This operation is CPU intensive. Once the analysis is done, we can scroll the video by unchecking the two buttons and obtaining the desired result. There are two ways of specifying the *input frames*. You can specify an *input frame rate* which is lower than the *project frame rate*. This causes input frames to be taken at even intervals. You can also specify *keyframe locations* as the positions of the *input frames*. In this mode the *output frame rate* is used as the input frame rate and you just create keyframes wherever you want to specify an *input frame*.

Inverse Telecine

This is the most effective deinterlacing tool when the footage is a video transfer of a film. This can be used to solve the problem, i.e., undo the damage caused by making film into a TV broadcast. That process came about because film is at 24 fps while TV is at 29.97 fps and *fields* are at 60. So the film was converted from 24 fps to 60 fps. Roughly speaking, converting every 4 frames into 5 frames plus a slight slow down in *speed*. Then the 60 fps was down-sampled to 30 fps by extracting *odd* and *even* lines and *interlacing* the lines. This process is referred to as three-two pull down (*3:2 pull down*) in filmmaking and television production for the post production process of transferring film to video.

The three-two pull down is where the telecine adds a third video field (a half frame) to every second video frame, but the untrained eye cannot see the addition of this extra video field.

The IVTC effect is primarily a way to convert *interlaced* video to *progressive* video. It reverses the effect of three patterns of interlacing. In the next lines A, B, and C represent fields.

A AB BC CD D
AB CD CD DE EF
Automatic

The first two options are fixed patterns and affected by the *pattern offset* and *odd field first* parameters. The last option creates several combinations of lines for each frame and picks the most progressive combination. It is a brute force algorithm that is trying to resample the lines. This technique does not rely on a pattern like other techniques and is less destructive but the *timing* is going to be jittery because of the lack of a frame rate reduction. In order to smooth out the timing, you need to follow *inverse telecine* with a *decimate* effect.

Invert video

Invert video is a method of reversing the colors of a video track. The four parameters refer to channels - *Red, Blue, Green, Alpha*. A very common use is to invert the *alpha* channel to change transparency.

Lens

Create the effect of looking through a *lens*.

R, G, B, A Field of View: quantity of deformation of the relative *fields*. Often used with *Lock* to simultaneously affect the 4 fields.

Aspect Ratio: determines the aspect ratio that you intentionally set.

Radius: radius of curvature of the distortion. At minimum, it is a sphere (*fish eye*) and at maximum, it is a rectangle (*no distortion*).

Center X, Y: determines the coordinates of the center of the sphere. It can be made visible with *Draw Center*.

Mode: determines the *type of distortion*. The choice is between *sphere shrink*, *sphere stretch*, *rectilinear shrink* and *rectilinear stretch*.

Interpolation: determines the interpolation algorithm; from the fastest and least precise *Nearest*, passing through *BiLinear* to the better *BiCubic*.

Linear Blur

This effect acts only in *one direction* which can vary up to an angle of 180° with these parameters:

Length - distance between original image and final blur step; corresponds to the distance of the *fields*.

Angle - angle of motion in one direction for linear blur

Steps - number of blur steps to be used in the calculation. Increasing the number takes more CPU.

Channels – RGBA

The linear blur image shown here has the parameters: Length=19, Angle=25, and Steps=2.



Fig. 88: For clarity of presentation only 2 fields are shown

Live Video

This effect reads video directly from the *capture card input*. It replaces any video on the track so it is normally applied to an *empty* track. Only one *Live Video* effect can exist at any time on the timeline. It can not be *shared* by more than one track. The configuration for the capture card is taken from the recording preferences. Go to **Settings->Preferences->Recording** to set up the capture card.

In the *Video In* section where it says *Record driver*, it should be set to either *Video4Linux2* or *IEC 61883*. Other video drivers have not been tested with *Live Video* and probably will not work. For *live video*, the selection for *File Format* and *Video* needs to be set to a format the timeline can use. The file format must be *Quicktime for Linux* and video recording must be enabled for it. Click on the *wrench* to set the video compression.

The video compression depends on the recording driver. For the *Video4Linux2* recording driver, the compression must be *Motion JPEG A*. For the *IEC 61883* driver, the compression must be *DV*. This gets the driver to generate *output* in a *color model* that the timeline can use. Some cards provide color and channel settings. *Live video* takes the color settings from the values set in the *Video In* window. Go to **File->Record** to bring up the *recording interface* and the *Video In window*. Values set in the *Video in* window are used by *Live Video*. Any channels the capture card supports need to be configured in the *Video In* interface since the same channels are used by the *Live Video* effect.

With the video recording configured, *highlight* a horizontal region of a video track or define *in* and *out points*. Then drop the *Live Video* effect into it. Drop other effects after *Live Video* to process the live video in realtime. For best results, you should use *OpenGL* and a video card which supports *GL shading* language. Go to **Settings->Preferences->Playback->Video Out** to enable the *OpenGL* driver.

Loop video

Sections of video can be *looped* by dropping a *loop* effect on them. Contrary to the **settings** → **loop playback option**, the loop effects can be rendered where the **settings** → **loop playback option** can not be. The loop effects are also convenient for short regions.

The loop effects have one option: the *number of frames* or *samples* to loop. This specifies the length of the region to loop starting from either the beginning of the effect or the latest keyframe. The region is replicated for the entire effect.

Every time a keyframe is set in a loop effect, the keyframe becomes the beginning of the region to loop. Setting several keyframes in succession causes several regions to loop. Setting a single keyframe causes the region after the keyframe to be looped throughout the effect, no matter where the keyframe is. The end of an effect can be looped from the beginning by setting the keyframe near the end.

Motion51

This plugin compensates for *unwanted motion* and *stabilizes* the picture. The *Motion51* Plugin simplifies motion stabilization so that without a lot of tweaking you can easily achieve reasonable results, either by using the *defaults* or varying a single *parameter*. Since the motion in every clip is specific, there are some additional parameters useful to adjust the settings accordingly. Alternatively, the *MotionCV* and *MotionHV* plugins can still be used as the originals, if more control over specific parameters is needed. The *Motion51* plugin uses different methods for tracking than the other motion plugins. Motion Stabilization is very useful if you have *jittery* video, for example when taken from a car window, or while walking.

The better results require more *samples*. Setting the sample set size is probably the most important setup change. Also, when computing motion compensation, the entire history of the image motion is important, and so it is desirable to enable the playback setting *play every frame* in order to get good results. When every frame has to be processed, it can be time-consuming. Reasonable results are possible with small sample sets. After setup, the sample size can be increased to produce a high quality rendered result.

Description of what the program is doing :

The motion is detected by “sampling” the video image in a circular field. This size and placement of the sample region defaults to most of the image area. When the “draw vectors” feature is enabled, the outer line trace encloses the searched region. The dotted circles define the target pixel set as the reference sample. The image is sampled using the circle pattern in a grid search. The best match is used to find the center and amount of rotation to transform the current image so that the reference area motion is canceled.

The amount of *sampling* does not significantly change for smaller or larger *search areas*. This means that a wide area can be searched just as easily as smaller areas. The main parameter which determines how *hard* it looks at the image is the *samples* parameter. It represents the number of possible *rotations*, as well as the search precision. More samples mean more precision, and less jitter, but the program will run more slowly.

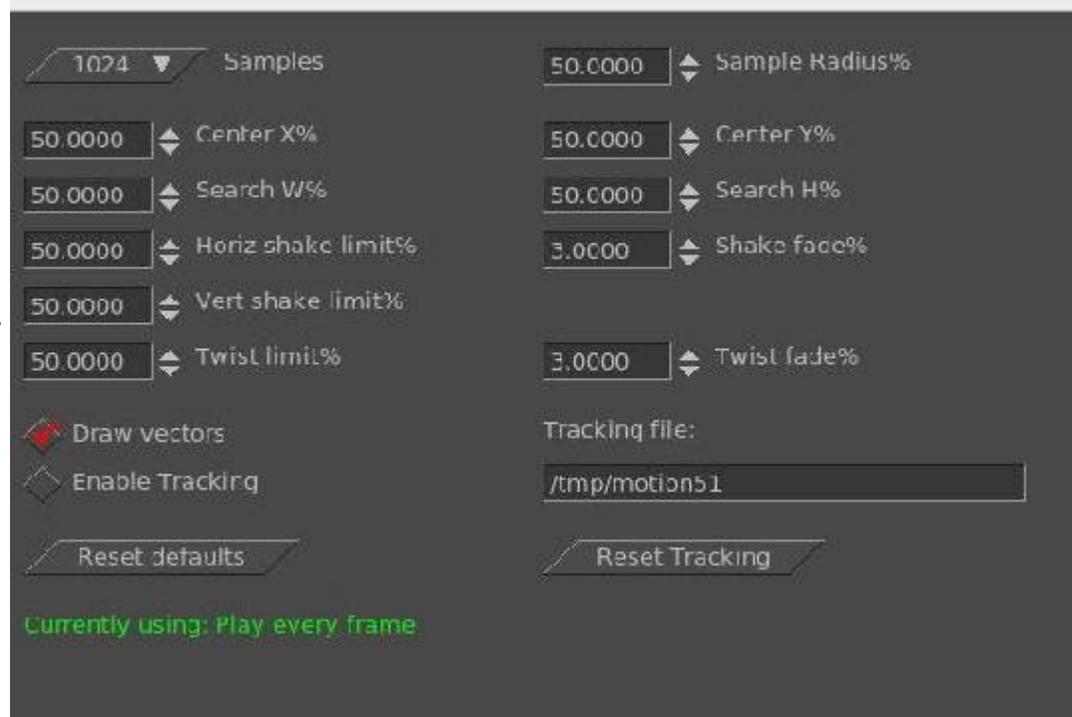
Figure 1 -

Note the faint circles shown in the center, the oval outside area boundary, and the line that is part of a box in the color orange.



Figure 2 -

Display of the Motion51 plugin window with its default options set. The Samples box at the top is most often the only parameter that you may want to vary.



Program Parameters Description follows:

X11-OpenGL - setting can speed up the computation significantly in some cases when hardware OpenGL is available.

Samples - is the number of pixels which the software will examine to stabilize the picture. The sample set is arranged in 4 equal concentric circular sets. Each sample *dot* represents content and position for a

pattern matching test. Setting the samples to larger values improves the match by adding lots of placement possibilities. The samples/pixels that will be utilized are distributed throughout the selected area – this is seen within the circles drawn when *Draw Vectors* is enabled. See Figure 1.

Draw vectors - demonstrates the search operation of motion stabilization. When enabled, the outer search boundary (oval), the search grid area (rectangle), and the reference sample (circles) are visible. 4 concentric circles show the reference sample set (target). You will also see an arrow in the center of the circle which shows each image displacement from frame to frame.

When you render the video using the motion plugin, these dots/lines/circles are drawn into the rendered output. *Draw vectors* helps to visualize the meaning of the parameters to aid in setup. You should disable *Draw Vectors* before the final rendering.

Sample Radius - is the radius of a circle that denotes the area of the sample locations. It is expressed as a percentage of the smallest image edge. For example, if it is set to 50%, then the circle will overlap about $\frac{1}{2}$ of the image. This does not change the number of samples. It does change the area from where the samples are gathered. If you have *Draw vectors* on, you can see the faint outline of a circle used for the radius.

Center X/Y - is the center position of the sample circle, as a percentage of image width and height. This is useful to reset the reference focal point in cases where the important feature target is off-center. When both X and Y are set to 50%, the samples will be used from around the center of the video.

Search W/H - determines the width and height of the rectangular area used for the grid pattern search. Samples are taken by moving the center of the circles in a grid pattern.

Shaking refers to image translation.

Horiz/Vert shake limit - determines translation constraints. If the motion determined by the search exceeds the limit, it is clipped to a value that is at the limit boundary. For example, if the match indicates that the motion is 60% off the reference target, but the limit is 50%, then the actual translation used will be limited to only 50%.

Shake fade - determines how fast the translation cancellation fades away, and the image resettles to its actual appearance. Every frame that is processed accumulates the motion of the past. The amount of the past motion which is applied is reduced by the fading factor. The current match is then added to the fading past motion history. Fading works **fast**. It is applied every frame. So if fade is 10%, and no new motion occurs in the input, the history will be 90%, 81%, 73%, 66%,... and in 30 frames only 4.2% of the past motion will be present in the effect. Fading insures that the image will eventually settle re-centered when the image motion ends.

Twisting refers to image rotation.

Twist limit - determines the rotation constraints. If the rotation determined by the search exceeds the limit, it is clipped to the limit boundary. Its operation is similar to the shake limits.

Twist fade - determines how fast the rotation cancellation fades away, and the image resettles to its actual appearance. Its operation is similar to the shake fade.

Enable Tracking - caches the search results in a file so that subsequent playback does not have to be recalculated. When tracking is enabled, before a frame is processed the frame number is used to look

for cached results. If cache data is available, it is used. If no data is available, the frame is processed by the motion tracking search, and the results are added to the cache file. If tracking is not enabled, the data is always sourced from the motion tracking search and the tracking file is not updated.

Tracking file - is the name of the file which will contain the calculated values to be saved. Note that the default is /tmp/motion51 which can be hazardous, since a system crash or a reboot can delete /tmp files.

Reset Defaults - button is used to revert to the initial defaults built into the program.

Reset Tracking - will delete the current Tracking file and disables tracking so that any previously calculated values are no longer available. However, because motion stabilization can often be cpu intensive, if the default file already exists, it will create a file name from the loaded asset.

Play Every Frame status shows if you are *Currently using: Play every frame* or *For best results, set play every frame*. This can be set in **Settings** → **Preferences**, the *Playback Tab* in the *Video out* section.

Motion

The *motion tracker* is almost a complete application in itself. The motion tracker tracks two types of motion: *translation* and *rotation*. It can track both simultaneously or one only. It can do *1/4 pixel tracking* or *single pixel tracking*. It can stabilize motion or cause one track to follow the motion of another track. Although the motion tracker is applied as a *realtime effect*, it usually must be rendered to see useful results. The effect takes a long time to precisely detect motion so it is very slow.

Motion tracker works by using one region of the frame as the region to track. It compares this region between 2 frames to calculate the motion. This region can be defined anywhere on the screen. Once the motion between 2 frames has been calculated, a number of things can be done with that motion vector. It can be scaled by a user value and clamped to a maximum range. It can be thrown away or accumulated with all the motion vectors leading up to the current position.

To save time the motion result can be saved for later reuse, recalled from a previous calculation, or discarded. The motion tracker has a notion of 2 tracks, the master layer and the target layer. The master layer is where the comparison between 2 frames takes place. The target layer is where motion is applied either to track or compensate for the motion in the master layer. Motion tracking parameters:

Track translation

Enables translation operations. The motion tracker tracks X and Y motion in the master layer and adjusts X and Y motion in the target layer.

Translation block size

For the translation operations, a block is compared to a number of neighboring blocks to find the one with the least difference. The size of the block to search for is given by this parameter.

Translation search radius

The size of the area to scan for the translation block.

Translation search steps

Ideally the search operation would compare the translation block with every other pixel in the translation search radius. To speed this operation up, a subset of the total positions is searched. Then the

search area is narrowed and re-scanned by the same number of search steps until the motion is known to 1/4 pixel accuracy.

Block X, Y

These coordinates determine the center of the translation block based on percentages of the width and height of the image. The center of the block should be part of the image which is visible at all times.

Maximum absolute offset

The amount of motion detected by the motion tracker is unlimited if this is 100. If it is under 100 the amount of motion is limited by that percentage of the image size.

Settling speed

The motion detected between every frame can be accumulated to form an absolute motion vector. If the settling speed is 100 the absolute vector is added to the next frame. If the settling speed is less than 100 the absolute vector is downscaled by the settling amount before being added to the next frame.

Track rotation

Enables rotation operations. The motion tracker tracks rotation in the master layer and adjusts rotation in the target layer.

Rotation block size

For rotation operations a single block is compared to equally sized blocks, each rotated by a different amount. This is the size of the rotation block.

Rotation search radius

This is the maximum angle of rotation from the starting frame the rotation scanner can detect. The rotation scan is from this angle counterclockwise to this angle clockwise. Thus the rotation search radius is half the total range scanned.

Rotation search steps

Ideally every possible angle would be tested to get the rotation. To speed up the rotation search, the rotation search radius is divided into a finite number of angles and only those angles compared to the starting frame. Then the search radius is narrowed and an equal number of angles is compared in the smaller radius until the highest possible accuracy is achieved. Normally you need one search step for every degree scanned. Since the rotation scanner scans the rotation search radius in two directions, you need two steps for every degree in the search radius to search the complete range.

Draw vectors

When translation is enabled, 2 boxes are drawn on the frame. One box represents the translation block. Another box outside the translation block represents the extent of the translation search radius. In the center of these boxes is an arrow showing the translation between the 2 master frames. When rotation is enabled, a single box the size of the rotation block is drawn rotated by the amount of rotation detected.

Track single frame

When this option is used the motion between a single starting frame and the frame currently under the insertion point is calculated. The starting frame is specified in the Frame number box. The motion calculated this way is taken as the absolute motion vector. The absolute motion vector for each frame replaces the absolute motion vector for the previous frame. Settling speed has no effect on it since it

does not contain any previous motion vectors. Playback can start anywhere on the timeline since there is no dependence on previous results.

Track previous frame

Causes only the motion between the previous frame and the current frame to be calculated. This is added to an absolute motion vector to get the new motion from the start of the sequence to the current position. After every frame processed this way, the block position is shifted to always cover the same region of the image. Playback must be started from the start of the motion effect in order to accumulate all the necessary motion vectors.

Previous frame same block

This is useful for stabilizing jerky camcorder footage. In this mode the motion between the previous frame and the current frame is calculated. Instead of adjusting the block position to reflect the new location of the image, like Track Previous Frame does, the block position is unchanged between each frame. Thus a new region is compared for each frame.

Master layer

This determines the track which supplies the starting frame and ending frame for the motion calculation. If it is Bottom the bottom track of all the tracks sharing this effect is the master layer. The top track of all the tracks is the target layer.

Calculation

This determines whether to calculate the motion at all and whether to save it to disk. If it is “Don't Calculate” the motion calculation is skipped. If it is “Recalculate” the motion calculation is performed every time each frame is rendered. If it is “Save” the motion calculation is always performed but a copy is also saved. If it is “Load”, the motion calculation is loaded from a previous save calculation. If there is no previous save calculation on disk, a new motion calculation is performed.

Action

Once the motion vector is known this determines whether to move the target layer opposing the motion vector or following the motion vector. If it is “Do Nothing” the target layer is untouched. If it is “Track...” the target layer is moved by the same amount as the master layer. This is useful for matching titles to objects in the frame. If it is “Stabilize...” the target layer is moved opposite to the motion vector. This is useful for stabilizing an object in the frame. The motion operations can be accurate to single pixels or subpixels by changing the action setting.

Secrets of motion tracking

Since it is a very slow effect, there is a method to applying the motion tracker to get the most out of it. First disable playback for the track to do motion tracking on. Then drop the effect on a region of video with some motion to track. Then rewind the insertion point to the start of the region. **Set Action → Do Nothing; Set Calculation → Don't calculate**; Enable *Draw vectors*. Then enable playback of the track to see the motion tracking areas.

Enable which of translation motion or rotation motion vectors you want to track. By watching the compositor window and adjusting the Block x,y settings, center the block on the part of the image you want to track. Then set search radius, block size, and block coordinates for translation and rotation.

Once this is configured, set the calculation to Save coords and do test runs through the sequence to see if the motion tracker works and to save the motion vectors. Next, disable playback for the track, disable

Draw vectors, set the motion action to perform on the target layer and change the calculation to Load coords. Finally enable playback for the track.

When using a single starting frame to calculate the motion of a sequence, the starting frame should be a single frame with the least motion to any of the other frames. This is rarely frame 0. Usually it is a frame near the middle of the sequence. This way the search radius need only reach halfway to the full extent of the motion in the sequence.

If the motion tracker is used on a render farm, Save coords and previous frame mode will not work. The results of the save coords operation are saved to the hard drives on the render nodes, not the master node. Future rendering operations on these nodes will process different frames and read the wrong coordinates from the node filesystems. The fact that render nodes only visualize a portion of the timeline also prevents previous frame from working since it depends on calculating an absolute motion vector starting on frame 0.

2 pass motion tracking

The method described above is 2 pass motion tracking. One pass is used just to calculate the motion vectors. A second pass is used to apply the motion vectors to the footage. This is faster than a single pass because errors in the motion vector calculation can be discovered quickly. This also allows the motion tracking to use a less demanding colormodel like RGB888 in the scanning step and a more demanding colormodel like RGB Float in the action step. The scanning step takes much longer than action. This has the disadvantage of not being practical for extremely long sequences where some error is acceptable and the picture quality is low to begin with, like stabilizing camcorder footage.

The slower method is to calculate the motion vectors and apply them simultaneously. This method can use one track as the motion vector calculation track and another track as the target track for motion vector actions. This is useful for long sequences where some error is acceptable.

Using blur to improve motion tracking

With extremely noisy or interlaced footage, applying a blur effect before the motion tracking can improve accuracy. Either save the motion vectors in a tracking pass and disable the blur for the action pass or apply the blur just to the master layer.

Using histogram to improve motion tracking

A histogram is almost always applied before motion tracking to clamp out noise in the darker pixels. Either save the motion vectors in a tracking pass and disable the histogram for the action pass or apply the histogram just to the master layer.

Tracking stabilization in action

This is an explanation of how to stabilize a video as in the case of a video taken from a vehicle.

First select on the timeline the part of the footage you want to stabilize, using the in and out points. Then apply the motion effect on that part of the video. Select the "Previous frame same block" option. That option is recommended for stabilizing jerky camcorder footage. Its goal is not to "follow" an object. The block stays exactly at the same place during all the effect length.

Enlarge the block and select almost half the size of the video. Select the "Stabilize subpixel" option as it will give a finer stabilization. Reduce the "Maximum absolute offset" value to limit the stabilization

amplitude. You probably prefer to get a non-perfect stabilization on some places on the video rather than having a very large black border on one side of the picture during big shakes. Set the "Translation search steps" value to 128. Increasing that value will not give a better result and only considerably increases the rendering time. Make sure the "Draw vectors" option is selected, and render the part of the video where the motion effect is applied.

If the result is good, deselect the "Draw vectors" option so that the block and vectors are not drawn anymore on the video. Then, render your video to a '.dv' file, and import it into your project. You will notice the video is stabilized but there are black borders which appear on sides of the frame. You have to zoom in and define projector keyframes to move the projector around the screen, in order to remove those black borders. The more your footage is jerky, the more you have to zoom in to discard the black borders. That is why the result is better with HDV footage than with DV footage.

An interesting side note about *add offset* usage is explained next (credit Pierre Marc Dumuid).

To stabilize video, the motion plugin uses a "tracking frame" to which to track to and a region within that frame to track (generally an object in the background) in the current frame. When the region is obscured, often by something in the foreground or by leaving the screen, then the motion compensation would fail, and the video jumps all over the place.

You set a second region to track, and then *add offset*.

This shows how it is used. It works very well:

```
-----k-----k-----  
      ^      ^^  
      A      BC
```

A - object1 is visible in the background up until C
B - (the frame before C) has both object1 and object2 visible
C - has only object2 visible

1. Make a keyframe and set to track object1.
2. Make a keyframe at C and track frame at B, set to track object2.
3. Set keyframe at C to add offsets that were calculated at B.

Motion 2 Point

Motion stabilization using 2 pass tracking.

Motion Blur

Uses *X/Y camera* automation vectors to apply a linear blur trailing camera direction due to movement.
Length - distance between original image and final blur step; corresponds to the distance of the fields.
Steps - number of blur steps to be used in the calculation. Increasing the number takes more CPU.

MotionCV

Motion tracking/stabilization from the community version of cinelerra.

MotionHV

Updated motion tracking/stabilization of 2017 from the original author of cinelerra.

Oil painting

This effect makes video tracks appears as a painting. It can be controlled by *Radius* slider and *intensity* of colors can be chosen as an option.

Overlay

This effect can *combine* several tracks by using the so called *Overlayer*. This is a basic internal device normally used by Cinelerra to create the dissolve transitions and for compositing the final output of every track onto the output bitmap. The *Overlayer* has the ability to combine one or several image layers on *top* of a *bottom layer*. It can do this combining of images in several different (and switchable) output *modes* such as *Normal*, *Additive*, *Subtractive*, *Multiply* (Filter), *Divide*, *Max* and *Replace*. For a detailed list refer to the chapter on *Overlay Modes - PorterDuff*.

The *overlay* plugin enables the use of this *Overlayer* device in the *middle* of any plugin stack, opening endless filtering and processing possibilities. It is only useful as a *shared plugin* (i.e. a multitrack plugin). To use the overlay plugin:

1. Add the effect to *Track A*.
2. Choose *attach effect* from the context menu of another track (*Track B*).
3. Choose *Track A:Overlay* as a *shared plugin*.
4. Manipulate the plugin parameters in *Track A*.

In the *Overlay* Plugin's parameter window you can choose the overlay order, that is, which track plays the role of the *bottom layer* and which plays the role of the *top layer*. For some overlay modes, this can make quite a difference, for example, the *top layer* is subtracted from the *bottom layer* for *Subtractive* mode. You can also choose on which of the tracks to overlay the combined output. In most cases, you will want to *mute* the other track and only retain this combined output.

Perspective

The *perspective* plugin allows you to change the perspective of an *object* and is used to make objects appear as if they are fading into the distance. Basically, you can get a different *view*. A transformation is used which preserves *points*, *lines*, and *planes* as well as ratios of distances between points lying on a straight line.

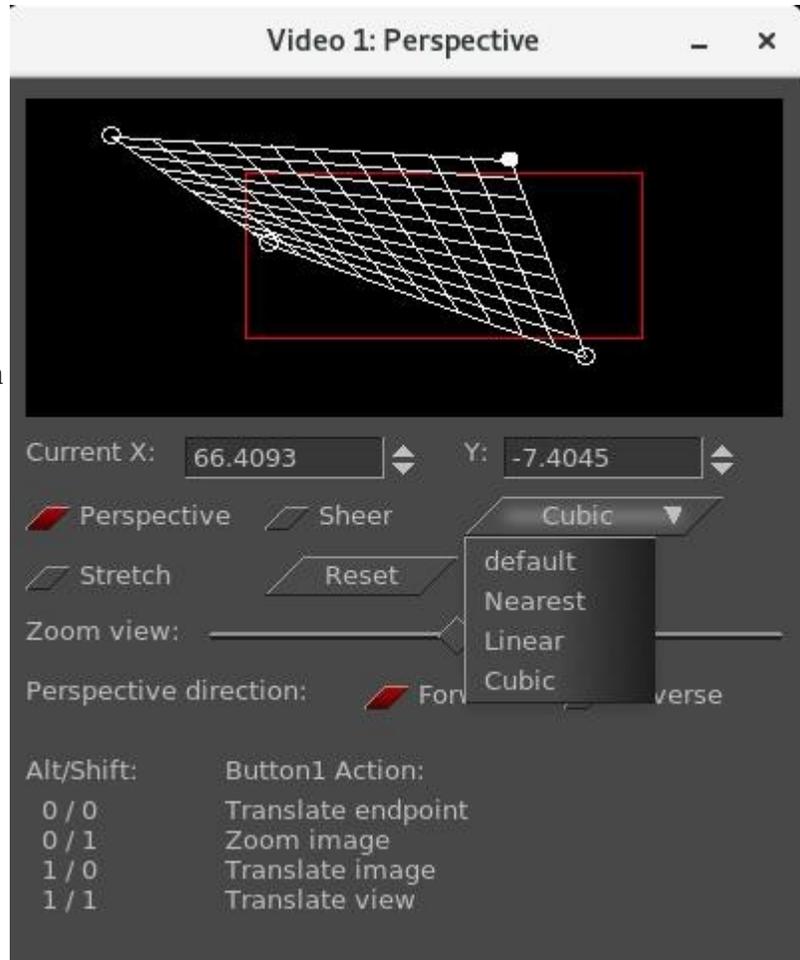
In the screenshot to the right, you can see that there are 4 options for the endpoints used for the edges.

Default – if OpenGL is being used with your graphics card, this will be the option in effect. If no OpenGL, then it will be Cubic.

Nearest – using software, nearest neighbor can look step-py.

Linear – software implementation of a linear algorithm.

Cubic – smoothest looking on the edges and considered the best.



Key Presses for using the *Perspective* plugin:

Left mouse button – drags the corner that is closest to current location

Alt key + left mouse – translates the perspective; drags the whole image

Shift key + left mouse – zooms the perspective

Alt+Shift + left mouse – translates view but does not change output

Note that the red color lines in the box show the *composer* boundary.

In order to see *endpoints* that go off the screen, you can use the zoom slider which changes only the zoom view and does nothing else. The slider uses a logarithmic scale ranging from 1/100 to 100.

The next screenshots show the results of the 4 different *smoothing* options.

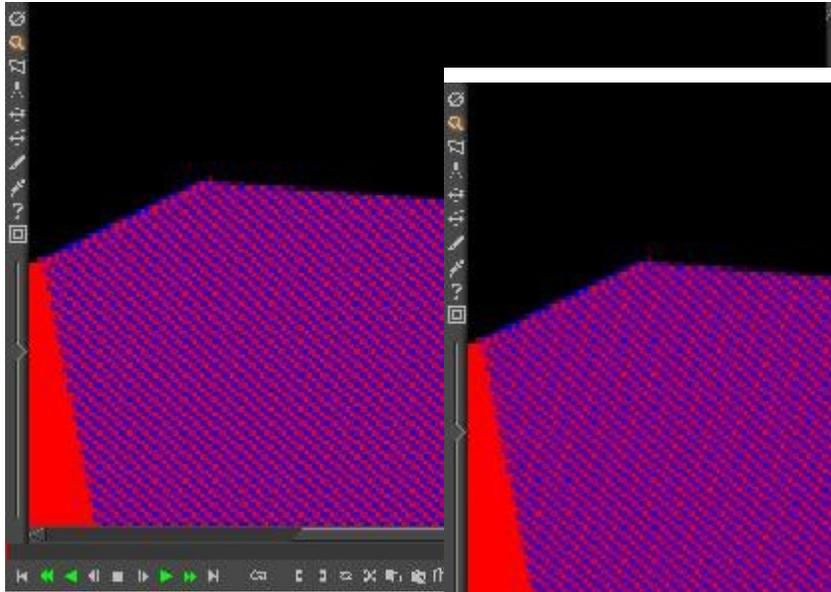


Fig. 89: Cubic

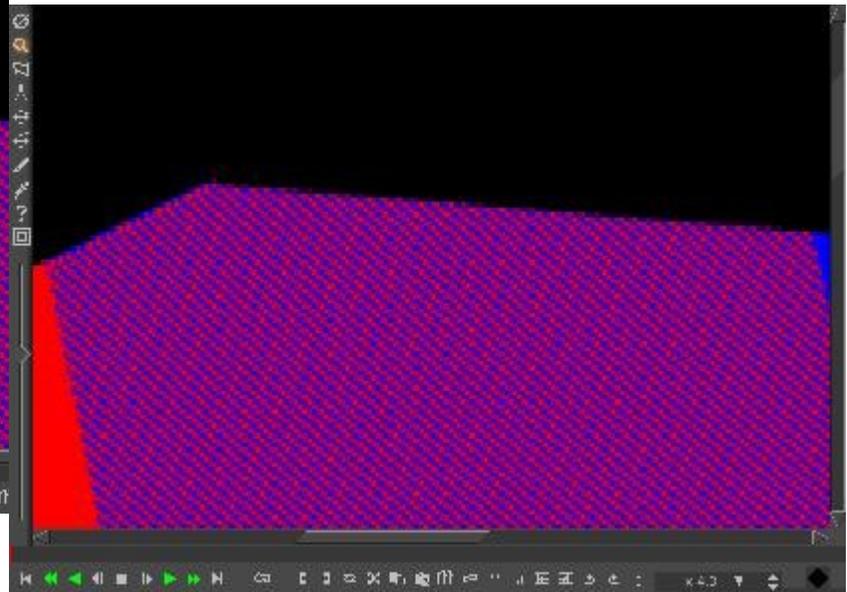


Fig. 90: Linear

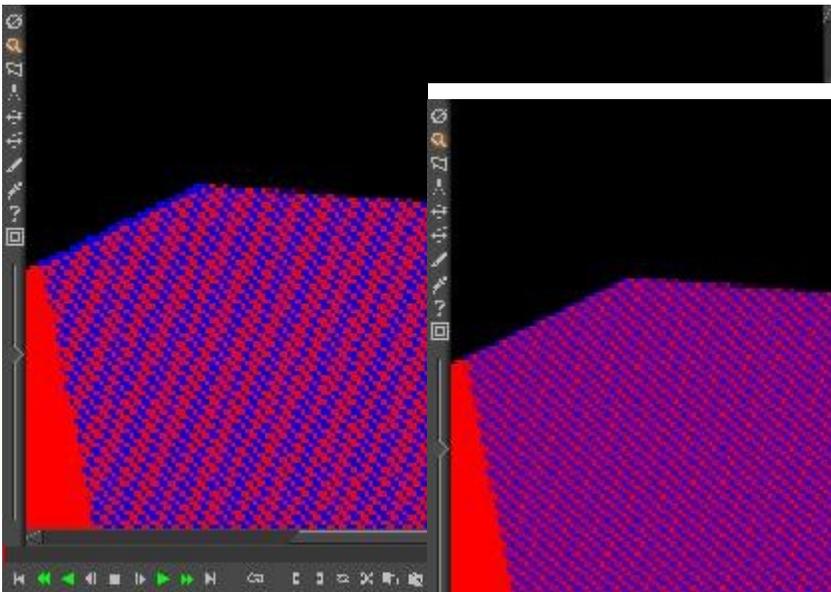


Fig. 91: Nearest Neighbor

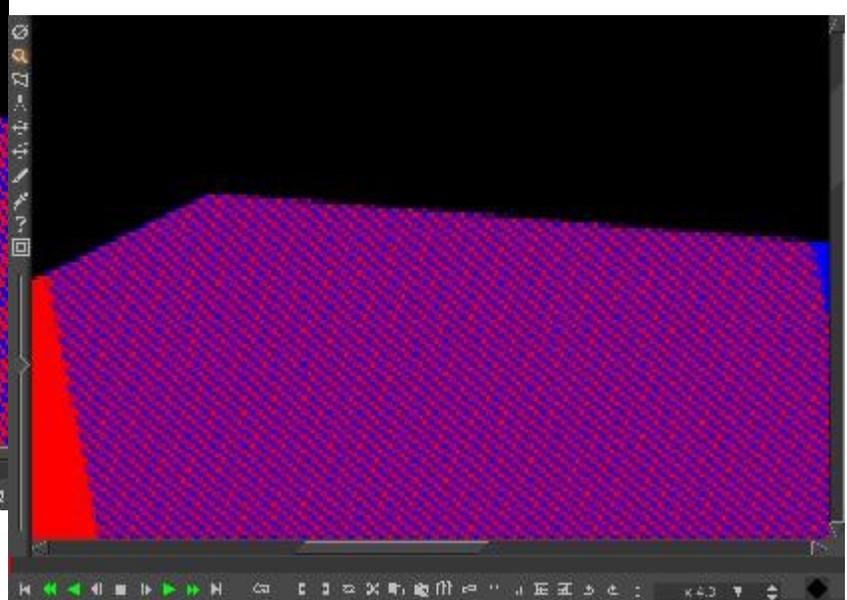


Fig. 92: Default - OpenGL with Graphics Card

Polar

The Polar effect *bends* and *warps* your video in weird ways. Mathematically, it converts your video from either *polar coordinates* to *rectangular coordinates*, or the reverse.

RGB-601

For *analog* video or *MPEG* (including *DVD*) output, the maximum range for R,G,B is **[16,235]** (8-bit). For *YUV*, the maximum range for *intensity* (Y) is **[16, 235]** (8-bit). This range corresponds to gray levels from **6%** to **92%**. When rendering, values outside of these ranges will be *clipped* to these limits.

To render to *MPEG*, add the *RGB-601* effect to all video tracks where material uses the *full* intensity scale (0-100%), and enable *RGB → 601 compression*. Consider adding the *Videoscope* effect after *RGB-601* to see how *RGB-601* affects your *dynamic range*. To preview how your rendered *MPEG* would look without *RGB-to-601* compression, instead enable *601 → RGB expansion* and you will observe a noticeable *contrast* increase. Although *RGB-601* will reduce contrast in your video tracks, the contrast will be restored during *MPEG* playback.

RGBShift

Most *cameras* take the light coming into the lens, and convert that into 3 sets of numbers, one for *Red* (R), one for *Green* (G), and one for *Blue* (B). Some of the older cameras were composed of sensors and originally the RGB sensors were on 3 separate planes and had to be aligned. If they were misaligned in the video, you can use *RGBShift* to get them realigned. To move a specific color *up/down*, modify the *dy* value using the slider bar in the *RGBShift* window. To move a color *left/right*, modify the corresponding *dx* value. Note that the current values of the *RGBShift* are maintained in the *.bcast5* defaults file and will be retained across sessions. If using the *y u v* color space, you will want to use *YUVShift* instead. Screenscasts showing RGB shift before and after.



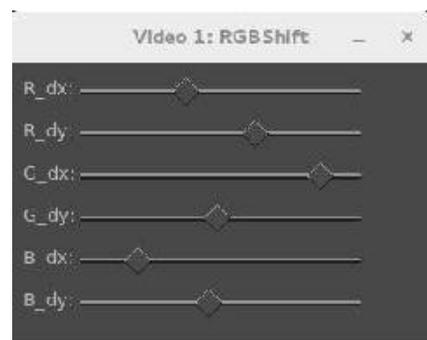
← Normal color aligned.

Misaligned color →



← RGB shift in a good video should be 0's.

You move the R, G, B sliders to align the colors correctly. →



Radial Blur

Radial blur is a *Bokeh* effect that creates a whirlpool which simulates a swirling camera. You can vary the location, type, and quality of the blur.

X,Y - center of the circle of movement.

Angle - angle of motion in one direction.

Steps - number of blur steps to be used in the calculation; increasing this number uses more CPU.

The radial blur image shown here has the parameters: Angle=-35 and Steps=2.



Fig. 93: For clarity of presentation only 2 fields are shown

ReframeRT

ReframeRT changes the *number* of frames in a sequence of video directly from the timeline. The faster method for getting the same results as this plugin is to use the *speed curve* which was a later addition. But if you need very precise results, *ReframeRT* is most useful. There are two ways to do this, which can be selected from the checkboxes in the configuration GUI. The first *Stretch* mode changes the number of frames in the sequence, and therefore its *length*, but not the *framerate*. The *Downsample* mode instead keeps the *length* of the movie by varying the *framerate*.

It is important to understand that the plugin works by varying the frames, the possible change of fps is only a side effect of the creation of new frames due to *interpolation*.

Stretch

Stretch mode multiplies the current frame number of its output by the *scale factor* to arrive at the frame to read from its input. The scaling factor is not entered directly but using a number of *input frames* to be divided by the number of *output frames*.

Scale factor = Input frames / Output frames

$1 / 8 \Rightarrow$ factor scale = 0.125 (slowmotion)

That is, one input frame of the original movie corresponds to 8 new output frames originated by interpolation. It is the opposite with regard to *fast play*.

The *stretch* mode has the effect of changing the length of output video by the inverse of the *scale factor*. If the *scale factor* is greater than 1, the output will end before the end of the sequence on the timeline. If it is less than 1, the output will end after the end of the sequence on the timeline. The *ReframeRT* effect must be *lengthened* to the necessary length to accommodate the *scale factor*. Change the length of the effect by clicking on the endpoint of the effect and dragging.

Although *stretch* mode changes the number of the frames read from its input, it does not change the *framerate* of the input. Effects before ReframeRT assume the same frame rate as ReframeRT. In stretch mode to create a *fast play* effect enter a value greater than 1 to get accelerated playback. For a *slow motion* effect, use ReframeRT in stretch mode with a value less than 1.

Example: you have a clip that you want to put in slow motion. The clip starts at 33.792 seconds and ends at 39.765. The clip is 5.973 seconds long. You want to play it at 4/10ths normal speed. You divide the clip length by the playback speed (5.973/.4) to get a final clip length of 14.9325 seconds. You create an in point at the start of your clip: 33.792 seconds. You put an out point 14.9325 seconds later, at 48.7245 seconds (33.792 + 14.9325). You attach a ReframeRT effect, set it to .4 and stretch. You change the out point at 48.7245 to an in point. You start your next clip after the slow motion effect at the 48.7245 out point. You can do this without making any calculations by first applying the effect and then lengthening or shortening the bar to where the stretched movie ends.

Downsample

Downsample mode does not change the *length* of the output sequence. It multiplies the frame rate of the output by the *scale factor* to arrive at a frame rate to read the input. This has the effect of replicating the *input* frames so that they only change at the scaled frame rate when sent to the *output*. It does not change the length of the sequence. If the scale factor is 0.5 and the output frame rate is 30 fps, only 15 frames will be shown per second and the input will be read at 15 fps. Downsample is only useful for scalefactors *below 1*, hence the name downsample.

Downsample mode changes the frame rate of the *input* as well as the number of the frame to read, so effects before ReframeRT see the *frame rate * the scale factor* as their frame rate. If the scale factor is 2 and the output frame rate is 30, the input frame rate will be 60 and the input frame number will be doubled. This will not normally do anything, but some input effects may behave differently at the higher frame rate.

Other important points:

- ReframeRT uses the fps indicated in **Settings** -> **Format** -> **fps project** and not the *fps* of the *assets*.
- It can be associated with *Nested Clips*.
- As an alternative to *ReframeRt* you can use the *speed curve*, or change the *framerate* in **Resources** → **info** and in the *Project*.
- It is keyframmable.

Reroute

The Reroute plugin enables you to selectively transfer the Alpha channel or the Components (RGB or YUV) or both from a source track to a target track, partially overwriting the target's contents. It works as a shared plugin. The typical usage scenario is to build up a possibly animated Mask in one track and then to transfer the Alpha channel to another content track.

Reverse video

Media can be reversed on the timeline in *realtime*. This is not to be confused with using the reverse playback on the *transport panel*. The reverse effects reverse the region covered by the effect regardless of the transport direction. The region to be reversed is first determined by what part of the track the effect is under and second by the locations of keyframes in the effect. The reverse effects have an enabled option which allows you to set *keyframes*. This allows many possibilities.

Every enabled keyframe is treated as the start of a new reversed region and the end of a previous reversed region. Several enabled keyframes in succession yield several regions reversed independent of each other. An enabled keyframe followed by a disabled keyframe yields one reversed region followed by a forward region.

Rotate

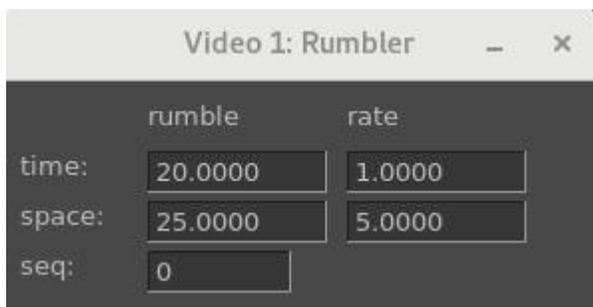
The Rotate filter can rotate the video in 90 degree increments or by any number of degrees through use of the *wheel* and about any pivot point. It can also *reverse* and *flip* the video.

Rumbler

The *Rumbler* plugin can be used to create *dream-like* or *earthquake-like noise* in the video. It creates noise by *jiggling* the corners through use of *perspective transformation* at the corners. The algorithm used is: **Rumbler(value) = (value at time) + amplitude * (random generator)**. The *random generator* varies from -0.5 to 0.5. The *rumble* perturbs the normal values at *time points* which occur rate times a second. The values used between the rumble points are interpolated, so that the value *jiggles* rate times a second, by as much as the rumble amplitude. The *time* unit is frames per second. The *corners* are in units of percent width/height.

Plugin variables:

- Rumble* gain applied to random rumbler.
- Rate* number of times per second.
- Time* number of frames per rate times a second; if 0 has no effect on results.
- Space* corners jiggling in percentage per rate times a second; if 0 has no effect on results.
- Seq* is a random seed number; any reasonable positive or negative number you choose.



Screencast shows:
Time jittering - 20 fps 1 time a second.
Corners jittering - 25 percent 5 times a second.
Using random seed 0 for a rumble sequence.

SVG via Inkscape

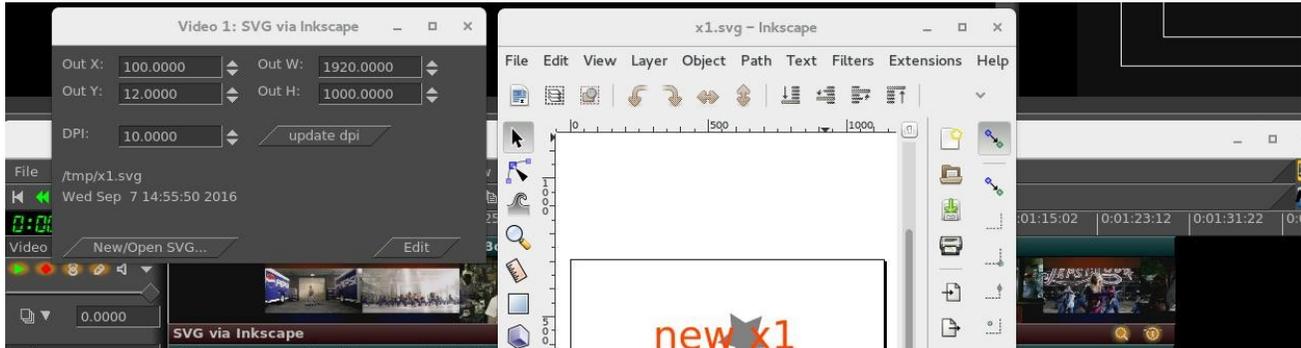
This plugin allows the user to manipulate an SVG (scalable vector graphics) image with *Inkscape* without having to leave the program. The associated cinelerra window provides the ability to change the *DPI*, the *Out x/y* coordinates, and the *Out w/h* values. For more information on use of inkscape, refer to: <https://inkscape.org/develop/about-svg/>

DPI - is *Dots per inch* and it represents the *resolution* of the SVG image. Since the image is scaled with interpolation mode *linear*, the edges will look blurry when the *input* resolution is lower than the *output* resolution. You can either set the desired DPI value in the window or use the tumbler on the integer text box, then use the *update dpi* button to have the change take effect. Changing DPI causes the entire

image to be *re-exported* via *inkscape*. DPI changes cause adjustments in the resolution, speed of re-import, and storage needed for the image data.

Out_x/Out_y parameters allow for changing the location of the SVG via the x or y coordinates. Scaling is controlled by the *Out_w/Out_h* parameters for width and height as they are normal parameters to *overlay*.

Screenshot below shows the menu options plugin window and the SVG image in the Inkscape window.



Scale

Reduce or expand the image size depending on the ratio.

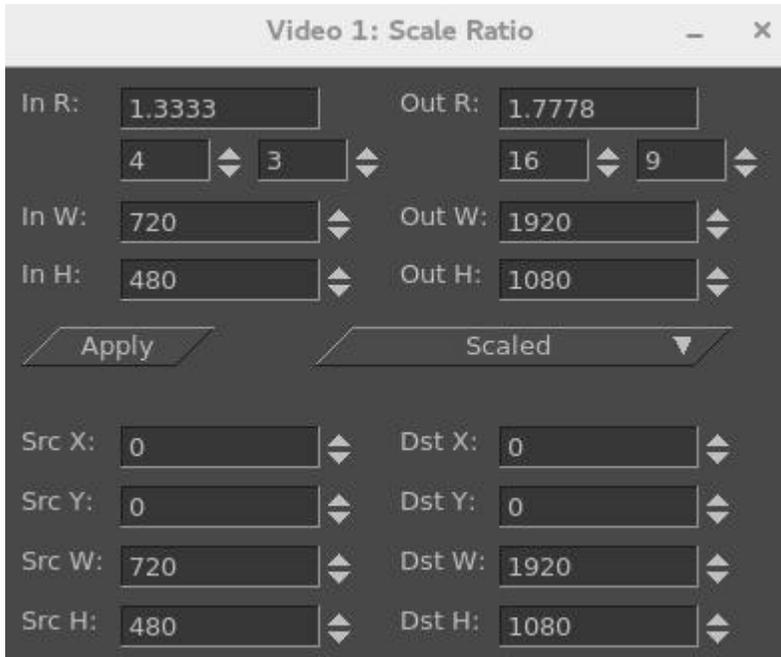
Size - Height and Width in pixel; plus pulldown menu with *preset*.

Scale - Aspect ratio

Constrain ratio - Lock height to width rate

Scale Ratio

With the *Scale Ratio* plugin you can manipulate your video to maintain the pixel aspect ratio (*proportional geometry*), possibly for a different output Display device.



On the first 2 lines, there is *In R* and *Out R* representing the current input and output aspect ratios. Use the *arrows* to change to your desired values. Next you have the *In W/H* and the *Out W/H* for Width and Height.

Left and right sides of the bottom portion show the *Source* and the *Destination X, Y, W, H* values. As you change the values on the left side, you can see how this will affect the *output* as you observe the results in the *compositor* window. For example, as you change the values for *SrcY* in a *cropped* Scale scenario, you see *up/down* movement.

In the middle of the plugin on the right-hand side, you can set the *Scale* type of *None*, *Scaled*, *Cropped*, *Filled*, *Horiz edge* and *Vert edge*. The *top* part (aspect ratio data) is used to compute the *bottom* part when the *Apply* button is pressed. The bottom part allows you to reposition the image *input* or *output* to customize the results. When the *in/out* aspect ratios are different, the output must be *cropped* or *filled* to fit the *output* and maintain *pixel square* appearance.

Selective Temporal Averaging

This plugin is designed to smooth out *non-moving areas* of a video clip. *Denoise* is generally done on a *spatial basis*, mediating the values of a group of adjacent pixels to achieve greater uniformity. The effectiveness of *Denoise* can be increased by also introducing a *time average* between a group of successive frames. The union of these two phases is the basis of the plugin. In fact the smoothing is performed by averaging the color component for each pixel across a number of frames. The smoothed value is used if both the *standard deviation* and the difference between the current component value and the *average* component value are below a *threshold*. The *standard deviation* is a mathematical index used to estimate the *variance* of a group of pixels: at high values corresponds more variation of the pixels and therefore more noise.

The *Selective Temporal Averaging* plugin plays on the homogenization of the values of a group of pixels in a group of frames, based on a *threshold* below which the original values are left and above which the average is performed and then the noise reduction.

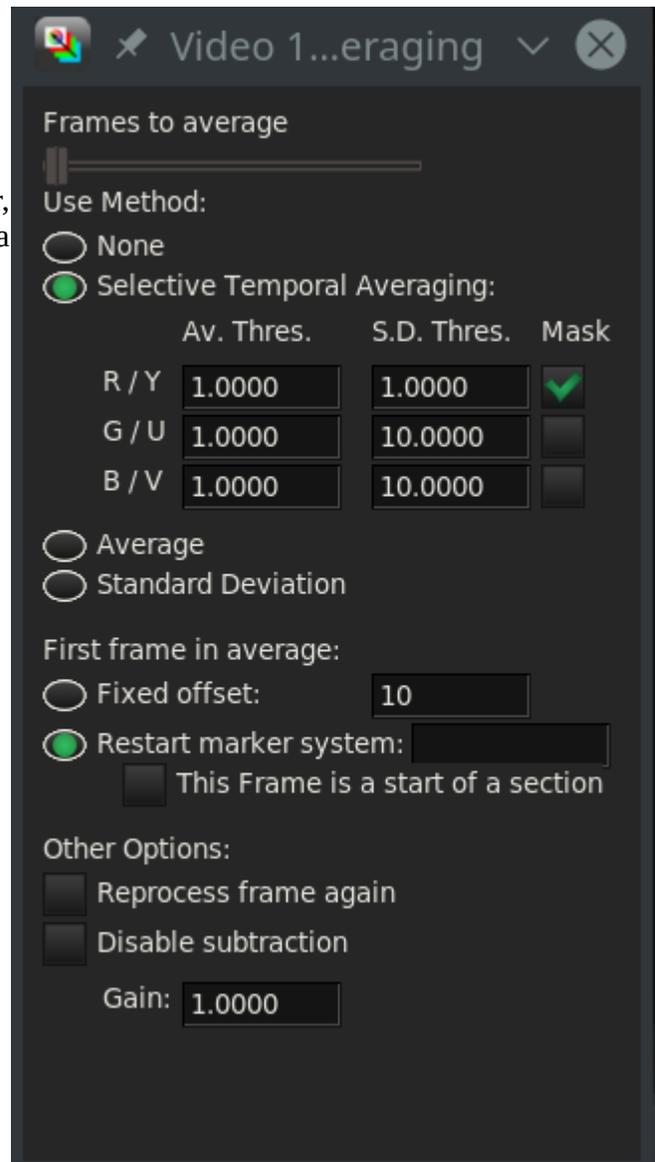
The *average* and *standard deviation* are calculated for each of the components of the video. The type of components averaged is determined by the *color model* of the entire project. The average and standard

deviation of the frames can be examined by selecting the specific radio button in the plugin options window.

The region over which the frames are averaged is determined by either a *fixed offset* or a *restart marker system*. In a restart marker system, certain keyframes are marked as *beginning of sections*. Then for each section, the frames surrounding the current frame are used as the frames to average over, except when approaching the beginning and end of a section, where the averaging is performed over the N beginning or ending frames respectively.

An example of common usage is to select the number of frames you wish to average.

1. Enter a reasonable number of *frames to average* (for example, 10).
2. Select the *Selective Temporal Averaging* method and enter 1 and 10 for all the *Av. Thres.* and *S.D. Thres.* respectively. This basically causes all pixels to use the average value.
3. Turn the *mask* for the first component on. This should make the whole frame have a solid color of that specific component.
4. Slowly reduce the *S.D. Thres.* value. As you do so, you will notice that the regions vastly different from the average will have a flipped mask state. Continue to reduce the threshold until you reach the point at which non-moving regions of the video have a flipped masked state. This value is known as the *noise-floor* and is the level of *natural noise* generated by the *CCD* in the *camera*.
5. Repeat the same procedure for the *Av. Thres.*
6. Turn off the *mask*.
7. Repeat this for all channels.



Sharpen

Sharpen the video, either the *luminance*, *horizontal*, or *interlace*.

Shift Interlace

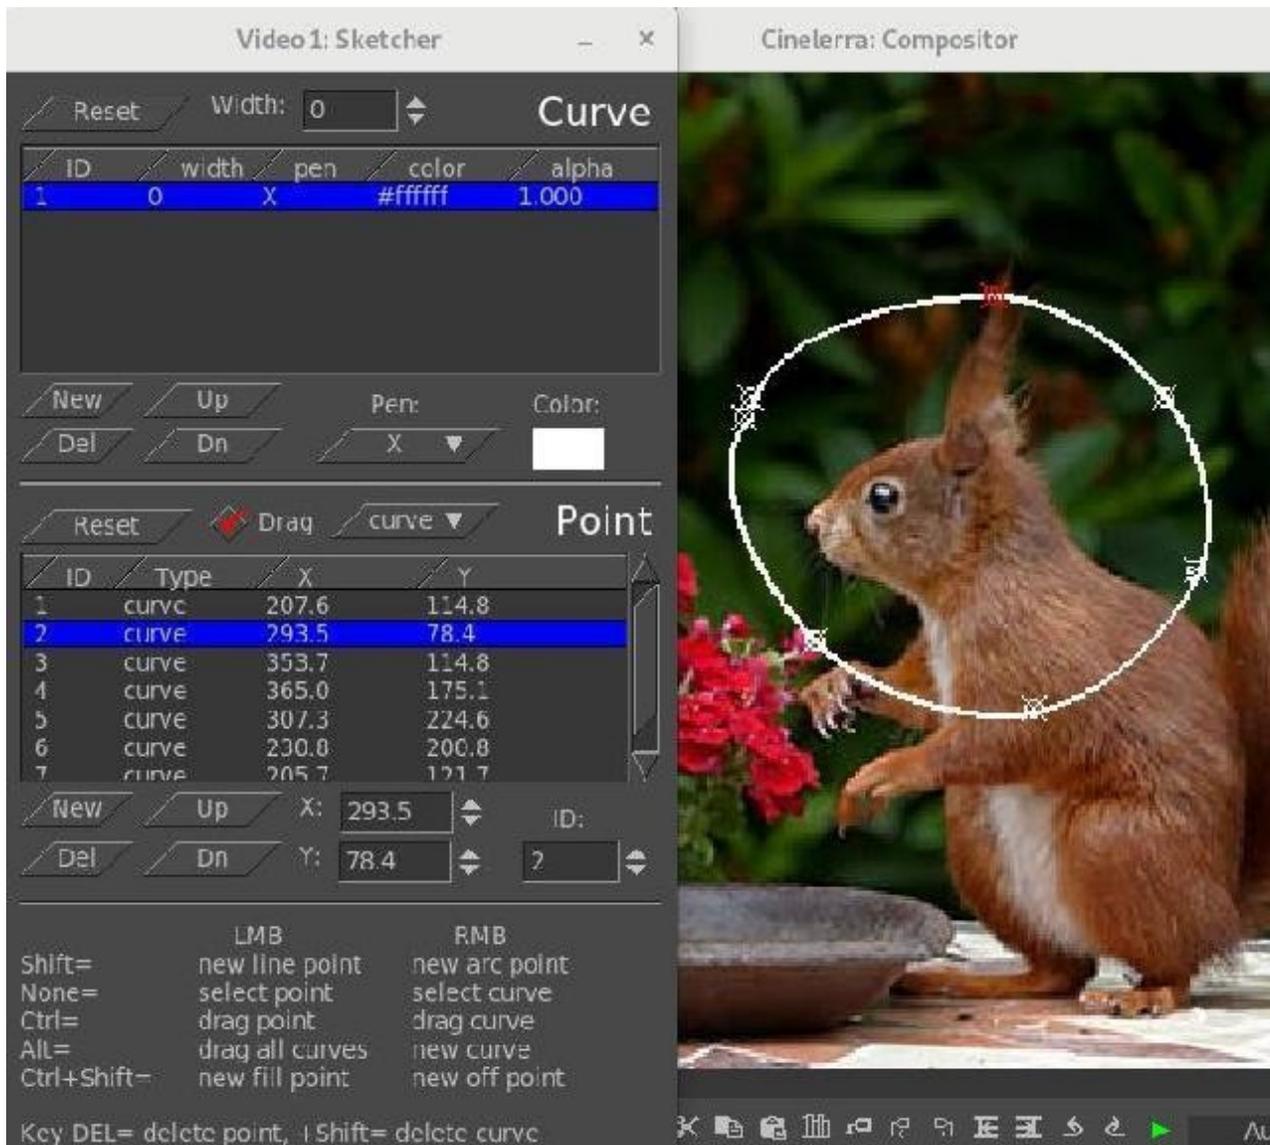
Shift the *interlace* lines using *odd* or *even*.

Sketcher

Now you can *sketch* lines, curves or points, on your video in different colors using various *pen* widths and *pen* type with the *sketcher* plugin. You can even *fill* them. Getting started is fairly easy - simply hold down the *shift* key while using the *left mouse* button to create a bunch of *points* in the *compositor*

window. However, to do more than that you will need to understand the buttons and options or you may end up with unexpected results.

Screencast below shows the *Sketcher* gui and the sketch lines/curves created in the *Compositor*.



In the above screencast, note the *Sketcher* window gui *Curve* top section and the *Point* bottom section. The *white circle* sketch is *id #1* in the curve section. Since *id #1* is highlighted in the *Curve* section, the *X/Y* coordinates in the *Point* section below show the points used to create the *loop*. The *red point* on the image is *id #2* point that is *highlighted*

Some basic rules:

Uncheck *Drag* before rendering; if you do not uncheck, the *drag points* will show in your output.

Drag must be checked on to edit the data. *Drag* must be checked **off** to use *Click to play*.

If the *drag* button flickers when clicked then another window has *drag focus*. Un-focus it first.

Turn *drag* **off** to see what the *sketcher* figure will look like when rendered.

Curves can be any number of points or just a single point.

You must create a new curve if you want a new string of points not connected to the current curve.

Highlight the curve id # for a specific curve to modify its values.

Left mouse click or right mouse click on an existing *point* on the highlighted line/curve id # will automatically highlight the selected point in the *Point* section of the gui and turn *red* in the image.

There will always be 1 empty curve automatically defined when you *start* or even *reset*.

You can not delete the empty default curve but you can use it for a *curve*.

Left mouse button click almost anywhere on the *compositor* screen will automatically show, in the currently highlighted curve, the closest point by turning it *red*.

There is no *undo* recorded between gui updates. Recommend using the option **b** to save a *backup* if you get to a place where you want to make sure you do not lose your sketch.

Compositor: Mouse usage	Compositor: Action	Plugin GUI
shift+left mouse button	create a new line point	New button in Point (line)
shift+right mouse button	create a new arc (curve) point	New button in Point (curve)
left mouse button	select a single line point	Click point listbox item
right mouse button	select a single arc (curve) point	Click curve listbox item
ctrl+left mouse button	drag point	
ctrl+right mouse button	drag curve	
alt+left mouse button	drag all of the curves together	
alt+right mouse button	create a new curve	New button in Curve section
ctrl+shift+left mouse button	create a new fill point	New button in Point (fill)
ctrl+shift+right mouse button	create a new off point	New button in Point (off)
delete key	deletes highlighted points	Del button in Point section
delete key+shift	deletes highlighted curves	Del button in Curve section

Other Button and Label Descriptions:

Color – refers to the current curve pen color. Click on the Color rectangle to bring up the Color window to change any of the color values, including alpha.

Drag – check to create curves and to be able to see any Off type curves. Uncheck for render/viewing.

Reset – in the *Curve section*, all of the curve lines will be deleted except for an empty default curve.

– in the *Point section*, all points for the highlighted curve in the Curve section will be deleted.

Width – integer width of line; width of 1 is the default; width of 0 is the smallest size of 1 pixel.

ID – is the label number of the curve or point.

Pen – value can be set with the pulldown to box (a square), +, /, x or **off**. The default is **x**.

Alpha – alpha value as set in the Color window. Default is 1 which is totally opaque.

Arc Type – value can be set with this pulldown, which is to the right of the Drag checkbox, to either *off*, *line*, *curve*, or *fill*. *Off* is the default. If the line type is *off* when you uncheck drag, you will no longer see the line/curve. Fill specifies a point be inside the region to be filled with the chosen color. Multiple fill points can be used within a single curve.

Up – moves the highlighted Curve or Point Up in the stacking order. Multiple points, but not multiple curves, can be highlighted and moved.

Dn – moves the highlighted Curve or Point(s) Down in the stacking order.

X – X is the point's X coordinate.

Y – Y is the point's Y coordinate.

ID number detailed information:

Points and curves are identified by numeric *ids*, and not the table position. This is so that deleting and inserting points/curves does not shift the interpretation of which points are associated for interpolation. For purposes of this discussion, “prev” refers to “previous”. The prev keyframe constitutes the draw list of curves and points. The prev ids are used to access the next keyframe interpolation input data. When “next” exists for prev.id the interpolation functions apply. When “next” does not exist, “prev” is returned for the interpolation value. The keyframes may be any id sequence; both the prev and next keyframe id sequences are arbitrary non-unique sets. This means that if an id is specified redundantly, it causes a redundant draw when it is the prev keyframe, and only the last item with the id in the next keyframe is used for interpolation.

Fill detailed information:

The *fill* point is a marker point, and not really part of the curve. That point identifies the *inside* of the loop. It can also be used to identify the *outside* of the loop in order to fill that. The loop is created by drawing a line draw from the last point to the first point in the curve. The line type of this last segment is from the last point of the curve. If there are isolated loops in the curve (it is self intersecting) then you will be able to use multiple fill points to fill these regions.

Sphere Cam

Converts multiple fisheye images into a panoramic projection.

Swap Frames

Swap frames with 0-1 2-3 4-5... or 1-2, 3-4, 5-6... .

Swap channels

Swap R G, B, Alpha with another color channel.

Threshold

Threshold converts the image to pure *luminance*, and replaces pixels with one of three colors based on the luminance. Pixels with luminance values in the *low range* are replaced with *black*, pixels in the *middle range* are replaced with *white*, and pixels in the *high range* are replaced with *black*. *Color* and *alpha* (transparency) for each range are configurable with the buttons and interpolate according to keyframes.

The threshold window shows a *histogram* of luminance values for the current frame. **Click** dragging inside the histogram creates a range (*blue*) to convert to *white*. **SHIFT-clicking** extends one border of this range. *Values* for the threshold range can also be specified in the text boxes.

This effect is basically a primitive *luminance key* to produce a *matte*. A second track above the track with the *threshold* effect can be *multiplied*, resulting in only the parts of the second track within the threshold being displayed.

Time average

Time average is one effect which has many uses besides creating *trail patterns* of moving objects. The main use is *reducing noise* in still images (or in the motionless parts of a video). Merely point a video camera at a stationary subject for 30 frames, capture the frames, and *average* them using *time average* and you will have a high quality print. In floating point colormodels, time average can increase the dynamic range of low quality cameras.

Inside the time average effect is an *accumulation buffer* and a *divisor*. A number of frames are accumulated in the accumulation buffer and divided by the divisor to get the average (for 10 accumulated frames the divisor is 10). Because the time average can consume large amounts of memory, it is best applied by first disabling playback for the track, dropping the time average in it, configuring time average for the desired number of frames, and re-enabling playback for the track.

Frames count: this determines the number of frames to be accumulated in the accumulation buffer. Ranges from 1 to 1024 frames.

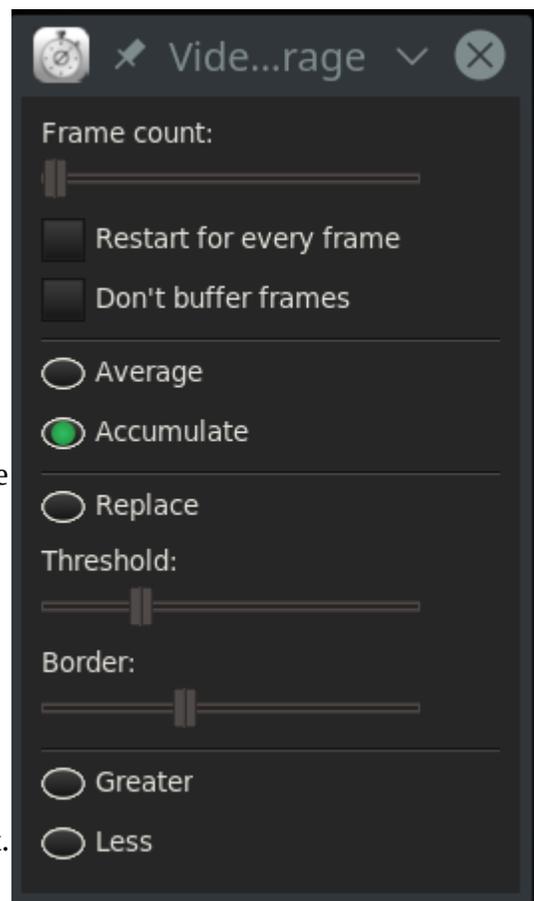
Accumulate: this outputs the accumulation buffer without dividing it.

Average: this causes the accumulation buffer to be divided before being output. This results is the average of all the frames. The result is similar to *Selective Temporal Averaging*, but not configurable.

Replace (Threshold, Border): causes the accumulation buffer to be replaced by any pixels which are not transparent. In combination with *motion tracking* it allows entire sequences to be combined to form panoramas.

Restart for every frames: if an effect before the time average is adjusted the *time average* normally does not reread the accumulation buffer to get the change. This forces it to reread the accumulation buffer when other effects change.

Don't buffer frames: in order to represent the accumulation of only the specified number of frames, the *time average* retains all the previous frames in memory and subtracts them out as it plays forward. It would run out of memory if it had to accumulate thousands of frames. By disabling *subtraction* the previous frames are not stored in memory and only the average function is affected by the frame count.



Timefront

Space-temporal warping enables time to *flow* differently at different locations in the video. This plugin divides the frame into segments (*bands*) whose *type*, *position* and *size* can be configured at will. Each segment will move with a different *speed* from the others, greater than the band that precedes it and less than the band that follows.

Type: allows you to choose a *Linear* or *Radial* segmentation; or to make it depend on the *alpha channel* or on another *track* (with *Shared Effect*). In the case of *Linear* you can choose the orientation of the strips (*Angle*); in the case of *Radial* you can choose the coordinates (*X,Y*) of the center. For a more precise adjustment we can make the bands visible with the button *Show Grayscale (for tuning)*.

With *Inner* and *Outer Radius* we can position the beginning and end of the *bands* in the frame (similar to the *Gradient* plugin), so you can choose the spatial *range* in which to apply the effect. With *Time Range* we decide the size (and therefore the number) of the bands.

Rate allows you to choose the type of algorithm to use when switching between the start and end bands. You can reverse the direction with the *Invers* button. The three modes are: *Linear*, *Log* or *Quadratic* (exponential). This is a *warping* framework plugin based on this article:

http://www.vision.huji.ac.il/videowarping/HUJI-CSE-LTR-2005-10_etf-tr.pdf



Title

The *Titler* allows you to add *text* from within Cinelerra. The titler has standard options for *font*, *size*, and *style* plus many options as described next.

Justify - justifies the text relative to the *entire frame*. Once justified, the *X* and *Y offset* is applied. This allows text to be justified while at the same time letting you push it within the title safe region.

Motion type - scrolls the text in any of the four directions. When using this, the text may disappear.

Make sure the *speed* is set to a reasonably high value (for example 150) and move the insertion point along the timeline until the text is far enough along the animation to reappear. The text *scrolls on* and *scrolls off*. Setting *loop* causes the text to scroll completely off and repeat. Without loop the text scrolls off and never reappears. The *speed* of the animation is determined by speed in pixels per second. Set it higher to speed up the animation.

Smooth - with *Chroma Key*, a shadow may show; remove by enabling this to *add/remove* the *shadow*.

Drop shadow - draws a black *copy* of the text to the bottom right of the original text. This is useful when drawing text over changing video to keep the border always visible.

Fade in/Fade out are a second type of *animation*. If the *fade* seconds are **0**, no fading is done.

Color and *Outline* – allows you to pick the color for using to draw the text or its outline.

Stamp timecode - replaces text with the current position on the timeline in *seconds*, *frames* or *samples*.

The title effect supports *keyframes* only for *Justify* and *Text*, with no *interpolation*. To add *subtitles* to your movie set a single title effect and then define keyframes. If you enable the *automatic keyframes* toggle, a new keyframe is created each time you *edit* the text. Check **View** → **Plugin autos** to make them visible on the timeline. In the text input box you will see the subtitle displayed under the *insertion point*. To correct an existing subtitle, the *automatic keyframes* toggle must be **off**. To adjust the *timing* of subtitles simply drag the *keyframes*.

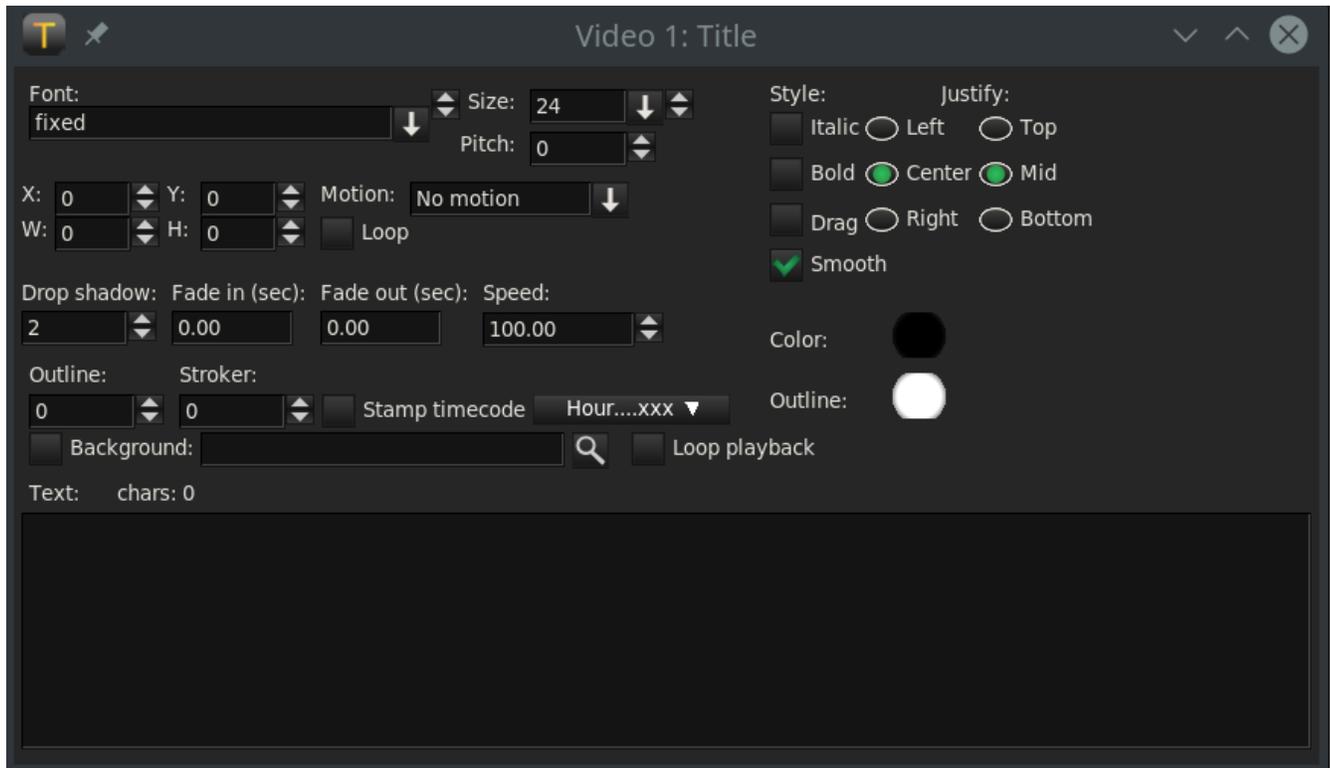


Fig. 94: GUI of the Title plugin

To create special effects for your title you can place it on a dedicated *track* and insert other realtime video effects just under the *title* effect and/or use camera and *projector*. With keyframing you can *animate* your title and make it change *position, size, color, transparency, texture, or shape* over time.

For improving *playback* performances of titles with effects, you can reduce the *size* of the dedicated track. Right click on the track and select *Resize track....* Enter the smallest resolution that still keeps the title visible.

To include graphical elements like *logos*, you may want to import your *title* as a *PNG image* (alpha channel *transparency* is possible), move it with *camera* and *projector* or add *effects*.

Adding fonts to the titler

The *X Window* system originally did not have a suitable font renderer for video. It also is restricted to the current *bit depth*. It does not have a convenient way to know which fonts work with the suitable font renderer in the desired bit depth. The easiest way we have found to support fonts in the *titler* is to have a directory for them at ``usr/lib/cinelerra/fonts'`.

The *titler* supports mainly TTF, true type fonts. It supports others but TTF are the most reliable. To add true type fonts, copy the ``.TTF'` files to the ``usr/lib/cinelerra/fonts'` directory. In that directory run **`ttmkfdir && mv fonts.scale fonts.dir`** and *restart* Cinelerra. The new fonts should appear. The usage of *ttmkfdir* changes frequently so this technique might not work.

If the video is displayed on a consumer TV, the *outer border* is going to be *cropped* by 5% on each side. To avoid text which is too close to the edge looking bad, you may want to enable the *title-safe* tool in the compositor window. The text should not cross the inner rectangle.

Some recently added options:

Drag – initial default checkbox is **off** so that the *Title* plugin will work as it always has.

Anchors - When you turn **on** the *Drag* feature, nine different *anchors/handles* will appear on *compositor* window. The *middle anchor* allows you to drag your title wherever you want in the *compositor* window (*X, Y coordinates*). The other 8 *handles*, drawn as arrows in each corner and in the middle of each side, let you change the *size* of the drag area box so that your title is within that area if it fits and as it is directed.

W/H – the values in these 2 boxes specify the *size* of the *drag area box* measured in pixels as shown in the *compositor* window. You can set these *manually* and if you can't see the location of your box or find your handles, set them to zero because 0 sets it to the same as the width/height of the media. The *Drag* effect ignores all *boundaries*, including the *Title Safe Region* of the *Compositor* so that if you drag your titles off the screen, it will look like they disappeared completely. Reset *X* and *Y* to reasonable values to have it reappear. The *Title* text, *background*, and *pngs* are applied on a *single layer* so that they will drag together as an entity. All of the *Title* capabilities work in conjunction with *dragging* so if you want to *justify* the title, you can still use the *Left/Center/Right/Top/Mid/Bottom* within the drag area. Be sure to turn **off** *Drag* when rendering or the box will show in the video; keep in mind that drag bars do not appear until there is some text in the text box and you can not actually *drag* until the *Title window controls* are available.

Attributes – in the *Text box* where you type your Title information, you can now change several *attributes* to give you plenty of flexibility. Each of these special *attributes* begin with an *open angle bracket* < and ends with a *closing angle bracket* >. Until the closing angle bracket is *keyed in*, the actual characters you type, will be seen in the *compositor* window. In addition, if you do not use the exact syntax or you keyin a filename that is not available, all of the characters will continue to show up. This helps to see what needs to be fixed or is missing. The attributes usage is described in the table below.

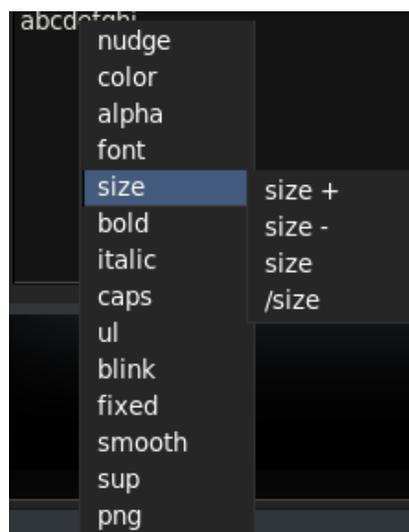


Figure 95: Pulldown Attributes

Format:

<Attribute_name Attribute_value> → Sets the attribute to value.

</Attribute> → Resets the attribute to previous value.

<Attribute> → Reverts to value set in controls config applies to color, alpha, size and font.

<i>Attribute name</i>	<i>Attribute value</i>	<i>Notes</i>
color	color name such as RED from <cin_path>/guicast/colors.h	Or use the hex value like #a000a0; color-hex.com shows examples
font	exact name from “Font” pulldown	When you set font, bold/size and italic will be as currently set up
alpha	floating-point number between 0 and 1	0 is transparent; 1 is opaque
size	+ increases the size by 5/4, - decreases 4/5 or use a number > 0 but less than 2048	Examples: <size +++>, or <size - or <size 14.5> (floating-point ok)
png	filename of a .png file	Example, add a logo. Full pathname needed if not in current directory
bold	for “on” (default) or use 0 for “off”	Font used must have bold available
italic	1 for “on” (default) or use 0 for “off”	Font must have italic already set up
blink	a number in seconds to flash on and off; a negative number for fade in and out	Can be a floating-point number; <blink> with no value is like 1
ul	underline 1 for “on” (default), 0 for “off”	
caps	1 for “on”, 0 for “off”, -1 for lower case	If no value set, treated same as 1
sup	1 for superscript; 0 ends; -1 for subscript; positive bumps up by 1/2 of font size; negative sinks by 1/2 of current size	Size of text will be 1/3 of font size; example: 12<sup 1>10<sup 0>; <sup> with no values, acts like 1
fixed	number of pixels between characters; fixed with no value uses ¾ current size	A fixed font is very useful for tables; /fixed reverts to previous fixed size
nudge	x,y to move the text by that amount; displacement is in number of pixels. Can be negative or positive	Example: <nudge 8,8>abc</nudge>
smooth	add anti-aliasing to smooth jagged edges	Turn off smooth for chroma key

These attributes stay in effect until you change them or reset them. Additional cpu time is needed for the *blink* attribute because it requires redrawing every frame and for the *background* option described below. Note that some *Title window controls* can not be set, such as *Justify* and *Outline color*. The lines below are examples for testing purposes. The accompanying screenshot displays the corresponding cinelerra output.

Examples of Title line keyins:

<size 15>Buddy, the <color red><bold 1>bad dog</bold><color white> ate my homework !<png /tmp/buddy.png><ul 1><size +>There are<ul 0><size +>2<sup 1>10<sup 0>cats to chase?

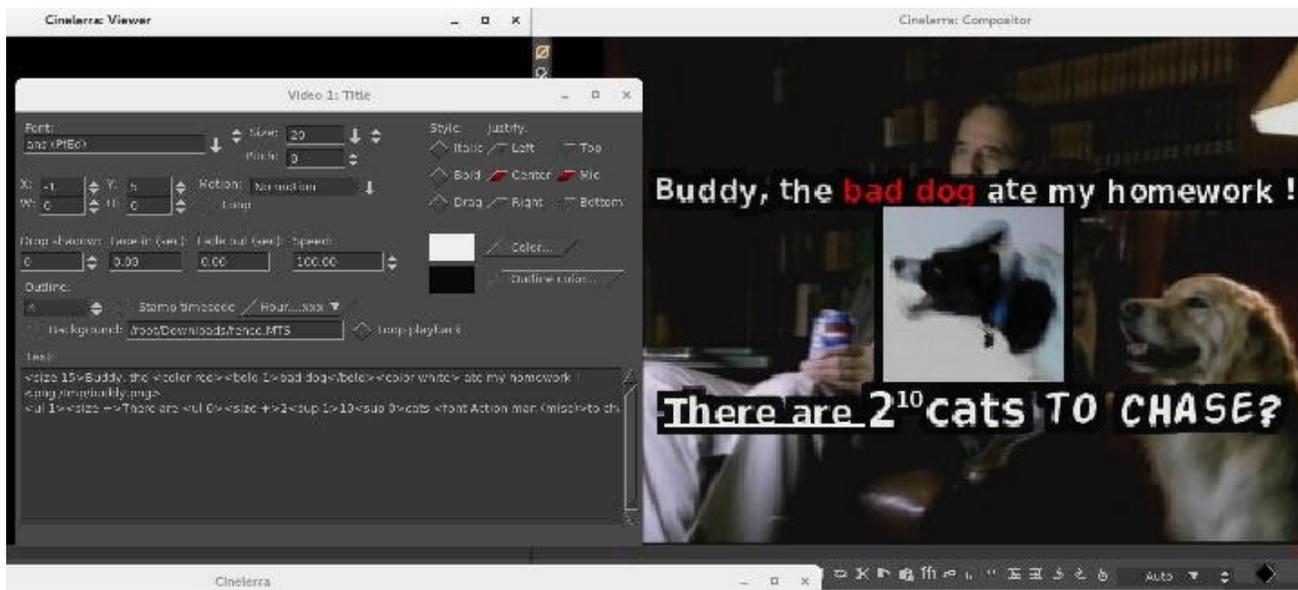


Fig. 96: Final result of the exemple

Special Characters – (< > / \ #)

Besides the previously described <, >, and / characters, there are two special characters: *backslash* “\”, and the *pound sign* “#”. The backslash character is used for two things. With the advent of the *attribute name* and *value*, your line may become quite long so you can use “\” followed by the *carriage return* to continue on the next line as if it is just a single line. It also can be used to designate that the following character does not represent the beginning of an attribute. For example, if you want to use the opening angle character “<” as a *title character*, precede it with the backslash character. The pound sign, “#”, designates the whole line as a *comment* or if in the middle of the line, then the rest of the line is a *comment* (this includes the *carriage return*).

Background – in this box you can keyin the name of a *file* of the type that cinelerra accepts and use that file as a *background* for your Title characters. This will be seen in the *compositor* window on top of the video that is loaded in the main track canvas. Besides typing in the filename, you must also check the checkbox. This makes it easy to turn it **on** and **off** to see what it looks like. Next to the background box is a *Loop* checkbox. If the *background file* takes less time than the main track canvas video to run, you can turn on the *loop* checkbox so that it runs over and over again to match the time size of your video.

Stroker – to add *pen strokes* to the text letters, adjust the stroke width numerically. This looks particularly nice on certain fonts and with a negative adjustment of the *Drop shadow*.

Unicode Insertion – if you want to enter a special character like the mathematical summation symbol, you can use the *unicode* equivalent to do so. **Ctrl-Shift-U** followed by **2022** and a *carriage return* is an example of the *bullet*. Refer to section 40.11 for details.

Popup Helper – put your cursor where you want to add an attribute, then right mouse will bring up a list of the available attributes for you to choose, along with a submenu to choose from. The program will insert that attribute for you and all you have to add is a value when required! (see Fig. 51).

Color Picker Usage:

The *Text Color* window has several enhanced features as listed here and seen in the next screenshot.

1. The *hex* value of the color you choose shows in the textbox and you can also keyin a value there.
2. A small square next to the hex text box, is a green *eyedropper* color picker. Use the left mouse button to click on the square to enable picking and you will see it turn red to designate that it is enabled. Your cursor will switch to a two-colored reticle. You can now move the mouse around to choose a color *anywhere on the screen* and then click there to have it picked. If you hold down the right or left mouse button while moving, you can see the color changing in the vertical bar in the color palette area as you move to give you a bigger view of the actual color. The *eyedropper square* is seen in the next screenshot.
3. A *history* of 16 of your last chosen colors is available to easily use again. Any time you choose a new color in any methodology, it will become the latest choice in the history either immediately or after checking OK/and leaving. History shows latest color starting left to right..
 1. Besides HSV, RGB, there is also a YUV color model to choose from.
 2. Several of these items may have associated Tool Tips.

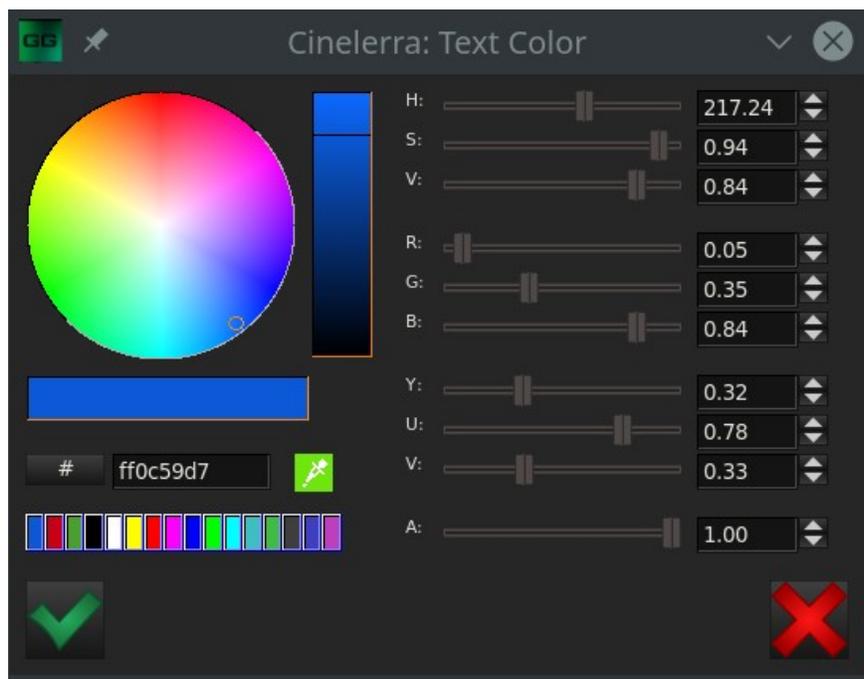


Fig. 97: Screencast showing the Color Picker menu.

Font Choice:

In order to choose a *font* faster, you can keyin the first few characters of the font name, being sure to use capital characters if used since it is *case-sensitive*. The steps to follow are:

- 1) next to the font box, click on the *down arrow pulldown*;
- 2) keyin the first character(s) of the desired font and you will see the first match become *highlighted*;

3) you can see the characters you keyed in the upper right corner of the fonts.

Font Addition / Font Subtraction:

Some of the *system fonts* are automatically included in the set of fonts being used by cinelerra. The easiest way to add additional fonts for the *Title* plugin's set, is to use ones available in specific directories on your computer as long as they have a `fonts.scale` file already set up. You can run **mkfontscale** to create this file within that directory. In order to include a specific directory you set an *environment variable* before starting cinelerra which stays in effect until it is unset or until the next reboot. Below is the method and an example.

```
export BC_FONT_PATH=<colon-separated-search-path-for-fonts>  
export BC_FONT_PATH=/usr/share/fonts
```

The current set of fonts in cinelerra's directory will be automatically included and will be the *default set* if this *environment variable* is not set. Keep in mind that if you add a lot of fonts, it will considerably slow down the *startup* every time you bring up the *Title* plugin.

If you want to only have a limited number of *fonts* set up, you can manipulate the cinelerra directory directly at `<cinelerra_install_path>/bin/plugins/fonts`. Here you will find the default set of fonts that come with the install. Copy any other fonts you would like to include here with read permission, delete any fonts you do not want to have, then execute **mkfontscale** which creates the file `fonts.scale` that cinelerra will read. However, the next time you install a new version of cinelerra, your changes will be written over so you will have to make sure to save them elsewhere and then re-establish.

If you have problems with a specific font or set of fonts, there is a debug option available to determine which font is an issue. When starting cinelerra, you should set up the variable:

```
export BC_FONT_DEBUG=1      (default is 0 for no debug)  
unset BC_FONT_DEBUG      (to remove debug messages)
```

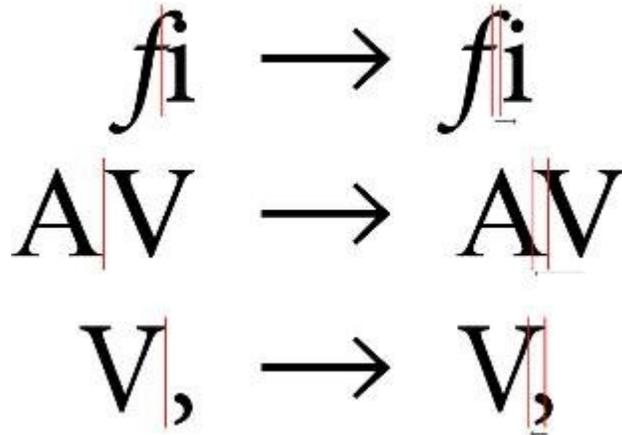
Then start cinelerra from a *terminal* window to see the fonts being loaded and previewed in the *Titler*. This should point out any issues. Another debug methodology is to *remove all fonts* from being used and subsequently add in the ones that you want. For example:

```
export BC_FONT_PATH=: (the : "colon" removes all automatic system and cinelerra fonts)  
export BC_FONT_PATH=/usr/share/fonts (remove all fonts and then add /usr/shar/fonts)
```

One last *item* of information about fonts that may lead to some confusion. The checkbox for *Bold* and *Italic* will occasionally be *ghosted out* if no bold or italic version of the selected font is available. This is no guarantee, but currently as good as it can get due to inconsistency in the creation of fonts. It is mostly just a *hint*. If boxes are checkmarked, but ghosted, you can not uncheck until you change to a font that does not ghost out the boxes. If you use the *popup helper* with the boxes checked, and attempt to keyin a font that does not have the bold/italic attribute as checked, the font will be considered *illegal*. *Text: chars* is output and updated to indicate the *number of characters* already used. The only limit to the number of characters based on a count of single 8-bit characters is the available resources on the user computer available for cinelerra use. Keep in mind that *unicode* or other special characters may consist of 2 to 4 8-bit bytes. Also, *newlines* are a character and any of the *attributes* you use count in the total. There is now a *horizontal scroll bar* as well as the *vertical* one in the textbox and they only appear when there are more lines or characters that can fit in the original sized textbox.

When using the Titler, kerning is applied in order to allow parts of a letter to go outside the standard sized letter box. Kerning is the process of adjusting the space between individual letters. It is not the same as proportional spacing. In kerning, a bounding box is allowed to overlay another bounding box. The philosophy here is to aim at the ability to have boxes overlap in order to make the letters look more visually appealing. Bounding box and escapement are tracked separately. In addition to adjusting individual letter spacing, the program will also expand the render box in order to have any parts of the letter extend outside the standard letter box. Kerning is applied to any and all fonts.

Screenshot shows how letters can be presented without and with kerning. Note this is just an example.



Translate

This effect allows *displacing*, *cropping*, and/or *scaling* the source video *horizontally* and/or *vertically*. The *In* and *Out* parameters operate similar to the *camera* and *projector* functions in the *Compositor*:
In X/Y specifies how many pixels from the left/top of the *source* you want to start (*camera*).
Out X/Y defines where on the screen you want the *output* to start (*projector*).
In W/H defines how many pixels of the *source* you want to include in each direction.
Out W/H defines how many pixels on the screen you want that source to take up.
Identical values for both *In* and *Out* that are *less* than the *source dimension* will simply *crop* the source. Different values will *stretch* (or *compress* if *Out* > *In*) the source in that direction (and *crop* if *In* is less than the *source dimension*).

This effect supports *keyframes* so these parameters can change smoothly over time. You can use this effect for many things such as having a cropped *insert clip* move across the screen, or have it change *size* or *stretch* while doing so. Be forewarned though, that for *interlaced* footage *horizontal displacements* are likely to destroy the *field order*, resulting in all sort of flickering and jumping movements.

Unsharp

This effect applies a traditional *darkroom* technique, the so called *unsharp mask* to every video frame. With different parameter values, this can be used to *soften* or to *sharpen* the image. Its parameters are:

Amount - moving the slider to the right makes *dark areas* get *darker* and *light areas* get *lighter*.

Radius - this slider controls how much *blurring* is used in the *edge-finding stage*. The practical effect of this is to specify how large a region is darkened or lightened.

Threshold - this slider controls how big a *difference* between a pixel in the *blurred copy* and the *original copy* is needed before any darkening or lightening will be applied.

VideoScope

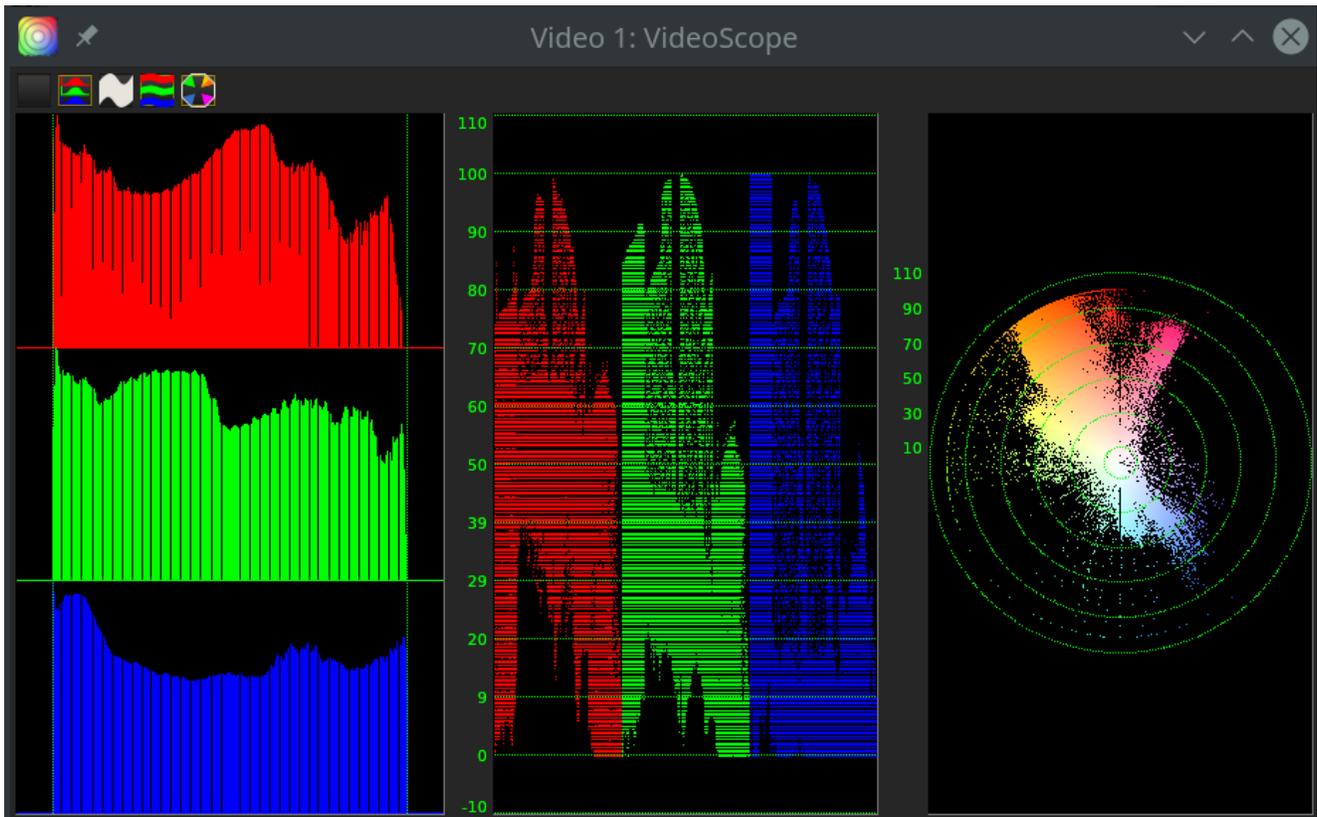


Fig. 98: GUI of the Videoscope. You see Histogram, RGB Parade and Vectorscope

Videoscope summarizes *intensity* and *color* on a *calibrated* display. The Videoscope can be used in conjunction with other Cinelerra plugins such as *YUV*, *HUE*, *Brightness*, *Color Balance* or *Histogram* to accurately correct video for *contrast*, *clarity*, *conformance* (to *normalize* various videos shot under different light settings), or for *cinematic* purposes. The *human eye* is not specialized to match precise level of *light* and *color*, but *Videoscope* is. Videoscope contains two displays: the *waveform* scope and the *vectorscope*, plus the histograms.

Waveform/RGB Parade

The *Waveform* Scope displays image *intensity (luminance)* versus image *X position*. The *RGB Parade* Scope displays image *RGB intensity* versus image *X position* (one graph per channel). The *Waveform* Scope appears on the left side or in the middle of the Videoscope window. The display is calibrated vertically from **0%** intensity (*black*) at the bottom up to **100%** intensity (*white*) at the top. Each *column* of pixels in the image corresponds to *one column* of pixels in the *Waveform* Scope. Note that the *height* of the values of a *waveform/RGB Parade* corresponds exactly to the values on the *x-axis* in the *histogram*. A vertical/horizontal correspondence is therefore obtained.

The next 3 images are described below.

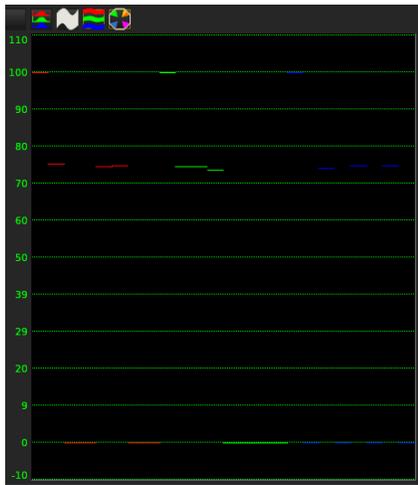


Fig. 99: RGB Parade of the colortest75%

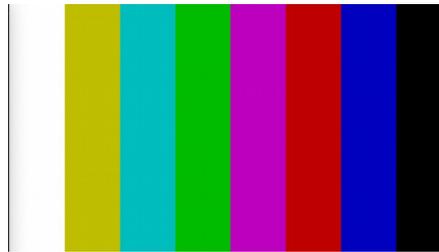


Fig. 101: Colortest75%

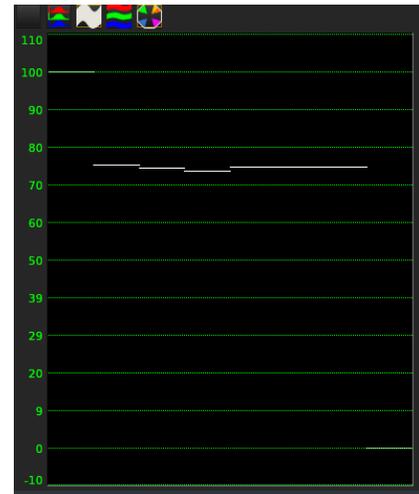
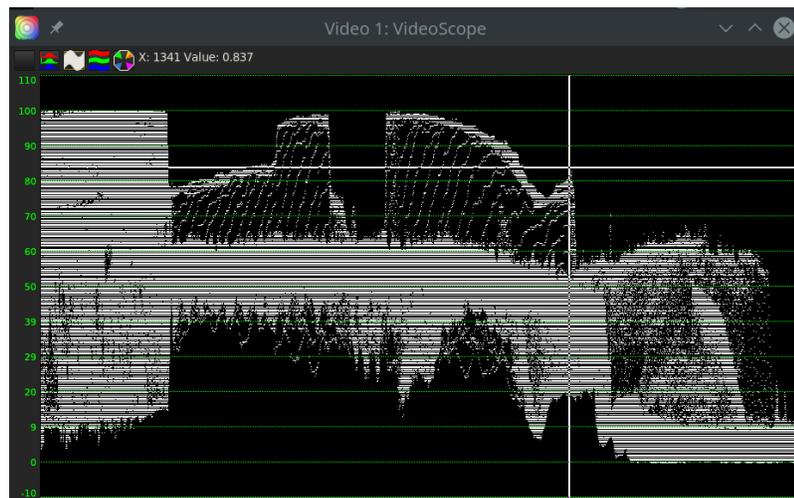


Fig. 100: Waveform of the colortest75%

On the left is shown *RGB Parade*: instead of the color shadows as in Fig. 54, we have lines representing the *color bar test at 75%*. They are *pure colors*, so all pixels have the same value. In fact, they are all at the level of 75 % except for the 100% *white band* and the 0% *black band*. In the *waveform* on the right, we have the same behavior with regard to *luminance*: the *white band* is 100%; the *black band* is 0% and all the others 75%.

The *Waveform* scope helps correct image *light levels* for *contrast range* or for *conforming* light levels on various scenes originally shot on different light settings.



Adjusting light levels (adjusting luminance):

1. Insert the *Brightness/Contrast*, *YUV*, *Color 3 Way* or another video adjustment effect on your track.

2. Insert the *Videoscope* effect on the track below. Make sure that it is placed below so it can see the adjustment effect's results. If it is not, right-click and move it *down*.
3. Show both the effect and *Videoscope*.
4. Adjust the effect while observing the *waveform* to match the desired light level.
5. Precise adjustments can be made by measuring the values on the *waveform* with the *crosshair* (by click with LMB, and reading numeric values on top left of the window) and reporting these numbers in the *effects* window (*Histogram Bézier/Curves*, for example).

For instance, if you are looking for maximum contrast range, adjust the *Brightness/Contrast* levels to align the darkest point on the scope with the 0% level and the brightest portion with 100%. Anything above 100% is over saturated. Limits which may be highlighted with checkbox controls:

HDTV or sRGB (ITU-R BT.709)

The maximum pixel range for *HDTV* or *sRGB* is **[0, 255]**. This range corresponds with levels 0% and 100%.

MPEG or Analog video (ITU-R BT.601)

For analog video or *MPEG* (including *DVD*), the maximum range for RGB is **[16, 235]** (8-bit). For YUV, the maximum range for intensity (Y) is **[16, 235]** (8-bit). This range corresponds to gray levels from 6% to 92%. Reference RGB-601 plugin.

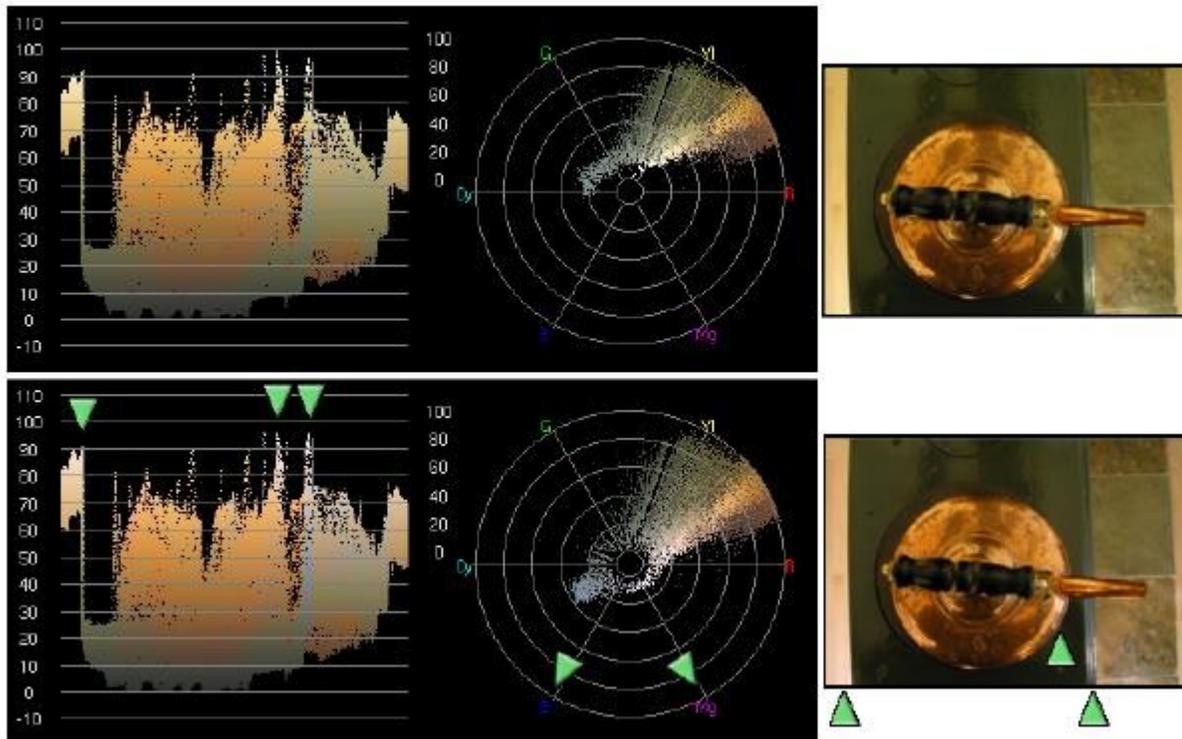
NTSC Television broadcast

If you are producing a video for *NTSC* television broadcast, keep the intensity between **7.5%** and **100%**. The minimum black value which can be broadcast is IRE 7.5% (indicated by the "7.5" level), and values below this level are no darker.

Vectorscope

The *Vectorscope* displays *hue* (angle on the *color wheel*) and *saturation* (radius). Each pixel in the *source* image is drawn as a point on the *color wheel*. The distance from the center is the color saturation. *Gray* values are close to the center, and high saturation values are near the perimeter (100%). In the center there is pure *white* (0%). By clicking with the mouse on the color wheel appear radius and circle whose values of *hue* and *saturation* are shown at the top left of the window, similar to the values of *X* and *luminance* of the *Waveform* and *RGB Parade*.

Vectorscope is used as monitor with other plugins to correct color, adjust image tint, and apply other effects for cinematic effects, image correction, or to conform images to look the same. For example, *skin* tones are found along an axis (+ *I-line*) between yellow and red, and between 0 and 50% saturation values. The blue of the sky is more or less along the opposite axis to that of the *skin* (- *I-line*), with a much wider saturation range.



In this example, the top image is *white balanced*. Vectorscope shows many pixels in the yellow region and few in the white region. To remove the yellow tint, the *Color Balance* effect is used to first shift the vectorscope plot towards *magenta* (Mg), and then towards *blue* (B) until the region previously near the center surrounds the *center*. In the bottom image, *yellow highlights* have become *white highlights* (arrows). Note that the corresponding features in waveform also appear whiter (arrows). The Vectorscope can also be used to verify that the video output will display properly on various monitors. Any points along the inner radius will be displayed as pure white and any points above the 100% radius, will probably not be correctly displayed on the screen.

Histogram:

You can also display the 4 *histograms* (*master* or *RGB*) on the left of the window. (see Fig. 54).

Wave

The *wave* effect adds waves on the image. *Amplitude*, *Phase*, and *Wavelength* parameters can be adjusted.

Whirl

Creates a *whirl* (spiral) of the video around the center.

YUV

Modify the Y, U, V settings.

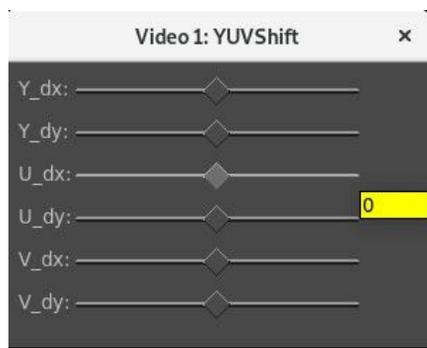
YUV411

Modify the 411 yuv to look like 420 color space instead.

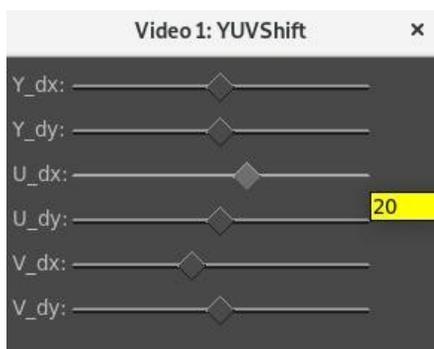
YUVShift

This effect is used for *YUV* input video from older cameras using 3 sensors. It is possible to have misalignment of the 3 sets of numbers: *Y*, which represents the *luminance* or brightness component, and for *U* and *V* representing the *chrominance* (color) components. If they were misaligned in the video, you can use *YUVShift* to realign. To move a specific component *up/down*, modify the *dy* value using the slider bar in the *RGBShift* window. To move a component *left/right*, modify the corresponding *dx* value. If you are using an RGB color space, you will want to use the *RGBShift* effect instead.

Screenshot below shows the *blue U* component aligned too far to the *left*. And the *red V* component is misaligned too far to the *right*. Note the *U_dx* current slider bar set to 0 as shown by the yellow box value in the *YUVShift* plugin window. All components are currently at zero.



A corrected video image is shown in the screenshot below. Now the red and blue colors are correctly aligned. Note how *U_dx* is now at +20 and *V_dx* is now negative to realign the image.



Zoom Blur

Blurs the video from the *center* outwards, like the sun's rays, and uses a *zoom* effect.

X,Y - center of the origin *field*.

Radius - Zoom on the *fields*.

Steps - number of blur steps to be used in the calculation. Increasing the number takes more CPU.

The zoom blur image shown here has the parameters: Radius=21 and Steps=3.



Fig. 102: For clarity of presentation only 3 fields are shown

10.10 Video Find Object and other plugins from OpenCV – Stylize,Puzzle, Flow, Gabor, Move

The *Find Object* plugin searches a *Scene* for an *Object*, and the *Object* can be overlaid with a *Replacement* object. It requires the thirdparty *OpenCV* (Open Computer Vision) library and you will have to do your own build. Also, be aware that the two useful algorithms of **Sift** and **Surf** have been removed from versions of *OpenCV* after version 4.

There are some potential pitfalls with using this to include:

1. *opencv* code is constantly changing, making it difficult to keep the plugin up to date
2. there is some confusion about *licensing/patents* for specific algorithms
3. some parts don't work all that well and it can be slow
4. internally the *colormodel* is converted to *greyscale*, which means the color information is not used

Because of the *build size* of *opencv*, it is not normally included in the thirdparty directory. Building *opencv* adds a lot of storage demand (4GB) to the *cinelerra* build tree, and the *opencv compile time* is significant. For these reasons, *findobject* is not normally built. You can however invoke a build which adds *findobject* to the *cinelerra* plugin library. Building *findobject* will *configure* and *build* *opencv* into the thirdparty library build area if needed (it does not rebuild if already done). System builds are possible, but the system compatibility of the newer *opencv* interfaces can vary between distros. The *opencv* interface is rapidly changing; we recommend the *static* library build and link to avoid problems.

10.10.1 How to Build OpenCV Plugins

To build *findobject* and the other plugins using *opencv*, access the **src** using *git*:

```
git clone --depth 1 "git://git.cinelerra-gg.org/goodguy/cinelerra.git" cinelerra5
```

then configure the build, but add the - **-with-opencv** configure parameter.

```
cd <path>/cinelerra-5.1
```

```
./autogen.sh
```

```
./configure <add std params, eg: --with-single-user> --with-findobject=sta
```

```
--with-opencv may be set to [<bld>][,<src>]
```

```
bld=typ sta,dyn,sys: sta=static, dyn=shared, sys=system libraries
```

```
src=typ git,tar,git=url,tar=url: git (default git_url), tar (default: tar_url)
```

For example, like:

```
--with-opencv=sta
```

```
--with-opencv=dyn
```

```
--with-opencv=sta,tar=https://cinelerra-gg.org/download/opencv/opencv-20180401.tgz
```

Once thirdparty/*opencv* is built, it will be reused. Use target **mrclean** to remove **thirdparty/opencv***. Network access is required to obtain the OpenCV source, at least once. You will need a minimum of 4 GB in the thirdparty build directory and more time to compile.

To get *opencv* built in the easiest way possible (need internet access because builds directly from the *opencv* github but this changes wildly): **./configure <params> --with-opencv=sta,git**

10.10.2 Using OpenCV Plugins from the Automatic Builds

The *OpenCV* plugins are built only in the *64-bit* tarball builds, both static and dynamic. However, due to size these plugins are not included with pkgs. But it is relatively easy to add the current 6 plugins for your distro via a simple procedure of copying the plugins from the tarball to the cin5 install plugin path. They are:

```
cin/plugins/opencv/findobj.plugin
cin/plugins/opencv/flowobj.plugin
cin/plugins/opencv/gaborobj.plugin
cin/plugins/opencv/moveobj.plugin
cin/plugins/opencv/puzzleobj.plugin
cin/plugins/opencv/stylizeobj.plugin
```

- 1) do your package install of the current build for your distro as usual;
- 2) look in <https://cinelerra-gg.org/download/tars> to see your distro name's static tar;
- 3) download the corresponding distro *static tarball*;
for example for *arch*: https://cinelerra-gg.org/download/tars/cinelerra-5.1-arch-{date}-x86_64-static.txz
- 4) create a temporary directory on your computer
- 5) **cd** that-directory
- 6) **tar -xf** location-of-the-tarball-you-downloaded
- 7) **cp plugins/*obj.plugin** <see below for your location>/. (note the period on the end!)
location for most User installs, this is: <cinlib_path>/plugins/
location for some System installs, this is:
 /usr/lib/cin/plugins/ (most ubuntu distros)
 /usr/lib64/cin/plugins/ (Leap distro)

Start cinelerra and look for the six plugins in *Video Effects*.

To reverse this, simply delete the six plugin files (eg. **rm /usr/lib*/cin*/*obj.plugin**).

10.10.3 Description of Find Object Plugin

As in the standard OpenCV *FindObj* program, there are 5 *detector* methods and 2 *matcher* methods which can be selected. They detect *features* and match them as a *rectangular projection*. The matched region will be overlaid with a *replacement image* if *replace object* is enabled. This is done using a variety of feature *detectors* and region *matches*. The *match* works by creating sets of *Feature points*. These points are generated for both the *source* and *reference* object frames. Then the two sets are *matched* for *Homography* (a regional similarity).

Matchers:

FLANN – Fast Library for Approximate Nearest Neighbors

BF – (flann not selected) Brute Force

Detectors Algorithm:

SIFT – Scale-Invariant Feature Transform

SURF – Speeded-Up Robust Features

ORB – Oriented fast and Rotated Brief

AKAZE - Accelerated Keypoint detector And descriptor Extractor

BRISK - Binary Robust Invariant Scalable Keypoints

Don't Calculate

The *cinelerra plugin* has several additional features for ease of use and more flexibility. A description of parameters is outlined below.

Mode:

Square
Rhombus
Rectangle
Parallelogram
Quadrilateral (default)

Reset – to get back to the default values.

Algorithm – *detector* selections as described above.

Replace object – checked if replacing the object is the desired outcome.

Use FLANN – use the *Matcher*, or if unchecked just use *brute force*.

Draw match – will display a blue line around the matching object only for the purpose of verifying the placement. The starting point for the box is designated by a white outlined circle in corner #1.

Scene/object/replace layer – number designating the attachment track of each findobject plugin element.

After *matching*, the resulting *projection corners* are used to do reshaping. The general procedure is:

- Calculate centroid, translate center to 0,0, check for bowties and rotation direction.
- Calculate replacement object corners using *mode*, *central angles*, *scale*, and *aspect* of the projection.
- Using the replacement, apply the inverse aspect, rotation, scaling, and translation as specified.

Selectable specifications:

Aspect – when set, the *matching* aspect ratio is applied, otherwise the original object aspect is applied.

Scale – when set, the match surface area is used for scaling, otherwise the original surface area is used.

Rotate – when set, the match orientation is used in the projection, otherwise the original orientation is used (orthonormal).

Translate – when set, the match replacement is centered to the projection center, otherwise the original object center is used.

Once the *replacement corners* are calculated, the replacement object is overlaid to the scene frame.

Scene/Object/Replace Drag – only 1 can be checked on at a time (and that includes counting the *Title* plugin *drag box*). The *Drag boxes* are faster alternatives to using the column dials.

Scene/Object/Replace columns of X/Y/W/H – x and y positions along with width and height; 0–100% + *Replace DX* and *DY* deltas for the *Replace* column.

Draw keypoints - when enabled shows the detection points created by the feature detection and used in the *matcher* Algorithm calculation.

Draw scene/object/replace border – displays each of the borders which helps to determine correctness.

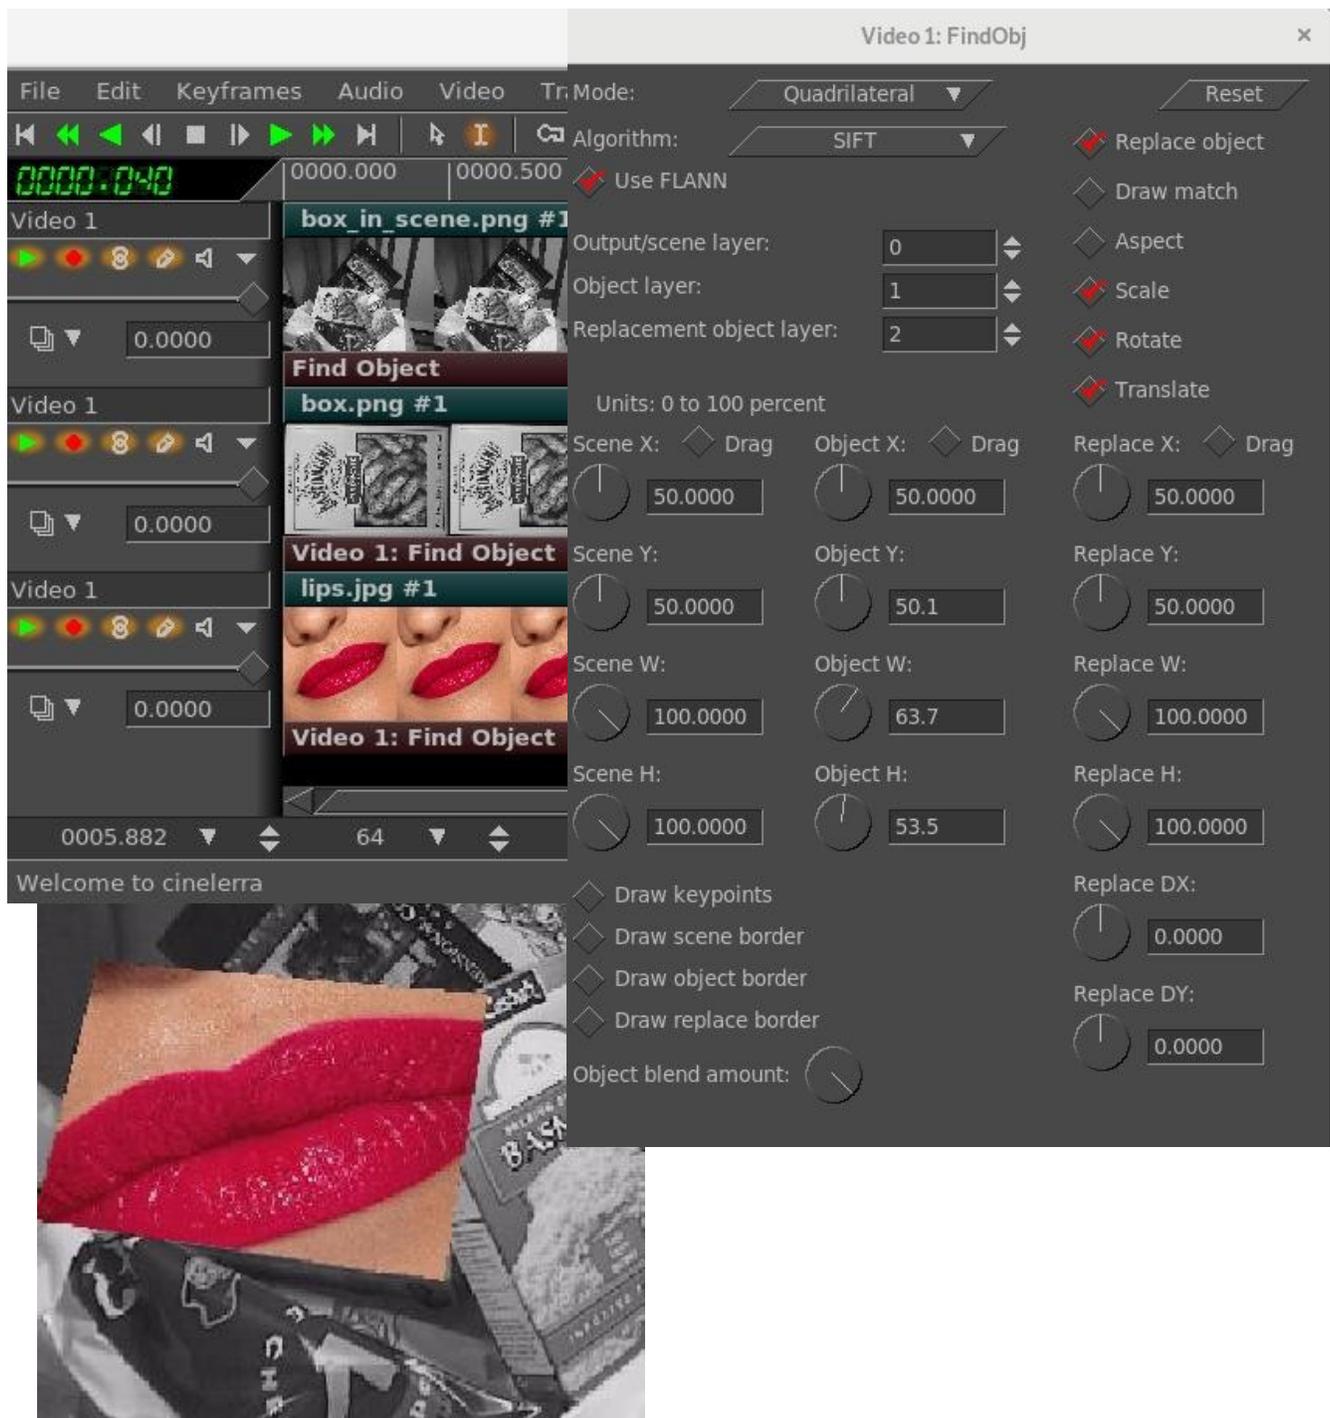
Object blend amount - can dampen (*smooth*) the corner motion from frame to frame to reduce *jitters*.

10.10.4 An example of how to use Find Object

The following steps were used to set up the example screencast shown next.

1. For best results, set *Play every frame* in **Settings** → **Preferences** → **Playback A**.
2. Load 3 tracks of *png/jpg* files – this is one of the more useful working cases:
 - i. 1st track should be the *scene*; that is the *output*
 - ii. 2nd track is the *object* to find
 - iii. 3rd track is the *replacement* object
3. Drag the *Find Object* plugin onto track #1.
4. On each of other 2 tracks, click the right mouse button; choose *attach effect*, highlight *Find Object* in the *shared effect* column, and click OK. All three tracks should now have a *findobj* plugin effect bar. One, (probably the first one) is the *master* plugin and the others are attached input tracks.
5. Check the plugin *show* icon on the *master* track to bring up the *controls* for the FindObj plugin. You will see that *Use FLANN* is checked for using *nearest neighbors*.
6. Set *Output scene*, *Object*, and *Replacement object layers*' track number accordingly (numbered from zero).
7. Check *Draw scene border* and you will see a *white* border in the *compositor* window surrounding the whole image of the scene. This assumes *default* settings for *Scene center X,Y*(at 50%), and *area W,H* (100%). Adjust these however you need to via the dials or more simply by checking *Drag* and dragging any of the 9 drag points. As shown above this in the *controls*, units are in 0-100%.
8. Turn off *Play track* in the *patchbay* on the first track so you can see track #2 in the *compositor* and then use the *Object X,Y,W,H* dials. You will see a *blue-dotted* line which should then be adjusted to surround the desired object. Since it can be a little tedious to use the dials, you will want to turn off the previous *Drag box* of track #1 and check this instead to do your dragging placement from drag points.
9. Turn off *Play track* in the *patchbay* on this track #2 so that track #3 is visible and if necessary, adjust the *dials/drag box* in the same manner. The border will have a *pink-dotted* line around it.
10. Turn back on all of the desired *Play tracks* in the *patchbays*.
11. Set the *detector* choice (should have *Don't calculate* set initially), for example use *SIFT*.
12. Turn off all of the *Draw borders* and whichever *Drag box* that may still be checked.
13. Check *Replace object* to see the replacement overlay. At this time you may still want to change the checkboxes of *Aspect*, *Rotate*, *Scale*, and *Translate* and adjust *DX/DY* in the *Replace* column.

Screencast showing the 3 tracks with *Scene*, *Object*, and *Replacement* + *Find Object* controls .

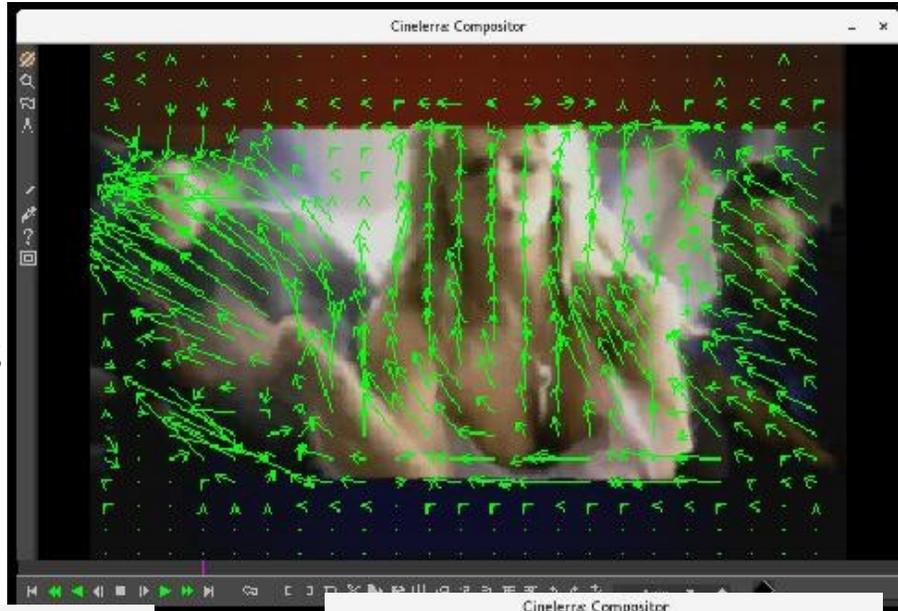


Screencast above shows the results in the *compositor* of *finding* the object on track 2 in the track 1 scene and then *replacing* it with the object in track 3 on the main track canvas.

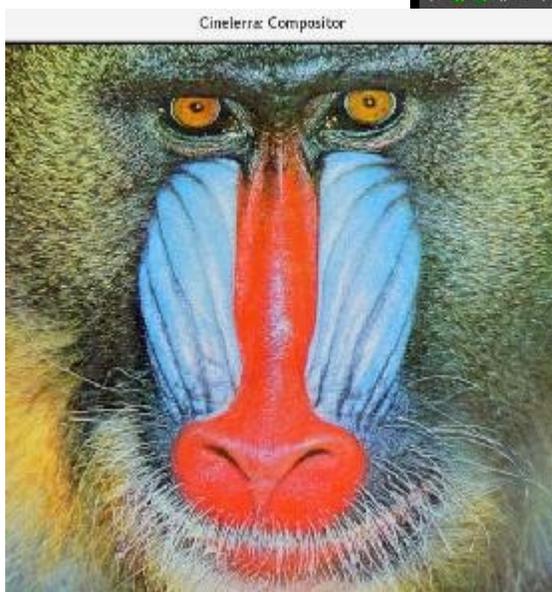
10.10.5 Additional OpenCV – MoveObj, FlowObj, GaborObj, StylizeObj, PuzzleObj

MoveObj plugin moves an object and stabilizes it.

FlowObj plugin retards image motion as shown with optical flow. You can turn on/off draw vectors. ● →

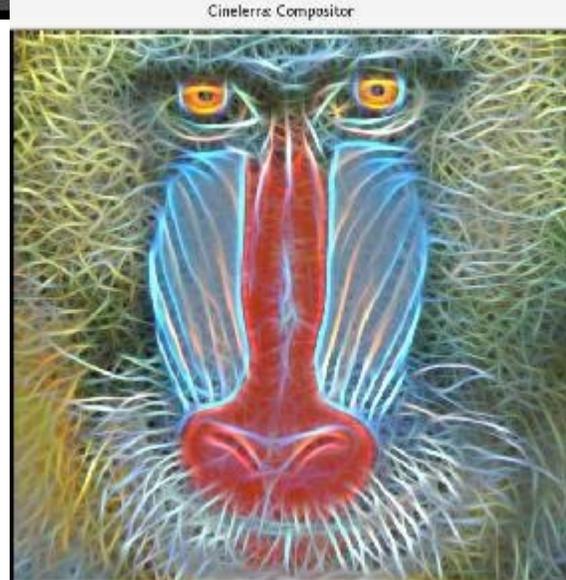


GaborObj creates an interesting fractalius-like image effect using the Gabor filter. ↓



← Before image

After →



PuzzleObj makes a puzzle out of an image. You can make the puzzle pieces smaller or larger with the *Pixels* slider bar. The *Iterations* slider bar allows for varying morphing distance. ● →



The *StylizeObj* plugin can be used to create some interesting edge effects using various options.

Edge Smooth / Edge Recursive

There are two edge preserving methods available – one to *smooth* the image *edges*, and the other to not smooth the edges/color boundaries, but instead replace the *color values* at a pixel by the average of the pixels around the area which have color similar to that pixel.

Detail Enhance

This option enhances the details of an image to make it look sharper.

Pencil Sketch / Color Sketch

Pencil like line drawings – either applied to a grayscale version of the images, which can give you thin pencil to charcoal like results, or applied to the color input image, which will look like a colored pencil drawing. For best results you can vary the Smooth-ing (the size of the surrounding area), the Edges (how dissimilar colors in the surroundings will be averaged; a larger value results in large regions of constant color) and the shading (simple scaling of the output image intensity with the higher the value, the brighter the result).

Stylization

Produces output that looks like the images were painted using water colors, which often gives the impression of cartoon-like pictures. It abstracts regions of low contrast while either preserving or enhancing features of high contrast.

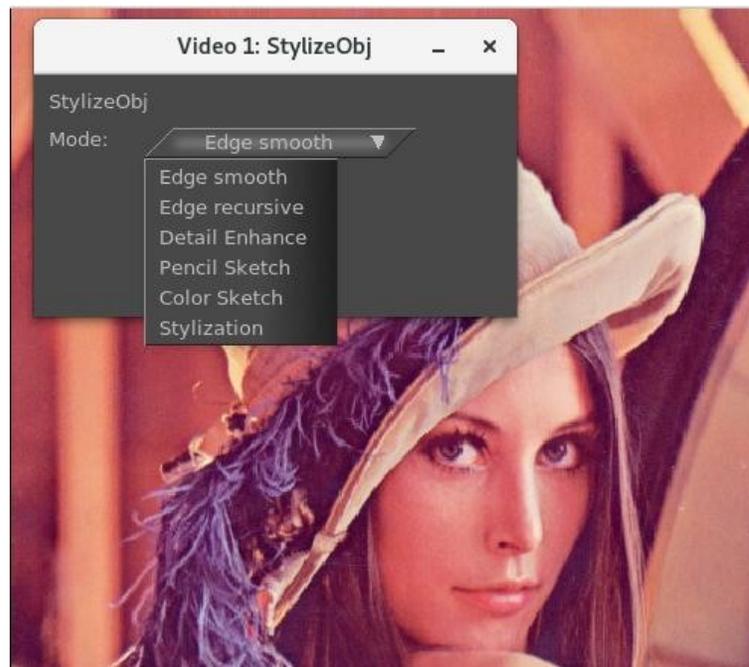


Fig. 103: Pulldown of Stylize mode

Screencast shows the original image and the StylizeObj plugin options as previously described.

The screencasts below show the images after adding the 6 various styles.



Fig. 106: Edge smooth

Fig. 104:



Fig. 105:

Detail Enhance



Fig. 108: Pen

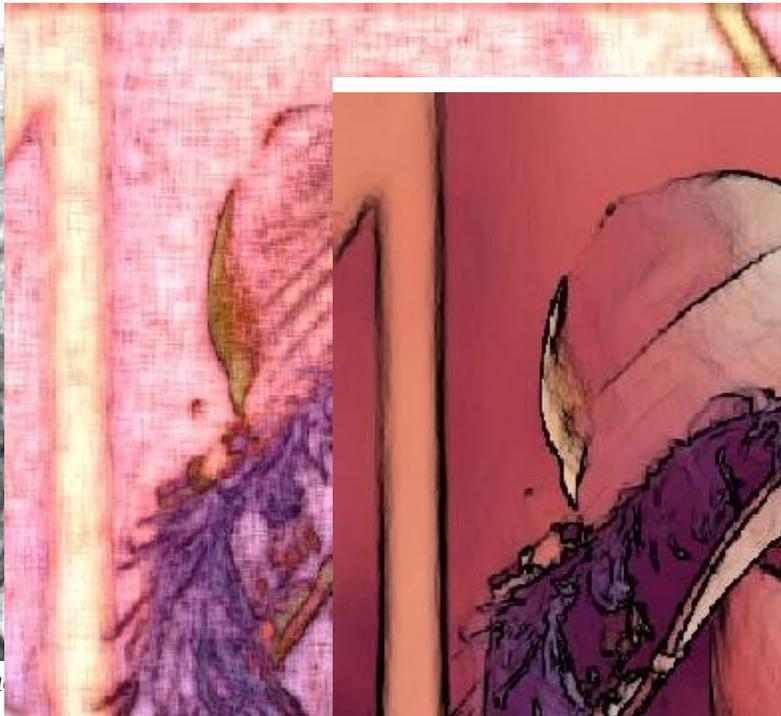


Fig. 107:



Fig. 109:

Stylization

10.11 FFmpeg Audio and Video Plugins

(credit to WPfilmmaker for the Ffmpeg info description lines taken from his contributed pdf).

Cinelerra-GG currently comes with more than 140 of the video plugins and 55 of the audio plugins developed by the FFmpeg project www.ffmpeg.org. These plugins do not have a *GUI* with buttons like the rest of plugins, therefore to change settings it is necessary to change the variables by hand by highlighting the *option*, typing a *value* in the *Range box*, and then hitting the *Apply* button. Many of these plugins provide *tooltips* at the bottom right corner of the window when the option is highlighted. A slider bar and a dial for numerical values can be used to easily vary the values which take effect immediately.

The screenshot below is what an FFmpeg video plugin looks like; example is F_chromakey.

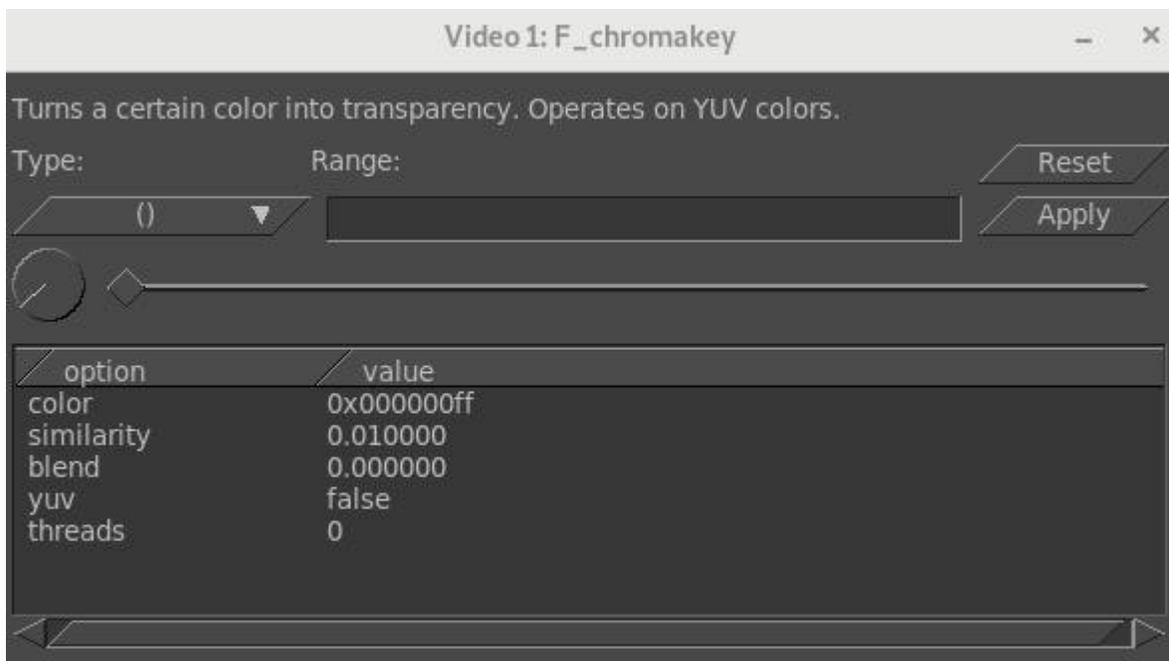


Fig. 110: GUI for FFmpeg plugins

10.11.1 FFmpeg Icons and how to Replace

Currently FFmpeg audio and video plugins that do not have a personalized icon use one of the default icons as shown below.



If you want to replace the icon with a more descriptive picture, see a previous section on *Updatable Icon Image Support* and *Details on where to put Plugin Icons*.

10.11.2 How to Use FFmpeg Audio/Video Plugins

Simply *drag and drop* the plugin on the *timeline*. To enter the settings option, once you have added the plugin to the timeline, right click on the colored bar that appears below the timeline track of the plugin and highlight the *show* option. Alternatively, you can left click the *magnifying glass* icon that appears on the right side of the colored bar. For the audio plugins, if the plugin is not working at all, you will hear a *trouble tone waveform* to indicate that. If the value is unspecified, it uses *default*. If you specify the default, the value becomes unspecified which means that nothing gets sent into the program.

Some of the ffmpeg plugins are not usable with cinelerra due to *input/output* requirements. Also, some do not come with *legal* initial supplied values for the parameters. These plugins get tested at least once and if they crash, cause problems, or are deemed unusable, they are *commented out* in the *plugin.opts* file in the cinelerra ffmpeg subdirectory. Generally they are not retested so if the ffmpeg software changes them making them usable, they will still not be accessible until the *plugin.opts* file is changed. You can easily retest these, by temporarily *uncommenting* the copy of the *plugin.opts* file in the bin subdirectory and using **Settings** → **Preferences** → **Interface** tab to *Reload plugin index*.

FFmpeg's plugin guide is at the link:

<https://ffmpeg.org/ffmpeg-filters.html>

10.11.3 FFmpeg Audio Plugins

F_abench: Benchmark part of a filtergraph.

F_acompressor: Audio compressor.

F_acontrast: Simple audio dynamic range compression/expansion filter.

F_acrusher: Reduces audio bit resolution.

F_acue: Delay filtering to match a cue.

F_adelay: Delays one or more audio channels.

F_aderivative: Compute derivative of input audio.

F_aecho: Adds echoing to the audio.

F_aemphasis: Audio emphasis.

F_aeval: Filters audio signal according to a specific expression.

F_afade: Fades in/out input audio.

F_aformat: Convert the input audio to one of the specified formats.

F_agate: Audio gate.

F_aintegral: Compute integral of input audio.

F_allpass: Applies a two-pole all-pass filter.

F_aloop: Loops audio samples.

F_anois-src: Generates a noise audio signal.

F_aperms: Set permissions for the output audio frame.

F_aphaser: Adds a phasing effect to the audio.

F_arealtime: Slows down filtering to match realtime.

F_aresample: Resamples audio data.

F_asetrates: Change the sample rate without altering the data.

F_astats: Shows time domain statistics about audio frames

F_atempo: Adjusts audio tempo.

F_atrim: Pick one continuous section from the input, drop the rest.

F_bandpass: Applies a two-pole Butterworth band-pass filter.

F_bandreject: Applies a two-pole Butterworth band-reject filter.

F_bass: Boosts or cuts lower frequencies.

F_biquad: Applies a biquad IIR filter with the given coefficients.

F_chorus: Adds a chorus effect to the audio.

F_compand: Compresses or expands audio dynamic range.

F_compensationdelay: audio compensation delay line.

F_crossfeed: Apply headphone crossfeed which blends the left and right channels of a stereo audio recording. It is mainly used to reduce extreme stereo separation of low frequencies in order to produce more speaker like sound.

F_crystalizer: Simple Expand Audio Dynamic Range filter.

F_dcshift: Applies a DC shift to the audio.

F_drmeter: Measure audio dynamic range where setting window length in seconds is used to split audio into segments of equal length.

F_dyaudnorm: Dynamic Audio Normalizer. When using this plugin, be sure to “attach effect” to all audio tracks by dragging the plugin to the 1st audio track and then right mouse clicking all subsequent audio tracks which brings up an menu. Highlight the effect shown in the middle section and click OK.

F_earwax: Widens the stereo image. When using this plugin, be sure to “attach effect” to all audio tracks by dragging the plugin to the 1st audio track and then right mouse clicking all subsequent audio tracks which brings up an menu. Highlight the effect shown in the middle section and click OK.

F_equalizer: Applies two-pole peaking equalization (EQ) filter.

F_extrastereo: Increases difference between stereo audio channels. When using this plugin, be sure to “attach effect” to all audio tracks by dragging the plugin to the 1st audio track and then right mouse clicking all subsequent audio tracks which brings up an menu. Highlight the effect shown in the middle section and click OK.

F_flanger: Applies a flanging effect to the audio.

F_haas: Apply Haas Stereo Enhancer for a more natural sounding pan effect or more clarity in the center of the mix. With this filter applied to mono signals it gives some directionality and stretches its stereo image.

F_highpass: Applies a high-pass filter with 3dB point frequency.

F_hilbert: Generate a Hilbert transform FIR coefficients.

F_loudnorm: EBU R128 loudness normalization.

F_lowpass: Applies a low-pass filter with 3dB point frequency.

F_mcompand: Multiband Compress or expand audiodynamic range. The input audio is divided into bands which is like the crossover of a loudspeaker, resulting in flat frequency response when absent compander action.

F_pan: Remix channels with coefficients (panning).

F_silenceremove: Removes silence.

F_sine: Generate sine wave audio signal.

F_stereotools: Applies various stereo tools. When using this plugin, be sure to “attach effect” to all audio tracks by dragging the plugin to the 1st audio track and then right mouse clicking all subsequent audio tracks which brings up an menu. Highlight the effect shown in the middle section and click OK.

F_stereowiden: Applies stereo widening effect. When using this plugin, be sure to “attach effect” to all audio tracks by dragging the plugin to the 1st audio track and then right mouse clicking all subsequent audio tracks which brings up an menu. Highlight the effect shown in the middle section and click OK.

F_treble: Boosts or cuts upper frequencies.

F_tremolo: Applies tremolo effect.

F_vibrato: Applies vibrato effect.

F_volume: Change input volume.

F_volumedetect: Detect audio volume.

10.11.4 FFmpeg Video Plugins

- F_amplify:** Amplify changes between successive video frames.
- F_atadenoise:** Apply an Adaptive Temporal Averaging Denoiser.
- F_avgblur:** Apply average blur filter.
- F_bbox:** Compute bounding box for each frame.
- F_bench:** Benchmarks part of a filtergraph.
- F_bitplaneoise:** Measure bit plane noise.
- F_blackdetect:** Detect video intervals that are (almost) black.
- F_blackframe:** Detect frames that are (almost) black.
- F_boxblur:** Blurs the input video. Through the settings you are able to change the power and the radius of the boxblur applied to luma, chroma and alpha.
- F_bwdif:** Deinterlaces the input image.
- F_chromakey:** Turns a certain color into transparency. Operates on YUV colors.
- F_ciescope:** Video CIE scope.
- F_color:** Provide an uniformly colored input.
- F_colorbalance:** Adjusts the color balance.
- F_colorchannelmixer:** Adjusts colors by mixing color channels.
- F_colorkey:** Turns a certain color into transparency. Operates on RGB colors.
- F_colorlevels:** Adjusts the color levels.
- F_colormatrix:** Converts color matrix.
- F_cover_rect:** Find and cover a user specified object.
- F_crop:** Crops the input video.
- F_cropdetect:** Auto-detect crop size
- F_curves:** Adjust components curves.
- F_datascope:** Video data analysis.
- F_dctdnoiz:** Denoise frames using 2D DCT.
- F_deband:** Debands video.
- F_deblock:** Deblocks video.
- F_deflate:** Applies deflate effect.
- F_deflicker:** Remove temporal frame luminance variations.
- F_dejudder:** Removes judder produced by pullup.
- F_delogo:** Removes logo from input video. When using this plugin a green box will appear on the screen, once the logo is inside the box the plugin will hide it. Through the settings you can specify the position of the logo to hide (on a X-Y axis) and the size of the box (so you can adjust it to the size of the logo).
- F_deshake:** Stabilizes shaky video.
- F_despill:** Remove unwanted foreground colors, caused by reflected color of green or blue screen.
- F_dilation:** Applies dilation effect.
- F_doubleweave:** Weave input video fields into double number of frames.
- F_drawbox:** Draws a colored box on the input video. Through the settings you are able to choose the position of the box on X/Y coordinates, the size of the box, the color and the thickness of the lines.
- F_drawgraph:** Draw a graph using input video metadata.
- F_drawgrid:** Draws a colored grid on the input video. Through the settings you can select the horizontal and the vertical offset, set the width and height of the grid cell, and the color and thickness of the lines. When using the Presets button on the plugin bar on the timeline, you can choose a preset of “rule_of_thirds”. The Rule of Thirds is a 3×3 grid on top of an image which is commonly used in filmmaking. The concept is that you align the key elements in the image using

this grid at the intersection of the lines or along and within the vertical/horizontal lines.

F_edgedetect: Detects and draws edge.

F_elbg: Apply posterize effect, using the ELBG algorithm.

F_entropy: Measure video frames entropy.

F_eq: Adjusts brightness, contrast, gamma and saturation.

F_erosion: Applies erosion effect.

F_fade: Fade in/out input video.

F_fftdnoiz: Denoise frames using 3D FFT.

F_fftfilt: Apply arbitrary expressions to pixels in frequency domain.

F_field: Extract a field from the input video.

F_fieldorder: Set the field order.

F_fillborders: Fill borders of the input video.

F_floodfill: Fill area of the same color with another color.

F_format: Convert the input video to one of the specified pixel formats.

F_framerate: Upsamples or downsamples progressive source between specified frame rates.

F_framestep: Select one frame every N frames.

F_fspp: Applies Fast Simple Post-processing filter.

F_gblur: Apply Gaussian Blur filter.

F_gradfun: Debands video quickly using gradients.

F_graphmonitor: Show various filtergraph stats.

F_greyscale: Estimates scene illumination by grey edge assumption.

F_haldclutsrc: Provide an identity Hald CLUT.

F_hflip: Horizontally flips the input video.

F_histeq: Applies global color histogram equalization.

F_histogram: Computes and draws a histogram.

F_hqdn3d: Applies a High Quality 3D Denoiser.

F_hqx: Scales the input by 2, 3 or 4 using the hq*x magnification algorithm.

F_hue: Adjust the hue and saturation of the input video.

F_idet: Interlace detect Filter.

F_il: Deinterleaves or interleaves fields.

F_inflate: Applies inflate effect.

F_interlace: Convert progressive video into interlaced.

F_kerndeint: Applies kernel deinterlacing to the input.

F_lenscorrection: Rectifies the image by correcting for lens distortion.

F_life: Generate a life pattern.

F_limiter: Limit pixels components to the specified range.

F_loop: Loops video frames.

F_lumakey: Turns a cerai luma into transparency.

F_lut: Compute and apply a lookup table to the RGB/YUV input video.

F_lut1d: Adjust colors using a 1D LUT.

F_lut3d: Apply a 3D LUT (lookup table) to an input video. LUTs are used to map one color space to another and are frequently supplied with high-end cameras as a .cube file to use as input.

F_lutrgb: Compute and apply a lookup table to the RGB input video.

F_lutyuv: Combine and apply a lookup table to the YUV input video.

F_mandelbrot: Render a Mandelbrot fractal.

F_mcdeint: Applies motion compensating deinterlacing.

F_mestimate: Generate motion vectors.

F_mpdecimate: Remove near-duplicate frames.

F_mptestsrc: Generate various test pattern.

F_negate: Negates input video.

F_nlmeans: Non-local means denoiser. Example usage is for recovery of VHS tapes which look bad.

F_noise: Adds noise to the video. Through the settings you can select the variables of the noise (strength, flag and seed).

F_normalize: Normalize RGB video.

F_oscilloscope: 2D video oscilloscope. Useful to measure spatial impulse, step responses, and chroma delays.

F_owdenoise: Denoises using wavelets.

F_pad: Add paddings to the input image, and place the original input at the provided x, y coordinates.

F_pal100bars: Generate PAL 100% color bars.

F_pal75bars: Generate PAL 75% color bars.

F_perms: Set permissions for the output video frame.

F_perspective: Corrects the perspective of video.

F_phase: Phases shift fields.

F_pixscope: Pixel data analysis for checking color and levels. It will display sample values of color channels.

F_pp: Filters video using libpostproc.

F_pp7: Applies Postprocessing 7 filter.

F_prewitt: Apply prewitt operator.

F_pseudocolor: Make pseudocolored video frames.

F_readeia608: Read EIA-608 Closed Caption codes from input video & write to frame metadata.

F_readvitic: Reads vertical interval timecode and writes it to frame metadata.

F_realtime: Slows down filtering to match realtime.

F_removegrain: Removes grain.

F_repeatfields: Hard repeat fields based on MPEG repeat field flag.

F_rgbtstsrc: Generate RGB test pattern.

F_roberts: Apply roberts cross operator which performs a simple/quick 2-D spatial gradient measurement on the video (usually a grayscale image). It highlights regions of high spatial frequency which most likely correspond to edges.

F_rotate: Rotates the input image.

F_sab: Applies shape adaptive blur.

F_scale: Scale the input video size and/or convert the image format.

F_separatefields: Split input video frames into fields.

F_setparams: Force field, or color property for the output video frame.

F_setrange: Force color range for the output video frame.

F_showpalette: Display frame palette.

F_shuffleframes: Shuffles video frames.

F_shuffleplanes: Shuffles video planes.

F_signalstats: Separates statistics from video analysis.

F_smartblur: Blurs the input video without impacting the outlines. Through the settings you can select the radius, the strength and the threshold of luma and chroma.

F_smptebars: Generate SMPTE color bars.

F_smptehdbars: Generate SMPTE HD color bars.

F_sobel: Applies sobel operator.

F_spp: Applies a simple post processing filter.

F_stereo3d: Converts video stereoscopic 3D view.

F_super2xsai: Scales the input by 2x using the Super2xSaI pixel art algorithm.

F_swaprect: Swaps 2 rectangular objects in video.

F_swapuv: Swaps U and V components.

F_tblend: Blend successive frames.
F_testsrc: Generate test pattern.
F_testsrc2: Generate another test pattern.
F_tile: Tile several successive frames together.
F_tinterlace: Performs temporal field interlacing.
F_tlut2: Compute and apply a lookup table from 2 successive frames.
F_tmix: Mix successive video frames.
F_transpose: Transposes input video.
F_unsharp: Sharpen or blur the input video.
F_uspp: Applies Ultra Simple/Slow Post-processing filter.
F_vaguedenoiser: Applies a Wavelet based Denoiser.
F_vectorscope: Video vectorscope.
F_vflip: Flips the input video vertically.
F_vfrdet: Variable frame rate detect filter.
F_vibrance: Boost or alter saturation.
F_vignette: Makes or reverses a vignette effect. Through the settings you can set the circle center position on a X-Y axis, choose the angle, the aspect and set the dithering of the vignette.
F_w3dif: Applies Martin Weston three field deinterlace.
F_waveform: Video waveform monitor.
F_weave: Weaves input video fields into frames.

F_xbr: Scales the input using xBR algorithm.
F_yadif: Deinterlaces the input image.
F_yuvtestsrc: Generate YUV test pattern.
F_zoompan: Applies Zoom & Pan effect.

10.12 Rendered Effects

(This capability is going to be deleted in the future unless receive notification of need to keep).

Besides the *Realtime* effects, as has been described in the previous sections, another type of effect is performed on a *section* of the track and the result *stored* somewhere before it is played back. The result is usually pasted into the track to replace the original data. The *rendered* effects are not listed in the resources window but instead are accessed through the **Audio** → **Render effect and Video** → **Render effect** menu options. Each of these menu options brings up a *dialog* for the rendered effect. In the *Select an effect* dialog is a list of all the realtime and all the rendered effects. The difference here is that the *realtime* effects are rendered to disk and not applied under the track. *Rendered* effects apply to only one type of track, either *audio* or *video*. If no tracks of the type exist, an *error* pops up.

A *region* of the timeline to apply the effect to must be defined before selecting *Render* effect. If no *in/out points* and no *highlighted* region exist, the entire region after the *insertion point* is treated as the affected region. Otherwise, the region between the *in/out points* or the *highlighted* region is the affected region. The tracks to apply the rendered affect to, need to be *armed*, other tracks are ignored. The rendered effect processes certain track *attributes* when it reads its input data but not others. *Transitions* in the affected track are applied, but *nudge* and *effects* are not. This allows the new data to be pasted into the existing position without changing the *nudge* value.

How to use a rendered effect:

- Highlight an effect in the list (*Select an effect*) to designate it as the one being used.
- Define a file to render the effect to in the *Select a file to render to* box. The *magnifying glass* allows file selection from a list.
- Select a file *format* which can handle the track type. The *wrench* allows configuration specific to the *file format*.
- There is also an option for *creating a new file at each label*. For example, if you have a *CD rip* on the timeline which you want to divide into different files, the *labels* would become dividing points between the files if this option were selected. When the timeline is divided by labels, the effect is *re-initialized* at every *label*. *Normalize* operations take the peak in the current file and not in the entire timeline.
- Finally there is an *insertion strategy* just like in the render dialog. It should be noted that even though the effect applies only to *audio* or *video*, the insertion strategy applies to all tracks just like a *clipboard* operation.
- When you click *OK* in the effect dialog, it calls the *GUI* of the effect. If the effect is also a *realtime* effect, a second *GUI* appears to prompt for *acceptance* or *rejection* of the current settings.
- After accepting the settings, the effect is processed.

10.12.1 Rendered Audio Effects

Resample

This multiplies the number of each *output sample* by a *scale factor* to arrive at the number of the *input sample*. The output file's sample rate is set to the *project* sample rate but its *length* is changed to reflect the scaled number of samples. It also *filters* the resampled audio to remove aliasing.

If the scale factor is 2, every 2 input samples will be *reduced* to 1 output sample and the output file will have half as many samples as the input sequence. If it is 0.5, every 0.5 input samples will be *stretched* to 1 output sample and the output file will have twice as many samples as the input sequence.

Normalize

Time Stretch

10.12.2 Rendered Video Effects

Reframe

This does exactly the same thing as *ReframeRT* in *Stretch* mode. It multiplies the *output* frame number by the *scale factor* to arrive at the *input* frame number and changes the *length* of the sequence. Unlike *ReframeRT*, this must run from the *Video menu* and *render* its output. Be aware *Reframe* does not write the scaled frame rate as the frame rate of the rendered file. It produces a file of scaled length and *equal* frame rate as the *project*. The new length is $1/\text{scale factor}$ as big as the original sequence.

To create a *slow-motion* of *fast moving* video:

- Select the video clip you wish to re-frame and put it on a video track.
- Select the area you wish to reframe.
- From the Video menu, select the *Render Effect* option.
- From the effect list, select *Reframe*.
- Enter the output *format* and *insertion strategy* for the new clip to be created.
- Press OK.
- At the *popup* menu, enter the *scale factor* 2 to run twice as fast, and 0.5 to run at half speed.

CD Ripper

720 to 480

This page intentionally left blank so that chapters begin on a right facing page.

11 - Transition Plugins

When one *edit* ends and another edit begins, the default behavior is to have the first edit's *output* immediately become the *output* of the second edit when played back. *Transitions* are a way for the first edit's output to become the second edit's output with different *variations*. The audio and video transitions are listed in the Resources window as shown below.

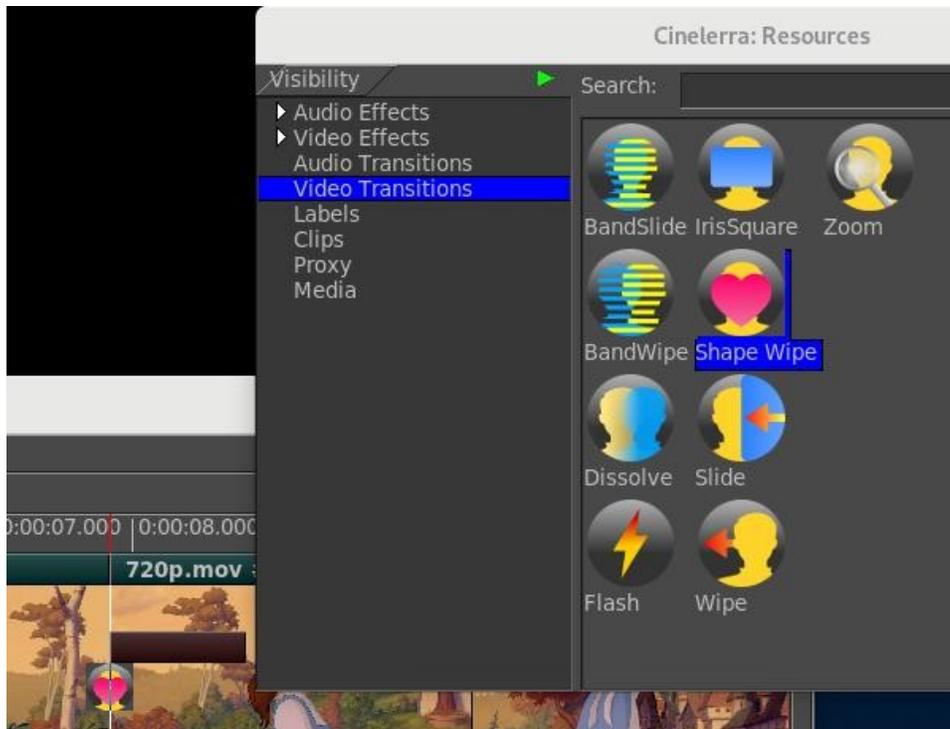


Fig. 111: Resources window displaying the Video Transitions

Note the *colored bar* ^ above the *Shape Wipe* transition.

Transitions may only apply to the *matching track type*. *Transitions* under *audio transitions* can only apply to audio tracks. *Transitions* under *video transitions* can only apply to video tracks.

An example usage of a transition follows:

1. Load a single video file and cut away a section from the center so that you make two *edits* out of a single file. Make sure the *edit boundary* between the two edits is visible on the *timeline*.
2. Go to the *Resources* window and click on the *Video transitions* folder. Drag a *transition* from the transition list onto the *second video edit* on the timeline. A *colored box* highlights over where the transition will appear. Releasing over the 2nd edit applies the transition between the 1st and 2nd edit.

Once the transition is in place, it can be *edited* similarly to a *plugin*. Move the pointer over the transition and right click to bring up the *transition menu*. The *show* option brings up specific

parameters for the transition in question if any. The *length* option adjusts the length of the transition in seconds. The *detach* option removes the transition from the timeline. If the *insertion point* or the *In point* is over an edit, the *beginning* of the edit is covered by the transition.

Dragging and dropping transitions from the *Resource* window to the *Program* window can be tedious so there are shortcuts to solve this issue. Once you drag a transition from the *Resources* window, the **U** and **u** keys will *paste* the same transition. The **U** key pastes the *last video transition* and the **u** key pastes the *last audio transition* on all the recordable tracks. Another easy way to add the same transition to multiple edits is to get into *Arrow mode* (*Drag and Drop* editing), select the *edits* you would like to add the transition to and use the *Video* or *Audio* pulldown to *Attach transition*. Choose which transition you would like and click the checkmark OK. You can also set your *default* transition at any time by doing so in the *Attach transition* popup box – highlight your choice, and then at the bottom, click on the button *Set Default Transition*. You will see that name appear.

Transitions make two edits *overlap* for a certain amount of *time*. Cinelerra does not move edits during transitions. Instead it uses *spare frames* from the *source* file to lengthen the first edit enough to make it *overlap* the second edit for the duration of the transition. The exact point in time when the transition takes effect is the *beginning* of the second edit. The transition lasts a set amount of time into the second edit. For example, if you set a duration of 1 second for a *dissolve* transition, it will not start at the last 0.5 second of the first edit and continue 0.5 second into the second edit. In fact, it will start exactly at the beginning of the second edit and last for 1 second into that second edit. On the timeline a *colored bar* over the transition symbol visually represents the position and the duration of the transition. The most important consequence of this behavior is that the first asset needs to have enough *spare* data after the *end boundary* to *fill* the transition into the second edit. Spare data duration should be equal or greater than the length of the transition effect set in the *Length* parameter of the transition popup menu. If the last frame shown on the timeline is the *last frame* of the source file, Cinelerra will lengthen the first edit using the last frame only, with the unpleasant result of having the first edit *freezing* into the transition.

It should be noted that when playing transitions from the timeline to a hardware accelerated video device, the hardware acceleration will usually be *turned off* momentarily during the transition and *on* after the transition, in order to render the transition. Using an un-accelerated video device for the entire timeline normally removes the disturbance.

11.1 Audio Transitions

Crossfade

Creates a smooth transition from one audio source edit to another. The crossfade has the first source fade out while the second fades in.

11.2 Video Transitions

In order to use a transition that you have dragged to the timeline, first right mouse click on the transition icon in the timeline. A menu will pop-up with the following controls:

Show - Pops up the transition specific menu if available (not available on the *dissolve* transition).

On - Toggle on/off the transition effect.

Transition length - Set the length in seconds, frames, samples, H:M:S:frm or H:M:S.xxx for the transition to complete. In addition you can use the mouse wheel to change the length in real time.

Detach - Remove the transition from the timeline.

BandSlide

Bands slide across video and you see the image slide.

BandWipe

Bands wipe across the video and you see the mask slides.

Flash

The video flashes when transitioning between segments.

IrisSquare

Video switches segments via a small rectangular view that gradually grows to full size.

Shape Wipe

Wipe a specific *shape* across the video. Currently available shapes are: *burst*, *circle*, *clock*, *heart*, *specks*, *spiral*, *tile2x2h*, and *tile 2x2v*.

You can add your own images to the *Shape Wipe* transition and there are some free ones available to download such as at assistcg.com. To include new images in the *Shape Wipe* Transition, simply copy the *{shape}.png* file to your location of cinelerra in the subdirectory **plugins/shapes**.

Slide

Image slides into view; you can set Left/Right/In/Out.

Wipe

Wipe the image across screen starting left or right.

Zoom

Zoom out video at X/Y magnification for some seconds.

This page intentionally left blank so that chapters begin on a right facing page.

12 - Overlay Modes, Alpha Blending, and PorterDuff

The purpose of the Overlay Modes is to control the foreground and background stacking and use blending to reshape image object boundaries. It normally makes use of a binary type alpha blending system for all in or all out. To use the available operations in Cinelerra, follow these steps:

- 1) In the main window, look at the Patchbay on the far left.
- 2) Click on the small arrow that points to the right, inside the panel to see it expand.
- 3) A down arrow shows up with a tooltip of “overlay mode”.
- 4) Use the down arrow to get to the popup menu and choose the desired effect from the 30 possibilities. “Normal” is the default.

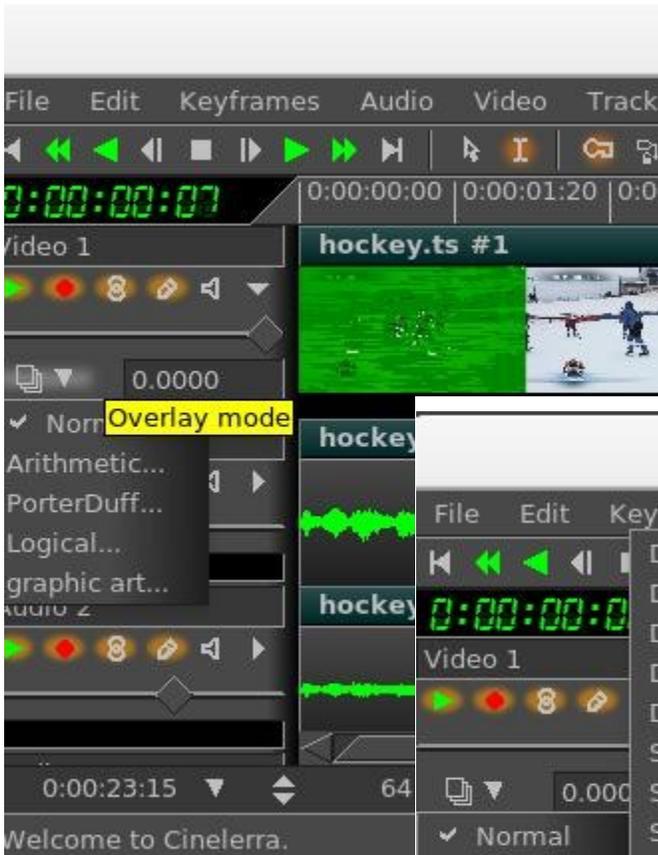


Figure 112: Patchbay pulldown

Screenshot on the left shows the pulldown in the patchbay and the tool tip “Overlay mode” which provides access to 4 types of overlays. Each will expand further as shown in the screenshots below for “PorterDuffs” and “graphic art”.

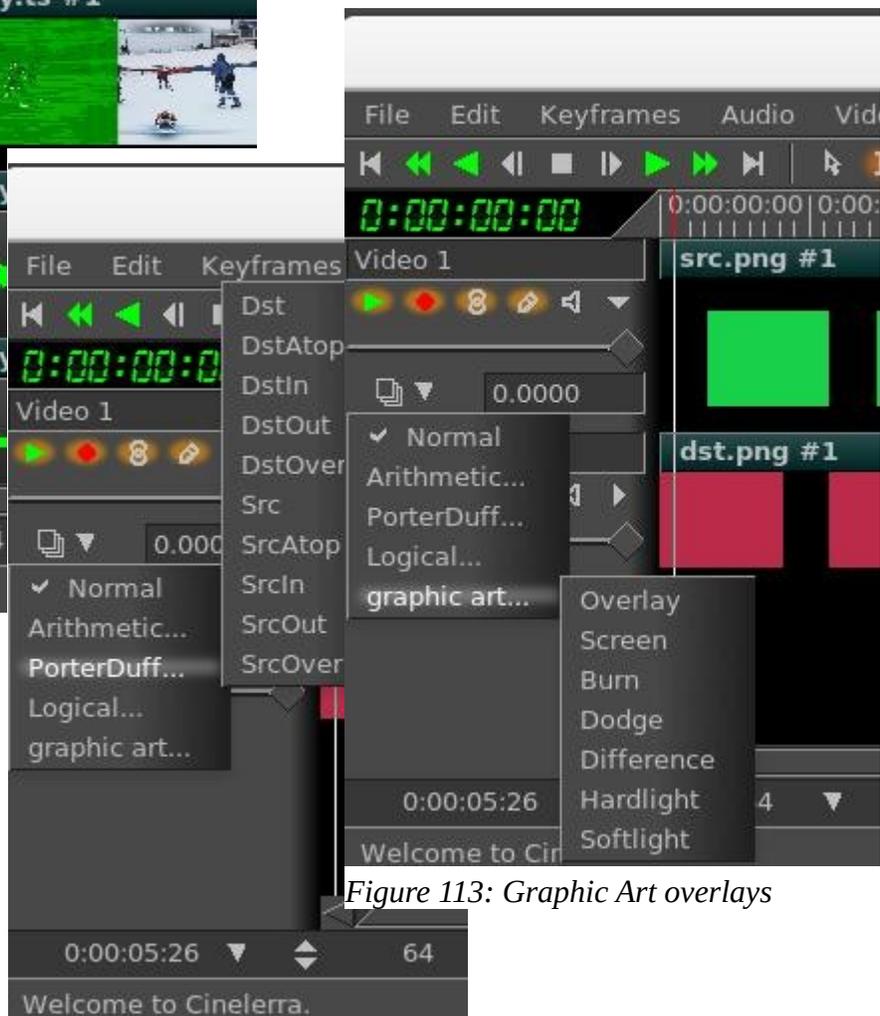


Figure 113: Graphic Art overlays

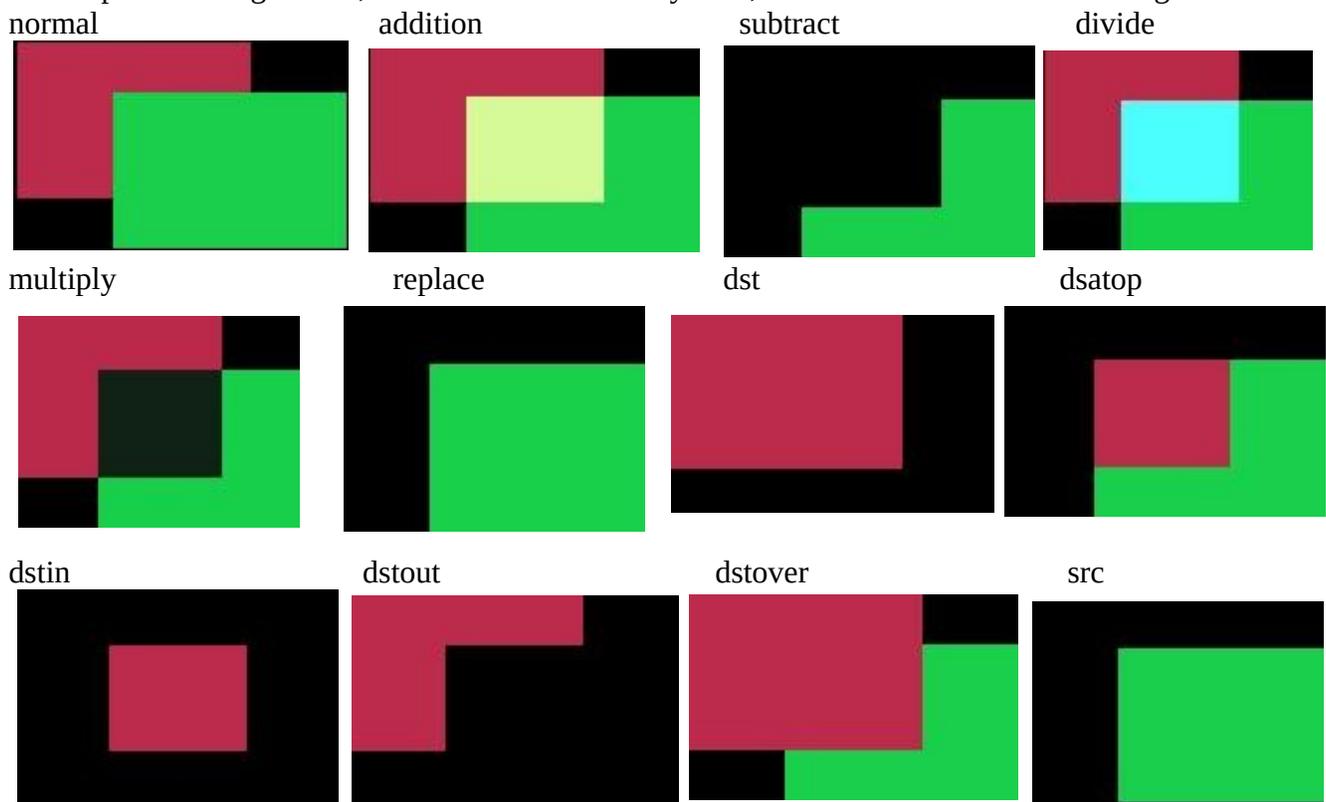
Figure 114: Porter Duff Overlays

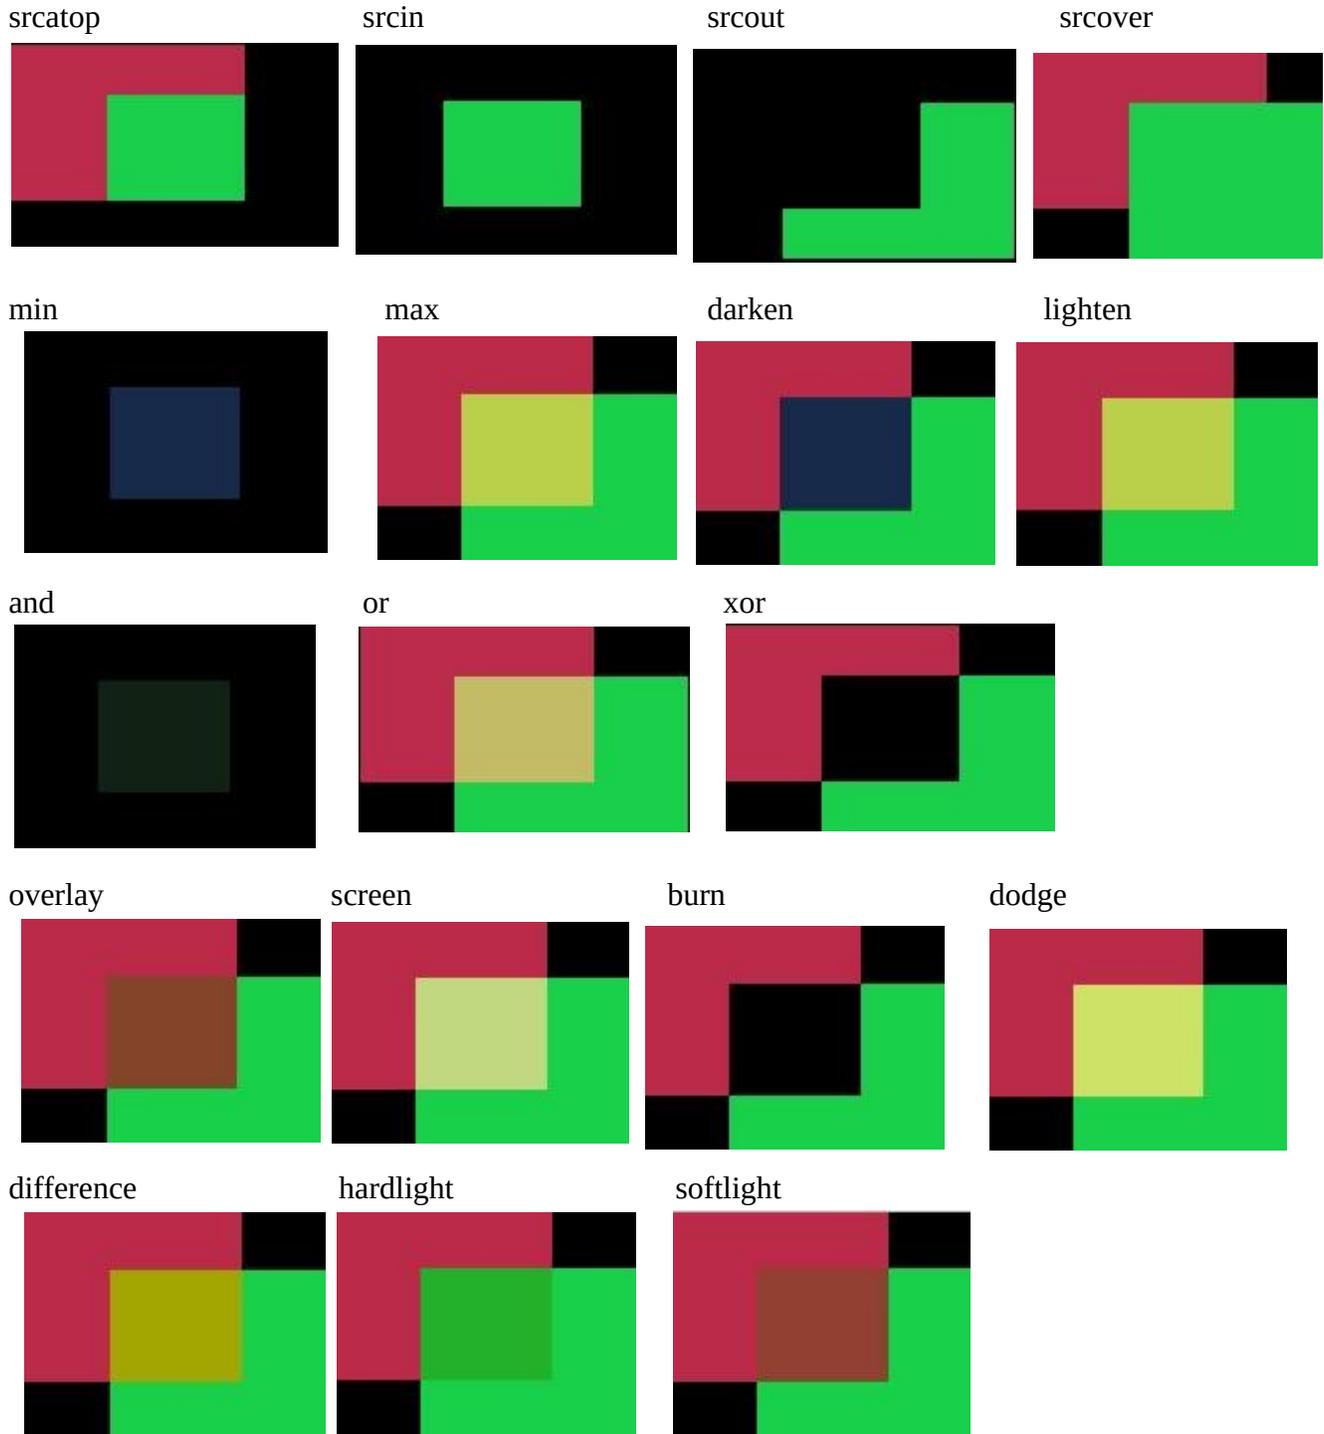
Porter-Duff is the industry standard for alpha blending operations. Only a short explanation follows here, but there is much more information to be found on the internet with complete descriptions and examples. Every pixel has 3 color channels (like RGB), and may have 1 alpha channel value. If there is no alpha defined for a color model, the alpha value is assumed to be 1. Regions of the image are created with the alpha image map. These regions are manipulated using the blending operations described below. Alpha blending is the process of combining a foreground color with a background color which produces a new blended color. The “alpha” channel describes how much opacity is present in a pixel. It may be completely transparent, completely opaque, or any range of translucency.

Conceptually, when the foreground color is completely opaque, the resulting blended color will be the foreground color. If it is transparent, the blended color will be the color of the background. When the value of the alpha channel is 1, the image is all there, if it is 0, there is no image at all, otherwise it is only partially there. In other words, the alpha value goes from 0 to 1, where full transparency is 0 and opaque is represented by 1. Alpha blending models opacity.

When blending source and destination shapes, the shape boundaries can be changed with the alpha blending effects. There are a total of 10 standard Porter-Duff operators, but there are 30 possible overlay modes used in Cinelerra-GG. Each is characterized by its value in the four regions: source, destination and both, with the “neither” region always being blank. The source and destination regions can either be blank or filled with the source or destination colors. A specific compositing math formula is used to calculate effect. This is only applicable to RGB; some effort has been made to accommodate YUV, but the effects are not as predictable, and may not be useful.

Below are the results of utilizing the 30 available operations within Cinelerra as listed on a following page. Src is the solid green rectangle and Dst is the solid red rectangle. There are better illustrations of what alpha blending can do, however for consistency sake, these are the results when using standards.





The implementation math forms are listed below, where D=Destination, S=Source, a=alpha, c=chroma (color). Each line describes a pair with the left one for alpha and the right one for chroma.

NORMAL $[S_a + D_a * (1 - S_a), S_c * S_a + D_c * (1 - S_a)]$

Arithmetic

ADDITION $[(S_a + D_a), (S_c + D_c)]$

SUBTRACT $[(S_a - D_a), (S_c - D_c)]$

MULTIPLY $[S_a + D_a - S_a * D_a, S_c * (1 - D_a) + D_c * (1 - S_a) + S_c * D_c]$

DIVIDE $[Sa + Da - Sa*Da, Sc*(1 - Da) + Dc*(1 - Sa) + Sc / Dc]$
 REPLACE $[Sa, Sc]$ (fade = 1)

PorterDuff

DST $[Da, Dc]$
 DST_ATOP $[Sa, Sc * (1 - Da) + Dc * Sa]$
 DST_IN $[Da * Sa, Dc * Sa]$
 DST_OUT $[Da * (1 - Sa), Dc * (1 - Sa)]$
 DST_OVER $[Sa + Da - Sa*Da, Sc * (1 - Da) + Dc]$
 SRC $[Sa, Sc]$
 SRC_ATOP $[Da, Sc * Da + Dc * (1 - Sa)]$
 SRC_IN $[Sa * Da, Sc * Da]$
 SRC_OUT $[Sa * (1 - Da), Sc * (1 - Da)]$
 SRC_OVER $[Sa + Da - Sa*Da, Sc + (1 - Sa) * Dc]$

Logical

MIN $[min(Sa, Da), MIN(Sc, Dc)]$
 MAX $[max(Sa, Da), MAX(Sc, Dc)]$
 LIGHTEN $[Sa + Da - Sa*Da, Sc*(1 - Da) + Dc*(1 - Sa) + max(Sc*Da, Dc*Sa)]$
 DARKEN $[Sa + Da - Sa*Da, Sc*(1 - Da) + Dc*(1 - Sa) + min(Sc*Da, Dc*Sa)]$
 AND $[Sa * Da, Sc * Dc]$
 OR $[Sa + Da - Sa * Da, Sc + Dc - Sc * Dc]$
 XOR $[Sa + Da - 2 * Sa * Da, Sc * (1 - Da) + Dc * (1 - Sa)]$

Graphic Art

OVERLAY $[Sa + Da - Sa * Da, Sc*(1 - Da) + Dc*(1 - Sa) + 2*Dc < Da ? 2*Sc*Dc : Sa*Da - 2*(Da-Dc)*(Sa-Sc)]$
 SCREEN $[Sa + Da - Sa * Da, Sc + Dc - (Sc * Dc)]$ (same as OR)
 BURN $[Sa + Da - Sa * Da, Sc*(1 - Da) + Dc*(1 - Sa) + Sc <= 0 || Sc*Da + Dc*Sa <= Sa*Da ? 0 : (Sc*Da + Dc*Sa - Sa*Da)*Sa/Sc]$
 DODGE $[Sa + Da - Sa * Da, Sc*(1 - Da) + Dc*(1 - Sa) + Sa <= Sc || Sc*Da + Dc*Sa >= Sa*Da) ? Sa*Da : Dc*Sa / (1 - Sc/Sa)]$
 DIFFERENCE $[Sa + Da - Sa * Da, Sc*(1 - Da) + Dc*(1 - Sa) + abs(Sc * Da - Dc * Sa)]$
 HARDLIGHT $[Sa + Da - Sa * Da, Sc*(1 - Da) + Dc*(1 - Sa) + 2*Sc < Sa ? 2*Sc*Dc : Sa*Da - 2*(Da - Dc)*(Sa - Sc)]$
 SOFTLIGHT $[Sa + Da - Sa * Da, Sc*(1 - Da) + Dc*(1 - Sa) + Da > 0 ? (Dc*Sa + 2*Sc*(Da - Dc))/Da : 0]$

The previous math forms are the only truly accurate description of each blending operation, but short descriptions are below where “Source” is the output from the next track and “Destination” is the output from the lower track stacking. Blending starts with new Source and combines it with the previous render stack output, which is referred to as Destination. The new output becomes the next Destination and the next up stack level becomes the new Source. Source is above; Destination is below.

Normal – Normal mode is the default layer mode. The result color is the source color. The layer on top covers the layers below it. If you want to see anything below the top layer when you use this mode, the layer must have some transparent areas. It is “stacked on top”. Math formula used is different than that used by Gimp; there is no SVG equivalent.

Arithmetic Group – standard numerical operations:

Addition – The source is added to the destination and replaces the destination. Addition mode is very simple - the pixel values of the upper and lower layers are added to each other. The resulting image is

normally brighter. The equation can result in color values greater than 255, so some of the light colors may be clipped to the maximum value of 255. Math formula is the same as that used by SVG but different than Gimp.

Subtract – Subtract mode reduces the pixel values of the upper layer by the pixel values of the lower layer. The resulting image is normally darker. You might get a lot of black or near-black in the resulting image. The equation can result in negative color values, so some of the dark colors may be clipped to the minimum value of 0. Math formula used originated from HV cinelerra and is different than that used by Gimp; there is no SVG equivalent.

Multiply – The source color is multiplied by the destination color and replaces the destination. The resulting color is always at least as dark as either the source or destination color. Multiplying any color with black results in black. Multiplying any color with white preserves the original color. Math formula is the same as used by SVG and Gimp.

Divide – Divides source color by destination color. If the source color is white, the result color is the underlying color. The resulting image is often lighter. Math formula used is different than that used by HV/CV and Gimp; there is no SGV equivalent.

Replace – Replace mode will cause any existing destination to be replaced by the source media. Mathematical formula used is the same as used by Gimp; there is no SVG equivalent.

PorterDuff Group – industry standard compositing operators:

DST - The destination is left untouched. Only the destination will be present in the output. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

DST_ATOP - The part of the destination lying inside of the source is composited over the source. The destination outside the source is dropped. You will see the existing canvas is only kept where the shapes overlap. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

DST_IN - The part of the destination lying inside of the source is displayed. This effectively allows for cutting out the background using the current layer/feature as a mask. You will see the new shape is drawn only where both the source and the destination overlap. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

DST_OUT - The part of the destination lying outside of the source is displayed. This is basically a reverse mask, compared to destination-in. You will see the existing content is kept where it doesn't overlap the source. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

DST_OVER - The destination is composited over the source. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

SRC - The source is copied. The destination is not used as input. Only the source will be present in the output. Math formula is the same as that used by Porter-Duff & SVG; there is no Gimp equivalent.

SRC_ATOP - The part of the source lying inside of the destination is composited over the destination.

The source outside the destination is dropped. You will see the existing canvas is only kept where the shapes overlap. Math formula is the same as that used by Porter-Duff and SVG; no Gimp equivalent.

SRC_IN - The part of the source lying inside of the destination is displayed. This effectively allows the destination to act as a mask for the layer/feature being shown. You will see the new shape is drawn only where both the source and the destination canvas overlap. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

SRC_OUT - The part of the source lying outside of the destination is displayed. This is basically a reverse mask, compared to source-in. You will see the new shape is drawn where it doesn't overlap the destination. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

SRC_OVER - The source is composited over the destination. Math formula is the same as that used by Porter-Duff and SVG; there is no Gimp equivalent.

Logical Group – Venn diagram operations of area combinations:

Min – The output color is the component-wise minimum value of the source and destination colors. There is no SVG or Gimp equivalent math formula.

Max – The output color is the component-wise maximum value of the source and destination colors. There is no SVG or Gimp equivalent math formula.

Lighten – Selects the lighter of the destination and source colors. The destination is replaced with the source when the source is lighter, otherwise it is left unchanged. Completely black layers have no effect on the final image and completely white layers result in a white image. Math formula is the same as used by SVG and Gimp's “lighten only”.

Darken – Selects the darker of the destination and source colors. The destination is replaced with the source when the source is darker, otherwise it is left unchanged. Completely white layers have no effect on the final image and completely black layers result in a black image. Math formula is the same as used by SVG and Gimp's “darken only”.

And – This operation intersects source and destination. If either image is not white, it down factors the other image. Usually you will end up with a lot of darker areas. There is no SVG or Gimp equivalent math formula.

Or – This operation is the union of the source and the destination. If either image is not black, it adds to the other image. Usually you will end up with a lot of lighter areas. Math formula is the same as that used by “Screen” operation in SVG; there is no Gimp equivalent.

Xor – The union of the source and destination with the intersection removed. Results in black if they are equal. Shapes are made transparent where both overlap, and drawn normal everywhere else. Math formula used is that used by Porter-Duff; there is no Gimp equivalent and SVG formula differs.

Graphic Art – typical operations from popular "paint" packages:

Overlay – Multiplies or screens the colors, dependent on the destination color. Source color overlay the destination while preserving its highlights and shadows. The destination color is not replaced, but is

mixed with the source color to reflect the lightness or darkness of the destination. Dark parts on the base layer become darker, and light parts become lighter. Math formula is same as SVG and Gimp.

Screen – The source and destination are complemented and then multiplied and then replace the destination. The resultant color is always at least as light as either of the two constituent colors. Screening any color with white produces white. Screening any color with black does not change the other layer. Using the Screen mode is similar to projecting multiple photographic slides simultaneously onto a single screen. A lighter picture is the result. Math formula is the same as used by SVG and Gimp. This is identical to “Or”.

Burn - Darkens the destination color to reflect the source color. Burning with white produces no change. In photography, burning is a technique used in a darkroom to increase the exposure in particular areas of the image. This brings out details in the highlights. Math formula is the same as used by SVG and Gimp.

Dodge - Brightens the destination color to reflect the source color. Dodging with black produces no change. The result color is a lightening of the source color to reflect the underlying layer color by decreasing the contrast. If the source color is pure black, the result color is the underlying color. In photography, dodging is a technique used in a darkroom to decrease the exposure in particular areas of the image. This brings out details in the shadows. Math formula is the same as used by SVG and Gimp.

Difference - Subtracts the darker of the two constituent colors from the lighter. This is the same as “Subtract” except that the absolute value is used. “Subtract” colors which would be black due to clipping become visible instead. Differencing with white inverts the destination color. Differencing with black produces no change. Math formula is the same as used by SVG and Gimp.

Hardlight - Multiplies or screens the colors, dependent on the source color value. If the source color is lighter than 0.5, the destination is lightened as if it were screened. If the source color is darker than 0.5, the destination is darkened, as if it were multiplied. The degree of lightening or darkening is proportional to the difference between the source color and 0.5. If it is equal to 0.5 the destination is unchanged. Using pure black or white produces black or white. The effect is similar to shining a harsh spotlight on the destination. This mode is useful for creating the appearance of shadows on a layer. You might use this mode to combine two photographs and obtain bright colors and sharp edges. Math formula is the same as used by SVG and Gimp.

Softlight - Darkens or lightens the colors, dependent on the source color value. If the source color is lighter than 0.5, the destination is lightened. If the source color is darker than 0.5, the destination is darkened, as if it were burned in. The degree of darkening or lightening is proportional to the difference between the source color and 0.5. If it is equal to 0.5, the destination is unchanged. Using pure black or white produces a distinctly darker or lighter area, but does not result in pure black or white. The effect is similar to shining a diffused spotlight on the destination. A layer with pure black or white becomes markedly darker or lighter, but does not become pure black or white. Soft light is not related to “Hard light” in anything but the name, but it does tend to make the edges softer and the colors not so bright. Math formula is the same as used by Gimp; SVG formula differs.

Note: the Graphic Art group operates principally on color, and the others operate principally on alpha.

This page intentionally left blank so that chapters begin on a right facing page.

13 - Capturing and Recording Media + DVD/Bluray Creation

This section covers the areas of capturing media from the web or television by the use of recording. It also covers how to save your media for long time storage on DVD or Bluray discs.

Access the Record function via File → Record... The recording application really only does one thing, capture media from some external source, and write it onto a file in a specified way.

Path:	output media file path	
Start time:	weekday / time of day	to begin capture
Duration:	hrs:mins:secs	until end of capture
Source:	channel/device input	selection of input subsorce
Mode:	timed	use start time / duration
	untimed	use transport controls

The media file will be written using the format and codec specified in the Settings → Preferences → Record tab, which you will need to set up first. See the Settings/Preferences section for parameters. Only ffmpeg can record both audio and video simultaneously, and some ffmpeg formats require too much cpu to do a realtime compression. A setting which is more likely to be usable (requires less cpu/memory for realtime encoding) is ffmpeg /qt/mp4.qt with bitrates like audio/256000 & video/6000000.

13.1 Record Web Media by real-time Audio Capture/Screen Capture

Below describes the necessary steps for recording freely available media from the internet for your own personal use. You have to be on a system using pulseaudio, such as ubuntu, fedora, centos.

- 1) Start cinelerra and select Settings->Preferences->Recording
From a terminal (with a wide text window) run: `pactl list`.
You will see all of the audio sources and sinks on your system.
Identify the source associated with the normal output your system uses; example: Source #1.
Locate the source which monitors your normal audio output. For example: front stereo.
- 2) Choose a File Format and File Type (for example: FFMPEG & qt).
Check Record audio tracks.
Click Audio wrench tool:
Select mp4, for example.
Set Bitrate (256000 is a reasonable number).
Check Record video tracks.
Click Video wrench tool:
Select mp4 video, for example.
Set Bitrate 1000000 - 4000000 (bigger=better, but more cpu)
Select Audio In->Record Driver: ALSA (or whatever you have instead)
Set the Bits: 16 Bit Linear (probably, because reasonable and fast)
Select Device: the source from above pactl list search, for example Source #1)

Samples read from device: aprox dev buffer size (2k-16k probably)

Samples to write to disk: 131072 (a good size)

Sample rate for recording: 44100 (will automatically change)

Channels to record 2 (probably)

Uncheck Map 5.1->2

Gain 1.0

Select Video In->Record Driver: Screencapture

Set the Display: leave blank (probably) or use ":0.0" (default screen/display)

Frames to record to disk at a time: 30 (a good number)

Frames to buffer in device: (2-6) (probably)

Positioning: Software Timing (important)

Uncheck: Sync drives automatically

Size of captured frame: 720x480 (this is your choice, actual capture size; 600x320 youtube)

This defines a "screen capture" rectangular area on the display.

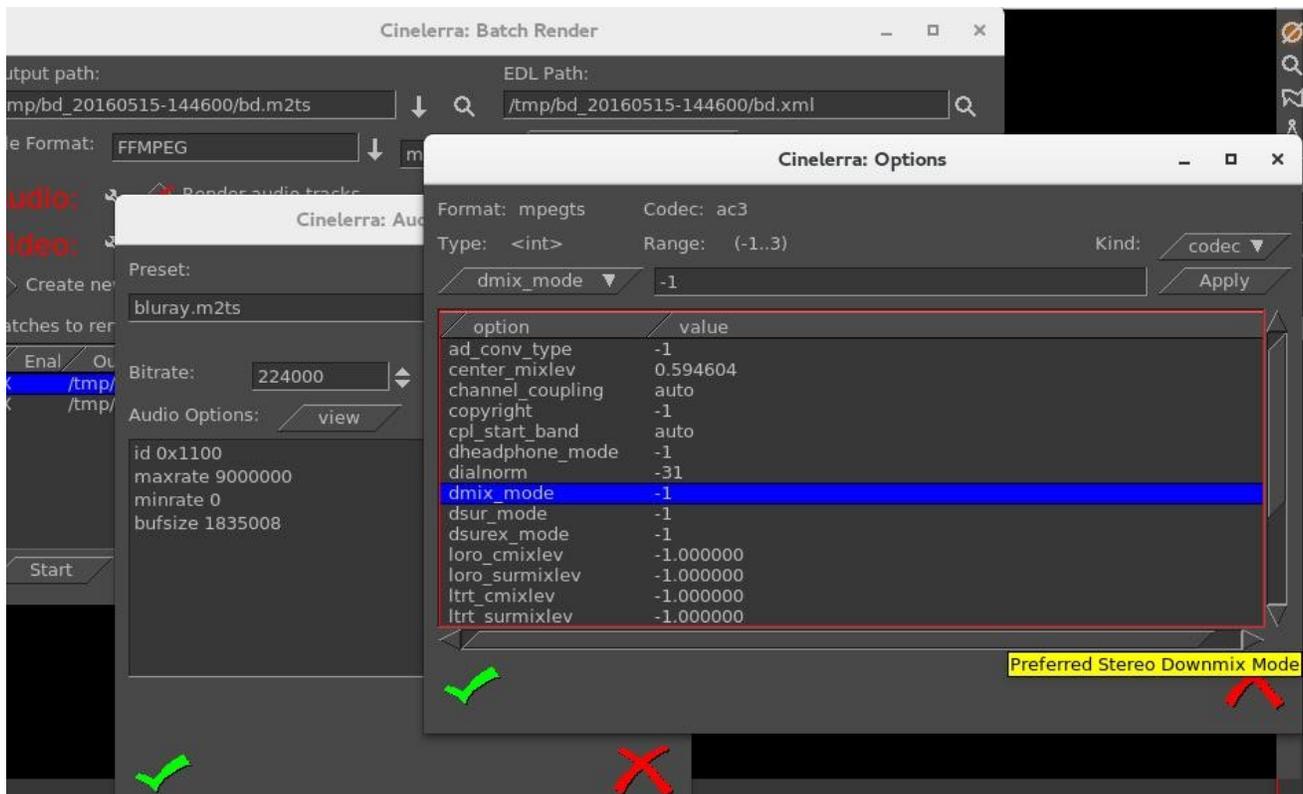
Try to make it the size you need to cover the screen playback area.

Frame rate for recording 23.97 (a good choice, not all choices work)

OK

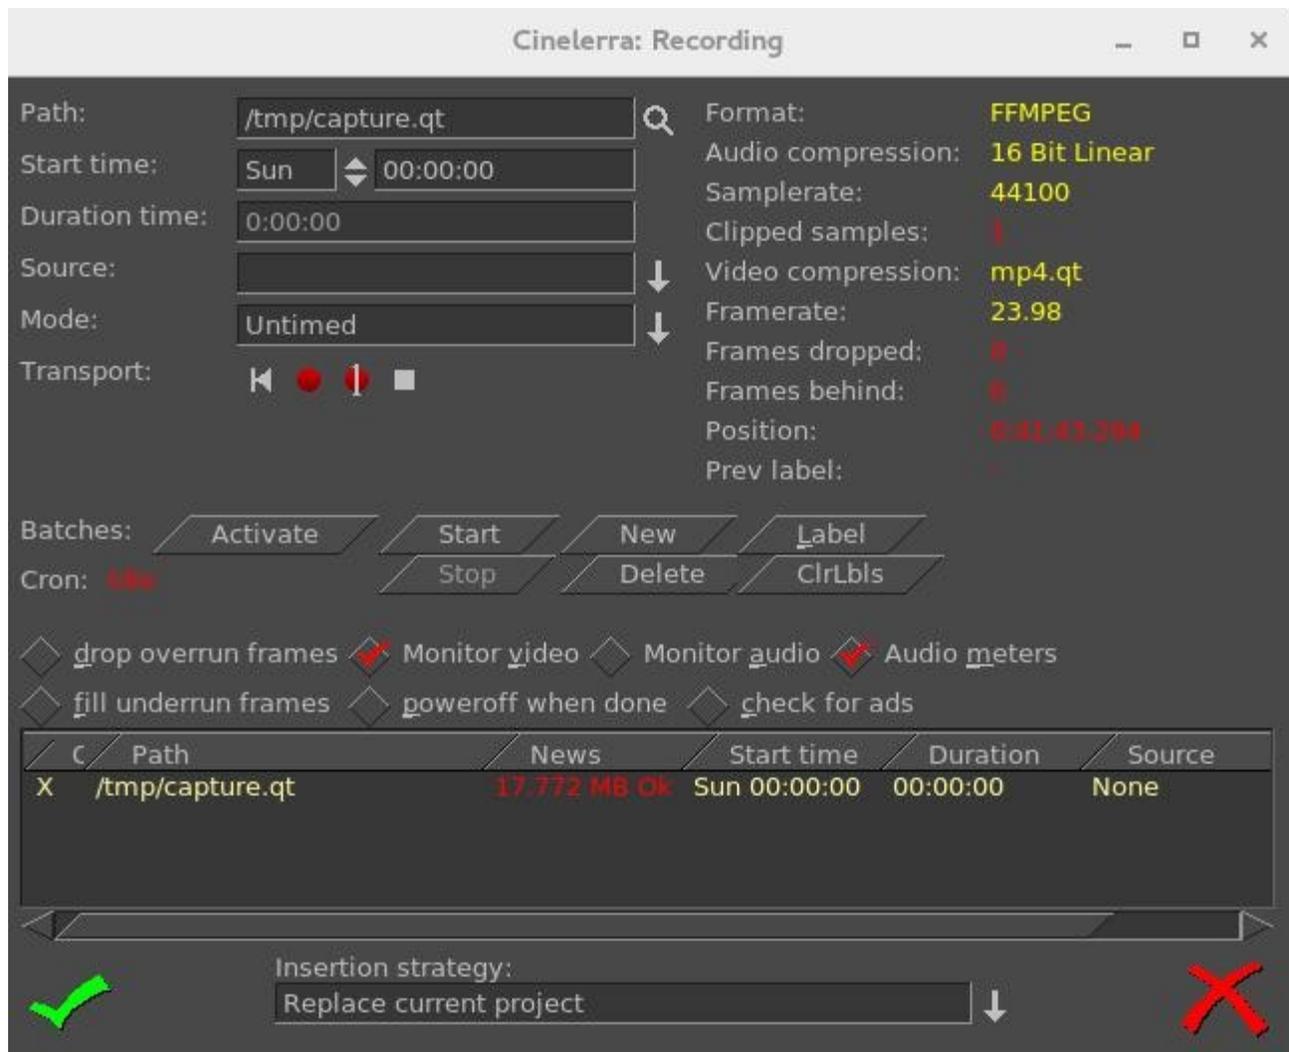
When recording from the screen, a large green-colored boundary box appears to allow you to easily frame the screen area to be recorded when you move it around to where you want to position it. In addition there are "record cursor" and "big cursor" check boxes which allow for also recording the cursor and for making the cursor bigger.

Screenshot to illustrate some appropriate settings which are described above/below the image.



- 3) Make sure the Compositor window is up since you will want to see the captured media later. From the main window press (lower case r) (move windows and resize if needed) Using the mouse pointer, mouse over the "Video In" Recording monitor display. Press and hold the left mouse button and drag the display until the desired target screen capture area is properly positioned inside the record monitor "Video In" display. At first this may seem confusing, but keep in mind that what you are actually doing is positioning the portion of your monitor display you want to capture over the "Video In" window.
- 4) From the "Recording window" select a Path file name for the recording.
- 5) Select mode untimed.
- 6) Make sure the "monitor audio" is unchecked else audio becomes looped and is very loud.
- 7) Press Transport "start recording" (red circle button).
- 8) Operate the playback of the media source, i.e. start playing the web video.
- 9) When the media ends, stop on the playback application.
- 10) Press Transport "stop recording" button (white square button).
- 11) Before clicking chk-Ok, be sure to set the Insertion Strategy.
- 12) Press chk-Ok to review the capture, or X to cancel / dismiss application.

Screenshot of what the "Recording" display will look like.



Two things are happening during recording, the data is being written, and it is being rendered. When recording, if video and audio are not in sync, it will usually be video behind audio as opposed to the other way around because video is more difficult to decode. During recording, video frame and audio sample positions are checked and efforts to maintain synchronization are used. Below is an explanation for some of the less obvious settings for recording as shown in the previous screenshot.

Frames dropped – more input than cpu(s) can process. This is the number of frames dropped when the system is too slow to keep up with the video. This value serves as a warning that it is behind and is dropping frames to keep up. You might want to consider stopping to make some parameter adjustments and start over to reduce drops.

Frames behind – active input buffers. This is the number of frames that the video is behind and it will drop frames to get caught up. “Frames behind” depends on the synchronization source, for example, time as the synchronization. This means that frames are not being written as fast as they are being captured and can occur when the system is too slow to keep up. This situation can happen with older systems or cpus where there is more input than can be processed in a timely manner. So that when there are “Frames behind”, frames will be dropped.

The algorithm for determining how many frames to drop is as follows:

If the number of frames behind > 3, only 3 frames will be dropped before it will then do 1 frame;
if the number of frames behind is 2, only drop 1 frame;
if only 1 frame behind, no frames will be dropped.

Drop overrun frames – checkbox when enabled, frames will be dropped. Drops occur when you capture a frame and there is no storage space available to save it, usually because buffer space is exhausted. This situation can occur when there are too many frames coming in too fast and the I/O system is not keeping up in emptying the buffers.

Fill underrun frames – checkbox if enabled, duplicate last frame when no input frame ready for output. As the tooltip states “write extra frames when behind”. This is desirable in order to keep audio and video synched with regards to sample and frame rates. Use underrun padding to fill frames when the input capture rate is lower than the output frame rate. Since demand is constant, you have to output something, so the last frame is duplicated as many times as needed to keep up. You should stop and adjust the frame/sample rates to match the input if not filling underruns, then start over. This situation can occur if the input stops for some reason like lightning interrupts the signal or the internet connection is too slow.

Transport controls - these control buttons mimic the functions of tape recorders from the old days.

Reverse button/left arrow – remnant from the past; does the same as the Start button really.

Red round button – this is the Start button to start recording.

White square – this is the Stop button to stop recording.

Red round button with white line through the center – start recording in single frame mode. The way this works is to just “take a picture now, take a picture now...”. The reason to use this mode is to get a still shot like you would with a camera. Sometimes input is continuous, for example looking at stars with a telescope -- surveillance goes on for hours, but you just want to take a picture now when something of significance interest comes up

Cron – The batch recording watcher, cron, is either Idle or Active. When you start or stop batch recording at specific times, there is a cron thread watching timers to perform the timed action. Idle/Active indicates whether the timers are running.

Position – this is a timebase which tracks frames/samples when obeying frame/sample rate. When a recording starts, it resets to 0. Timing is against audio (when available). Audio time and video time are based on position.

You can select synchronization time source “Positioning” in Settings → Preferences under the “Recording” tab. Possible choices for time base are:

Presentation Timestamps – use time code which is in both the audio and video media input stream. Uses these timestamps to sync the 2 streams.

Device Position – this is the device hardware position of where you are. It is usually only on the audio side.

Sample Position - Sample # or frame # divided by frame rate tells you where you are.

Software Timing – usually used for things like YouTube; it just will “take a picture now”...“take a picture now”... over and over again until you tell it to stop.

Positioning & Timing needs more detailed explanation for complete understanding and application. The overall goal is to maintain media stream timeline synchronization. The reason for providing different “Positioning” options is that different input media may have different timebase standards. Additionally, the input may be damaged. Damaged data can skew the timeline during presentation. The timebase standards make it possible to correctly resynchronize the media presentation to the original time position. For example, the transport layer may have timestamps provided in it. These timestamps record “audio time” and “video time” and are called “presentation timestamps”.

The kind of positioning used depends on the input device and the media format. In the transport stream example, the media has already been coded at least once as a stream and contains position information. In other situations, sample position or device position can provide timeline position. If the input device provides data at a variable or asynchronous rate, then software timers can provide timeline position. Sample position is position from the perspective of software, and device position is position from the perspective of hardware. The difference is buffer time. Device position is usually more accurate, but may not be available, or may contain errors.

Examples of media which usually has a particular type of positioning is:

- 1) mpeg has timestamps therefore Presentation Timestamps is a good recording choice;
- 2) raw media streams/screen capture have no timestamps, so tell it time via Software;
- 3) a web cam may supply data at variable rates, so again you would use “Software timing”;
- 4) an example of Device Position is audio (timestamps can come off the device);
- 5) if Software Positioning working for you, try Sample Position for really bad audio.

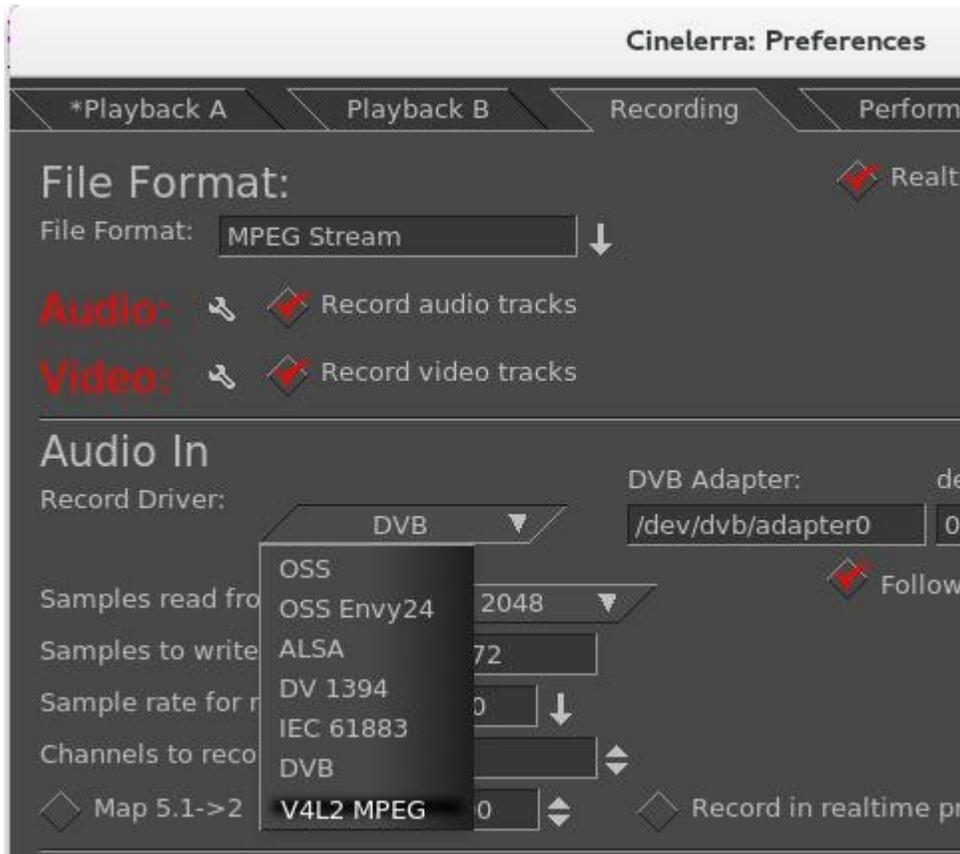
Other “Recording” settings are more pertinent to capturing and editing broadcast television with all of its many commercials. Their usage will be explained in more detail in a later section.

Label – create a label. During capture, put a time marker at the designated spot. When the recording gets pulled in later for review, you will see a green arrow marking the spot that was clicked.

Ctrl+L - clears all of the previously set labels.

Check for ads - check for commercials. In real time, labels are added when a commercial that matches a previous ad is already in the database (the TDB, Traveling Data Base).

One other noted new feature is a new choice for recording - "V4L2 MPEG". Some digital tuners now stream mpeg data for broadcast data streams. V4L2 has been upgraded to be able to process data from these device streams. Screenshot below shows the available options to include v4l2 mpeg.



13.2 Digital Video Broadcasting (DVB) Capture Configuration

You can do real-time capture of a full transport stream of Broadcast TV. Record television programs using Cinelerra by setting up ahead of time to start recording later at the specified time. Then you can use Cinelerra to watch later and easily fast forward through the commercials or edit the program, deleting the commercials, to watch uninterrupted. Some details may be slightly different than what it is in the United States. There are 3 requirements you must have to take advantage of this capability.

- 1) You need to have a suitable DVB Adapter such as the Hauppauge WinTV-HVR 950Q usb adapter. Make sure it is the linux version, not the new version.
- 2) Your DVB adapter on your PC should connect to a broadcast antenna input.
- 3) For now, you may have to run as the “root” user.

13.2.1 Overview

There is a lot of detailed explanation below to be thorough, but an overview is presented here.

- 1) Set up Record Preferences.
- 2) Scan the Channels.
- 3) Watch TV by selecting the desired channel.
- 4) Or set up a batch job to record a program and watch later via Ctrl-Alt-s in the main window.
- 5) Or record a currently running program by using the buttons in the Record window.
- 6) Load a previously recorded program.

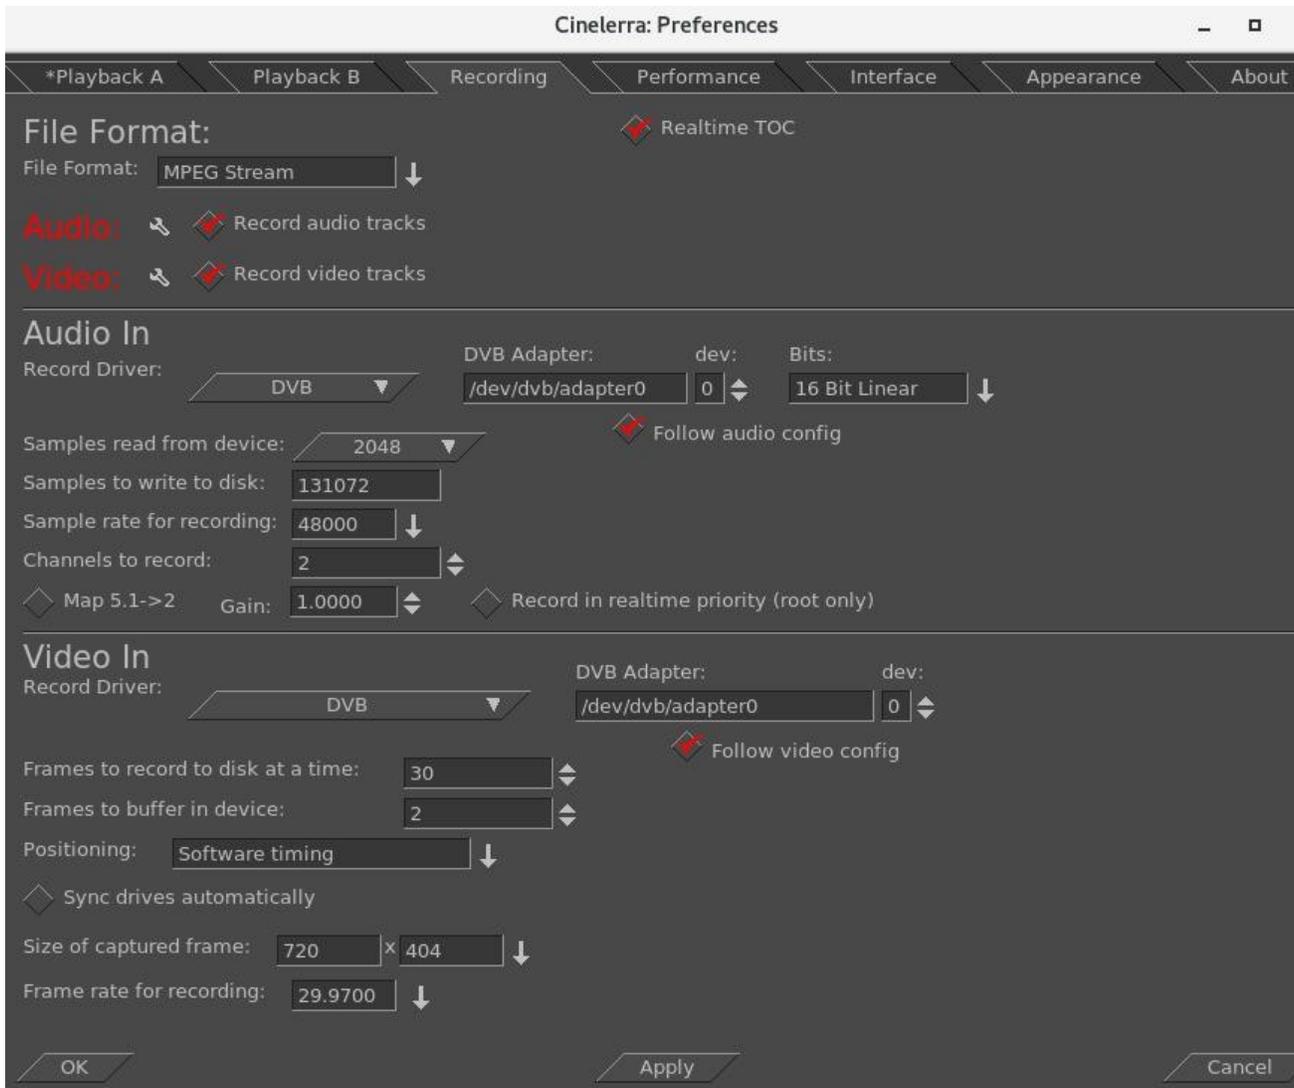
13.2.2 Initial Setup for Digital Video Broadcast Capture

Listed below are the steps for the initial DVB setup. Setup will be preserved until you make changes.

- 1) Start cinelerra, and click Settings->Preferences->Recording
 - Select File Format checkbox: Realtime TOC (Table of Contents)
 - It is not necessary to set File Format at this point as it will automatically be set later.
 - Check Record Audio Tracks.
 - Check Record Video Tracks.
 - Select Audio In->Record Driver: DVB
 - Set the adapter path = /dev/dvb/adapter0 (probably)
 - Set the adapter dev = 0 (probably)
 - Set the Bits: 16 Bit Linear (probably)
 - Check: Follow audio config
 - Samples read from device: 2048 or approximate dev buffer size (2k-16k probably)
 - Samples to write to disk 131072 (a good size)
 - Sample rate for recording: 48000 (will automatically change, probably to 44100)
 - Channels to record: 2
 - Uncheck Map 5.1->2
 - Gain: 1.0
 - Uncheck: Record in real-time priority
 - Select Video In->Record Driver: DVB (note, File Format above changes to MPEG Stream)
 - Set the adapter path = /dev/dvb/adapter0 (probably)
 - Set the adapter dev = 0 (probably)
 - Check: Follow video config

Frames to record to disk at a time: 30 (a good number)
Frames to buffer in device: 2 (2-6 probably)
Positioning: Presentation Timestamps (**this is important and may not be already set**)
Uncheck: Sync drives automatically
Size of captured frame 720x480 (will automatically change)
Framerate for recording 29.97 (will automatically change)
Click OK

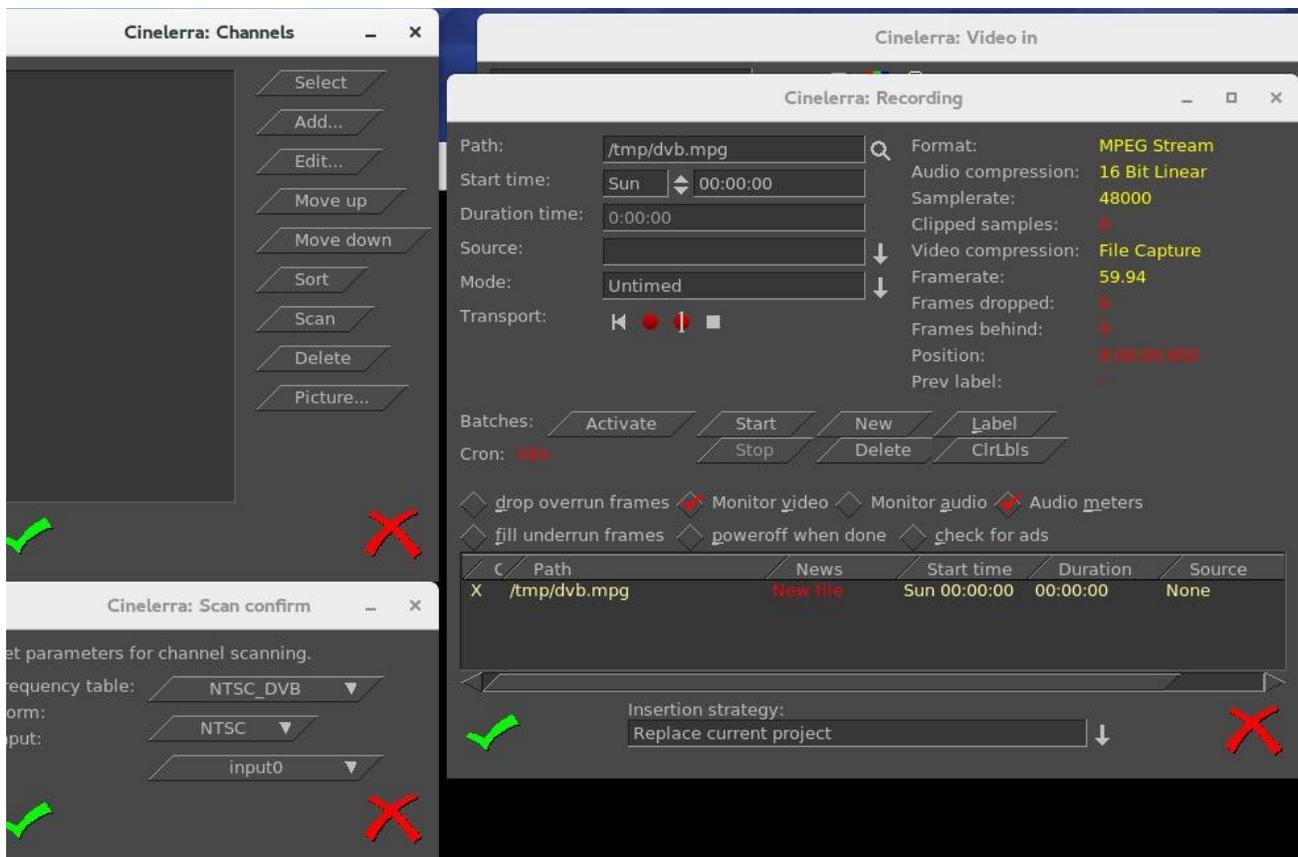
Screenshot below shows the Preferences window as a good example of what you should see.



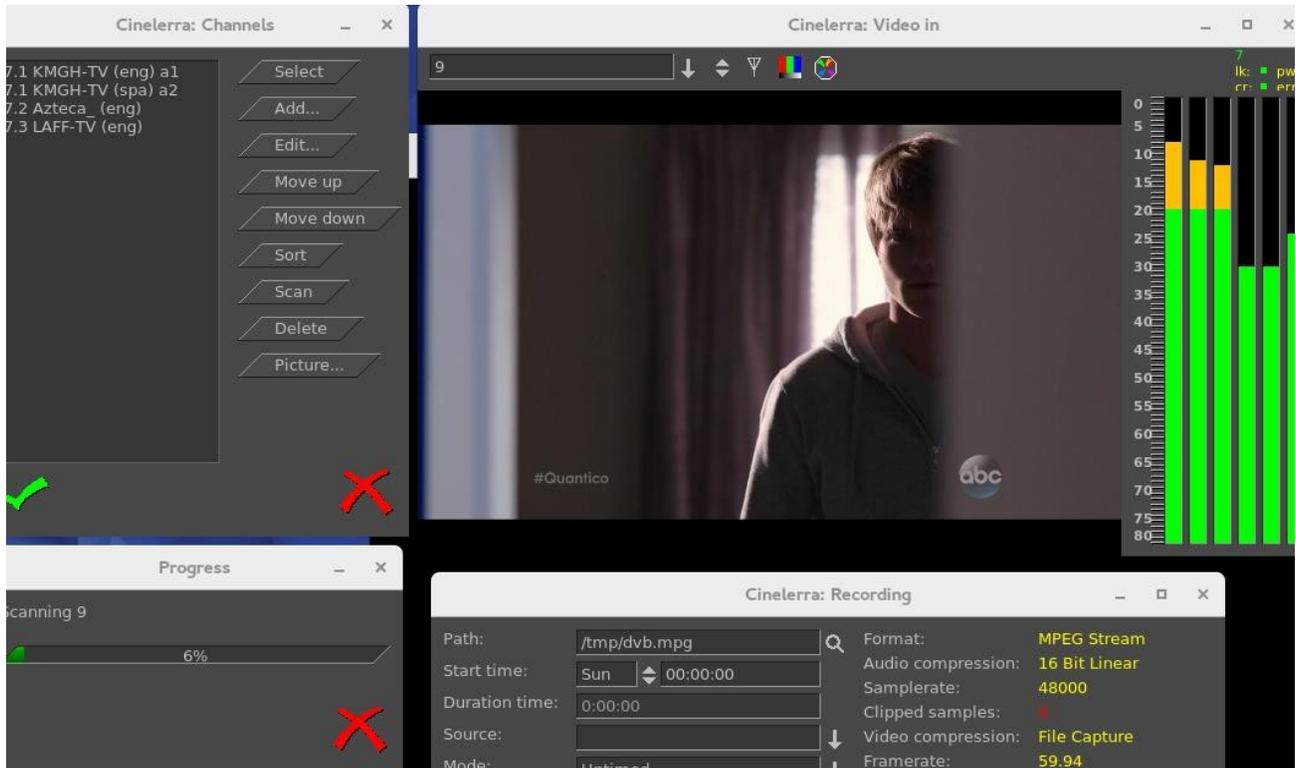
Scan the channels once the devices are setup by performing the following.

- 2) From the main window pulldown use: File → Record (lower case r) & resize windows. Two windows appear - the “Recording” and the “Video In”. The “Recording” window should have “Monitor video” checked and “Monitor audio” unchecked. Remember to check the Monitor audio later when you want to listen to a TV program. On the “Video In” toolbar, click the antenna button (Edit channels tool tip). When the Channels window opens, click scan - “Scan confirm” appears. Make sure Freq table NTSC_DVB, Norm NTSC, and input0 are chosen automatically. Click OK and then wait for scan to end which may take several minutes.

Screenshot below on the left shows the Channels window and the popup waiting for scan confirmation. On the right side is the usual Recording window with the Video in window hidden behind.



Screenshot below shows scanning in progress. Note the “Video in” with the antenna (3rd symbol).



Any time after the channels have been scanned, you can edit them from the “Recording” window channel selection dialog (shown on left of screenshot above, such as “Sort”). You do not have to scan the channels every time as they are saved in a database referred to as the channel-db.

Click chk-OK when done scanning channels and making any edits you want.

After the channel-db is ready, you can watch any of the selectable channels.

13.2.3 Watching TV using Cinelerra

If you have setup for DVB capture as outlined in the section above, just start here to watch TV.

From the main window use **File → Record** to request the “Recording” application.

In the “Video in” window, press the down arrow on the left to see the channel pulldown.

Select the desirable channel, and watch TV. You will have to enable the audio by checking the “Monitor audio” in the Recording window for sound.

Helpful hints: for any computer that you might choose to use that is a little slower, you can improve the available resources for watching TV by unchecking “Audio meters”, checking “drop overrun frames”, and checking “fill underrun frames”. Also, uncheck “Realtime TOC” in the Preferences → Recording window.

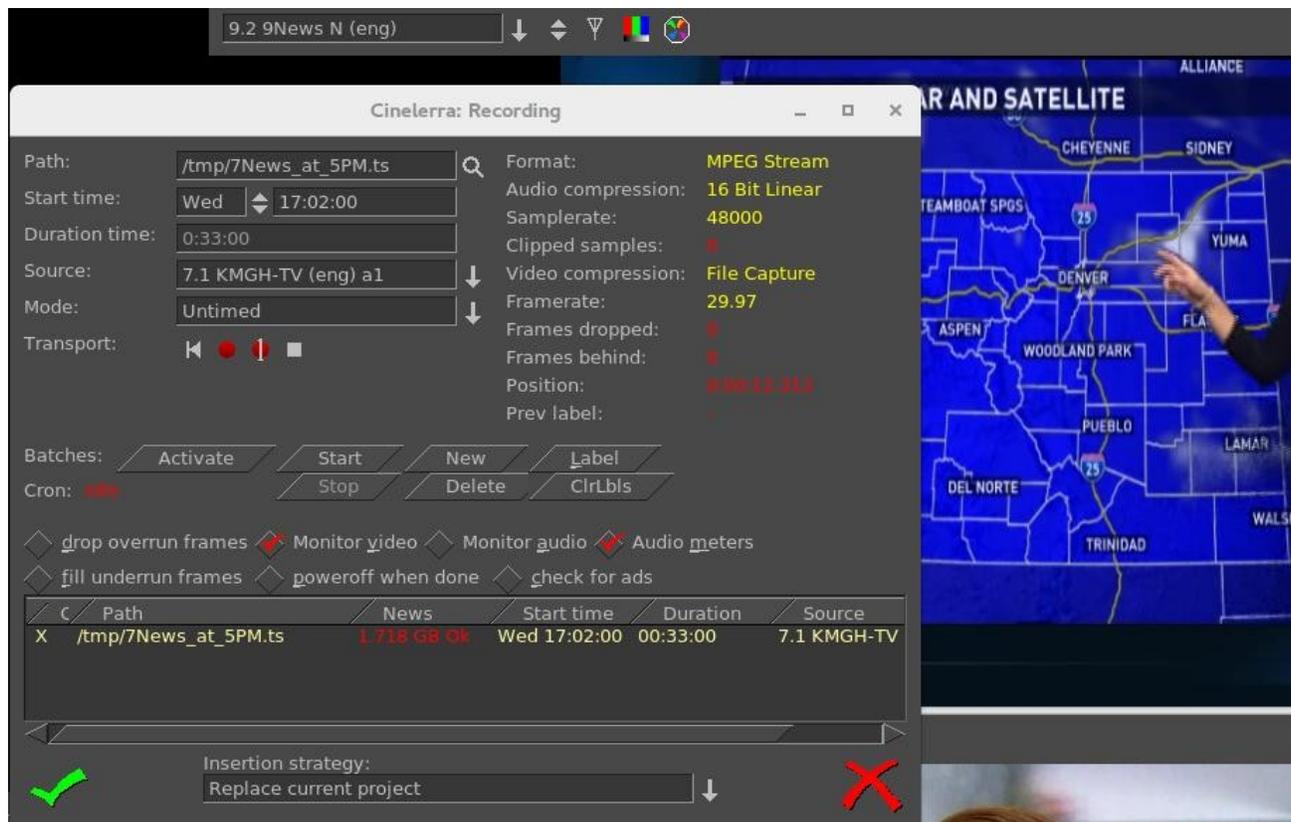
Dismiss the “Recording” window with chk-OK, or X to cancel when finished.

13.2.5 Recording Immediately to Capture a current TV program

If you have setup for DVB capture as outlined in the first section above, just start here to record a TV program that is currently playing and stop it whenever you want. Media files can be directly captured from the “Recording” window to include Start, Stop, and then loaded via chk-OK.

- Bring up the “Recording” window via the main window File → Record (lower case r).
- Set the “Path” to the write file Path; be sure there is enough disk space (9.5Gb/hr).
- Set “Mode” pulldown to “Untimed” found in the upper left corner area of the window.
- Press transport "start recording" button (red round button).
- Path will flash and update in batch window and remain flashing while recording.
- Press transport "stop recording" button (white square button).
- Before clicking chk-OK, be sure to set the Insertion strategy.
- Pressing chk-OK attempts to load any newly recorded media data for review..
- Or press X-cancel to terminate and dismiss the “Recording” application.

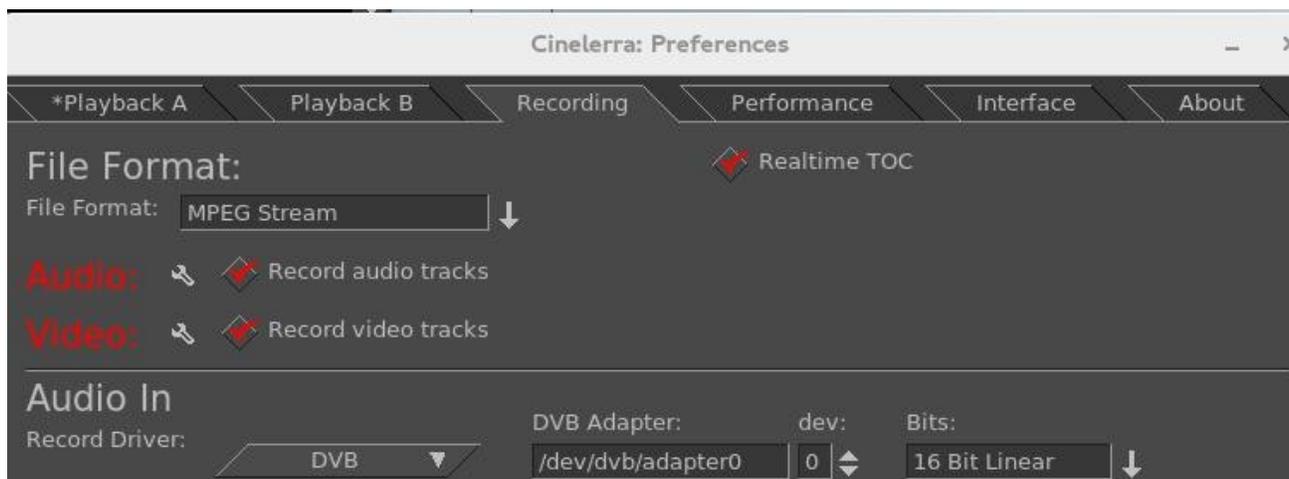
Screenshot showing the recording of a currently running broadcast TV program weather report.



13.2.6 Record “on the fly” Table of Contents (TOC)

If the Settings → Preferences → Recording tab is selected, and the “video Record Driver” is changed to DVB, the file format is switched to MPEG Stream automatically (note that the “audio Record Driver” must also be set to DVB for broadcast recording). This kind of record stream can be captured by simply writing the hardware capture stream directly to the record capture asset file path. But by checking the “Realtime TOC” in the Preferences Recording window, it is possible to scan the stream data “on the fly” on its way to being written while the asset is being captured. The scan is used to generate a TOC – Table of Contents file in the \$HOME/.bcast5 directory. This file is used by the mpeg library to view the stream. This avoids the need to scan the file on first open since the TOC builder can run in realtime.

Screenshot showing the red-checked “Realtime TOC” for the “Record Driver” of DVB.



13.2.7 Watching a Previously Recorded TV program

If you have setup for DVB capture as outlined in the section above, just start here to watch a previously recorded TV program. Use the following steps.

In the main window choose File → Load files and select the desired media file to watch.

Broadcast transport streams may consist of multiple programs that all run together on one channel (for example 9-1, 9-2, and 9-3 are all on channel 9). Therefore when you play the recorded program back all of these programs will play simultaneously and that is a jumble.

In order to select just the one you want, press Alt-1 in the main window to select Asset Program 1 (or Alt-1...8 to select the desired asset program number) which removes all of the unrelated audio/video tracks except for the selected program. It also resets the session to match the media.

Move timeline to where you would like to start watching.

Be sure the Compositor window is up and you will see the pre-recorded program there. Use either the keyboard, a remote control that has been set up, or Android device for controls. To toggle between fullscreen or original size, press “f” in the Compositor window.

To watch on a large TV screen or monitor which you have connected, you will have to setup for Dual Screen Mode as explained in the next section. In addition, you can either use a Remote Control hardware device or special key setup for manipulating the playback.

13.2.8 Dual Screen Mode

Dual screen mode is useful for displaying cinelerra video/project on a large screen digital Television set or for watching a previously recorded broadcast TV program via Cinelerra. The physical configuration results in one screen (an Xwindow) on one monitor and the other screen on the other monitor or TV screen. Before the mode is usable for Cinelerra, X has to be configured in “dual screen mode” - that is, you must first split the configuration into 2 displays to operate Dual Screen Mode. On some systems, you may be able to get to the dual screen mode by using Xinerama allowing for use of 2 physical displays to act as 1 virtual display.

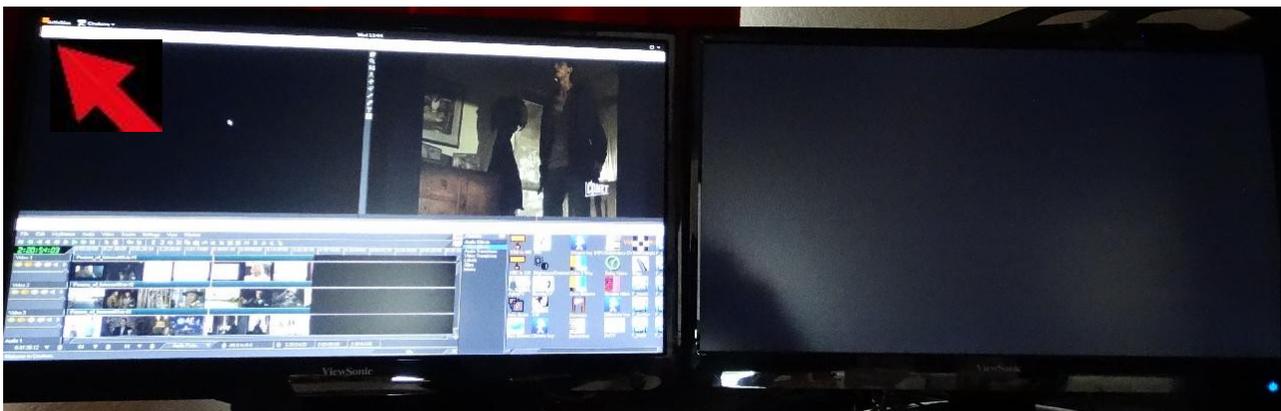
The main menu pulldown, Settings → Preferences → Playback A and B Tabs, must be properly configured in order to take advantage of the Dual Screen Mode capability.

In the main menu Window pulldown there are “Tile left” and “Tile right” options for the 2 monitors. If you have only 1 monitor, you will see no change upon activating these options. Tiling is possible due to the existence of one long horizontal screen display buffer. Tile left takes Cinelerra and moves all of its windows to the 2nd monitor/TV and Tile right puts it back to the 1st monitor. The left and right monitors are the left part or right part of a Playback A/B Xwindow host referenced in the Video Config Playback A/B for separate X screens. Shortcuts for Tile left and Tile rights are the letters a and b on the keyboard once you have set Remote Control mode. Below is a summary of the configuration letters:

- a - select playback a and tile windows left (all on monitor or laptop)
- b - select playback b and tile windows right (all on 2nd monitor or TV)
- c - select playback c and tile windows left, but composer right (TV mode)
This option has the effect of taking the Compositor window and moving it to the 2nd monitor/TV without the surrounding borders interfering with the video.

Since Dual Screen Mode is most widely used for either a demonstration or watching TV, you will want to take advantage of Remote Control mode as described in a following section. In the camera shots below, note the square red box that shows up in the upper left hand corner of the Viewer window (arrow pointing to it but it looks more orange than red) which denotes that the Application/Menu key was used to get into Remote Control mode.

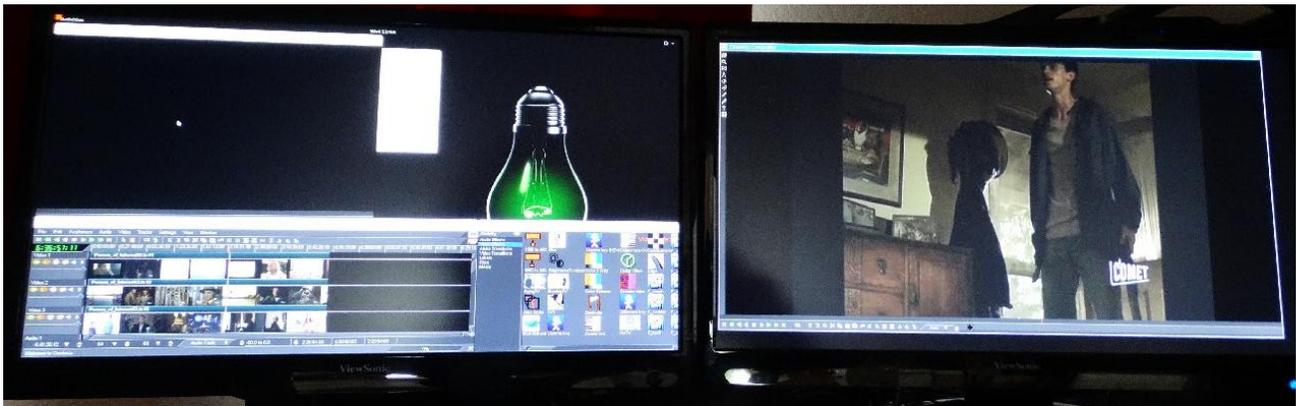
Camera shot below shows 2 monitors (could just as well be a big-screen TV for the second monitor) with Tile left where all of the cinelerra windows are on the left most monitor. The big red arrow points to the little red square box (actually may look orange) indicating that Remote Control mode is in effect.



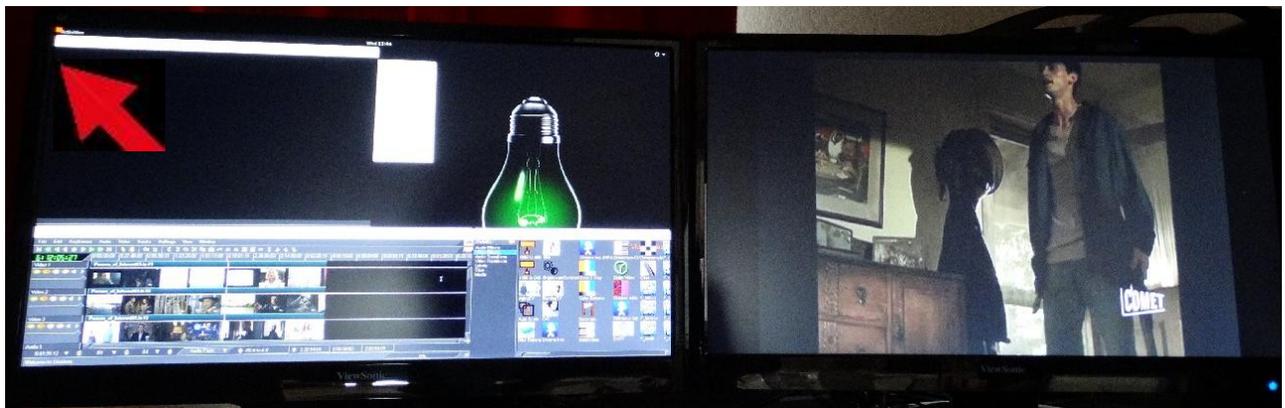
Camera shot below shows 2 monitors with Tile right (b key) where all of the cinelerra windows are on the right monitor. The left monitor shows the Suse distro logo and a couple of xterm type windows.



Camera shot below shows 2 monitors with Compositor window on the 2nd monitor which was obtained through use of the "c" keypress available when in Remote Control mode.



Camera shot below shows 2 monitors with Compositor window on the 2nd monitor in FULLSCREEN mode so that no cinelerra borders are visible to distract from the picture. The big red arrow points to the little red square box (may actually look orange) indicating that Remote Control mode is in effect.



13.2.9 Remote Control for DVB

Cinelerra DVB recording/playback can now easily be done in “couch potato” mode using a remote control. This comes in handy when you want to playback on a big TV screen for multiple person viewing. You can use the Application/Menu key on the keyboard to toggle between standard Cinelerra usage or a Dispatcher methodology through use of a remote control, an Android programmed device such as a tablet, or keyboard. The Application/Menu key on most keyboards is between the Alt and Ctrl keys on the right hand side bottom and depicts a menu on it either with or without a pointer (see image below). An ati-x10 Remote Control device is currently working with Cinelerra and other remote controls may work but have never been tried. The red/orange box on the top left main screen of Cinelerra indicates Dispatcher Method access is activated which allows for keyboard grab which routes all key strokes until toggled off. Below are the currently defined operations.

Remote Control Keys (Application/Menu key toggle for ati-x10 remote)

Ati-x10 Remote

- up arrow - forward 1 minute
- down arrow - backward 1 minute
- left arrow - back 10 seconds
- right arrow - forward 10 seconds
- d - toggle channel scan menu
- f - toggle full screen
- a - select playback a and tile windows left (all on laptop)
- b - select playback b and tile windows right (all on TV)
- c - select playback c and tile windows left, but composer right (TV mode)
- 0-9 - select timeline by 1/10s from beginning to end
- a - toggle audio mute
- m - toggle audio meters
- keypad / - display mpeg ts stream program data
- keypad * - display mpeg ts stream channel schedule date
- keypad + - change channel, next
- keypad - - change channel, prev
- keypad . - start channel change by numeric entry
- keypad .. - clear channel entry
- book key - toggles channel scan (always available)
- hand key - terminates cinelerra (always available)

Plus usual Transport keys

- e - ??
- f - toggle full screen
- << - fast reverse
- > - toggles normal playback
- >> - fast forward
- red o - slow reverse
- square - stop
- 2 lines - fast reverse



Application/Menu key
right bottom of keyboard



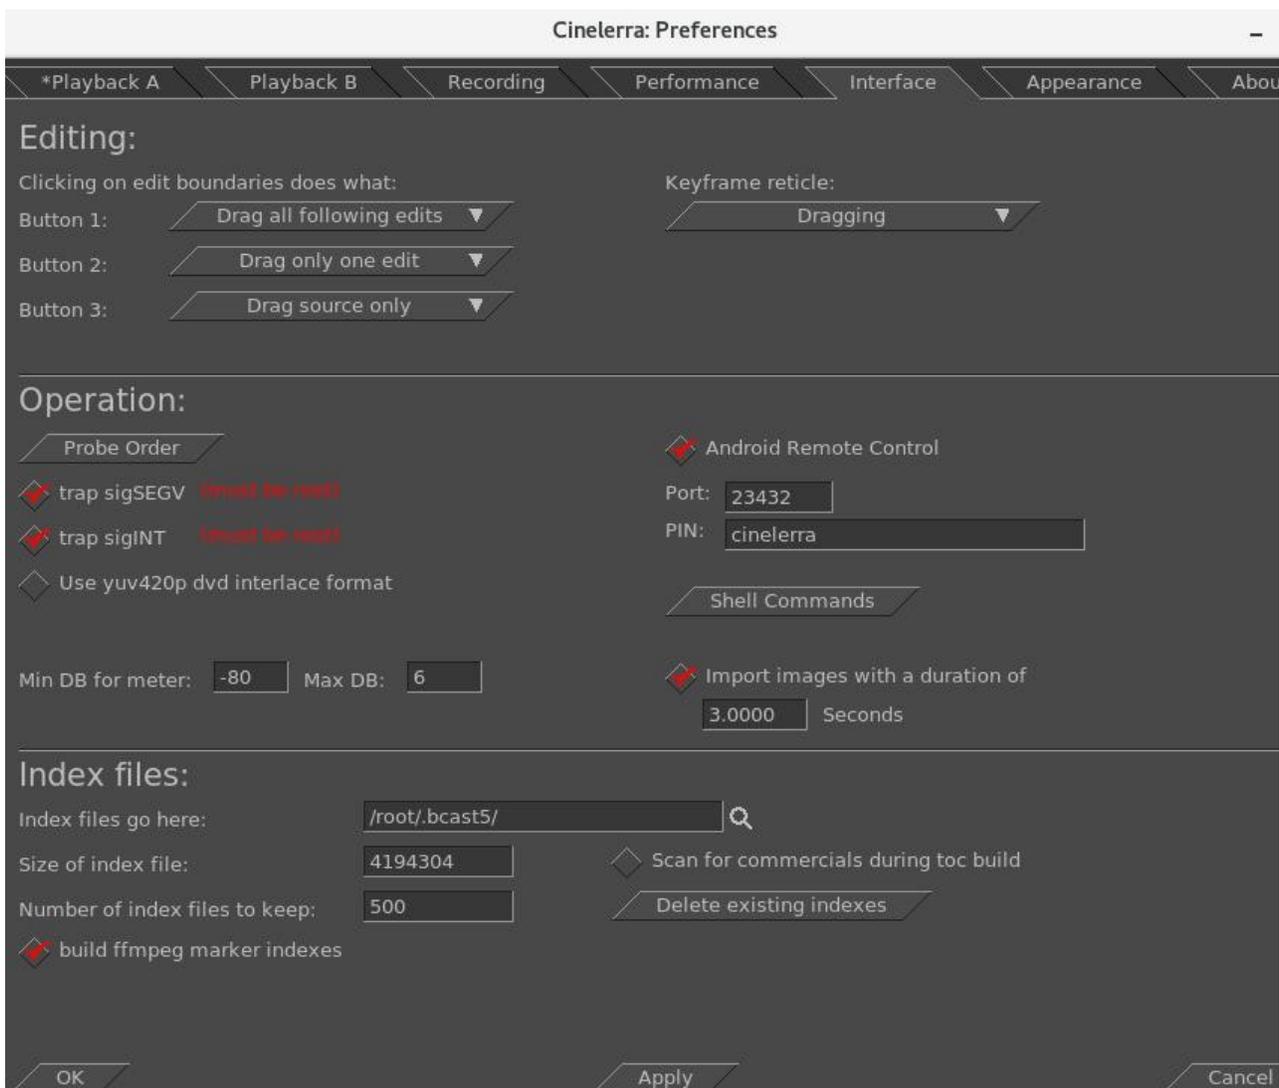
The Application/Menu key is used to get into remote control mode, even if there is no remote control

hardware device in use, allowing for use of the same defined remote control keys to work on the keyboard. Once the Application/Menu key enables remote mode, the remote control keys don't allow the standard Cinelerra keys to work.

13.2.10 Android Remote Control for DVB

Cinelerra 5.x has android remote interface code. Any device, such as a tablet or a phone, can be used as long as it is running the Android operating system. Programming an Android Remote Control for DVB is a bit complicated at first, but becomes pretty simple after a bit.

- 1) The Android Remote Control requires you to download and install the android CineRmt app package (apk). To install it simply download it (any way you can) to your phone or tablet and click it to install, if it doesn't run automatically.
- 2) You should be able to open Settings → Preferences → Interface tab. And see Android Remote Control in the Operation section. Enable it using the checkbox. The default port 23432 and default PIN cinelerra are OK.
- 3) After you install the new phone/tablet app in step 1, start it up from your apps menu.



Before you can use it, the Android device must have the ip address of your computer entered into the

configuration menu. The PC you are running Cinelerra on and the Android device have to be on the same network. You can tell it is in communication by tapping the Power (menu) button. When the remote is operated with the Power (menu) button, there is a little orange box on the upper left corner of the display to indicate the remote state is active. If you do not see the little orange box when the menu key is tapped, the communication is broken and must be fixed before proceeding. Follow the directions in the paragraph below.

BIG NOTICE: the firewalls in your computer and wifi router can stop this thing dead. Re-configuring a firewall is tricky and varies from distro to distro. If possible, during setup it is suggested that you disable the firewall temporarily. For those with advanced skills, use tcpdump or wireshark to look for udp messages from the IP address of the device.

SECTION NOT COMPLETE – (manualandroid)



Screenshot of an Android Tablet that can be used as a Remote Control for Cinelerra Recording and Playback.

13.2.11 How to Setup a Remote Control Device to work with Cinelerra

SECTION NOT DONE !! This section explains how you can take your own brand of remote control and set it up to easily work the way you want with Cinelerra.

13.3 The commercial DB... What is that? A TDB – Traveling Data Base

While watching broadcast television, it is possible to use Cinelerra to automatically mute commercials being aired during a broadcast capture. Visual muting as well as sound muting is done! You can NOT use ffmpeg for the format when working with the commercial DB. This database that was added is a “novel architecture” shared memory database, a Traveling Data Base (TDB). With TDB you have direct access to the shared database presence (instead of having to go to the server).

Basically it works as follows, but there are a few differences in just how it goes about it:

- 1) Identify a zone in a timeline as a commercial.
- 2) Characterize it concisely, in the audio and video timeline.
- 3) Store the characterization in a database.

Then, when capturing media or building a table of contents for an mpeg transport stream (broadcast capture DVB):

- 1) Characterize each I-Frame (a few frames per sec).
- 2) Search the database in real-time to identify the image.
- 3) If a match is identified, automatically press the mute button.
- 4) Unmute if the pattern match fails, or the commercial ends.

It mostly works, but it takes more time to setup and maintain the database as the advertisements change frequently or vary by a small amount. It would take a small set of personnel to keep the database current and some finessing of the code.

Currently, to activate the commercial database (db) methods you must do the following as root:

```
cd <cin-path>/bin
mkdir /cinelerra          (yes, currently only in the / filesystem)
./cin_db /cinelerra/media.db (to create a new empty commercial capture database)
export BC_USE_COMMERCIALS=1 (in the command shell that is starting cinelerra)
```

This activates the gui tools and database access. As of yet, all of the requisite permission changes necessary to run as non-root have not been worked out (initially just easier to run as root).

When a commercial is "cut", the following occurs:

- 1) transmits the capture timeline fragment to /tmp
- 2) scans the snip to characterize it statistically
- 3) creates a miniature b/w movie of the beginning and end
- 4) stores all of this into the media.db

There are 2 ways currently to identify a commercial:

- 1) Load mpeg3 transport media, highlight select the commercial in the timeline, and press "cut commercial" on the toolbar.
- 2) Use the remote control interface to indicate the boundary in/out cut points in the realtime capture timeline, and operate a cut command.

The cut command creates a popup to let you know that it is processing the captured data. This can take several seconds depending on the size and media format. The results from the scan are stored in a new record in the “media.db”.

The stats are somewhat disappointing. Image recognition is way past zero, it is not 100 percent. Broadcast streams have tons of minor variances which hamper digitally perfect reproduction from instance to instance of commercial presentation. This puzzle remains to be solved.

To review media stored in the “media.db”, and delete old expired commercial data records, press (Shift-M) from the main window. This will create a media-db window which shows the various captured samples stored there.

Currently, only a customized ati-x10 (kernel mod required) radio remote control works. But this can quite readily be modified for other remote control models.

13.4. DVD and Bluray Creation

This section describes how to create a blu-ray DVD, or **BD** referring to a **B**lu-ray **D**V**D**, and a regular DVD, or **SD** referring to a **S**tandard **D**V**D**. The DVDs (plural usage of DVD means either BD or SD) created are unencrypted, unlike commercially available movie DVDs. This Cinelerra version conceivably can create different variations of DVD/Blu-ray media but for the casual user the most standard usages are readily usable and will be described here.

Some preliminary information follows. For NTSC, SD media is almost always 720x480 interlaced (the format in the United States, US). For PAL, SD media is almost always 720x576 interlaced (Europe, EU, and most of the world). An SD can conceivably be created with a lower resolution – for example 352x240 MPEG-1 – but it is not useful. Aspect ratio for either NTSC or PAL can be 4:3 or 16:9.

Frames per seconds is usually 29.97 for NTSC and 25 fps for PAL. The standard SD dvd generally uses the MPEG-2 program stream with a filename extension of m2v. BD blu-ray media is not normally interlaced, but you can leave it interlaced. Blu-ray uses the MPEG transport stream which contains 1 or more program streams with a default filename extension of m2ts.

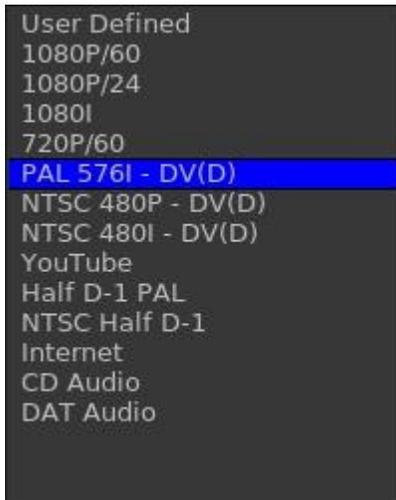
Requirements for creating DVDs is the hardware device to write the media on and, obviously, the blank media for either BD or SD. When generating SD media, you will have to install dvdauthor and for BD media, install udftools if they are not on your system. Also, keep in mind that **to mount filesystems for creating files and to burn DVDs, you will have to be root since you have to have privileges**, unless special permissions have been provided for a non-root user. It is also highly recommended to run the cinelerra startup from a terminal window, instead of the icon, in order to see informative messages of how to actually write the output (at prompt: keyin <install_directory_path/bin/cin).

A warning here – writing blu-ray BDs and regular SDs can take a large amount of clock time. Keep in mind that a blu-ray can contain 25 hours of viewing, so would take multiple hours to just write one.

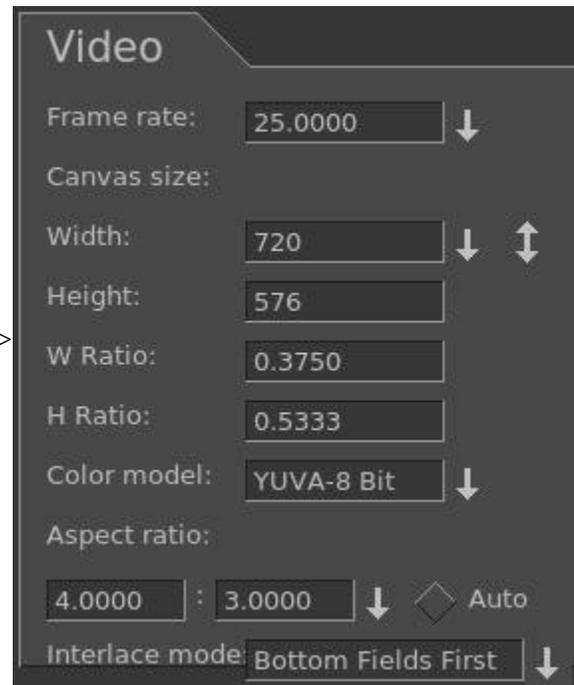
The max disk space needed is a little over 100GB for 50GB Double Layer (DL) media or 50GB for a single layer BD blu-ray. A standard SD of 4.7GB needs about 10GB disk space. You probably can get by with much less if the render for blu-ray is less than 25 hours of media. You will need twice as much disk space as the media holds to ensure you have sufficient space for working and copying.

The most important thing you need to know about in order to get started is "Format" and "Scale" in the "Create dvd or bd" window. Format settings shown in the Set Format window are set accordingly to an algorithm. Basically, it will take whatever you say in the asset format. It matches that against the known Presets available so that is what will be shown. If the asset format doesn't match any of the Presets default formats, then that will be shown as the "User Defined" format. In addition when you load media, the format is initially set to "same as source" so matches the source input and if that matches a known preset, then that is what is shown. The PAL versus NTSC only comes into play when there is no known correct matching format when you attempt to create a DVD Render batch job. It is applied when you click OK. For example, if you load up a YouTube video, it will not match any known format and will choose PAL or NTSC based on time zone.

Screenshot of Presets in Set Format window



When choose PAL, values get changed in window to reflect PAL -->



A quick set of **basic steps to create DVDs** is

immediately below and usually just using the defaults will get you something. However there is a serious issue with interaction between the Operating System and bdwrite when creating a BD/blu-ray that requires automount to be turned off. Refer to the details and more specific explanations below the following steps for how to do this.

- 0) If not logged in as root, you will get an error message in order to avoid doing a lot of work and then failing out because root is required for automount and to write on DVD hardware.
- 1) Load your input source media via: File → Load files .
- 2) Choose PAL or NTSC for SD/dvd or 1080P/24 for blu-ray in Settings → Format .
- 3) For blu-ray, choose BD Render or for PAL/NTSC, choose DVD Render in File → .
- 4) Designate a “work path” with sufficient disk space and then Chk-OK .
- 5) When the Batch Render window comes up, click on Start and the batch jobs will run.
- 6) Read the final messages echoed to the screen to see the command for burning... OR
- 7) Use the provided directory name to: cd /<target directory>bd_(or dvd_)<date-time> .
- 8) Load your media, format if needed, note device name to substitute for <bd> or <dvd>.
- 9) If rewritable blu-ray, use dd if=./bd.udfs of=/dev/<bd> bs=2048000
If write-once blu-ray, use growisofs -dvd-compat -Z /dev/<bd>=./bd.udfs
If any DVD media, use growisofs -dvd-compat -Z /dev/<dvd> -dvd-video ./iso

Any problems encountered will require that you read more information in this section to include specific details, helpful hints, and problem resolution.

Details and specific explanations to create blu-ray or regular DVD are provided here. It is very advantageous to startup cinelerra from the command line prompt instead of the icon. Also, please be root or your hard work will be lost when the automount is issued and fails for bluray udfs mounting.

The general design of the DVD/blu-ray generation operations is to first render media using batch rendering and then terminate Cinelerra to start a script which creates the target device filesystem data. These scripts are the dvd.sh and bd.sh scripts written into the target directory. For DVD, the general plan is to write a directory <target>/iso with the dvd filesystem via dvdauthor and then generate an iso9660 filesystem and write it to a dvd via growisofs.

For blu-ray, the filesystem generation is slightly harder. First, it creates an empty filesystem image <target>/bd.udfs using mkudffs which makes a big hole for the filesystem data. The hole is made just a little bigger than the data written by bdwrite so that you don't have to write an entire 25GB or 50GB disc even if no data exists. This empty filesystem is loopback mounted to make it writable, and the linux kernel manages the filesystem image. The bdwrite program applies the blu-ray structure to the UDF filesystem by creating the needed BDMV blu-ray filesystem, which the kernel stores onto the image file bd.udfs. When udfs is unmounted, the kernel finalizes the disk image on bd.udfs. The bd.udfs image can be written directly to a blu-ray disk via dd or growisofs.

NOTE of IMPORTANCE: there is a serious situation with the interaction between the Operating System (OS) and bdwrite when creating blu-ray, that requires automount to be turned off. The blu-ray automatic script unmounts the blu-ray/UDF filesystem but the system has not finalized the directories so the OS creates a new loop file device and the data is loaded and cached for use by the new loop but it is stale. Consequences is that not all of the data is written where it should be. The solution is for the OS not to mount this second mount so we have to make sure it doesn't. There are 2 methods to fix this. The first and easiest is by using the following command to disable automount:

```
gsettings set org.gnome.desktop.media-handling automount false
```

This can be reversed when you have completed the blu-ray generation via:

```
gsettings set org.gnome.desktop.media-handling automount true
```

A different and more complicated method you can use to turn off automount is to download and install the dconf-editor. Automount is a system parameter and only needs to be done once unless you do not want automounts to always be disabled.

- 1) run: dconf-editor
- 2) select: org → gnome → desktop → applications->media-handling
- 3) uncheck: automount
- 4) close dconf-editor window

Immediately below are the detailed steps with explanations for creating SD or BD media.

1) Construct a session with the desired presentation:

Format frame rates choices are 29.97 or 25 for SD, based on the user's timezone, with NTSC 29.97/US or PAL 25 /EU. For BD, the media input will be analyzed to automatically pick the default format or if unknown, the user's timezone will be used to default to 1920/29.97i for US or 1920/25i for EU. Be sure to set the rendering parameters in the settings -> format menu. Choose audio stereo or 5.1, again depending on your media.

Target Geometry will be 720x480 (US) or 720x576 (EU) for SD.

2) From the main window, select file → BD Render or select file → DVD Render. Then:

In the "Title" textbox there will automatically be a unique name based on the current day and time but you can change that to anything you want that is NOT already created.

Choose a "Work path" Target Directory in the work path textbox. Another big warning here is to use a path that is NOT /tmp because a lot of linux distros wipe out /tmp on a regular basis via cron, crash, or reboot. Furthermore, they tend to use ram based filesystems, which will be way too small. You will lose your hard work so use disk space that is permanent. An icon in the form of a magnifying glass is next to the work path which shows you file names/paths.

Disk space message will be green if you have at least 100GB/BD or 10GB/SD available

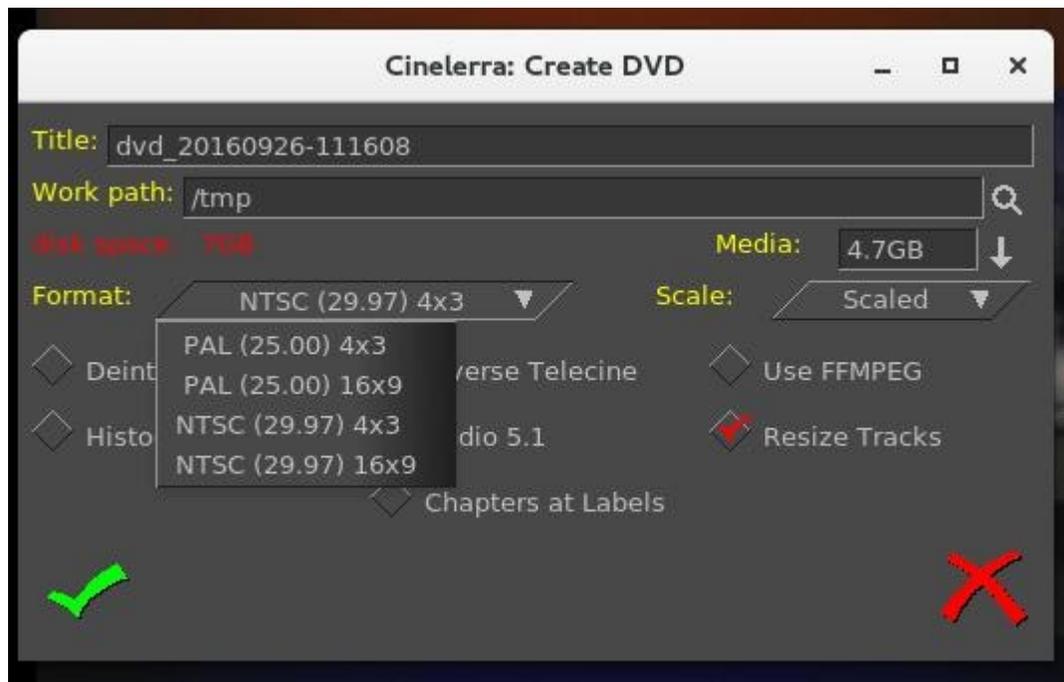
space (the disk space message will be red to warn you if you have less than these amounts).

Screenshot showing sufficient available disk space in green for blu-ray creation is below.



← Note the Scaled options in the pull down menu which will be described later.

Screenshot showing insufficient available disk space in red letters for SD dvd creation is below.



Note the pull --> down displaying the PAL and the NTSC formats.

(Continuation of Step 2)

Select desired features, check / un-check as appropriate.

Click OK check button. It is **very important to realize** that when you check OK, the EDL will be saved and that will be used for batch job rendering. If you bring up the Batch Render and then change some parameters, they will not take effect UNLESS you remember to either check "Save to EDL Path" or "Use Current EDL" in the Batch Render window. You will get a reminder automatically if "warn if jobs/session mismatched" is checked.

Explanation of the choice boxes as seen in the previous 2 screenshots for both SD and BD menus is given below. Many of them are plugins which allow you to further manipulate the settings for best results. They are just suggestions set by the program automatically based on your input media, and can be reset to suit your needs. These are listed in the next 4 points.

- 1) If the media does not match the DVD target geometry, and the **scale** plugin is not already in use, then the scale plugin is applied with scaling set to fit the media dimensions to the DVD format target geometry.
- 2) If the video height is at least twice the DVD height and the input media is interlaced, then the **deinterlace** plugin is applied with odd line sampling.
- 3) **Audio 5.1** will automatically be set to the wide-audio if you have 6 tracks of audio.
- 4) To allow video data to be accessible and overlay properly, the track buffers are “resized” to the largest track frame size in use (**Resize Tracks**). The theory behind this is to make sure to have enough memory to cover the entire presentation for transcoding.

All of the current choice boxes are further defined immediately following.

Deinterlace – remove the interlace. Interlacing is a video scanning system in which alternating lines are transmitted so that half a picture is displayed each time the scanning beam moves down the screen. You lose a lot and the quality is bad when you view interlacing on a progressive TV. You might not really want to use deinterlace, because if you deinterlace non-interlaced media, it will look awful.

Scale – alter the spatial mapping of an image to increase or reduce the size; modifies the picture. When some programs scale from 4:3 to 16:9 they will automatically cut off the appropriate section of the image for you. It is necessary to keep in mind, that square pixels is the true end goal of scaling, not the aspect ratio which could result in squished or stretched output. More information about scaling will be provided on a subsequent page with usage of the “Scale Ratio” plugin.

Histogram – remaps the color space. The color space ranges from 0-255 for 8-bit color values. You can use this tool to remap the color space to use the entire space or for stretching the contrast. Also, it lets you perform global color-correction on the image. You can use this to correct for color screens that are 'too blue', or for color Televisions that produce 'brownish' output, or whatever. In addition to color-correction, you can use the RGB modification tool to add color to images that didn't have color to begin with. For instance, you can “pseudo-color” greyscale media.

Inverse Telecine - the reverse of 3:2 pulldown where frames, which were duplicated to create 60-fields/second video from 24-frames/second film, are removed. MPEG-2 video encoders usually apply an inverse Telecine process to convert 60-fields/second video into 24-frames/second encoded video. The encoder adds information enabling the decoder to recreate the 60-fields/second display rate. Telecine, i.e. 3:2 pulldown, is used to transfer film to video. That's where the 3:2 ratio comes in. To ensure that there will consistently be 60 frames per second, the first frame is displayed on the TV screen 3 times and the second frame is displayed 2 times. The following frame is repeated 3 times, the next one 2 times, etc. throughout the film. For inverse telecine, you show 2 of the film frames for 3 of output frames. Only check “Inverse Telecine” if you have film or something that is 24fps and want to project to 30fps (most likely this will never be necessary).

Audio 5.1 – 6 channel surround sound. For most home systems, uses five full bandwidth channels and one low-frequency effects channel. Could automatically get set as explained previously.

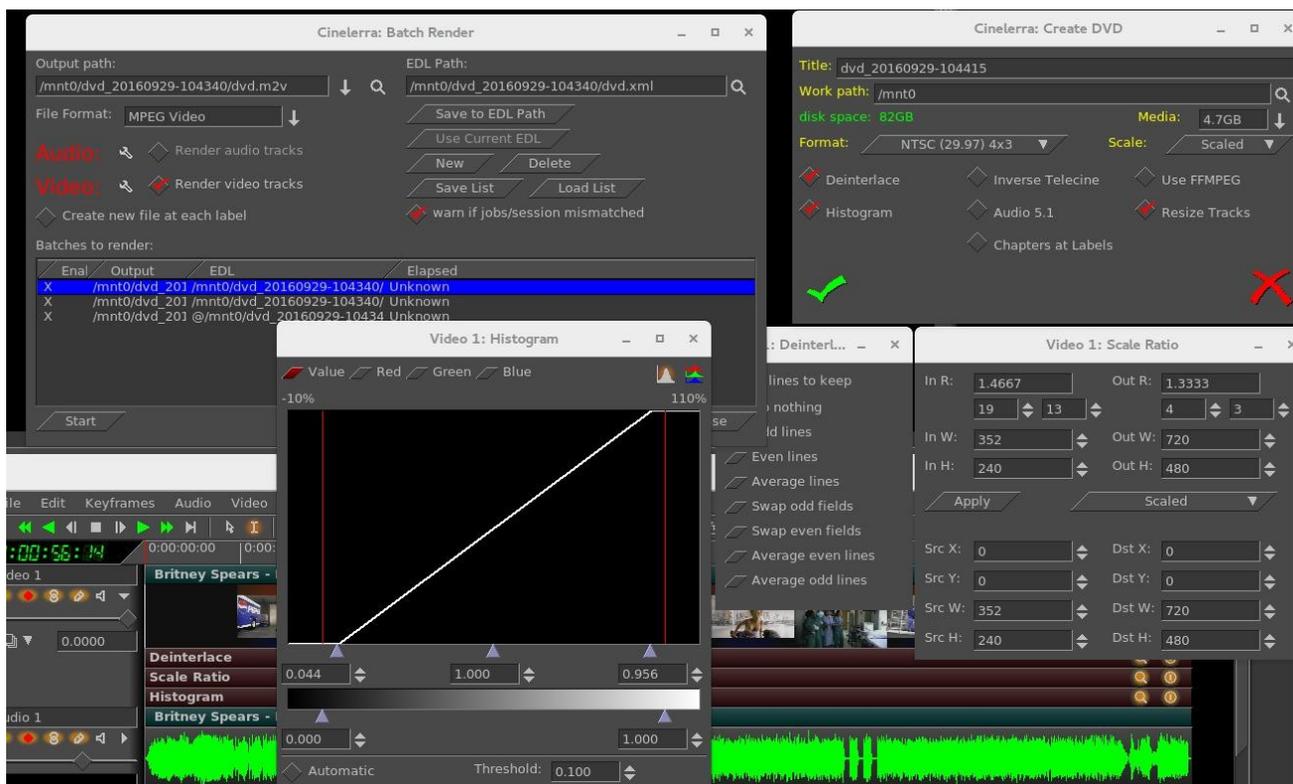
Aspect Ratio – aspect ratio may be automatically set to 4:3 or 16:9. Aspect ratio would better be defined as the size of the display, monitor, or TV which will be used to view the output. If you measure your old TV, which supposedly is 4:3 and your latest digital TV, which is supposedly 16:9, you will see that those ratios aren't always correct anyway. Then measure your laptop monitor, your desktop monitor, and your neighbor's, and lo and behold, the ratios don't fit either of the purported “standard” aspect ratio. Maintaining square pixels via scaling is more important in the long run.

Use FFMPEG – this is user's choice; it is recommended and faster but more difficult to modify due to numerous options. For blu-ray, ffmpeg must be used and is not an available option.

Resize Tracks – change track width and height as explained previously. The size is adjusted to the largest frame size needed.

Chapters at Labels – without this checked, chapters markers are automatically inserted every 5 minutes. The chapter labels can then be “skipped to” when playing the DVD. If instead, you want to put labels in at opportune times, you will have to run dvdauthor outside of Cinelerra and mark the chapter labels yourself by hand. In that case, you should check “Chapters at Labels” so that the automatic 10-minute labels are not created. This checkbox is not currently available for blu-ray.

The screenshot below shows on the upper right side the “Create DVD” window with Deinterlace, Histogram and Resize Tracks checked. Also Scale is set to “Scaled”. Once the green checkmark is clicked, the “Batch Render” window comes up and in the main window you will see the 3 plugins below the video track on the main window. By clicking on the magnifying glass that appears on the rightmost side, the controls for each will popup and you can make any necessary adjustments. Note the numerous choices for Deinterlace; the Value, Red, Green, and Blue for color adjustments in the Histogram window; and Scale Ratio popup menu for numerical settings control.



The Scale parameter gives you a lot of flexibility. A default based on your input media is provided for you but possible choices are None, Scaled, Cropped, Filled, Horiz Edge, and Vert Edge. You will have the opportunity to manipulate the desired results in the “Scale Ratio” window. Values for W (width), H (height), and X/Y coordinates are the number of pixels. For example, if video is 720 x 432, that is obviously 720 pixels by 432 pixels and this would be the values for Dst W and for Dst H. So if you have some media that is off center you can crop by changing the SRC Y value AND then change DST X/Y to non-zero. It will become the output origin. To see what it does, change them from 0.0 to 400.0 and you will see big changes in the compositor window.

For example, if you have the Cropped choice for the Scale, you will want to manipulate the ScaleRatio plugin (the magnifying glass on the main window video track) which brings up the Scale window. For cropped top instead of crop both top and bottom, modify the "Src Y". As you change the Y scale, you will see the cropping take place in the Compositor.

Scaling options are provided in order to preserve image aspect ratio. To determine which scaling option to use, it is important to correctly identify your source/destination video aspect ratios. Next is a short explanation of possible options.

None – do not do any scaling. The destination output matches the source input. There is no resizing.

Scaled – do not use uniform scaling on X and Y. Just make it fit and you will might end up with squishing/stretching in one direction or the other in order to do that. This happens when the input aspect ratio is different than that displayed on the output.

Cropped – removes outer edges of a source image in order to fit the image on the output display. This is done in order to maintain uniform geometry scale on the destination display without being affected by the particular factor of original media aspect ratio. Since cropping omits image area, and the areas which you wish to view may not be viewed when the image is centered, you can pan the image source using the Src X/Y to modify which area is in view. You can also think of this as “scaled up”.

Filled – the entire output display will be filled with image content but in order to do so, some of the image may have had to be cut. Any mismatch between the two will be filled with black. Centering of the image can be modified by using the Dest X/Y variables in the Scale Ratio controls. You can also think of this as “scaled down”.

Horiz Edge – this indicates preserving of the Horizontal edge.

Vert Edge – this option preserves the Vertical edge.

Horizontal and Vertical are duplicates or restatement of same functionality as Cropped or Filled but are provided as options to accommodate different ways of thinking. In any case, you can choose which outer edges of the image to crop by using the “Show Controls” of the Scale Ratio plugin. For example, you can ensure that no “action” is lost by displaying the center of the screen only or making sure that any textual information on the bottom is not lost by cutting only the top off.

Step 3:

3) Batch render menu appears with “m2ts” format selected for blu-ray or “dvd” format selected for regular/standard DVD when File Format selected is ffmpeg in the previous “Create DVD” menu. It will work just fine without selecting ffmpeg for DVD and may be advantageous not to. Using the audio /

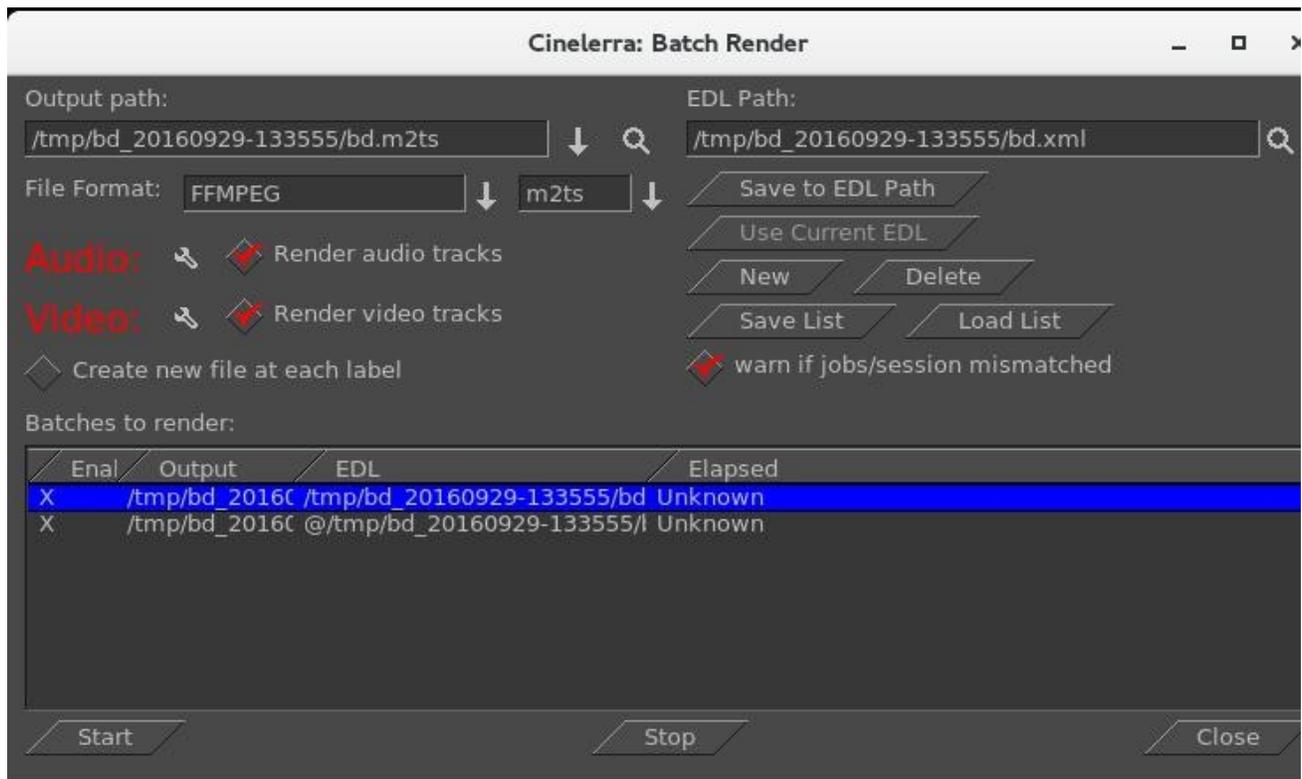
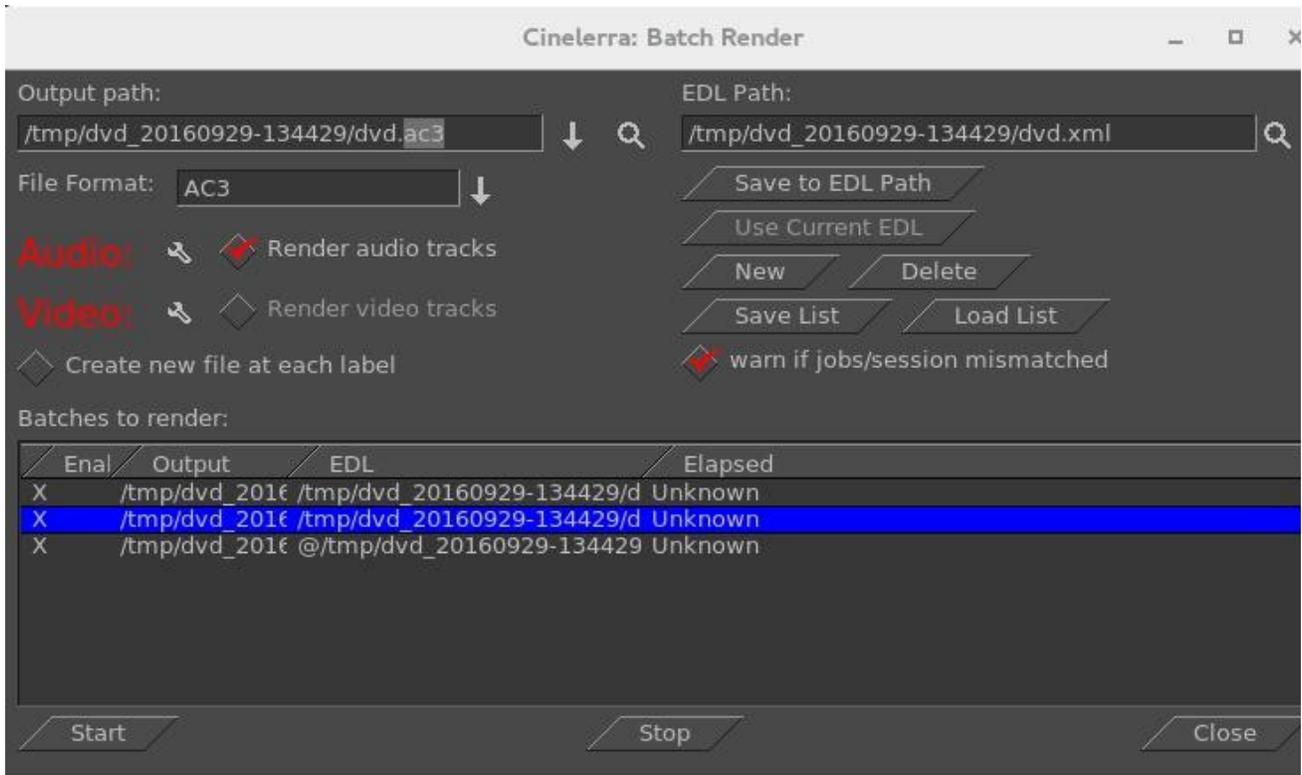
video wrench tools (you will have to have the video batch job highlighted to manipulate the video or audio batch job for audio).

Setup the audio bitrate (192000 recommended). Data rate is 192K default.

Setup the video bitrate (6000000 - 12000000 recommended). 10Mb/sec is the current default. Click OK check button.

The default bitrate is the largest value possible. The actual 'target' bitrate is calculated based on a formula from the blu-ray/DVD code. It divides the media size (in bits) by the video time (in seconds) to find the bitrate that will "just fit" on the target media. This could create a weird bitrate if the media is large, and the video time is small, so the default/target bitrate is limited to 10Mb/s. Batch jobs are then built and appended to the job list. Once these batch jobs are built, if you make any changes, you have to start over. You will see listed the batch jobs that are created to perform the rendering/tasks – 2 jobs for blu-ray, one for audio/video processing and one for the scripts. There will be 3 batch jobs created for DVD when not using ffmpeg to include one each for audio and video and then one more for the scripts. When you click on “start”, it fires off those jobs and you will see the rendering main window progress bar in the bottom corner reflecting the fact that it is currently processing. Be aware that the Cinelerra program will be totally shutdown AFTER the batch jobs finish. Screenshot below shows BD blu-ray creation with 2 batch jobs queued up and ready to go.

Screenshot below shows a DVD creation render ready to be started. Note that because it is not ffmpeg, the processing of video is done separately from audio. If you need to modify the video tracks, which you can see is ghosted out, you need to highlight the first batch job listed. Be sure to highlight the first batch job before pressing Start so it runs all of the jobs and be sure the X for job enabled is set.



When you click on “start”, it fires off those jobs and you will see the rendering main window progress bar in the bottom corner reflecting the fact that it is currently processing. Be aware that the Cinelerra program will be **totally shutdown after the batch jobs finish** and you will be at the command line prompt.

This will produce a new directory in your target path which contains a filesystem image file.

For example: /TargetDirectory/bd_20150820-093747

Directory and file names should not be changed at this time because the scripts and programs rely on the given names in order to proceed. You can change them later for your own purposes.

If bluray to test the filesystem you just created, use the command line interface; loopback mount the filesystem image which was generated in the target directory. For example if blu-ray:

```
cd /TargetDirectory/bd_20150820-093747/  
mount -o loop,ro ./bd.udfs ./udfs  
(check the media using a compatible media rendering tool like ffplay)  
umount ./udfs
```

To burn blu-ray media you will need to run from the command line interface. In the examples below, /dev/bd represents your blu-ray writer device (for example: /dev/sr1) and /dev/dvd represents your dvd writer device (for example: /dev/sr0).

Blu-ray Media

For rewritable blu-ray (recommended) (BD-RE):

Note - unwritten (virgin) media must be formatted first using:
dvd+rw-format /dev/bd (only done once and does not take very long)

To write or rewrite rewritable blu-ray media:

```
cd /TargetDirectory/bd_20150820-093747/  
dd if=./bd.udfs of=/dev/bd bs=2048000 (the growisofs command below also works)
```

To write blu-ray write-once media: (BD-R) (no pre-formatting needed):

```
cd /TargetDirectory/bd_20150820-093747/  
growisofs -dvd-compat -Z /dev/bd=./bd.udfs
```

DVD Media

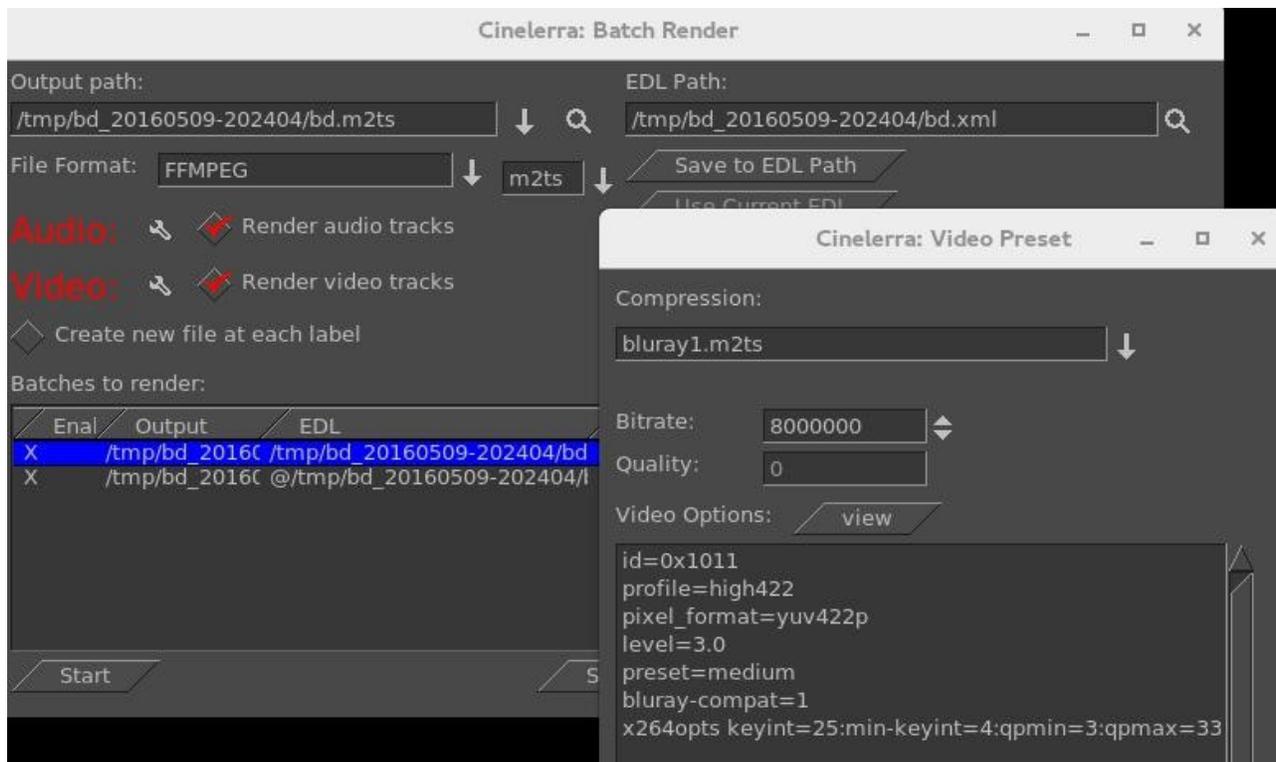
For rewritable DVD (DVD+RW):

Note – unwritten (virgin) media must be formatted first using:
dvd+rw-format /dev/dvd (only done once and does not take very long)

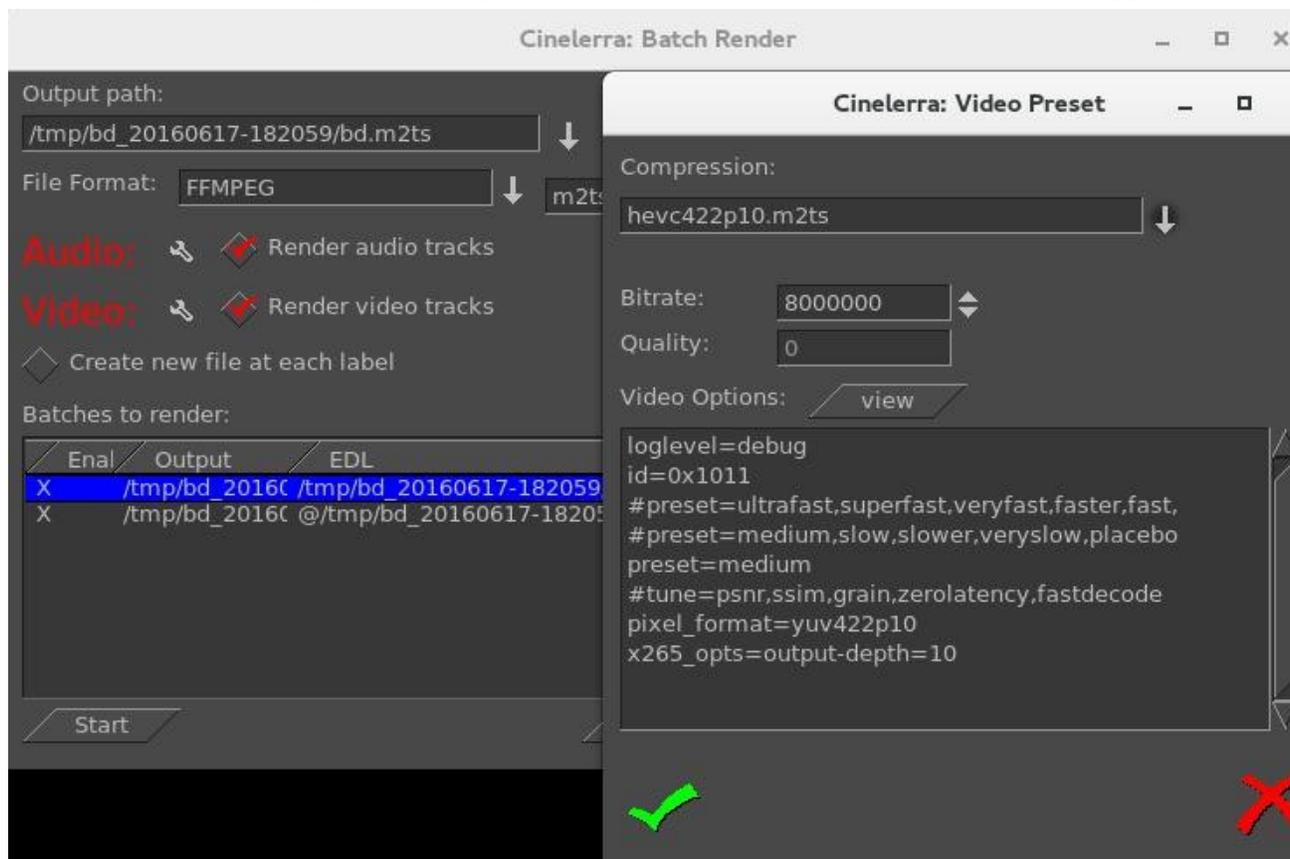
To write a DVD, load blank media and run the following from the command line (requires dvdauthor):

```
cd /TargetDirector/dvd_20160404-175416  
growisofs -dvd-compat -Z /dev/dvd -dvd-video ./iso
```

Screenshot showing the availability of 4:2 :2 for a Batch Render seen by clicking on wrench icon.



Screenshot showing the availability of 10-bit high quality 4:2 :2 for a Batch Render seen by clicking on wrench icon. You need specially compiled cinelerra in order to use the x265 10-bit as opposed to 8-bit.



13.4.1 Output Terminal Messages from Creating DVDs

Below are examples of what the batch jobs generate and you will see on the terminal screen if you started the cinelerra program in the recommended manner from a terminal window. It is just informational but will let you know if errors. In looking at any of the output, you can safely ignore the errors that read “Unsupported codec with id 100357 for input stream 0” - this comes from “nav-data” (navigation data). The first 2 examples are seen from running the batch jobs; the last 2 are from the single line execution which records the media output to the DVD hardware.

SD Example – Partial Output during Cinelerra run

```
...
FFMPEG::open_decoder: some stream times estimated
Render::render_single: Session finished.
FFMPEG::open_decoder: some stream times estimated
Render::render_single: Session finished.
running /tmp/dvd_20160407-113530/dvd.sh 1 /tmp/dvd_20160407-113530
  INFO: [mplex] mplex version 2.1.0 (2.2.7 $Date: 2012/11/17 01:55:16 $)
  INFO: [mplex] File /tmp/dvd_20160407-113530/dvd.m2v looks like an MPEG Video stream.
...
  INFO: [mplex] MUX STATUS: no under-runs detected.
DVDAuthor::dvdauthor, version 0.7.1.
Build options: gnugetopt graphicsmagick iconv freetype fridibi fontconfig
Send bug reports to <dvdauthor-users@lists.sourceforge.net>

INFO: default video format is NTSC
INFO: dvdauthor creating VTS
STAT: Picking VTS 01

STAT: Processing /tmp/dvd_20160407-113530/dvd.mpg...
STAT: VOB 32 at 15MB, 1 PGCs
INFO: Video pts = 0.133 .. 22.789
INFO: Audio[0] pts = 0.133 .. 22.789
STAT: VOB 46 at 21MB, 1 PGCs
CHAPTERS: VTS[1/1] 0.000
INFO: Generating VTS with the following video attributes:
INFO: MPEG version: mpeg2
INFO: TV standard: ntsc
INFO: Aspect ratio: 16:9
INFO: Resolution: 720x480
INFO: Audio ch 0 format: ac3/6ch, 48khz drc, 'en'

STAT: fixed 46 VOBs
INFO: dvdauthor creating table of contents
INFO: Scanning /tmp/dvd_20160407-113530/iso/VIDEO_TS/VTS_01_0.IFO
To burn dvd, load blank media and run:
growisofs -dvd-compat -Z /dev/dvd -dvd-video /tmp/dvd_20160407-113530/iso
```

BD Example – Partial Output during Cinelerra run

```
...
FFMPEG::open_decoder: some stream times estimated
Render::render_single: Session finished.
+ PATH=/usr/lib64/ccache:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/root/bin:/mnt0/build5/
cinelerra-5.1/bin
+ mkdir -p /tmp/bd_20161224-162059/udfs
++ du -sb /tmp/bd_20161224-162059/bd.m2ts
++ sed -e 's/[ ].*//'
+ sz=19811904
+ blks=13769
+ mkudffs /tmp/bd_20161224-162059/bd.udfs 13769
start=0, blocks=16, type=RESERVED
start=16, blocks=3, type=VRS
start=19, blocks=237, type=USPACE
start=256, blocks=1, type=ANCHOR
start=257, blocks=16, type=PVDS
start=273, blocks=1, type=LVID
start=274, blocks=13238, type=PSPACE
start=13512, blocks=1, type=ANCHOR
start=13513, blocks=239, type=USPACE
start=13752, blocks=16, type=RVDS
start=13768, blocks=1, type=ANCHOR
+ mount -o loop /tmp/bd_20161224-162059/bd.udfs /tmp/bd_20161224-162059/udfs
+ bdwrite /tmp/bd_20161224-162059/udfs /tmp/bd_20161224-162059/bd.m2ts
+ umount /tmp/bd_20161224-162059/udfs
+ echo To burn bluray, load writable media and run:
To burn bluray, load writable media and run:
+ echo for WORM: growisofs -dvd-compat -Z /dev/bd=/tmp/bd_20161224-162059/bd.udfs
for WORM: growisofs -dvd-compat -Z /dev/bd=/tmp/bd_20161224-162059/bd.udfs
+ echo for RW: dd if=/tmp/bd_20161224-162059/bd.udfs of=/dev/bd bs=2048000
for RW: dd if=/tmp/bd_20161224-162059/bd.udfs of=/dev/bd bs=2048000
```

SD Example – Partial Output during writing disc media

```
# growisofs -dvd-compat -Z /dev/sr0 -dvd-video /tmp/dvd_20161224-160756/iso
WARNING: /dev/sr0 already carries isofs!
About to execute 'mkisofs -dvd-video /tmp/dvd_20161224-160756/iso | builtin_dd of=/dev/sr0
obs=32k seek=0'
I: -input-charset not specified, using utf-8 (detected in locale settings)
75.62% done, estimate finish Sat Dec 24 16:09:51 2016
Total translation table size: 0
Total rockridge attributes bytes: 0
Total directory bytes: 4096
Path table size(bytes): 42
Max brk space used 1a000
```

```
6624 extents written (12 MB)
/dev/sr0: "Current Write Speed" is 4.1x1352KBps.
builtin_dd: 6624*2KB out @ average 0.7x1352KBps
/dev/sr0: flushing cache
#
```

BD Example – Partial Output during writing disc media

```
# growisofs -dvd-compat -Z /dev/sr0=/tmp/bd_20161224-155658/bd.udfs
Executing 'builtin_dd if=/tmp/bd_20161224-155658/bd.udfs of=/dev/sr0 obs=32k seek=0'
/dev/sr0: "Current Write Speed" is 2.0x4390KBps.
 1605632/24524800 ( 6.5%) @0.0x, remaining 1:39 RBU 100.0% UBU  0.0%
 1605632/24524800 ( 6.5%) @0.0x, remaining 2:22 RBU 100.0% UBU 100.0%
 1605632/24524800 ( 6.5%) @0.0x, remaining 3:05 RBU 100.0% UBU 100.0%
 5865472/24524800 (23.9%) @0.3x, remaining 0:54 RBU 100.0% UBU  33.3%
11829248/24524800 (48.2%) @0.4x, remaining 0:21 RBU  75.8% UBU  37.5%
17858560/24524800 (72.8%) @0.4x, remaining 0:08 RBU  39.8% UBU  79.2%
23789568/24524800 (97.0%) @0.4x, remaining 0:00 RBU  4.5% UBU  4.2%
builtin_dd: 11984*2KB out @ average 0.2x4390KBps
/dev/sr0: flushing cache
```

13.4.2 Helpful Hints for Problems/Debugging DVDs Creation

This section contains helpful hints, how to initially check the results, and some information on determining what might have gone wrong and how to address it.

- 1) For first time users, taking the defaults seem to work very well when running as root.
- 2) You may want to use rewritable media to see how it goes before using permanent media.
- 3) Until you are familiar with the procedure, start with shorter input in order not to waste time.
- 4) Test the generated output with a compatible media rendering tool before burning DVDs.
- 5) Check the list of files and file sizes after the batch jobs are complete before burning DVDs.

For blu-ray creation, cd /workpath/bd_date-time directory and look for similar files:

```
bd.jobs      bd.m2ts      bd.sh        bd.udfs      bd.xml
```

udfs directory which is used as a loopback mount point

Note that the size of bd.udfs should be larger than bd.m2ts because this is the final file which is actually going to be written to the disc media. It contains contents of bd.m2ts and all of the required disc structure.

For DVD creation, cd /workpath/dvd_date-time directory and look for similar files:

```
dvd.ac3      dvd.jobs     dvd.m2v      dvd.mpg      dvd.sh      dvd.xml
```

iso directory with VIDEO_TS and AUDIO_TS subdirectories of non-zero size

Note that there will be no files in the actual AUDIO_TS directory.

- 6) The bd.sh and dvd.sh files are script files that you can carefully run manually from some start point to determine where the failure occurred. You must BE CAREFUL and know what you are doing and what directory you are in because dvd.sh contains an “rm” command and will delete files. The script

takes a command line parameter of the directory where the file was rendered to and which is usually the directory where dvd.sh or bd.sh was created.

7) There is also a file in the same directory, called bd.jobs. It was the information that was used in creating the batch jobs and may be helpful in determining what parameters were actually used if there are any resulting problems. With enough background knowledge, you can make changes and rerun.

8) For blu-ray check to make sure you do not have any spurious loopback disks mounted that may interfere with the correct generation. Use the df command to check this and then the umount command to unmount these. Also, check to make sure you have used the gsettings command to disable automount.

9) For blu-ray loopback mount the <target>/bd.udfs image, and see if it has the BDMV filesystem written to it, and in particular a subdirectory named STREAM. Look at the results in ./udfs and check for the stream file which should exist in ./udfs/BDMV/STREAM/00000.m2ts and should have the same size as ./bd.m2ts.

```
mount -o loop <target>/bd.udfs <target>/udfs
```

```
ls -lR <target>/udfs
```

```
du -sc <target>/udfs
```

```
umount <target>/udfs
```

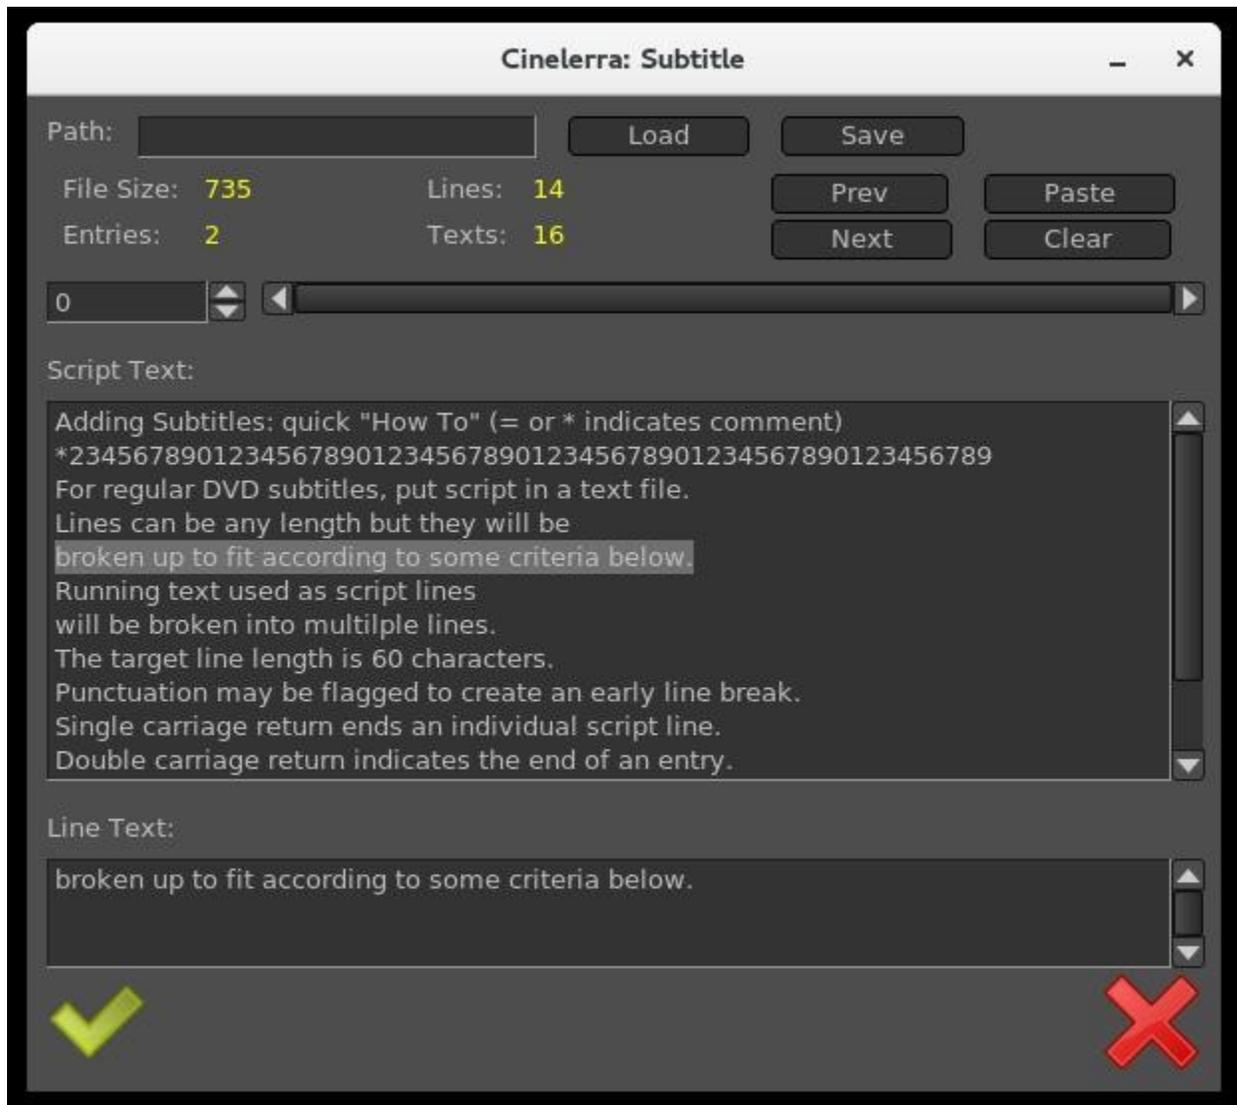
Checklist for Troubleshooting BD and DVD Render:

- ___ Are you logged in as root? This is required in order to loopback mount files for bluray and to write media on /dev/hardware. See section 8.3.5 for a workaround for normal user mode.
- ___ Did you startup cinelerra from a terminal window so you can see informative messages?
- ___ Is udftools installed for BD and dvdauthor installed for SD?
- ___ Do you have loopback not enabled for bluray? At least temporarily, disable automount via:
gsettings set org.gnome.desktop.media-handling automount false
- ___ Did you have sufficient disk space for working/writing files? In the Create window, the disk space will be displayed in green if sufficient, but in red if less than what fits on the disc media.
- ___ Did you use /tmp as the work device, then rebooted the computer, which deleted files on /tmp?
- ___ If the input media is interlaced, did you check the Deinterlace option to eliminate interlacing?
- ___ Did you change the output name in the Batch Render window after the batch jobs were already created? These filenames have already been written to disk. If you want to change either the Title or the Work_path, you have to start over.
- ___ Have you selected a Title in the Create window that is a directory that already exists? The program attempts to create that directory and will give you an error message if it exists.
- ___ Did you replace the /dev/bd or /dev/dvd on the command line with your hardware device name?
- ___ If a warning was issued in the Create BD/SD window of “* non-standard format” and your bluray reader could not play the disc, did you change to a standard format instead?
- ___ Did you correctly interpret the frame rate if using interlaced format to be ½ due to interlacing?

13.4.3 Subtitles

DVD (not blu-ray... yet) subtitles are added by using the main window pulldown File → Subtitle which brings up a window allowing you to type the filename of a previously generated text file containing the desired words/lines, the script. After entering the filename, click “Load” to read in your script. By creating a script file ahead of time, it lets you easily add dialog that was already written out and carefully edited for spelling and proper grammar.

Screenshot below shows the Subtitle window you will see.



The format of the script/text input file has specific requirements as listed below:

- Lines can be any length but they will be broken up to fit according to some criteria below.
- Running text used as script lines will be broken into multiple lines.
- The target line length is 60 characters.
- Punctuation may be flagged to create an early break.
- Single carriage return ends an individual script line.

Double carriage return indicates the end of an entry and helps to keep track of where you are. The “=” sign in column 1 indicates a comment seen in the script text to assist you in location. An “*” at the beginning of the line is a comment and not a script line. Whitespace at either the beginning of a script line or the end will be removed.

To put the subtitles onto your media, first add a subtitle track via the pulldown **Track → Add Subttl**. In the Subtitle window, note that there are 2 major textboxes. There is the “Script Text” textbox showing the current entry of text from your input file and there is the “Line Text” textbox showing the currently active text. In your subtitle track, select a timeline region (in/out or drag select with hairline cursor/highlight) to indicate the region where you want the active Line Text to be pasted. Then click the “Paste” button in the Subtitle window to paste the line onto the subtitle track. Silence will be added to the subtitle track in the places in the media where there are gaps.

Editing in the Line Text box can be used to change the active script line. By double clicking the timeline over the subtitle track, you can reselect the active script line. The subtitle text will be reloaded into the Line Text box and can be edited and re-pasted as the new active subtitle text. You can also highlight multiple lines in the Script Text box and paste them (using the usual window paste methodology) into the Line Text box. After pasting to the timeline, the Line Text box will be updated with the next script line. In addition, if you triple click a line in the “Script Text” box, it will automatically become the current line in the “Line Text” box.

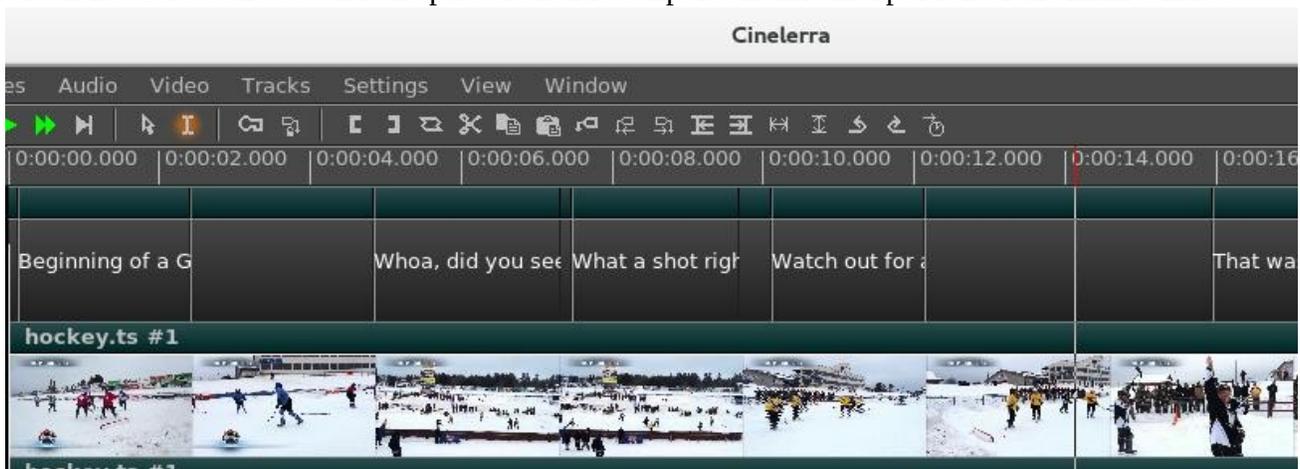
When you are finished, before clicking “Save”, you must supply a legitimate filename in the “Path” box; your current directory will be used if only a filename but no directory path is supplied. The filename used will automatically have a “-” after it followed by the “track label” and then “udvd” extension added; any extension in the filename will be removed. If you click OK before saving, the subtitle script position is saved with the session. This is convenient for continuing where you left off.

To reposition the script, use the slider or tumbler buttons:

Slider bar to move through the text entries quickly.

Prev or Next buttons to go to the previous or next script line.

Screenshot below shows what the pasted subtitle script looks like in a portion of the main window.



13.4.4 Dvd Interlaced Chroma

Cinelerra uses 4:4:4 colorspace to edit, so it is necessary to convert interlaced 4:2:0 video to 4:4:4. But you can run into problems, referred to as the “chroma bug”, which you see in DVD media displayed on higher resolution monitors – streaks or spiky horizontal lines are visible in the chroma channel, especially on diagonal edges. The Chroma Bug is specific to MPEG and 4:2:0 encoding. Now you can use the “YUV420P DVD Interlace Mode” when rendering DV directly to mpeg2 through a yuv4mpeg stream and when using video effects on HDV video.

With this option enabled, improved chroma results will be obtained from your DV or HDV source. Editing DV or HDV and rendering it back to the same format does not require any special handling. In order to perform colorspace conversions correctly in Cinelerra and avoid Chroma errors for interlaced 4:2:0 video, check the box as follows:

Settings → Performance → YUV420P DVD Interlace Mode

This option maintains the interlacing in Chroma sample addressing, which ordinarily would be deleted because the upsampling of interlaced chroma fields is normally done using a progressive algorithm. With this mode enabled, the MPEG decoder uses a different algorithm for interlaced frames so that the 4:2:0 format chroma interlacing is preserved.

13.4.5 MPEG utility programs

There are 2 utility programs that come in handy when dealing with DVD media for creating or reading previously written DVDs.

1) `zmpeg3cc2txt` – convert closed captioning data to subtitle data.

This program can be used to scan captured broadcast data streams and convert the closed captioning text data into subtitle track data. The result can be used to add subtitles to DVDs created from the captured media using the text collected from the closed captioning.

```
usage: ./zmpeg3cc2txt [-c cc_service ] [-s start:length, ...] [ -t track ] [-v verbose ]
        [-w wdw mask ] [-x file.xml ] [-o file.udvd ] file.ts
```

<code>cc_service</code>	= closed caption service id
<code>start:length</code>	= start:length frames (comma separated list)
<code>track</code>	= video track number
<code>verbose</code>	= verbose level
<code>wdw mask</code>	= bit mask for windows (-1 = all)
<code>file.xml</code>	= filename for edl xml format subtitle data
<code>file.udvd</code>	= filename for udvd format subtitle data
<code>file.ts</code>	= filename for transport stream

To use this program, the input file must be a transport stream (broadcast video) which contains closed captioning services. The service id defaults to one, and the default video track is zero. Either -o or -x must be specified to indicate the output file format desired. If the output file name is a '-' then stdout is selected as the output file.

For example:

```
zmpeg3cc2txt -o - /dvb_data/channel5.ts
```

2) zmpeg3ifochk – check DVD ifo file for usable features.

For some time, DVD manufacturers have been employing a variety of measures to make reading a DVD difficult on a computer. One technique which is widely deployed is to add a bunch of extra program data, so that correct playback is only likely if you can read the DVD virtual machine data and decode a maze of program data to find the undamaged stream definitions. Only a few streams are created which are machine usable, and dozens are created as decoy streams. The decoy streams fail or introduce errors. This program scans the IFO (info file) playlists and verifies the contents of the stream that does not contain obvious damage. The result is a list of program ids which can be entered into the playback preferences to select a program which qualifies.

13.4.6 HDV-sourced Video Recording on a Blu-ray Disc Without Re-encoding

An MTS file is a video file saved in the high-definition (HD) MPEG Transport Stream video format, commonly called "AVCHD". It contains HD video compatible with Blu-ray disc format and is based on the MPEG-2 transport stream. MTS files are often used by Sony, Panasonic, Canon and other HD camcorders. Legal input for Video – MPEG1VIDEO, MPEG2VIDEO, H264; Audio – MP1, MP2, AC3, AC3PLUS, DTS, TRUHD.

For creating a blu-ray disc, if you have HDV MPEG-2 media that is in blu-ray format, you can save the original quality of your work, rather than rendering it to another format. Follow the steps below directly instead of going through cinelerra. It has been tested on 10 different MTS files.

- 1) `du -sb /yourHDVfile.MTS` ! Determine the size of your file in bytes.
- 2) `blocks=((size-in-bytes/2048 + 4096))` ! Convert bytes into blocks + a little more.
- 3) `mkudffs /tmp/newfilename.udfs blocks` ! Create a file with that # of blocks + some extra.
- 4) `mount -o loop /tmp/newfilename.udfs /mntX` ! Use a mount point like mntX that is not in use.
- 5) `<cinelerra_installed_path>/bin/bdwrite /mntX /tmp/yourHDVfile.MTS` ! Substitute cinelerra path.
- 6) `umount /mntX` ! You must unmount the udfs filesystem
- 7) `growisofs -Z /dev/bd=/tmp/newfilename.udfs` ! Replace /dev/bd with your bluray hardware device.
OR `dd if=/tmp/newfilename.udfs of=/dev/bd bs=2048000` ! if using rewritable blu-ray; replace bd.

13.4.7 User-mode Blu-ray Render Workaround for Mount/Umount

Creating BD images to be written to media requires usage of “mount” and “umount” which typically can only be done by the root user due to security. If you want to avoid running cinelerra as root, you can implement a workaround by adding a line in /etc/fstab (must be root to edit the file initially) and by creating a directory in your home area, called “bluray”. You only have to do this once unless you upgrade the Operating System and it wipes out the line in /etc/fstab. Now the cinelerra program will

automatically do the mount and umount for you each time you execute BD Render and you can run as an ordinary user.

The line to add to /etc/fstab will look something like the following, assuming your username is “name” and your groupid may be “users” or “name”. If you do an “ls -l” in your home directory, the 3rd and 4th fields shown will be your uid or name and gid or groupid which you must substitute in the line below.

```
/home/name/image /home/name/bluray udf noauto,loop,rw,user,uid=name,gid=groupid 0 0
```

Also, be sure to do a “mkdir bluray” in your /home/name directory as this is a requirement (owned by you; uid=gid=name). When the actual image to be written to disc media is created, it will first delete any current /home/name/image file. Warning – make sure you do not already have a file called “image” that you want to save as it will be automatically deleted every time you initiate a BD Render. So you will want to burn a bluray disc after cinelerra creates the “image” since it will be written over on the next rendition. The actual writing to your bluray burner (something like /dev/sr0) is done outside of cinelerra at a terminal prompt and requires root privilege usually. You can either use “sudo” for 1 line or create user wheel group to get around this.

13.4.8 How to Create Bluray Media from Multiple Cinelerra Output

Writing prepared multiple cinelerra output files, “bd.m2ts”, to a single bluray disc is relatively easy to do but is not done automatically. You can render all of the desired files via the Create BD menu, save each individual bd.m2ts file with a unique name, construct a Menu Title that reflects the contents of each of these files, then manually use a few commands to create a udfs file to be written to BD.

Usage of the final preparation taken from the bdwrite program comments:

```
./bdwrite <tgt_dir_path> <playlist-0> <sep> <playlist1> <sep> ... <sep> <playlist-n>
```

```
<sep> == -<pgm_pid> | --<pgm_pid> | ---<pgm_pid>
```

```
<pgm_pid> may be empty string, or a numeric pgm_pid for current title clip
```

```
<pgm_pid> defaults to first pgm probed
```

```
<playlist-x> == <clip-0.m2ts> <clip-1.m2ts> ... <clip-n.m2ts>
```

One title is built for each playlist; playlist-0 is used as first-play item. The separators (<sep> represented by the dash character) have unique roles. The double “--” means stop after playing, and the triple “---” means pause.

For example:

```
./bdwrite /tmp/dir /path/menu_titles.m2ts --- /path/clip0.m2ts -- /path/clip1.m2ts -- /path/clip2.m2ts
```

The basic idea is to use playlist-0 as a menu or directions to use the bluray player remote control to select the desired Title and start the play, avoiding the need for a menu system. Planning in advance to get the desired results is necessary. The following steps provide an outline to get started.

- 1) Create all of the bd.m2ts files that you want to put on the Bluray.
- 2) Using cinelerra, design your Title page using a few seconds of video and the Title plugin.

Screenshot to the right demonstrates an example of setting up a Title menu on a 5 second video. There is a list of 4 menu title options that can be searched via remote control using the Title search option for your player.



3) Use BD Create to render your short Title video.

4) Next is the most complicated part which is to run mkudffs with a sufficient amount of disk space to hold all of the bd.m2ts files “plus a little more!” To calculate this, you can record the sizes from having run BD Create mkudffs. This number is displayed on the terminal screen when using the command line interface each time and add them together. Or recalculate the size of each bd.m2ts using the formula below and adding them all together. This is the number of blocks used to make a bluray image space for bdwrite to use. For many files, this could require a huge amount of space, so check first.

Total size = File#0 size in bytes / 2048 + 4094 “+” File#1 size in bytes / 2048 + 4094 “+” ...

Now create the image file via: mkudffs image <Total size> where image or udfs is the image name.

5) Loop mount the disk image (refer to Sections 8.3.4 and 8.3.5).

6) Then actually write your multiple bd.m2ts type files onto the “image” where <cin_path> is the location of the cinelerra binary “bdwrite” file and <path> is your directory path. Below is a single line that wrapped around with 4 Titles.

```
<cin_path>/bin/bdwrite image /<path>menu_titles.m2ts --- /<path>/bd1.m2ts -- /<path>/bd2.m2ts --  
/<path>/bd3.m2ts -- /<path>/bd4.m2ts
```

Note that the 3 dashes after the “menu_titles.m2ts” lets the bluray player know to “pause” after playing the few seconds video which contains the index to the rest of the files. The 2 dashes after each of the bd.m2ts signify “stop”. That is when you will have to use your remote to “Search Titles” in order to play the next one you want to see. In addition, if for some reason you just want to “play all”, you will have to add another line to the Title menu as a choice and list all of the 4 files in a row at the end of your bdwrite line without any dashes in between.

7) Umount the loopback disk.

8) Use your favorite dd or growisofs tool to write to a formatted bluray disc.

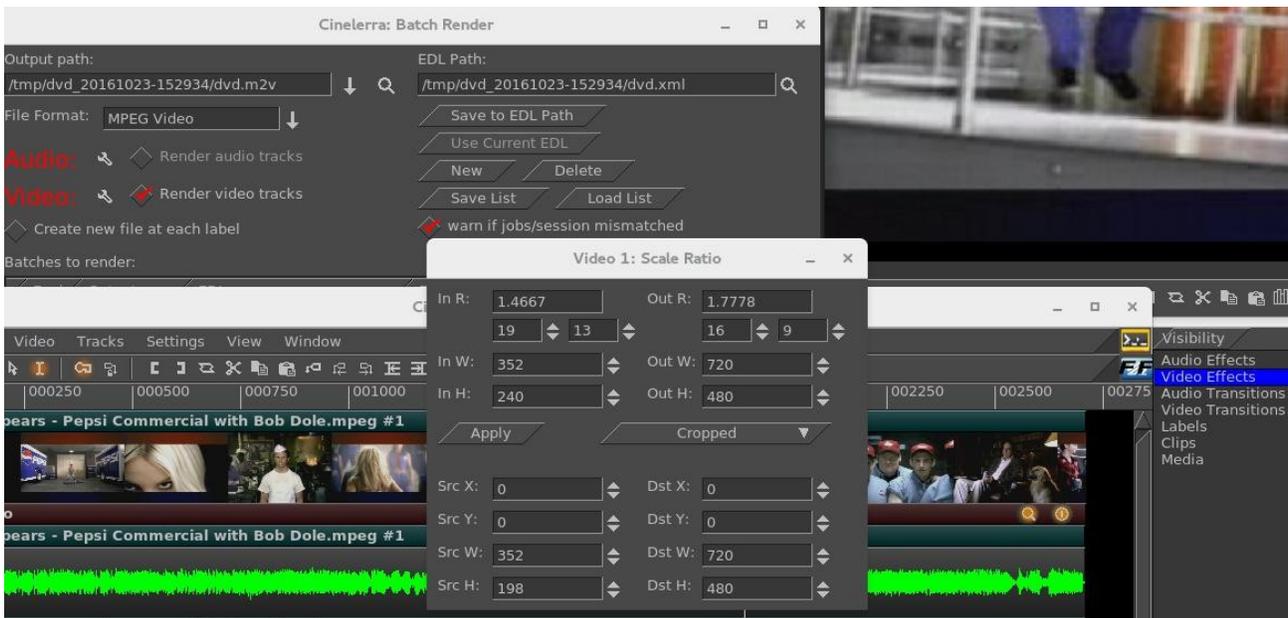
13.5 Complete Example of Video Source with 4:3 Aspect Ratio, Being Transcribed to 16:9 and Creating a DVD to be Displayed on a Digital TV

Illustrated steps to take source input with 4:3 aspect ratio and convert to 16:9, with the bottom of the image being cropped in order to preserve top of video so nobody's head gets cut off are provided here.

- 0) In order to write to a DVD writer hardware device, you must be “root”!!
- 1) Start Cinelerra-GG to bring up the 4 usual screens with main track canvas in the lower left corner.
- 2) Load media via the pulldown File → Load files... by choosing the directory path with the desired file.
- 3) Bring up the “Create DVD” window using File → DVD Render. This window is shown below.
- 4) Choose Format: PAL or NTSC with 16x9 aspect ratio for today's digital TV as in below screenshot.
- 5) Modify the “Work path:” parameter to a disk that has sufficient disk space and you will see the amount of disk space in green letters (/tmp is default, but is often deleted so may be a bad choice).
- 6) Note that in the following screenshot, Scale of “Cropped” has been chosen.
- 7) Click the green checkmark on the bottom left side of the window to close it and proceed.



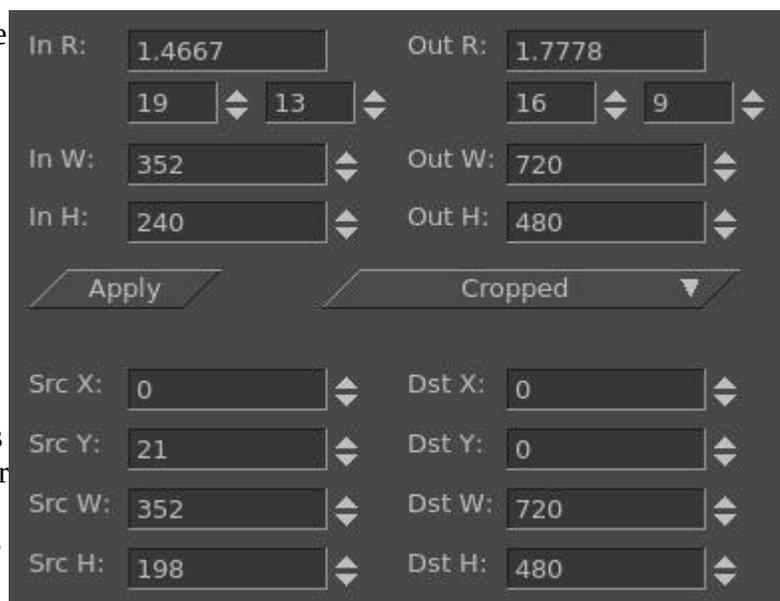
- 8) Now the “Batch Render” window will appear along with the Scale Ratio brown-colored line below the video in the main track canvas as seen in the screenshot below. Note that in this screenshot the top right most corner of the shot displays the bottom portion of the Compositor window.
- 9) Next, right mouse click the gold-colored magnifying glass, which is on the right hand side of the brown-colored line. This will bring up the Scale Ratio window that you can see below. Note in the Compositor window, the blue legs are only showing up to the knees.



With the Scale Ratio plugin you can manipulate your video so that it will look the way you want it on a different output Display device. In this case we are going to create a DVD for playing on a Digital TV screen.

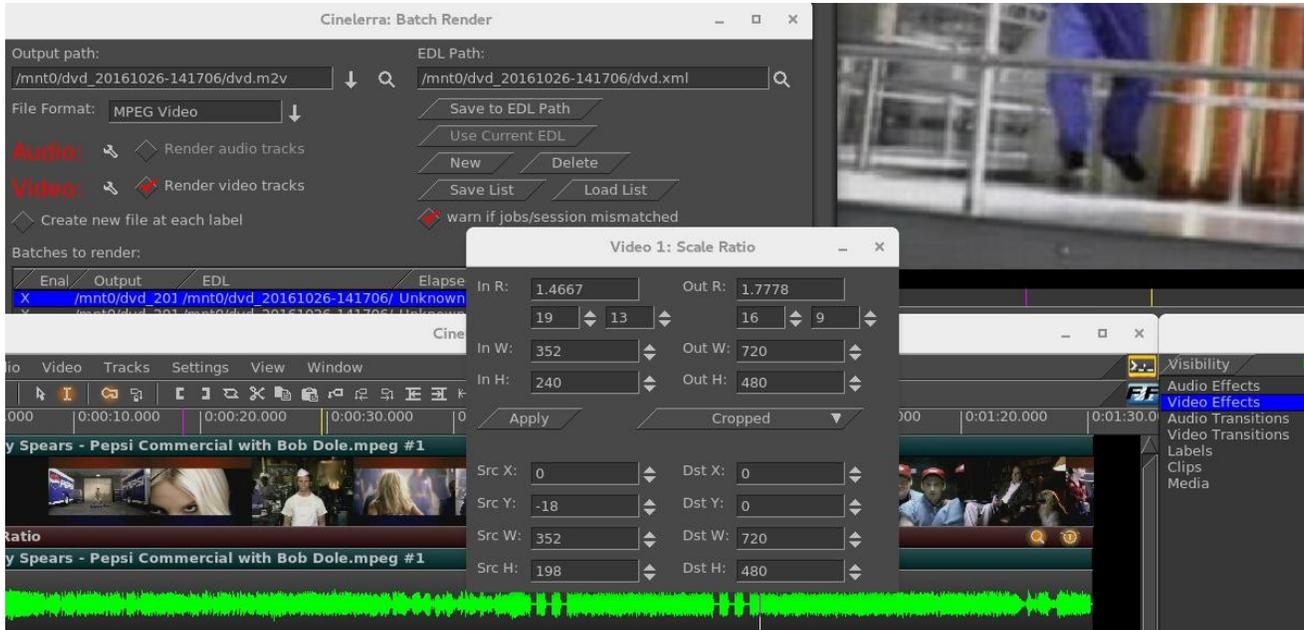
In this screenshot, the left side shows the Input Ratio, Width, and Height of input. The top right half shows desired output values. In this particular case, the input was a YouTube video which was not quite 4:3 aspect ratio.

Left and right sides of the bottom portion show the Source and the Destination X, Y, W, H values. As you change the values on the left side, you can see how this will affect the output as you observe the results in the compositor window. For example, as you change the values for SrcY in a “cropped” Scale scenario, you see up/down movement.



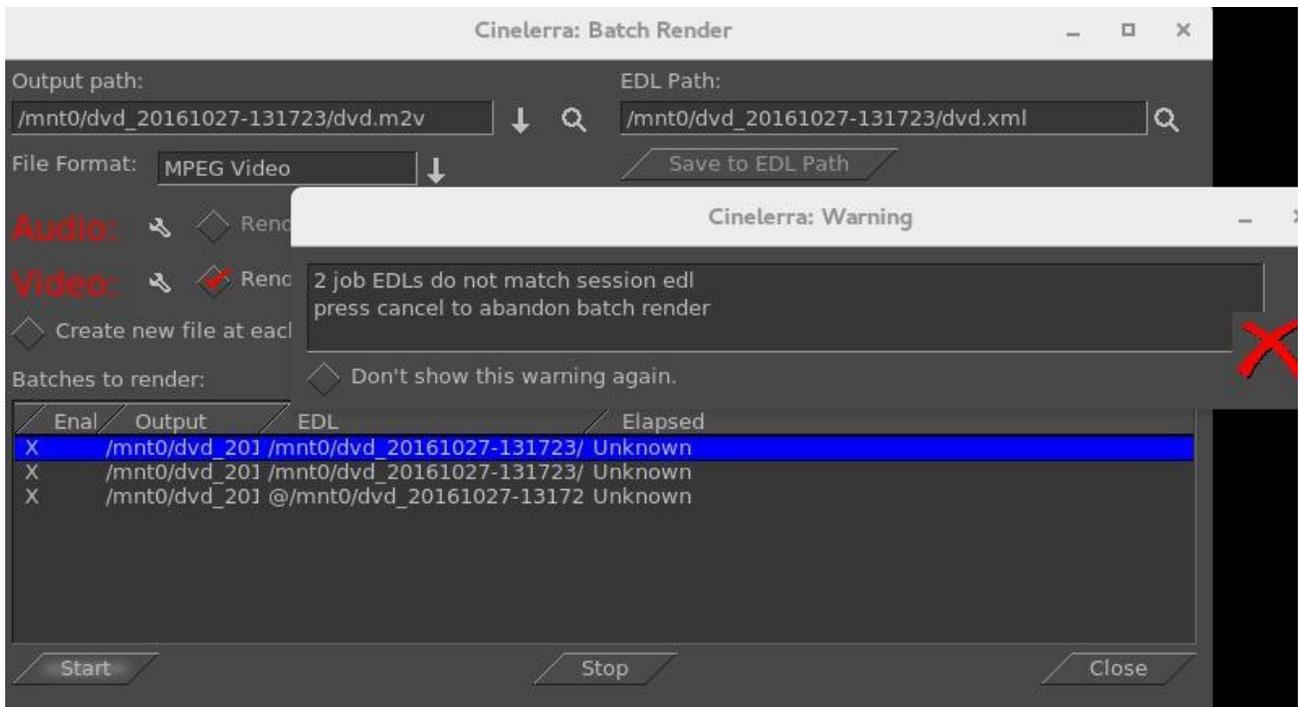
Keep in mind that the monitor you are using is NOT the intended output display device – your digital TV is, which most likely will have different looking aspect/pixels, etc.

9) In order to “crop” the bottom of the video in order to preserve all of the image on the top, modify the Src Y value on the bottom of the left hand side in the Scale Ratio plugin. Src Y which was 21 has now been changed to -18. You will see in the Compositor window how the bottom dark colored border is now gone so that none of the top portion which contains a person's head will be chopped off. Compare the next screenshot to the previous and note the blue legs can be now seen to the waist.



10) Click the “Apply” box in the Scale Ratio window.

11) Click the “Save to EDL Path” in the “Batch Render” window for creating a DVD. If you do NOT do this, you will get a Warning box as seen in the screenshot below, to remind you to Save because you have changed the EDL by modifying the scaling parameters in the Scale Ratio window.



- 12) Next, make sure you have the Timeline set in the Main window at the beginning of where you want to start rendering. Also, make sure the first line in the “Batches to render” section is highlighted as you can see above by the blue highlighting. Click on the “Start” box in the Batch Render window and you will see the video playing in the Compositor window.
- 13) Cinelerra program will be stopped when done rendering; you will be at the terminal prompt where you will see it has printed out some informational messages (or errors if problems), the last 2 are:

To burn dvd, load blank media and run:

```
growisofs -dvd-compatible -Z /dev/dvd -dvd-video /mnt0/dvd_20161027-131723/iso
```

- 14) Load a blank or rewritable DVD into your DVD writer device, which will be similar to /dev/dvd as in the “growisofs” line above – something like /dev/sr0 on your computer.
- 15) Keyin the “growisofs” line, substituting your actual writer device name. Again, you must be root.
- 16) When back at the terminal prompt, and if there are no errors, keyin “eject /dev/dvd” substituting.
- 17) Play it on your DVD player connected up to your Digital TV screen.

This page intentionally left blank so that chapters begin on a right facing page.

14 – The 5 Multi-s

This section is simply a conglomeration of multiple usage possibilities.

14.1 Multi-Camera / Mixer

Use the Mixer Viewer to see multiple media playing simultaneously in re-sizable mini-viewers. This can be used in various ways and is useful to edit videos shot by multiple cameras from different viewpoints that were simultaneously recorded in order to create a single good video. Everything will have to be initially synced so you can decide which one of the camera angles is best suited at any time.

The number of cameras/mixers you can have is generally limited to the available resources on your computer. Currently, the number of File Descriptors available in the O/S limits cameras to about 50. If you have many “mixer viewers” you will probably want to use proxy mode whenever possible. Also, in the Settings → Playback A tab “Video Out” section, uncheck “play every frame” and choosing a Video Driver of X11 with “use direct X11 render if possible” checked, will provide better performance.

Screencast below shows 9 media sources in the left corner, the composed video in the right corner, the timeline with the top video track with pieces of the 9 overwrites, and the choice in Resources of Mixed.



Figure 115: Using Mixer capability in Cin for multiple cameras

Easiest Method to Getting Started

- 0) This method assumes all of your media or cameras are aligned the way you want them already.
- 1) From the File pulldown, create a “New project” with the desired format for Audio and Video output (or you can just use the default).

- 2) File->Load the media files you want to work with using “Create new resources only”.
- 3) In the Resources window, with the Media folder, highlight the list of media you want to “Mix”. This is done using a ctrl or shift mouse button press as you would in a standard listbox selection.
- 4) Right click the mouse on the media selection and choose “Open mixers”. This opens multiple mixer viewer windows, one for each media item that was highlighted. You can do them 1 at a time instead. This also adds the source media tracks to the main window.
- 5) Now use the timeline to play and you will see all viewers/cameras playing. Stop when you get to the end of the “good” camera playback.
- 6) Simply double click the “good” mixer viewer and from where you first started playing to the playback insertion pointer is the source section, which will be pasted in the destination video/audio tracks at the top of the new project.
- 7) Repeat steps 5-6. Start playing again, stop when you want, double click the desired mini-viewer!

Some Hints and Caveats:

You can easily overwrite a section of the new track by “selecting” a section on the timeline, then double click on one of the mini-viewers to overwrite/replace that section.

If you edit the output tracks, it only edits output tracks, and the input tracks may no longer be lined up.

You can add a silent section by selecting past a section and start overwriting that section from then on.

If you use the cursor hairline to create the selection endpoint, it must be past the end of the destination.

The compositor shows composed media. This is the media that will be rendered.

The program always uses overwrite as the paste operation.

Use the timeline edit handles to move the start and end points of that section.

Only middle mouse drag handle operations should be used normally.

Other drags will displace the media source/destination timeline correspondence.

To re-tile the mixer windows after you have resized and moved them around, you can use the Window pulldown of “Tile mixers” or the shortcut of Alt-t.

But, I want to use only the first set of audio tracks...

There are many cases where you may want to compose using media from several different tracks while using the the same audio tracks as associated from a specific viewer. Since mixer source tracks can be updated any time by using a mixer toggle, this makes it possible to do this.

Procedure to update the mixer audio source track list:

- 1) Single click to highlight the mixer window you want to re-associate to the audio track.
- 2) In that audio track’s patchbay click the expand toggle, the arrow on the right side.
- 3) In the expanded pane that appears, there is another arrow on the left side. This icon has the tooltip “Mixer”. Click this and because in step #1 you highlighted the mixer window, it will now be toggled on. Once you click the mixer icon it will then point up.
- 4) Now, disassociate any audio that is unwanted by expanding its patchbay and toggling off the mixer.

Expert Usage:

When you double click a mixer viewer window, it operates an “overwrite” paste operation. This moves src (source) track edits to dst (destination) track edits over the same selected timeline region.

Src tracks should be not playable and not armed in the main window patchbay gui.

Dst tracks should be playable and armed in the main window patchbay gui.

Each mixer viewer maintains a list of the tracks which will be used as src. This list is made visible selecting the window with the left mouse button. When the mixer viewer is selected, a highlight is drawn around the media image. All track patchbay “mixer” toggles are updated to reflect the src tracks included in the selected viewer src track list. The track patchbay toggles can be used to manage the list.

Turning on a toggle (pointing up) includes the track in the src track list.

Turning off a toggle (pointing right) removes the track from the src track list.

New Mixer viewers can be created using the main menu Window → Mixer Viewer, or with a shortcut of Shift-M. When a new viewer is created, the currently enabled patchbay “mixer” toggles are used to create the viewer source track list. The toggles are cleared after the window is created. This is to improve the work flow. Use the following list of steps to create individual mixer viewers.

To create a list of mixer viewers:

- 1) Setup the session settings → format, width, height, frame rate, color model, aspect ratio.
- 2) Create dst tracks using the a/v track pulldowns (or use shortcuts ‘t’ / ‘T’), armed and playable.
- 3) Append src tracks using file → open → append tracks, or the resource window using pasting.
- 4) Using the track patchbay, disarm editing and disable playback of the audio/video src tracks.
- 5) Using the track patchbay, mark the new tracks as “mixer” source to be added to the viewer.
- 6) Create a mixer viewer using the main menu pulldown, or the shift “M” shortcut.
- 7) Repeat steps 3 – 6 for each mixer viewer needed for the session editing.

When you single click a mixer window, it becomes selected and highlighted and all of the patchbay mixer toggles are updated to reflect the state of the viewer’s src tracks. Tracks that will be src are shown as enabled. If you change a toggle, the src tracks for the selected window will be modified. This means you can associate or dis-associate any media track to any mixer window.

When you double click a mixer window, an overwrite paste is invoked. The mixer viewer’s src tracks are overwritten to the dst tracks. The timeline region for both the source and destination are the same for the overwrite paste function. The selection region is used if it is active. If the selection is empty, that is it is a hairline, the selection region is from the end of the destination playable edits to the selection cursor hairline. The hairline must be past the end of the playable edits on the destination tracks.

The mixer viewer configuration is saved with the session data. When a saved session is loaded in “replace project” or “replace project and concatenate tracks”, the mixer viewer will be reopened.

Using Proxy with “Open Mixers”

The best way to use proxy with your multiple cameras is to follow the steps below:

- 1) Load media with insertion strategy of “create resources only”.
- 2) Highlight the media in the Resources window and right click on this to choose “open mixers”.
- 3) Use the Settings pulldown and choose “Proxy settings...” to bring up the proxy menu.
- 4) Choose the size and other options you want and click the checkmark OK.
If you choose the option “Beep when done” you will hear a short beep if all media is already proxied or a longer beep when all proxies have been created.
- 5) When your editing is complete, use Settings pulldown and proxy to “original size”.

Instead of Open Mixers, you can Insert Mixers with new tracks at the timeline insertion point.

14.1.1 Mixer Align by Audio

Multi-camera footage of a single event can have various shots starting and ending at different times. So when the footage start times are different, you can use the mixer audio to synchronize the clips on the timeline. The program algorithm attempts to find and align automatically the waveforms of the media.

Synchronizing multiple camera videos based on audio tracks can be done with Cinelerra-GG easily enough with the Window pulldown Mixers... → Align mixers option. *Align mixers* brings up a window displaying your mixers, the currently selected Master Track, and a list of all of the Audio Tracks. There is a limit of 32 audio tracks per each mixer (that should be enough!)



Figure 116: A 2 camera Align mixer window using audio initial condition before any match applied

Two different methods of aligning the audio for mixers are available. They are most easily referred to by the button that is pressed – Match and Match All. There are also 2 methods of activating the alignment – Apply and Move.

Match consists usually of the next set of steps to take advantage of this feature:

- 1) Load your camera media with insertion strategy of “resources only” .
- 2) Highlight in the Resources window, all of the media you want to mix.
- 3) Right mouse button on one of the media and click on Open Mixers; all mixer windows come up.
- 4) Provide a small target audio pattern on the Master Track for syncing by marking with the In/Out points ([and]).
- 5) Make a selection on the timeline in which to look for the pattern. Left mouse click, then drag select and highlight a search time domain.
- 6) Use the Window pulldown, Mixers... → Align mixers to bring up its dialog window.
- 7) Highlight in the first listbox, the Mixer number you want to align. Click on “Match”. This will take a few seconds so watch the rendering time percentage on the lower right hand side zoom panel. The buttons will be ghosted out until finished. Now note the changed values in the Audio Tracks listbox.
- 8) If you are satisfied with the calculated Nudge values – that is they are very close to 1.0 - in the Audio Tracks listbox and the audio track selected as the Master Track in the Master Track listbox,

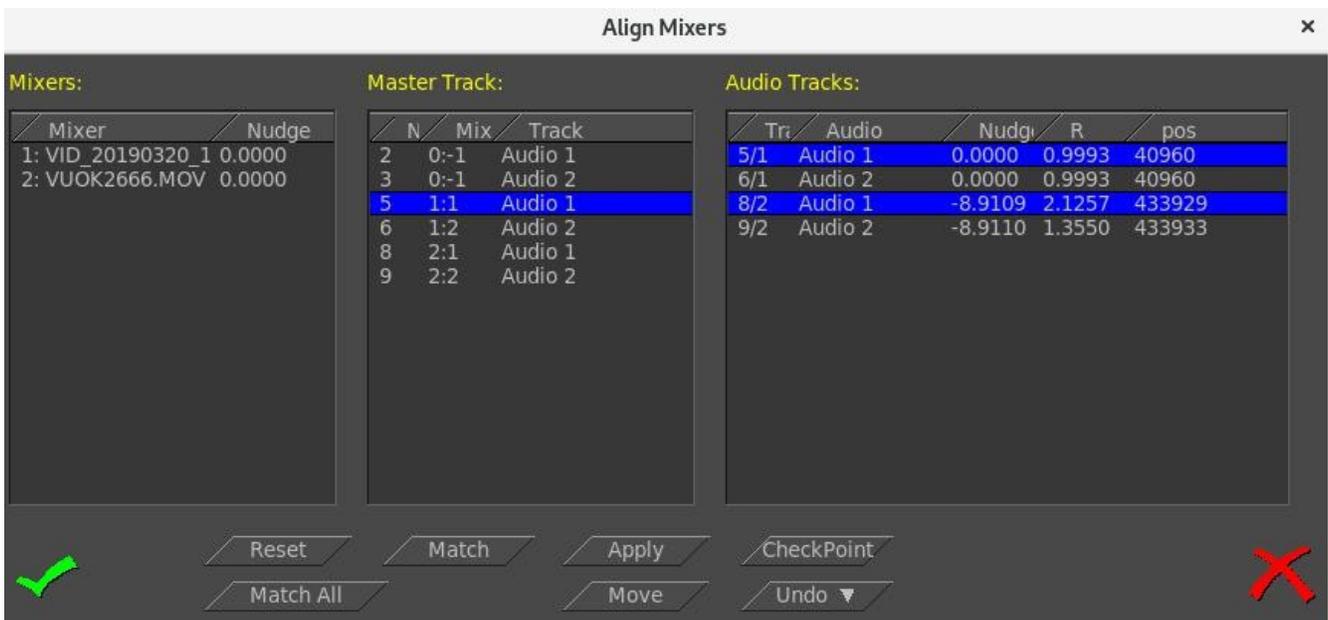


Figure 117: Align mixer window after a Match has been applied. Note the Nudge amounts above.

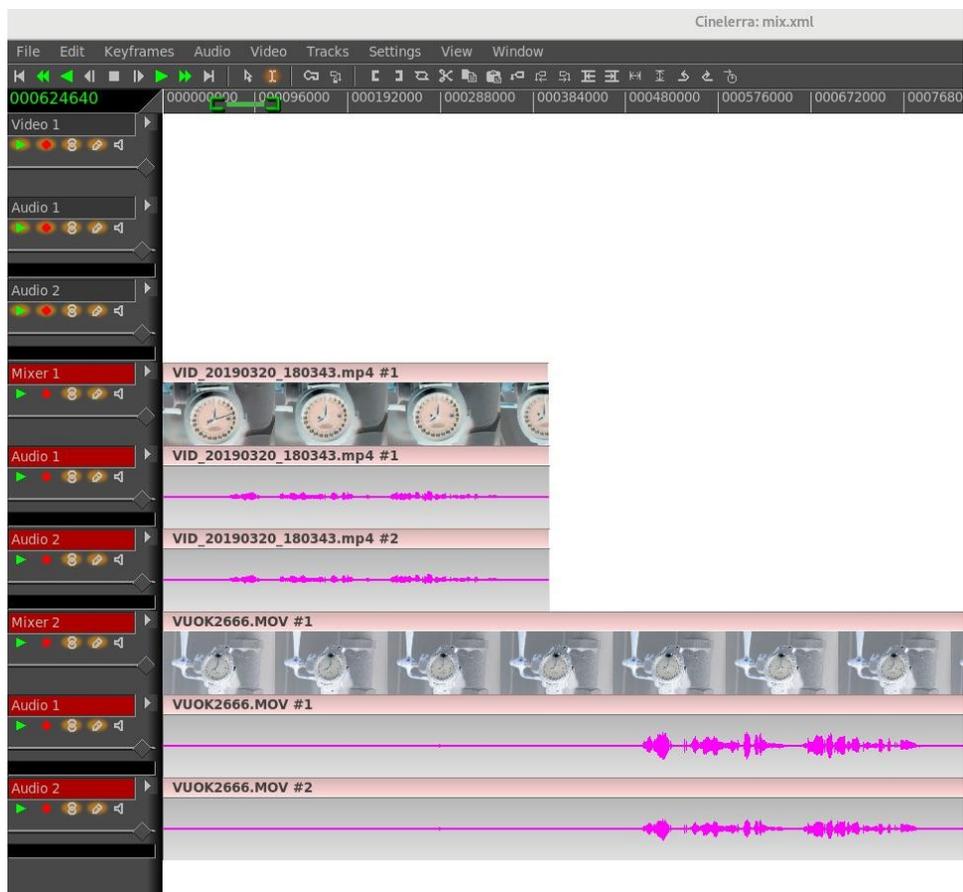


Figure 118: Setup for aligning by audio. Note that [] are set over a small sample waveform and highlighted selection that includes that.

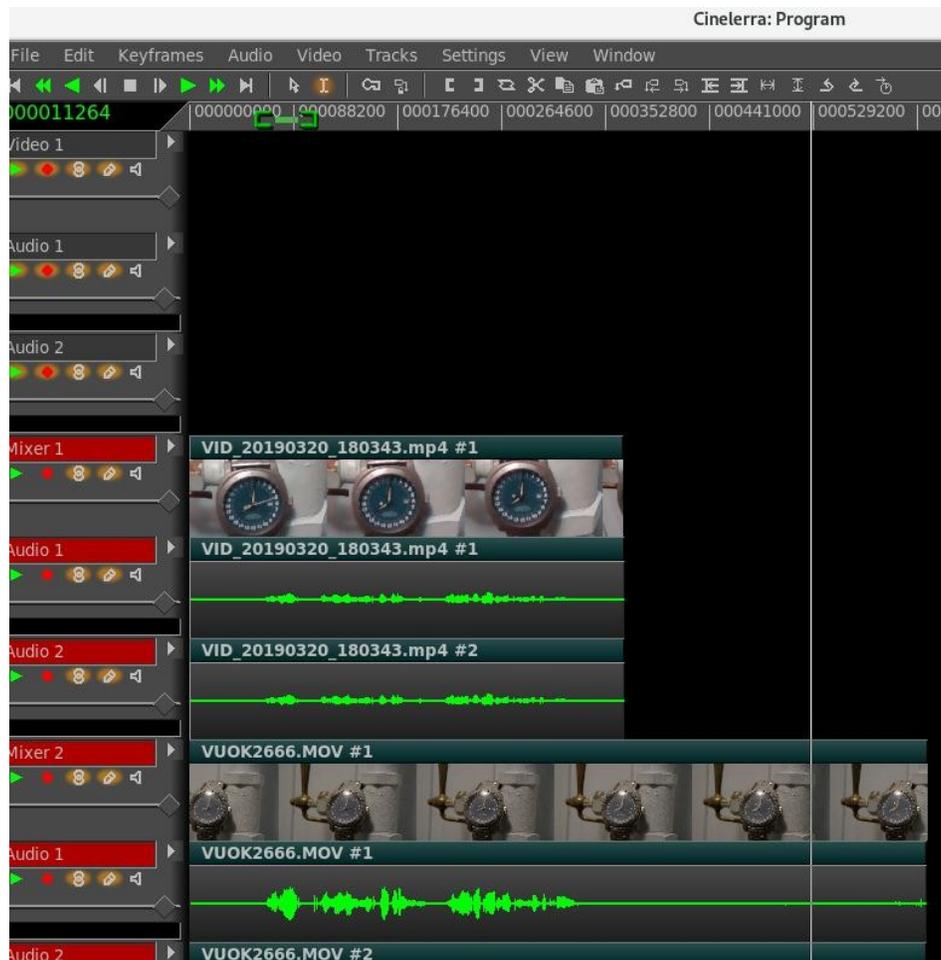


Figure 119: An Audio match is complete. Note the waveform is aligned.

hit the “Apply” button.

- 9) If you plan on performing more alignment tasks, click on Checkpoint so you can go back to a previous step in case you make a mistake.
- 10) Last, click on the OK green checkmark or to cancel click on the red X. Or just close the gui.

Reset is used to start over with the current session data, not an undo. This means you can use the match repeatedly to refine alignments. All of the Audio Tracks listbox values will be reset.

Checkpoint provides a method to create checkpoints that save the current state. This is especially helpful while learning or doing more complicated operations where you might make a mistake or do not like the results and need to get back to a previous state.

Undo is used to put the media back to a previous state on the timeline. If you choose “start over” the session will reload with the original, before any changes were applied. You can also go back to any of your previous checkpoints that you created earlier which are listed there, such as “chkpt 1”.

Match All is used when you have several mixers, instead of only picking 1 to match, it picks the best match for EACH of the mixer tracks based on a single master track. So when you hit Apply, each track might move differently. You do not set In/Out points but you have to make a selection within which to match.

Apply button will apply the nudges that were generated during the Match or Match All execution.

Move is very handy when you are using mixers, if you have an edit somewhere that you need to fix specifically without moving any of the other track pieces. In this case you have to select a section (like you do a group, but do not make a group), generate a match, and then you can just "Move" that set only - everything else stays where it is at its current location.

More detailed information follows about how this all works and the information in the dialog window. It is important to know that the result of the calculation is "best match" but you can still override the selections if you decide there is a better one. The dialog window is split into 3 sections:

- 1) *Mixers* lists the mixers that are active by highlighting them all initially. You can decide that you do not want 1 or more mixers to be used in the correlation calculation by un-highlighting the one(s) that should not be used. However, obviously you have to have at least 2 in order to align audio.
- 2) *Master Track* lists each of the audio tracks currently loaded for all of the mixers. You can decide to highlight a different audio track to be used as the master for correlation, but only 1 can be used.
- 3) *Audio Tracks* lists each of the mixer audio tracks. Again, you can highlight a different set of which mixer audio tracks that you want to use for the waveform correlation. Any audio tracks that are not highlighted, that is "turned off" will not be considered in the correlation calculation.

The corresponding input position is determined by track input correlation.

The letter "R" in the Audio Tracks listbox represents the correlation value. R=1.0 designates that if both the pattern and the matching section were in the highlighted area, they are completely correlated – this is a good self-test to check. Nudge=0.0 means just that!

The header "pos" stands for the timeline position.

When the "Apply" button is pressed, only the Mixers listbox is relevant at that time.

The Mixer with the master track will not move, everything else will be lined up.

Match All option basically consists of the following steps:

- 1) Highlight the Mixer to use in the Mixer listbox.
- 2) Highlight the Master Track you want to use in the Master Track listbox.
- 3) On the timeline, mark your selection on the Master Track.
- 4) Click on the Match All button.
- 5) Note the nudge values to see if they make sense, and if so, press Apply.

Match option basic steps (just for comparison with Match All):

- 1) Set the In/Out points [] of the target.
- 2) On the timeline, mark your selection.
- 3) Click on the Match button.
- 4) Note the nudge values to see if they make sense, then press Apply (or Move when doing a group).

14.1.2 Recover Mixer Windows

It is a hazard that you might accidentally undo (z) too far and lose your mixer windows. Here are the steps to recover. It is recommended that you make a backup of your project before performing the recovery steps just in case there are other problems.

In the patchbay box to the left of the main timeline, there is a right pointing arrow on the right side. This is called the Expander. When you click on each expander, another line appears below that expander arrow and the timeline track height is slightly increased. If you Shift-click on a single expander, the patchbay will expand for all of the tracks.

- 1) Expand all of the patchbay lines, either one by one, or Shift-click on one to do them all. This is so you can see the “mixer” right pointing arrow on the second expanded line.
- 2) Use the Window pulldown and choose “Mixer Viewer” to bring up a new mixer window. Now you will be making an association between the mixer viewer and the track’s video.
- 3) Click on the new mixer window to make sure it is highlighted with a white border. This designates it as the “in use” mixer viewer.
- 4) Set your play to the beginning of the video using the Home key or Home transport button.
- 5) In the patchbay for a video track click on the “mixer” arrow on the expanded 2nd line which is a right facing arrow. Now the arrow will point up. If there are audio tracks with that video, click on each of its audio tracks “mixer” arrow until they point up also.
- 6) Next move your insertion pointer on the timeline where there is video. Some of the time this just helps so that the new mixer viewer window gets redrawn and you can see that the images appear; but the image may not appear until the program does a redraw later. Now the mixer viewer should be correctly associated. Note if you have large video, give it some time to update. You may have to click on the mixer viewer window if the image does not show. You can always start over with that mixer if you encountered any problems.
- 7) Click the arrows that are pointing up in that video and its audio so they go back to pointing right. That mixer viewer is complete so you need to do this to make sure the “mixer” arrows are off.

Repeat steps 2 through 7 for each of the mixer viewers you need going down the patchbay starting on step 2 first with Mixer 1, then 2 to 7 steps for Mixer 2, then again run 2 to 7 for Mixer 3 and so on.

Sometimes the association does not stick initially. If not, highlight the mixer viewer with the problem, change the mixer arrows to point up, and reassociate.



Figure 120: Mixer patchbay

14.2 Multi-Pane Support

The main Cinelerra edit window holds the Track Canvas which can be divided into 4 panes of track data: 1 or 2 vertical panes and/or 1 or 2 horizontal panes. To split the track, use the Window pulldown, and then click on Split X or Split Y depending on how you wish to split the track. Alternatively, the canvas pane types can be changed using keys <Ctrl-1> for toggle split horizontal or <Ctrl-2> for toggle split vertical. Or the track can be split into panes by using the + widget in the lower right hand corner of the track canvas. Once the track has been divided, you can use the + widget shortcut or the drag bars to change the size of the panes.

Multi-Pane, or split screen, allows you to look at the first part of a movie at the same time as a part that is a long ways away on the timeline which would have been off the screen. By having multiple panes, you can see the 2 parts you want to look at simultaneously and drag/drop easily between the 2. Also, the X pane split is extremely convenient for laptop users and computer monitors with small screens since it can be used with horizontal scrolling with the mouse wheel + Ctrl. The Y-pane split makes it easy to see 2 simultaneous drag and drop zones when you have lots of tracks.

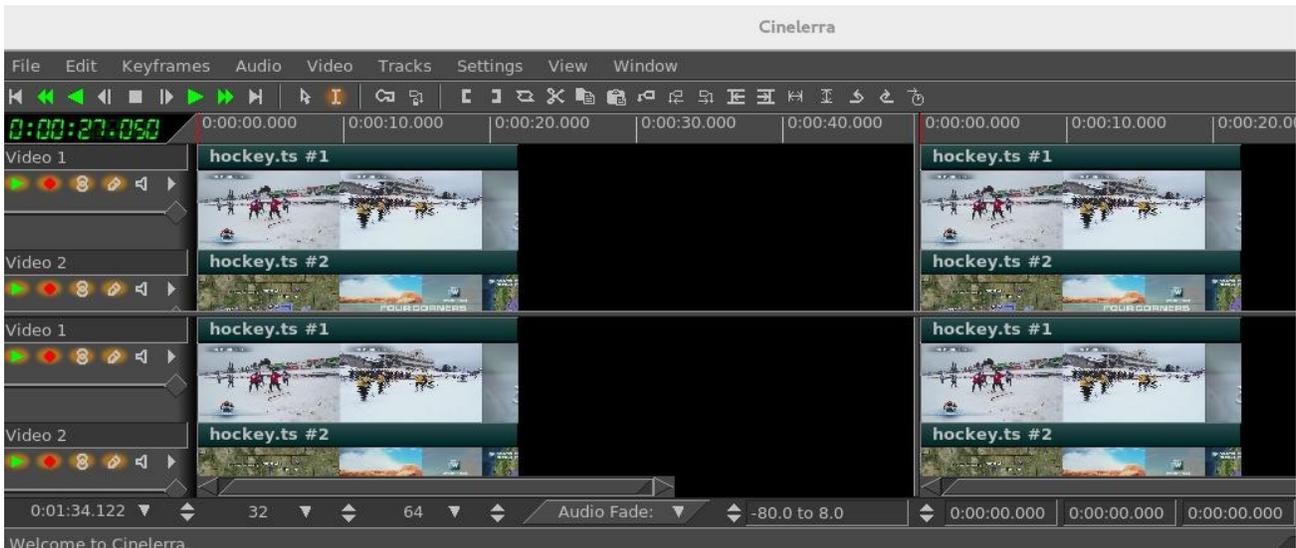


Figure 121: Shown are 4 panes that have split the main track canvas for some editing scenarios

14.3 Multi-screen / Playback Configuration

Cinelerra-GG supports 2 separate preferences for the playback configuration. Cinelerra can be operated in a single or dual screen configuration, both by using Xinerama or dual screen configuration of X windows. It will take some setup using Xconfig to make this work.

The Settings → Preferences menu has “Playback A and Playback B” tabs. The target display and audio device configuration can be separate, to support up to 2 display and/or audio device stations. The active configuration displays an asterisk (*) in its selection tab and the selected tab will be made active when OK is pressed. For example: you may have a dual screen monitor system with the left screen showing the cinelerra main window and the right screen showing the composer. Another setup might

use a monitor for the left screen and an HDTV as the right screen displaying the composer. When a playback configuration is selected, the audio/video device configuration is switched to the playback selection. The active playback setup can be changed through use of the menu pulldown of Settings → Preferences or via the remote control menu selection (see the section “Remote Control for DVB” for more detail).

Yes, you can watch TV on cinelerra instead of cinelerra on TV.

Screenshot showing partial window of *Playback A selected and the second tab for Playback B. Note that on the bottom right of the window, “Default B Display:” is set to :0.1, representing the setting for Screen 1. On the unseen “Playback A” window, the “Default A Display:” will be set to :0.0 meaning for Screen 0. Otherwise, the default would be nothing there or just <empty>.

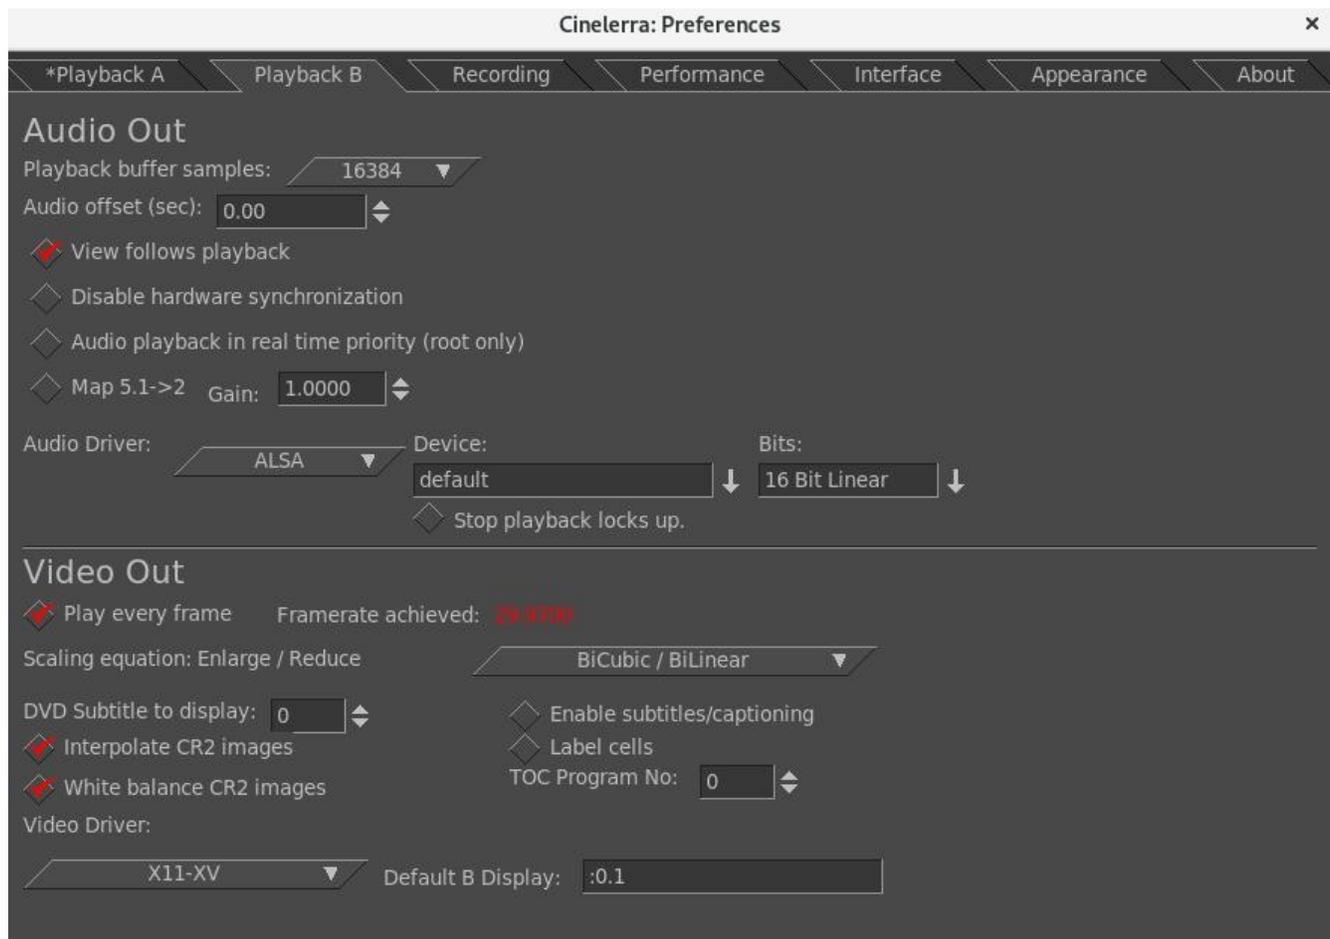


Figure 122: Multi-screen Playback example useful for watching Cinelerra run on the “big screen”

14.4 Multi-Session

You can run as many sessions of Cinelerra as your computer resources allow. However, if you are using the same \$HOME/.bcast5, changes you make for one may impact the others.

14.5 Multi-Viewer Window Support

You can create as many Viewer windows as you want in Cinelerra. These are handy for users who are adept at working with a lot of different clips simultaneously. By bringing up multiple Viewer windows, each clip can be edited in its own area, making it easy to see all of the separate pieces. After you have loaded some media files, to start another Viewer window, right click on one of the pieces of media in the Resources window. This brings up a menu of several options, one of which is “view in new window”. Choose this option and that media will come up in a new Viewer window for you to work.

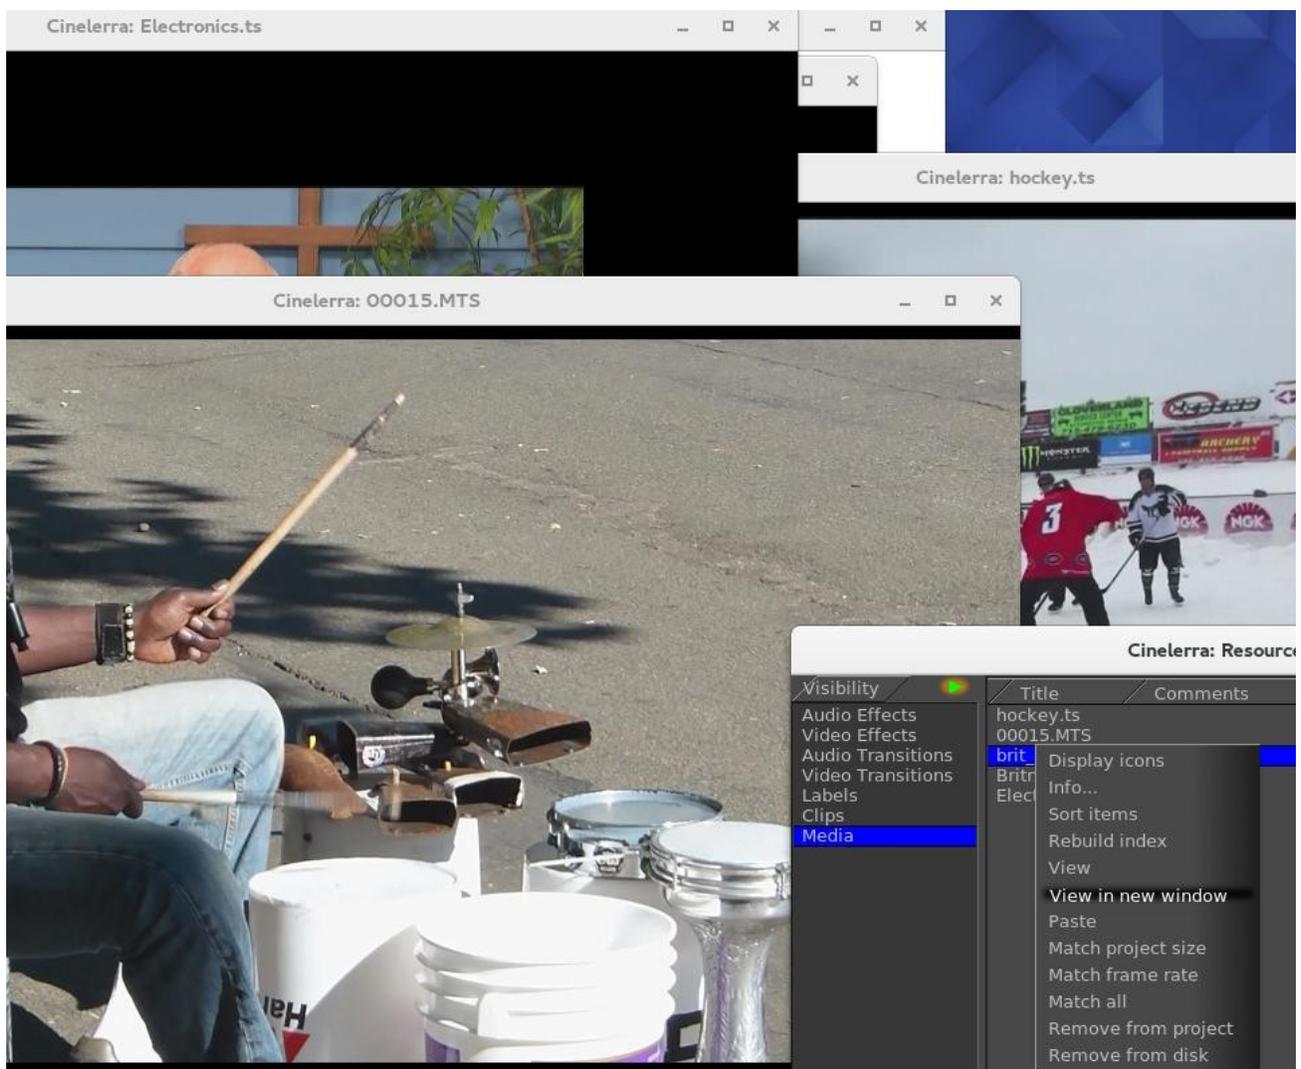


Figure 123: Shown here are 3 Viewer windows and the "View in new window" popup

This page intentionally left blank so that chapters begin on a right facing page.

15 - Shortcuts with Keyboard and Mouse

Main Menu (- is a checkbox)

01/19/2019 update

PULLDOWNS **Menu Items**

Shortcuts

Description

File	New	'n'	Create new project
Pulldown	Load files...	'o'	Load file(s)
	- load tool	'Ctrl-a'	Selects all files
	- load tool	'Ctrl-z'	Deselects any highlighted files
	Save	's'	Save project
	Save As...	'Shift-S'	Save and name project
	Export Project	'Alt-s'	Save a project for easy moving
	Record...	'r'	Open record window
	Scan...	'Ctrl-Alt-s'	Open dvb scan window
	SubTitle...	'Alt-y'	Open subtitle script window
	Render...	'Shift-R'	Open render window
	Export EDL...	'Shift-E'	Open export EDL window
	Batch Render..	'Shift-B'	Open batch render window
	BD Render...	'Ctrl-d'	Open create bluray disk window
	DVD Render...	'Shift-D'	Open create dvd disk window
	Quit	'q'	Quit the program
	Dump EDL		Write current EDL to stdout
	Dump Plugins		Write defined Plugins to stdout
	Load Backup		Load auto-backup file "backup.xml"
	Save Backup	'b'	Save auto-backup file "backup.xml"
Edit	Undo clear	'z'	Undo from history previous operation
Pulldown	Redo	'Shift-Z'	Redo from history next operation
	Cut	'x'	Copy selection to cut buffer and delete
	Split	'x'	Split edit into 2 sections at hairline cursor
	Copy	'c'	Copy selection to cut buffer
	Paste	'v'	Insert selection from cut buffer
	Clear	'Del'	Delete selection
	Paste Silence	'Shift-Space'	Paste empty edit at selection or frame if none
	Mute Selection 'm'		Replace selection with empty edit
	Trim Selection		Replace track with selection
	Select All	'a'	Toggle select all/none
	Shuffle Edits		Randomly reorder track edits
	Reverse Edits		Reverse track edits
	Edit Length...		Change selected track given duration
	Align Edits		Change selected corresp. track edits to start/end same
	Trans. Length...		Change selected track transitions to given duration
	Detach trans.		Remove selected track transitions
	Clear labels		Delete selected timeline label markers
	Cut ads		Remove selected region, add to commercial trking db
	paste subttl	'y'	Replace selected subtitle edit text w/ subtitle script text
Keyframes	Cut Keyframes	'Shift-X'	Copy selected keyframes to cut buffer and delete
Pulldown	Copy Keyframes	'Shift-C'	Copy selected keyframes to cut buffer
	Paste Keyframes	'Shift-V'	Paste keyframes in cut buffer
	Clear Keyframes	'Shift-Del'	Delete selected keyframes

	Change to linear Change 2 smooth Create curve type		Change keyframe auto curves to piecewise linear Change selected keyframes auto curves to bezier Set new keyframe type: smooth,linear,tangent,disjoint
	Copy default key Paste default key	'Alt-c' 'Alt-v'	Copy default keyframe to clipboard Paste clipboard keyframe into default keyframe
Audio Pulldown	Add track Default Trans. Map 1:1 Map 5.1:2 Attach Trans. Attach Effect... Render Effect...	't' 'u'	Append new empty audio track Insert default audio transition at selected edit boundary Setup audio mixing output = input Setup audio mixing stereo = 5.1 Open audio transition select for insert at edit boundary Open audio selection menu for insert at edit boundary Open audio render select to render from select as pcm
Video Pulldown	Add track Default Trans. Attach Trans. Attach Effect... Render Effect...	'Shift-T' 'Shift-U'	Append new empty video track Insert default video transition at selected edit boundary Open video transition select for insert at edit boundary Open video effect selection for insert at edit boundary Open video render select to render from select as pcm
Tracks Pulldown	Move tracks up Move trks down Delete tracks Delete last track Concatenate trks Append to Proj. add Subttl	'Shift-Up' 'Shift-Down' 'd' 'Shift-N' 'Shift-Y'	Circulate tracks up Circulate tracks down Delete all tracks Delete last track Concatenate tracks together Add set of tracks Add subtitle track
Settings Pulldown	Format... Preferences... Proxy Settings Align cursor... - Edit labels - Edit effects - Keyfrs fol. edits - Typeless keyfrs Save settings Loop Playback Set bkg render	'Shift-F' 'Shift-P' 'Alt-r' 'Ctrl-a' 'Ctrl-s' 'Shift-L' 'Shift-G'	Open render format setup menu Open preferences setup menu Bring up Proxy settings window Align cursor on frames Toggle labels follow edits Toggle plugins follow edits Toggle keyframes follow edits Toggle typeless keyframes mode Save Cinelerra_rc Set loop playback region to selection/all Toggle background rendering
View Pulldown	- Show assets - Show titles - Show trans. - Fade - Mute - Overlay Mode - Pan - Plugin Autos - Mask - Speed - Camera X	'0' '1' '2' '3' '4' '5' '6' '7' '8' '9' 'Ctrl-Shift-X'	Toggle show asset data Toggle show asset title Toggle show transition effects Toggle show fade auto curves Toggle show mute auto curves Toggle show mode auto curves Toggle show pan auto curves Toggle show plugin keyframes Toggle show mask auto keyframes Toggle show speed auto curves Toggle show camera X auto curves

- Camera Y	'Ctrl-Shift-Y'	Toggle show camera Y auto curves
- Camera Z	'Ctrl-Shift-Z'	Toggle show camera Z auto curves
- Projector X	'Alt-Shift-X'	Toggle show projector X auto curves
- Projector Y	'Alt-Shift-Y'	Toggle show projector Y auto curves
- Projector Z	'Alt-Shift-Z'	Toggle show projector Z auto curves
Window		
Pulldown		
- Show Viewer		Toggle show view window
- Show Resrs.		Toggle show resources window
- Show Compos.		Toggle show compositor window
- Show Overlays		Toggle show auto overlay window
- Show Levels		Toggle show audio levels window
- Split X pane	'Ctrl-1'	Toggle ½ horiz track timeline window panes
- Split Y pane	'Ctrl-2'	Toggle ½ horiz track timeline window panes
Mixer Viewer	'Shift-M'	Bring up a Mixer Viewer window
Tile mixers	'Alt-t'	Tile mixer windows to original position/size
Default Positions	'Ctrl-p'	Reset window positions/size to defaults
Tile Left		Set window positions/sizes to tile left screen
Tile Right		Set window positions/size to tile right screen
>_._	shell cmds	Top rt. Corner Run scripts (setup in interface preferences)

Main Menu

BUTTONS	<u>Menu Items</u>	<u>Shortcuts</u>	<u>Description</u>
Transport	Rewind	'home'	Jump to beginning
Buttons	Fast reverse	'KP +' / Alt-p	Double speed reverse play
	Normal reverse	'KP 6' / Alt-o	Normal speed reverse play
	Frame reverse	'KP 4' / Alt-u	Render previous frame
	Stop	'KP 0' / Alt-m	Stop playback
	Frame forward	'KP 1' / Alt-j	Render next frame
	Normal forward	'KP 3' / Alt-l	Normal speed forward play
	Fast forward	'KP Enter'/A-;	Double speed forward play
	Jump to end	'end'	Jump to end
	(No button)	'KP 5' / Alt-i	Slow speed reverse play
	(No button)	'KP 2' / Alt-k	Slow speed forward play
	(No button)	'Shift-KP1'	Render next frame and audio
	(No button)	'Shift-KP2'	Slow speed forward play without audio
	(No button)	'Shift-KP3'	Normal speed forward play without audio
	(No button)	'Shift-KP4'	Render previous frame and audio
	(No button)	'Shift-KP5'	Slow speed reverse play without audio
	(No button)	'Shift-KP6'	Normal speed reverse play without audio
		'Ctrl-KP#'	If [] set, "KP2,3,5,6,+,enter" between in/out
		+Shift w/Alt	Shift with Alt+x above, adds/removes audio
		+Shift w/Ctrl	Loops play, all or between In/Out
Edit Modes	Auto Gen. Keys	'j'	Generate keyframes while tweaking
	Lock Labels		Lock labels from moving
Edit	In point	'[' or '<'	Toggle In point timeline marker
Operations	Out point	']' or '>'	Toggle Out point timeline marker

To clip	'Ctrl-t'	Clear both In and Out pointers
Split	'i'	Copy selection and create clip
	'x'	Split edit into 2 sections at hairline cursor
Cut	'x'	Copy selection to cut buffer and delete
Copy	'c'	Copy selection to cut buffer
Paste	'v'	Insert selection from cut buffer
Toggle label	'l'	Toggle label at current position
Previous label	'Ctrl ←'	Move to label before cursor
Next label	'Ctrl →'	Move to label after cursor
Next label select	'Ctrl-Shift →'	Select from cursor to next label
Prev label select	'Ctrl-Shift ←'	Select from cursor to previous label
Next keyframe	'k'	Move to next keyframe
Previous keyframe	'Ctrl-k'	Move to previous keyframe
Next keyframe select	'Shift-K'	Select from cursor to next keyframe
Previous keyfr select	'Ctrl-Shift-K'	Select from cursor to previous keyframe
Previous edit	'Alt ←'	Move to previous edit
Next edit	'Alt →'	Move to next edit
Cut left label	Ctrl-alt-shift- '<'	Cut from insert pointer to previous label
Cut right label	Ctrl-alt-shift- '>'	Cut from insert pointer to next label
Cut previous edit	Ctrl-alt- ';'	Cut from insert pointer to previous edit
Cut next edit	Ctrl-alt- ']'	Cut from insert pointer to next edit
Fit	'f'	Fit time displayed to selection
Auto Fit	'Alt-f'	Fit the max/min range of all automation types
Auto Fit	'Ctrl-alt-f'	Fit the max/min range of current automation types
Undo	'z'	Undo from history previous operation
Redo	'Shift-Z'	Redo from history next operation
Manual Goto	'g'	Jump to time selected by popup

Drag/Drop Edits

Clear Select	Ctrl-Shift-A	Deselect all selected edits
Copy	Ctrl-c	Copy selected edits into copy buffer
Cut	Ctrl-x	Delete selected edits/put in buffer/collapse
Mute	Ctrl-m	Delete selected edits/put in buffer/insert space
Copy Pack	Ctrl-Shift-C	Copy selected edits and pack together
Cut Pack	Ctrl-z	Delete selected edits/put in buffer/collapse
Mute Pack	Ctrl-Shift-M	Delete selected edits/put in buffer/insert space
Paste	Ctrl-v	Paste (splice) buffer at insertion point or reticle
Overwrite	Ctrl-b	Paste (overwrite) buffer at insertion pt/reticle
	Ctrl / LMB	Toggles selected of group/edit
	LMB	If "Clear before toggle", only select this edit/group
	LMB	If not "Clear before toggle", toggle edit/group
Double click		Only select this edit/group
Double clic	Ctrl	Selects column and highlights
LMB	Shift	Group/ungroup selected edits/group
LMB	Ctrl + Drag	Column drag
LMB	Drag	Clip drag
	Drag	Drag unselected edit under cursor

FFMpeg probes - FFmpeg early early/late

Top rt. Corner Run ffmpeg open probes first / last

Main Window Key

Qualifier **Description**

Key		
'e'		Toggle editing mode: from cut&paste or drag&drop
'1 through 9'		Toggle auto mode settings as in main menu shortcuts
'1 through 8'	Alt	Select asset pgm 1-8; reset format (remove other trks)
← (left arrow)		Scroll window timeline display left (not insertion pt)
← (left arrow)	Ctrl	Move cursor to label left of cursor (main shortcuts)
← (left arrow)	Alt	Move cursor to edit boundary left of cursor position
← (left arrow)	Shift-ctrl	Move cursor to label left of cursor, expand selection
← (left arrow)	Shift-alt	Move cursor to edit boundary left of cursor&expand
→ (right arrow)		Scroll window timeline display right (not insertion pt)
→ (right arrow)	Ctrl	Move cursor to label right of cursor (main shortcuts)
→ (right arrow)	Alt	Move cursor to edit boundary right of cursor position
→ (right arrow)	Shift-ctrl	Move cursor to label right of cursor, expand selection
→ (right arrow)	Shift-alt	Move cursor to edit boundary right of cursor&expand
, (comma)		Scroll window timeline display left (not insertion pt)
, (comma)	Ctrl	Move cursor to label left of cursor (main shortcuts)
, (comma)	Alt	Move cursor to edit boundary left of cursor position
. (period)		Scroll window timeline display right (not insertion pt)
. (period)	Ctrl	Move cursor to label right of cursor (main shortcuts)
. (period)	Alt	Move cursor to edit boundary right of cursor position
UP (up arrow)		Increase timeline duration (zoom out timeline)
UP (up arrow)	Ctrl	Increase audio sample waveform scale
UP (up arrow)	Alt	Increase auto curve limits (zoombar selected curve)
UP (up arrow)	Ctrl-alt	Increase all auto curve limits
DOWN arrow		Decrease timeline duration (zoom in timeline)
DOWN arrow	Ctrl	Decrease audio sample waveform scale
DOWN arrow	Alt	Decrease auto curve limits (zoombar selected curve)
DOWN arrow	Ctrl-alt	Decrease all auto curve limits
PGUP (page up)		Scroll current track vertical window pane up
Mouse wheel	Up	Scroll current track vertical window pane up
PGUP (page up)	Ctrl	Increase track scale geometry (zoom in track data)
Pgdn page down		Scroll current track vertical window pane down
Mouse wheel	Down	Scroll current track vertical window pane down
Pgdn page down	Ctrl	Decrease track scale geometry (zoom out track data)
Home		Move insertion point to beginning of timeline
End		Move insertion point to end of timeline
Mouse wheel	Shift + up	Zoom out view without resizing the window
Mouse wheel	Shift + down	Zoom in view without resizing the window
Mouse wheel	Ctrl + up	Move left on the timeline without resize
Mouse wheel	Ctrl + down	Move right on the timeline without resize
'Shift+click'		Over edit causes highlight section to extend to cursor
'Shift+click'		Over boundary of effect, trims only that effect
'Shift+click'		Over Hard Edge of Blade Cut, toggles marker
Double click		Between labels, highlights selection
Toggle single trk	'Tab'	Toggle single track arming status
Toggle other trks	'Shift-tab'	Toggle all of the other tracks arming status
	Double click	On plugin title bar, selects that area
	Double click	On an edit, selects that area
	Double click	On subtitle track, displays text in Line text
	Double click	On fade/speed, synch video/audio ganged

	Shift	While dragging auto, turns of gang
	DoubleMMB	On fade/speed, select keyframe position
'U'		Paste the last Video transition
'u'		Paste the last Audio transition
'r'	Ctrl	Proxy quick switch
F1	Shift	Toggle on/off all XYZ of camera
F2	Shift	Toggle on/off all XYZ of projector
F1	'Ctrl+Shift'	Use window layout #1
F2	'Ctrl+Shift'	Use window layout #2
F3	'Ctrl+Shift'	Use window layout #3
F4	'Ctrl+Shift'	Use window layout #4

Timebar

<u>Key</u>	<u>Qualifier</u>	<u>Description</u>
Left mouse btn	Ctrl	Changes to the next time format
Middle mouse btn	Ctrl	Change to the previous time format
Left mouse btn		Moves cursor position on the timeline
Click label	Hold LMB	Drag label to elsewhere
Click In/Out	Hold LMB	Drag In/Out pointer elsewhere

Patchbay

Shift/Move	Hold LMB	On Fade slider bar, sets gain to 100% or 0db
------------	----------	--

Compositor

BUTTONS

<u>Menu Items</u>	<u>Shortcuts</u>	<u>Description</u>
Protect video	F1	Disable changes to the compositor output
Zoom view	F2	Zoom in view without resizing the window
Edit mask	F3	Brings up the mask editing tool
+ Edit mask	Shift-drag pt.	Hold down shift+left mouse to drag a point
+ Edit mask	Ctrl-drag pt.	Hold down Ctrl+left mouse to move control pointers
+ Edit mask	Alt-drag mask	Hold down Alt+left mouse to translate mask
+ Delete mask	Shf+delete btn	Hold down Shift + click Delete button to delete mask
Ruler	F4	Brings up the ruler for usage
+ ruler hotkey	Ctrl	Lock ruler to nearest 45 degree angle
+ ruler hotkey	Alt	Translate the ruler
Adjust camera	F5	Brings up the camera editing tool
Adjust projector	F6	Brings up the projector editing tool
Crop layer/output	F7	Crop a layer or output
+ crop tool	Click-drag	In video, starts a new rectangle
+ crop tool	Click-drag	Over rectangle to reposition
+ crop tool	Alt-click	In cropping rectangle to translate
Get color	F8	Detects color at current spot & stores temporary
Show tool info	F9	Works with editing/cropping buttons
Show safe regs.	F10	Draws safe regions in the video output
	'p'	Turn on/off Click to Play
	Shift+w/Ctrl	With transport btn. loops play; all or with []

Keypresses

<u>Key</u>	<u>Qualifier</u>	<u>Description</u>
Mid. click+drag		If zoom in, pans the view
Right click		Brings up zoom%/auto;reset cam/proj;hide ctrls
Lt clk Zoom		Zooms in
Lt clk Zoom	Ctrl	Zooms out

Lt clk Zoom '+' or '=' '-'	Shift-alt	Zooms out Zooms in video output Zooms out video output
f F11 F12 Click label Click In/Out	Hold LMB Hold LMB	Toggle on/off fullscreen Reset camera Reset projector Drag label to elsewhere Drag In/Out elsewhere

Timebar

<u>Key</u>	<u>Qualifier</u>	<u>Description</u>
Left mouse btn	Ctrl	Changes to the next time format
Middle mouse btn	Ctrl	Change to the previous time format
Left mouse btn		Moves cursor position on the timeline
Double click		If between labels, selects that section
Click label	Hold LMB	Drag label to elsewhere
Click In/Out	Hold LMB	Drag In/Out pointer elsewhere
Fat arrow end	Hold RMB	Preview region drag on blue-colored bar

Transport & Edit Buttons (plus 3 below)

	Same as in Main window
	Same as in Main window
'Left click'	Only Protect enabled, start/stop forward play
Middle click	Only Protect enabled, start/stop play reverse
Middlewheel	Only Protect enabled, plays forw/rev 1 frame
'p'	Turns on/off Click to Play button

Viewer

BUTTONS

<u>Menu Items</u>	<u>Shortcuts</u>	<u>Description</u>
Rewind	'home'	Jump to beginning
Fast reverse	'KP +' / Alt-p	Toggle double speed reverse play
Normal reverse	'KP 6' / Alt-o	Toggle normal speed reverse play
Frame reverse	'KP 4' / Alt-u	Render previous frame
Stop	'KP 0' / Alt-m	Stop playback
Frame forward	'KP 1' / Alt-j	Render next frame
Normal forward	'KP 3' / Alt-l	Toggle normal speed forward play
Fast forward	'KP Enter' / A-;	Toggle double speed forward play
Jump to end (No button)	'end'	Jump to end
(No button)	'KP 5' / Alt-i	Toggle slow speed reverse play
	'KP 2' / Alt-k	Toggle slow speed forward play
	'KP#-Shift'	Shift added to KP #, adds or subtracts audio
	'Ctrl-KP#'	If [] set, "KP2,3,5,6,+enter" play between ptrs.
	+Shift w/Alt	Shift with Alt+x above, adds/remove audio
	+Shift w/Ctrl	Loops play; all or between In/Out
	'Left click'	Start or stop play forward
	Middle click	Start or stop play reverse
	Middle wheel	Play forward or reverse 1 frame
	'p'	Turns on/off Click to Play button
In point	'[' or '<'	Toggle In point timeline marker
Out point	']' or '>'	Toggle Out point timeline marker
	'Ctrl-t'	Clear both In and Out pointers
Splice	'v'	Create splice
Overwrite	'b'	Overwrite

End	'Shift'	Selects to end of line
Home		Goes to the beginning of the line
Home	'Shift'	Highlights from ibeam to beginning of line
Backspace		Deletes the character in back of the ibeam
Delete		Deletes the next character
Ctrl-c	(also+'Shift')	Copies selection to the 2nd cut buffer
Ctrl-v	(also+'Shift')	Pastes selection from the 2nd cut buffer
Ctrl-x	(also+'Shift')	Copies selection to 2nd cut buffer and delete
	Double click	On word in textbox, highlites that word
	Double click	On highlighted Probe order, enables/disables
	Triple click	In subtitle script text, loads line into line text

Unicode	Ctrl-U	'Shift'	Go into Unicode enter mode
	Enter		Fires the Unicode
	Esc		Cancels Unicode state
	Backspace		Deletes previous character
	A-F, 0-9		Enters the hex code

Resources Window	Key	Qualifier	Description
	Drag Mouse	'Shift'	Hold down Shift and drag to move text icon
		Double click	Draw vicons off- loads media in Viewer
		Single click	Draw vicons on- bigger thumbnail plays or stops
		Middle click	Draw on – brings up Inter-View modeling
		Wheel move	Draw on – increase/decrease preview size
		Mouse over	Draw on – instantly brings up thumbnail
	'v'		Switch to next Display option
	'o'		Opens Load files... menu
	- load tool	'Ctrl-a"	Selects all files
	- load tool	'Ctrl-z'	Deselects any highlighted files
	DEL key		Resources/Media, deletes highlighted from project
	DEL key	'Shift'	Resources/Media, deletes highlighted from disk + OK
	DEL key	'Ctrl+Shift'	Resources/plugin, deletes highlighted from disk + OK

Overlays Window	Key	Qualifier	Description
	LMB	'Shift'	Turns off all else, or reverts to previous state

Record Window	Menu Items	Shortcuts	Description
	Start/pause	'space bar'	Start and pause recording of current batch
	Label	'l'	Toggle label at current position
		Double click	On channel in window, pops edit ch window

Channels BUTTONS	Menu Items	Shortcuts	Description
	Select		Select a channel
	Add...		Add another channel
	Edit...		Edit the wording associated with the channel
	Move up		Move the channel up in the selection
	Move down		Move the channel down in the selection
	Sort		Sort all of the channels numerically
	Scan		Scan the available channels
	Delete		Delete the selected channel

Picture...

Double click Select the device picture
Double click In Find window, cursor moves to hilite program
Double click On batchbay line, selects and loads ch info

**Remote
Control**

MENU
KPTV
KPBOOK
KPHAND

Toggle remote control state (grabs/ungrabs keys)
Open or close the record window
Start/stop dvb channel scan
Quit cinelerra program

Commercial

'Shift-A' Edit commercial

16 - Configuration, Settings and Preferences

The user's default settings and preferences are retained across sessions in a file called `.bcast5` in the user's `$HOME` directory. Initially when cinelerra is launched there is an empty project and there are program default settings, and from then on the `.bcast5` directory will contain the settings that were set when quitting. If you need to revert to the default settings, delete the `.bcast5` directory contents and restart Cinelerra. Or you may want to rename it temporarily if you think you might want it back later.

Although the location defaults to `$HOME/.bcast5`, you can use the `CIN_CONFIG` variable to override this location. For example: `export CIN_CONFIG=/tmp/.bcast5` will use a temporary setup for testing purposes. It is also useful for multiple users sharing the same home directory who would like to have different configuration/preferences settings data. And if you are experiencing inexplicable errors or crashes in cinelerra, they may be due to a problem with `.bcast5` in which case taking it out of the picture can at least eliminate this as the cause.

Several ways exist to change Cinelerra's operational characteristics. A lot of variations can be made to settings and preferences by using the Settings pulldown from the main window and choosing Preferences.

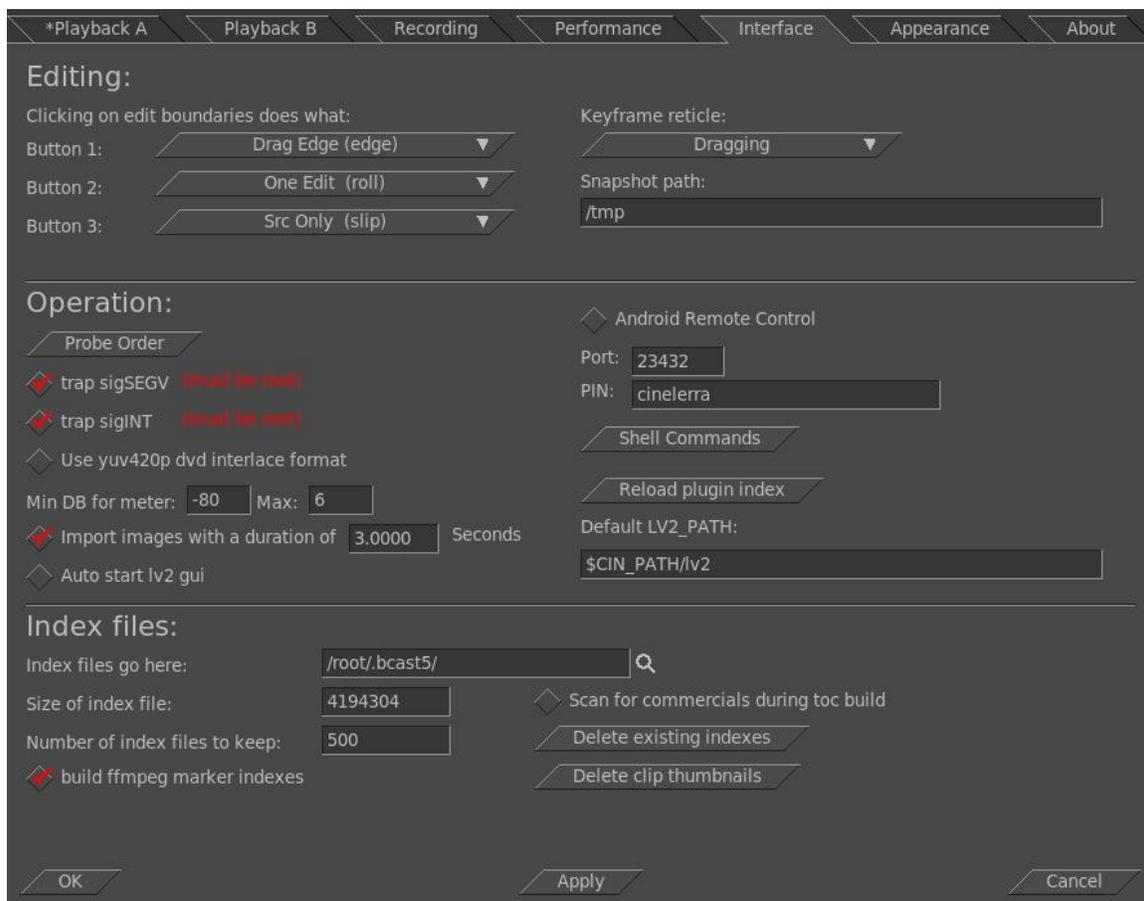


Figure 124: Settings->Preferences window with Interface tab displayed

16.1 Playback A / Playback B

Audio Out section

The audio drivers are used for both recording and playback. The Audio Out settings affect the outcome when you play sound on the timeline.

Playback buffer samples - for playing audio, small fragments of sound are read from disk and processed sequentially. A larger value here causes more latency when you change mixing parameters but yields more reliable playback. Some sound drivers do not allow changing of the fragment, so latency is unchanged no matter what the value. Since different stages of the rendering pipeline can change the rate of the incoming data, it would be difficult to disconnect the size of the console fragments from the size of the fragments read from disk.

Audio offset (sec) - the ability to tell the exact playback position on Linux sound drivers is poor. The audio offset allows users to adjust the position returned by the sound driver in order to reflect reality. The audio offset does not affect the audio playback or rendering at all. It merely changes the synchronization of video playback. The easiest way to set the audio offset is to create a timeline with one video track and one audio track. Expand the audio track and center the audio pan. The frame rate should be larger than 24 fps and the sampling rate should be greater than 32000. The frame size should be small enough for your computer to render it at the full framerate. Highlight a region of the timeline starting at 10 seconds and ending at 20 seconds. Drop a gradient effect on the video track and configure it to be clearly visible. Drop a synthesizer effect on the audio and configure it to be clearly audible. Play the timeline from 0 and watch to see if the gradient effect starts exactly when the audio starts. If it does not, expand the audio track and adjust the nudge. If the audio starts ahead of the video, decrease the nudge value. If the audio starts after the video, increase the nudge value. Once the tracks play back synchronized, copy the nudge value to the audio offset value in preferences. Note: if you change sound drivers or Disable hardware synchronization, you will need to change the audio offset because different sound drivers are unequally inaccurate.

View follows playback - this causes the timeline window to scroll when the playback cursor moves. This can slow down the X Server or cause the timeline window to lock up for long periods of time while drawing the assets.

Disable hardware synchronization - most sound cards and sound drivers do not give reliable information on the number of samples the card has played. You need this information for synchronization when playing back video. This option causes the sound driver to be ignored and a software timer to be used for synchronization.

Audio playback in realtime priority (root only) – for really old computers, this setting allows uninterrupted playback during periods of heavy load. It forces the audio playback to the highest priority in the kernel. Today, it is most useful for achieving very low latency between console tweaks and sound card output. You must be root to get real-time priority. Only experts might want to use this because it interferes with ordinary time-share scheduling and can lock up the system. When this is enabled, audio gets the first shot and burns audio until audio lets go. To explain, there are 2 kinds of scheduling, “time-sharing” which is the default, and “real time” where the scheduled task must explicitly request scheduling to allow other tasks to execute. Time-share interrupts when you use up your allocated time slice. Realtime priority audio will execute audio decode until it finishes, which may slow down other types of processing like video decoding. Most decoders use a policy that video may be downsampled to accommodate scheduling, but will never skip audio because it creates a much more obvious defect. This feature helps to make sure audio gets priority over video during decode. Be sure to check apply in order for this feature to take effect.

Map 5.1 → 2 – playback 5.1 → 2 driver downmix maps 6 tracks to 2 channels when checked, that is mixes 5.1 down to stereo on the output device side. This is different from the patchbay and menubar functions which reset the pan/mix levels of the input channels. In this way, you can render 5.1 media, and use stereo speakers to listen in the same session setup. This downmix only occurs if the playback is 5.1 (6 channels) and the device config is stereo (2 channels).

Gain – set audio gain to a different value than the default of 1.0

This feature, device level gain, corrects for hardware conditions which some devices may need to be useful. For example, you may need to increase the gain for a weak microphone or a noisy speaker, since it affects rendering when you crank up or down the audio via use of the patchbay. With the audio H/W gain support, you have the ability to fine tune the audio volume by some numerical value for the scale. You are adjusting the scaling of data into the audio driver - H/W scaling is done before it goes into or out of the driver. This is a one time linear multiplication of the sample values, and may offer better control than the logarithmic DB gain controls of the application.

Audio driver - there are many sound drivers for Linux. This allows selecting one sound driver and setting parameters specific to it. The currently available possibilities are listed next.

ALSA - is the most common sound driver these days and supports almost all sound cards. ALSA is frequently updated but is very stable.

OSS – was one of the first Linux sound drivers and has an open source implementation with many sound cards supported.

OSS Envy24 – is the commercial version of OSS with a variant for 24 bit 96 KHz sound cards. This variant required changes to the way the sound drivers were used and so needed a different driver.

Raw 1394, DV 1394, IEC 61883 are older audio drivers used by camcorders and not much else.

Device – with the down arrow, you can see the device choices on your computer.

Bits - 8, 16 or 24 Bit Linear are the current choices for the number of bits of precision Cinelerra should set the device for. The meaning of the number of bits can be misleading. Some sound drivers need to be set to 32 bits to perform 24 bit playback and will not play anything when set to 24 bits. Other sound drivers need to be set to 24 bits for 24 bit playback.

Stop playback locks up - this ALSA only checkbox is needed if stopping playback causes the software to lock up. This has worked some time ago, but may no longer work as expected.

Video Out section

The video drivers are used for video playback in the compositor and the viewer. These determine how you will see video on the timeline.

Play every frame - this causes every frame of video to be displayed even if it means that the playback of the video tracks fall behind. Most likely you will want this enabled because, after all, in order to edit you want to see each frame. However, if you are just watching a big video, you can switch to not play every frame so that you can at least not be distracted by slowness.

Framerate achieved - the number of frames per second being displayed during playback. This is updated during playback only. The goal is to get as close to the frame rate as possible, even if Play every frame is not enabled.

Scaling equation: Enlarge / Reduce - this algorithm is used when video playback involves scaling or translation. This does not affect 1:1 playback. Choices available are:

Nearest Neighbor / Nearest Neighbor - low quality output with fast playback. Often produces jagged edges and uneven motion.

Bicubic / Bicubic - Bicubic interpolation is used for both enlarging and reducing, enlarging blurs slightly but does not show stair step artifacts.

Bicubic / Bilinear - High quality output with slow playback. Bicubic interpolation is used for enlarging, which blurs slightly but does not show stair step artifacts. A bilinear interpolation is used for reduction, which produces very sharp images and reduces noise. Bilinear reduced images can be sharpened with a sharpen effect with less noise side effects than a normal sized image.

Bilinear / Bilinear - when slight enlargement is needed, a bilinear enlargement looks better than a bicubic enlargement. Bilinear uses less CPU than either Bicubic or Lanczos.

Lanczos / Lanczos - is not necessarily a general purpose upscaler, but is intended for low resolution sources. However many people like the sharpening effects. More quality from Lanczos does take more CPU.

DVD subtitle to display - DVD IFO files usually contain subtitle tracks. These must be decoded with the MPEG decoder. Select Enable subtitles to enable subtitle decoding. There are usually multiple subtitle tracks indexed by number and starting from 0. Enter the index number of the subtitle track to be decoded in the "DVD Subtitle to display" text box or use the tumbler to increase the index value. Go to the asset corresponding to the MPEG file in the Resources window and right click. Click on Info. The number of subtitle tracks is shown at the bottom.

Enable subtitles/captioning – for broadcast TV ??

Label cells - ??

TOC Program No – Table of Contents program number used in DVB ??

Interpolate CR2 images - enables interpolation of CR2 images. Interpolation is required since the raw image in a CR2 file is a Bayer pattern. The interpolation uses dcrw's built-in interpolation and is very slow. This operation can be disabled and the Interpolate Pixels effect used instead for faster previewing.

White balance CR2 images - this enables white balancing for CR2 images if interpolation is also enabled. This is because proper white balancing needs a blending of all 3 primary colors. White balance uses the camera's matrix which is contained in the CR2 file. Disabling white balancing is useful for operations involving dark frame subtraction. The dark frame and the long exposure need to have the same color matrix. If you disable Interpolate CR2 Images and use the Interpolate Pixels effect, be aware the Interpolate Pixels effect always does both interpolation and white balancing using the camera's matrix, regardless of the settings in Preferences. Dark frame subtraction needs to be performed before Interpolate Pixels.

Video driver - normally video on the timeline goes to the compositor window during both continuous playback and when the insertion point is repositioned. Instead of sending video to the Compositor window, the video driver can be set to send video to another output device during continuous playback. However, this does not affect where video is routed when the insertion point is repositioned. Options are listed next.

X11 - this was the first method of graphical display on Unix systems. It just writes the RGB triplet for each pixel directly to the window. It is useful when graphics hardware can not handle very large frames. And when X11 is used with the associated checkbox enabled of "use direct x11 render if possible" it can be a really good playback method to speed up playback for large frames.

X11-XV - this was an enhancement to X11 in 1999. It converts YUV to RGB in hardware with scaling.

In some cases it may be the preferred playback method, but it can not handle large frame sizes. Maximum video size for XV is usually 1920x1080.

X11-OpenGL - the most powerful video playback method is OpenGL. With this driver, most effects are done in hardware with the graphics board installed in the computer. OpenGL allows video sizes up to the maximum texture size, which is usually larger than what XV supports, depending on the graphics driver. OpenGL relies on PBuffers and shaders to do video rendering. Plugins or transitions that do not have “handle OpenGL” in the code will use software instead of hardware and this will slow down playback.

OpenGL does not affect rendering. It just accelerates playback. X11-OpenGL processes everything in 8 bit color models, although the difference between YUV and RGB is retained. The scaling equation set in the preferences window is ignored by OpenGL – it always uses linear scaling. Camera and projector operations use OpenGL, but some of the effects may not support OpenGL acceleration.

Raw 1394, DV 1394, and IEC 61883 are for old camcorders.

Default A/B Display - the interface is intended for dual monitor displays. Depending on the value of Display, the Compositor window will appear on a different monitor from the rest of the windows.

16.2 Recording

The parameters here expedite the File->Record... function by allowing the user to pre-configure the file format and the hardware used for recording, since the hardware generally determines the supported file format. Once set, the file format is applied to all recordings.

File Format section

File Format - this determines the output file format for recordings. It depends heavily on the type of driver used. The menu selections are the same as those of the rendering interface.

Record audio tracks - toggle must be enabled to record audio.

Record video tracks - toggle must be enabled to record video. The wrench button left of both the audio and video tracks toggle, opens a configuration dialog in order to set the compression scheme (codec) for each audio and video output stream. The audio and video is wrapped in a container format defined by the File Format menu. Different wrappers may record audio only, video only, or both. Some video drivers can only record to a certain container. If the video driver is changed, the file format may be updated to give the supported output. If you change the file format to an unsupported format, it may not work with the video driver.

Realtime TOC – setup for DVB recording to automatically generate a Table of Contents. This will scan the stream data "on the fly" on its way to being written while the asset is being captured. ??

Audio In section

Record Driver - is used for recording audio in the Record window. It may be configured the same as the Record Driver for video if the audio and video are wrapped in the same stream. Available parameters vary depending on the driver. The drivers are the same as described in Playback A/B with the addition of DVB and V4L2 MPEG but no Raw 1394.

DVB Adapter – name of a suitable DVB adapter for linux that is usb connected to your computer and has a connected broadcast TV antenna. ??

dev – your DVB adapter device number, which is usually 0.

Bits – same as described in Playback A/B audio section.

Follow audio config - ??

Samples read from device - a good value is 2048 or approximate dev buffer size (2k-16k probably).

Samples to write to disk - at a time. First, audio is read in small fragments from the device. Then, many small fragments are combined into a large fragment before writing to disk. The disk writing process is done in a different thread. The value here determines how large the combination of fragments is for each disk write. A good starting value is 48000 but this will most likely automatically change, probably to 44100 if necessary.

Sample rate for recording - regardless of what the project settings are, the value set here will be the sample rate used for recording. The sample rate should be set to the highest value the audio device supports.

Channels to record – usually set to 2.

Map 5.1 → 2 – leave unchecked to record all possible channels.

Gain – usually leave at default of 1.0, but this device level gain corrects for hardware conditions on some devices which need help. This gives you the ability to fine tune the audio volume by some numerical value for the scale. It is useful as better explained for the Gain in the Playback A/B Audio section discussed previously.

Record in realtime priority (root only) - only experts might want to use this because it interferes with ordinary time-share scheduling and can lock up the system. When this is enabled, audio gets the first shot and burns audio until audio lets go.

Video In section

Record driver - used for recording video in the Record window. It may be configured the same as the Record Driver for video if the audio and video are wrapped in the same container. Available parameters vary depending on the driver. The drivers available are as follows.

Video4Linux2

JPEG webcam

YUYV webcam

Video4Linux2 JPEG

Video4Linux2 MPEG

Screencapture

DV1394

IEC61883

DVB

DVB Adapter – name of a suitable DVB adapter for linux that is usb connected to your computer and has a connected broadcast TV antenna. ??

dev – your DVB adapter device number, which is usually 0.

Follow video config - ??

Frames to record to disk at a time - frames are recorded in a pipeline. First, frames are buffered in the device. Then, they are read into a larger buffer for writing to disk. The disk writing is done in a separate thread from the device reading. For certain codecs the disk writing uses multiple processors. The value set here determines how many frames are written to disk at a time.

Frames to buffer in device - the number of frames to store in the device before reading and this determines how much latency there can be in the system before frames are dropped.

Positioning – if set to *Software timing*, use software for positioning information. Video uses audio for synchronization, but most sound cards do not give accurate position information so selecting this option makes Cinelerra calculate an estimation of audio position in software instead of hardware for synchronization. You can also choose *Presentation Timestamps*, *Device Position*, or *Sample Position*.

Sync drives automatically - for high bitrate recording, the disk drives you use may be fast enough to store the data but your operating system may wait several minutes and stall as it writes several minutes of data at a time. This forces the operating system to flush its buffers every second instead of every few minutes to produce slightly better real-time behavior.

Size of captured frame - is the size of the recorded frames in pixels. It is independent of the project frame size because most video devices only record a fixed frame size.

Frame rate for recording - the frame rate recorded is different from the project settings. This sets the recorded frame rate.

16.3 Performance

The main focus of the performance section is rendering parameters not available in the rendering dialog with the obvious gain of perhaps better performance.

Cache size - to speed up rendering, several assets are kept open simultaneously. This determines how many are kept open. A number too large may exhaust your memory rapidly. A number too small may result in slow playback as assets need to be reopened more frequently.

Seconds to preroll renders - some effects need a certain amount of time to settle in. Checking this option sets a number of seconds to render without writing to disk before the selected region is rendered. When using the render farm, you will sometimes need to preroll to get seamless transitions between the jobs. Every job in a render farm is prerolled by this value. This does not affect background rendering because background rendering uses a different preroll value.

Force single processor use - Cinelerra tries to use all processors on the system by default, but sometimes you will only want to use one processor, like in a render farm client. This forces only one processor to be used. The operating system usually uses the second processor for disk access. The value of this parameter is used in render farm clients.

Project SMP cpus – to restrict the number of processors utilized, change the count number. This number will be used for the plugin per load balance operation `cpu limit`, which uses `smp-cpus` to stripe your data. It does not affect the number of cpus used in any other cinelerra operation besides plugins. On large cpu systems, it can come in handy to downgrade the number of cpus used for some plugins; otherwise it uses all of the processors and splits up the program into too many pieces which may add considerable overhead in high cpu count systems.

Background Rendering section

Use background rendering – checking this box, enables automatic background rendering. This works in conjunction with the interactive function Settings menu -> Toggle background rendering which sets the point where background rendering starts up to the position of the insertion point.

Frames per background rendering job - this only works if a render farm is being used; otherwise, background rendering creates a single job for the entire timeline. The number of frames specified here is scaled to the relative CPU speed of rendering nodes and used in a single render farm job.

Frames to preroll background - the number of frames to render ahead of each background rendering job. Background rendering is degraded when preroll is used since the jobs are small. When using background rendering, this number is ideally 0. Some effects may require 3 frames of preroll.

Output for background rendering - background rendering generates a sequence of image files in a certain directory. This parameter determines the filename prefix of the image files. It should be on a disk, accessible to every node in the render farm by the same path.

File format - the file format for background rendering has to be a sequence of images. The format of the image sequences determines the quality and speed of playback. JPEG is a good choice usually.

Video wrench – this has the single option of “use alpha”. It is by default unchecked.

Render Farm section

In the Render Farm Section are many options that are explained in detail in the Rendering chapter of this manual. Just make sure if you do not intend to use a render farm, that *Use render farm* is not checked.

16.4 Interface

Editing section

Clicking on edit boundaries does what - Cinelerra not only allows you to perform editing by dragging edit boundaries, but also defines three separate operations that occur when you drag an edit boundary. Here you can select the behavior of each mouse button. The usage of each editing mode is described in great detail in the Editing chapter.

Keyframe reticle – the options are *Newer*, *Dragging*, or *Always*. This is used to help in checking edit alignment across tracks. *Always* renders a line over all plugins, and *dragging* only over the drag icon. *Never* draws nothing.

Snapshot path – designates the default directory path for snapshot and grabshot generated output.

Operation section

Probe Order – clicking on this box brings up a popup allowing you to change the probe order usually for media that is raw camera output but it is also helpful if you want to ensure that a specific driver is used for certain media; for example you may want tiff files to be read natively instead of by ffmpeg.

trap sigSEGV – always enable this so that if cinelerra crashes, a dump will be generated for analysis.

trap sigINT – always enable this so that you can use Ctrl-c to interrupt the program if it appears to be hanging. This will often generate some useful information for analysis.

Use yuv420p dvd interlace format – for DVD media this option maintains the interlacing in Chroma sample addressing, which ordinarily would be deleted because the upsampling of interlaced chroma fields is normally done using a progressive algorithm. With this mode enabled, the MPEG decoder uses a different algorithm for interlaced frames so that the 4:2:0 format chroma interlacing is preserved.

Min / Max DB for meter – Min DB is useful because some sound sources have a lower noise threshold than others. Everything below the noise threshold is meaningless. This option sets the meters to clip below a certain level. Max DB sets the maximum sound level represented by the sound meters. This value is presented merely to show how far over the limit a sound wave is. No matter what this value is, no sound card can play sound over 0 dB.

Import images with a duration of # seconds – when you load single images, like png or jpeg, automatically load for # number of seconds. This makes it easier to see an image on the timeline.

Auto start lv2 gui - some lv2 plugins display a "glitzy" UI (User Interface); for example the Calf plugins. For these LV2 plugins, if you want that to automatically come up without having to click on the UI button on the simplified ui interface, this is the flag to enable that.

Android Remote Control – check this to enable using an android device as a remote control for broadcast TV.

Port - default port 23432 is used for the android remote control.

Pin - default PIN cinelerra is used for the android remote control.

Shell Commands – this button brings up the controls for setting up your own shell commands or editing previously set up commands. See the section on Menu Bar Shell Commands for information.

Reload plugin index – execute this reload command when you have modified plugins and want to make sure your changes take effect.

Default LV2_Path - when there is no system LV2_PATH set, if you want lv2 plugins loaded, you must set the correct directory path name here. When you change this field, cin will automatically restart and load the newly specified lv2 plugins. I

Index Files section

Index files go here - index files exist in order to speed up drawing the audio/video tracks. This option determines where index files are placed on the disk.

Size of index file - determines the size of an index file. Larger index sizes allow smaller files to be drawn faster, while slowing down the drawing of large files. Smaller index sizes allow large files to be drawn faster, while slowing down small files.

Number of index files to keep - to keep the index directory from becoming very large, old index files are deleted. This determines the maximum number of index files to keep in the directory.

build ffmpeg marker indexes - improves ffmpeg seeks in certain cases although not clear which ones.

Scan for commercial during toc build – used for working with broadcast TV commercial removal.

Delete existing indexes - when you change the index size or you want to clean out excess index files, this deletes all the index files.

Delete clip thumbnails – as clip thumbnails accumulate over time, you may want to delete them to get the disk space back.

16.5 Appearance

Layout section

Theme - Cinelerra supports 10 different themes to suit the preferences of different users. When you change the theme, Cinelerra automatically saves your session and restarts exactly where you were. The Themes User Interface are described in more detail next. Akirad themes are all available in Cin-GG with major thanks to the Cinecutie project for all of their hard work. These modifications provide alternative User Interface environments and allows you to choose your favorite look. Be aware that when you change your theme, Cinelerra will automatically shutdown and restart for it to take effect. The theme you choose is preserved across sessions.

Theme choices available:

Blond
Blond-cv
Blue
Blue Dot

How To Change your Theme:

To change a “Theme” in the main window pulldown, select Settings → Preferences then in Preferences window, click on the “Appearance” tab. In the Editing section in the lower left hand corner,

Bright
Hulk
Neophyte
PinkLady
S.U.V. (default)
Unflat

click on the down arrow next to Theme to see your choices.
Click on your desired choice from the list given.
Check OK, cinelerra will automatically shutdown and restart.

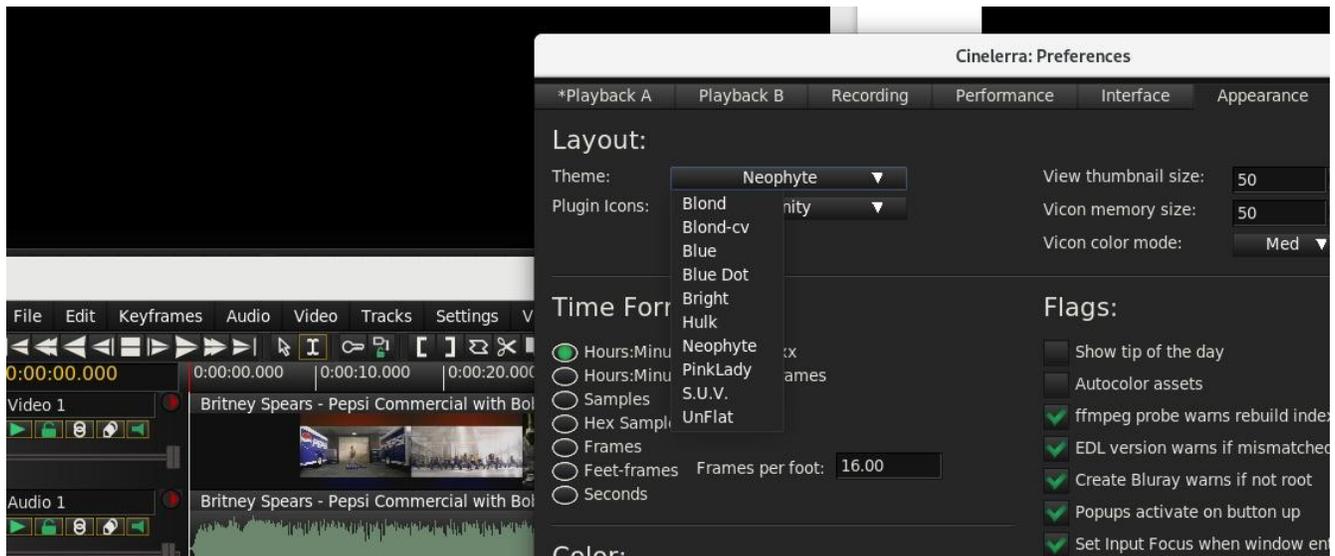


Figure 125: Shows the Neophyte theme (courtesy Olaf)+ Preferences window with list of themes

Plugin Icons – there are currently 4 choices for different plugin icons to include the old original.

View thumbnail size - you can increase or decrease the thumbnail size to use more or less cpu.

Vicon memory size - you can increase the memory used for thumbnails to get more clarity of pixels.

Vicon color mode – modify the color mode to Low, Medium, or High for the thumbnails.

Time Format section

Various representations of time are given so that you can select the most convenient one for you. The time representation can also be changed by Ctrl-clicking on the timebar in the main window.

Color section

Highlighting Inversion color – modify the selection area color; default is ffffff which is white. When you make a selection, that area becomes an inverse image which by default becomes a whitish color. You can set it to a different color by modifying the hex value in the box next to "Highlight inversion color". Keep in mind that if you set the value to a low value, you will not be able to see the outlined selected area (for example the hex value "f" is not readily visible and leads to confusion). A leading 0 or blank is not allowed and will be automatically changed to ffffff.

YUV color space – default is BT601; others BT709 (high definition), BT2020 (ultra high definition).

YUV color range – JPEG and MPEG

Flags section

This section contains many useful options to cater to the various preferences of individual users.

Show tip of the day – if checked, a tip will be displayed in a popup box when start up cinelerra.

Autocolor assets – to make it visually easier to see your clips on the timeline that are from the same media file, you can have them automatically colored. Use of this feature requires additional memory and cpu on every timeline redraw, therefore smaller computers may not want this checked on.

ffmpeg probe warns rebuild indexes - ??? ask bill

EDL version warns if mismatched – in the case of a Batch Render, it is often helpful to be warned if the EDL has been changed so that you are aware that what is going to be rendered is different than your current EDL session.

Create Bluray warns if not root – if checked and you are not logged in as root, you will get an error message in order to avoid doing a lot of work and then failing out because root is required for automount and to write on DVD hardware.

Popups activate on button up – this is the default but if unchecked, popups activate on button down.

Set Input Focus when window entered – this is checked on by default because on some operating system distros, when you move your mouse to a different window, nothing happens and you are left wondering why you can not enter information. When checked this causes the input focus to shift to any cinelerra window when the cursor enters an exposed region of the window which eliminates the need to switch input focus by tabbing.

Click to activate text focus – if checked, you will have to click to activate text focus.

Click to deactivate text focus – if checked, you will have to click to deactivate text focus.

Always show next frame – in this mode the insertion pointer reflects the same as the Compositor so that for playing forward, the result is what looks like 1 was added to the frame displayed in the Compositor window. This is fully explained in another section.

Use thumbnails in resource window - the Resource Window displays thumbnails of assets by default, but drawing asset thumbnails can take more time and CPU so you may want to uncheck this.

Perpetual session - is very useful for working on a project over many days so you can just quit before shutting down and the next time you start up Cinelerra you will be right back where you left off. You will retain all of your undo's and redo's.

Clears before toggle – when using copy/paste in drag and drop mode some users prefer to resort to the addition of the Ctrl key for adding multiple selections. By checking this flag, the user retains usage as is commonly done for listbox operations.

16.6 About

This section gives you information about the Cinelerra program and version you are running. The original author's copyright and name are first and foremost. Next is a textbox with additional information and a summary of the monthly new features of note. Below that is a summary of the GPL License and the fact that it is provided without any warranty. Then the licensing verbage is the item that you may need to refer to most often - the "built" date and time in case you need to know which version you are currently running.

16.7 Environment Variables for Customization

Environment variables are global variables in the shell which all applications can read. They are set with a command like "set VARIABLE=value" or export VARIABLE=value. Environment variables can be viewed with a command like "env". The values set can be removed with "unset VARIABLE".

The following exported variables can be set to customize your environment. The CIN_CONFIG variable could be extremely useful for testing purposes or for multiple users sharing the same home directory who would like different default preference settings.

CIN_BROWSER	name of browser to use by Shell Cmds options
CIN_CONFIG	configuration data; defaults to \$HOME/.bcast5
CIN_DAT	location of data files, such as documentation, models, tip of the day
CIN_LADSPA	LADSPA directory path
CIN_LIB	location of library programs, such as bdwrite
CIN_LOCALE	locale text domain path to use for translating text
CIN_PLUGIN	plugin directory path
CIN_RENDER	complete filename with path for use in Shell Scripts
CIN_XSYNCH	(for Developers only) set to 1 helps debugging for windows primitives to execute immediately and not be buffered up so you can see what is happening

One example: `export CIN_BROWSER=chrome` would override default firefox for Shell Cmds.
An example: `export CIN_CONFIG=/tmp/.bcast5` use a temporary setup for testing purposes.

`BC_USE_COMMERCIALS=1` to activate the commercial database (db) methods.
`BC_TRAP_LV2_SEGV` to get a dump of the failure of an LV2 plugin for help in debugging.

`BC_FONT_SCALE=2.0` for changing the default size of the characters to be twice as big.
`BC_ICON_SCALE=1.5` for changing the default size of the icons to be bigger; can change # to any. You can increase the size of the characters in the fonts and icons on your Cinelerra system. This will make it easier to read the characters if you have trouble seeing the default small letters, which have been auto-scaled based on the window geometry. The user-friendly font/icon scaling default is 1 but you can set it to any decimal value. To defeat default auto scaling and get any size characters/fonts, override the setting via the previous 2 listed shell environment variables.

`BC_FONT_PATH=<colon-separated-search-path-for-fonts>` to add additional font sets for the Title plugin or to remove all fonts set to `:` (colon). An example: `export BC_FONT_PATH=/usr/share/fonts`.
`BC_FONT_DEBUG=1` debug for determining which font is causing problems. `0` for no debug.

`LADSPA_PATH` specify an alternate set of ladspa plugins or include the default with the use of a colon separated list of directories to search for LADSPA plugins.
`LV2_PATH` specify a certain set of LV2 plugins to use. Separate multiple paths with colons.
`GLOBAL_PLUGIN_DIR` directory in which Cinelerra should look for native plugins. The default is `/usr/lib/cinelerra` but you may need an alternate directory if you share the same executable directory among many machines via NFS. Plugins of different binary formats need to be in different directories.

`LANG` and `LANGUAGE` Cinelerra can be localized to display menus and messages in many languages. Language settings are normally read from your linux O/S language settings. To run on a language different than the one selected on your system just change the `LANG` and `LANGUAGE` environment variables. For example, open a shell and type: `export LANG=es_ES` or `LANGUAGE=es_ES`, then run Cin from the same shell and you will have translations in Spanish.

`SHUTTLE_CONFIG_FILE` alternate shuttle configuration file.

17 - How some stuff works

This section describes in detail some areas of Cinelerra to help explain how things work.

17.1 Copy/Paste and Highlight Usage with Cinelerra

There are 3 types of copy/cut and paste methods which exist in X windows, and most modern programs use 2 of them. The 3 cases are:

- cut_buffers 0-7 - these are obsolete but they still work and are the simplest to use.
- highlighting - called primary selection; almost all clipboard programs only use this.
- copy and paste - called secondary (or clipboard) selection. Some more modern programs use ctrl-C/ctrl-X and ctrl-V for this (some use other keys or qualifiers, like shift).

How it works:

All of the methods use window "properties" to attach data, called a selection, to a source window. The program advertises the selection by using the X server. The window property used determines which selection type is set/advertised by the new selection.

When a paste is used in a target window, the target program requests the advertised selection data. This may access one of two buffers depending on which type of load/paste action is used. The user loads a cut buffer via drag select or ctrl-C/ctrl-X, and pastes a cut buffer via middle mouse press or ctrl-V.

Cinelerra cut and paste:

1) Text cut and paste operations:

To use a text selection, create a drag selection in textboxes by pressing and holding mouse button1 with the pointer over the beginning of the text selection, then move the pointer over the desired selection to the selection end, then release the mouse button. This continuously reloads the "primary" clipboard buffer and highlights the text selection. It can then be pasted to most programs by pressing the middle mouse button with the pointer over the text insertion point. Some examples of these programs are xterm, gnome-terminal, cinelerra, and browser input text boxes. After you create a text selection, if you then press ctrl-C the selection will also be copied to the secondary (or clipboard) selection buffer. This second paste buffer can be used for a more lasting save effect, since it will not be lost until you again press ctrl-C (copy). Using ctrl-X (cut) will also copy the selection to the secondary clipboard buffer, and then delete the selection from the textbox. If you press ctrl-V (or paste) in a target window, the secondary selection will be inserted at the target window cursor. If a text selection exists in the target window, it is replaced by the pasted text.

2) Media cut and paste operations:

To create a media selection, highlight a region on the cinelerra media timeline, then use the main menubar or compositor/viewer edit panel to operate the clip cut, copy, or copy-keyframe menu buttons. This selection can then be pasted to a target selection on the timeline using the main menubar or compositor/viewer edit panel to operate the clip paste or paste-keyframe operation. Also, by using the resource window you can select the "Clips" folder and right mouse the resources list box, then use the

"Paste Clip" menu item to paste the selection to a named clip. Additionally, these methods work between running instances of cinelerra, which means you can move media clips between the cinelerra program instances. The clip data is also copied to the secondary clipboard buffer. This makes it possible to examine the clip content directly if so desired.

3) Finally, the older `cut_buffer` method is used:

3.1) For text, if there is an active selection when a window closes, it uses `cut_buffer0`. Normally when a paste is performed, the target window "notifies" the selection owner to "send it now" when you do a paste, but if the window has closed there is no source window, so no pasting. Some programs, like cinelerra, use `cut_buffer0` as a fallback. This makes it possible to paste data from a closed window.

3.2) To move media clip, data `cut_buffer2` is used because it does not require the selection owner interface, and works simply and reliably. This buffer is not normally in use by other programs.

One final note, when a text selection is set, the selected text is redrawn using selected-highlight color when the textbox loses focus. This convenience feature shows the active text selection as you move the pointer to the new target window. When a new selection is set anywhere else on your screen, the current text selection will be redrawn using the inactive-highlight color as the textbox loses selection ownership. In most cinelerra themes, the drag selection text-highlight color is BLUE (#0000FF), the selected-highlight color is SLBLUE (#6040C0) - really sort of purple, and the inactive-highlight color is MEGREY (#AFAFAF).

17.2 Playing is Different than Seeking/Positioning!

Seeking targets and displays the next frame. The next frame is targeted because frame zero has no previous. When you seek, you reposition to just before the target frame, and since the play direction has not been established (there is no direction when seeking) it shows you the next frame. This produces the expected behavior when you seek to frame zero; you see the first frame. Seeking displays in the compositor what you are getting ready to work with/edit/etc.; always showing the next frame in relation to the cursor. Technically, since seeking just resets the position, it would be correct not to update the compositor, but it is best to seek and show the next frame to confirm that it is the frame you expected to see.

Playing shows you what has just been played in the compositor window. It is not the same as seeking. When you use keypad 1 to play frame forward, then the 1st frame that is played and shown in the compositor window is frame zero (which was already displayed). The position is incremented to 1. Press keypad 1 yet again, and the next frame displayed is 1, and the new position is 2, and so on. According to the implemented strategy, the insertion point moves to track playback. When playback stops, the insertion point stays where playback stopped. Thus, with playback you change the position of the insertion point.

Simple explanation of what you will be seeing in the compositor when playing:

Play forward - the frame to the right of the cursor in the timeline gets displayed.

Play backward - the frame to the left of the cursor in the timeline gets displayed.

The reason behind this "play" methodology is that you want to know what you just played so that you know what matches what you just saw/heard in case that is the desired stuff. You don't want the

compositor to show you what you have not yet played - you need to see this frame to analyze/check to see if it is what you want. This behavior applies to any playing operation, such as the keypad or Frame

forward/Frame reverse buttons. You can still easily see the actual insertion point in the zoombar at the bottom of the timeline - sixth button over or 3rd button from the right side. Also note the following:

Blinking insertion point on the timeline - seeking/positioning was the last operation.
Solid *non-blinking* insertion point on timeline - playing was the last operation.

Example and explanation based on the code:

- 1) open a small example of 10 numbered frames (or use the Title plugin to add a timestamp)
- 2) press f to "fit" the timeline
- 3) make sure settings->align_cursor_on_frames is set
- 4) seek to frame 4 by clicking on the timeline at position 4, compositor shows the 5th frame, since the media counts from 1 and the timeline counts from 0. This is correct behavior.
- 5) press KP1 to play next frame. According to playback strategy:
When play is forwards, the next unit is displayed, and the position is advanced one unit.
So the next frame is 4, (shows the 5th frame) it is displayed.
The position is advanced from 4 to 5. This is correct behavior.
- 6) press KP4 to play the previous frame. According to playback strategy:
When play is in reverse, the previous unit is displayed, and the position is reduced one unit.
So the previous frame is 4, (shows the 5th frame) it is displayed.
The position is reduced from 5 to 4. This is correct behavior.

If you watch the zoombar (bottom of main window) position, it shows the current position is just before the next frame to be displayed when going forwards, and just after the frame to be displayed when going backward.

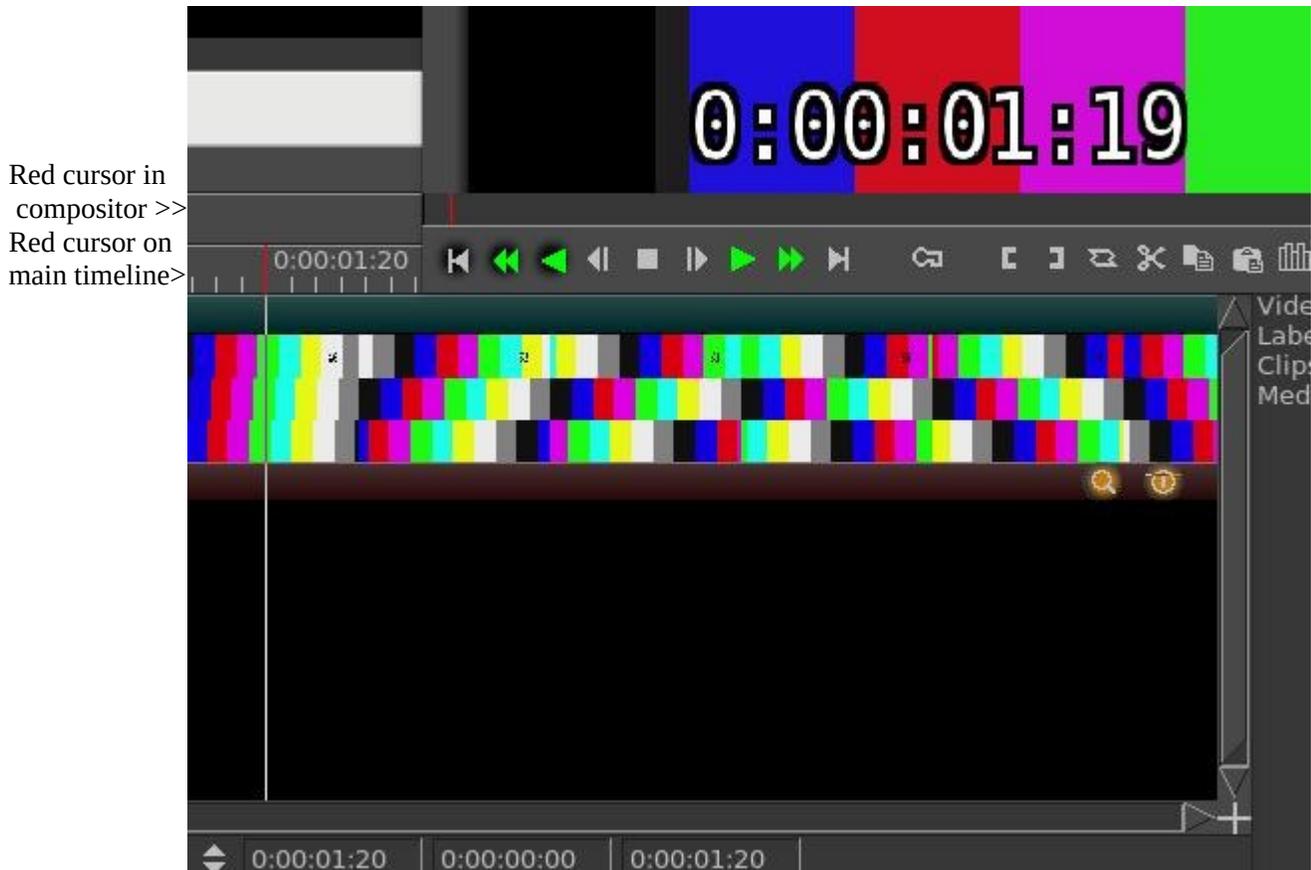
To recap, position is usually set in the program as a location that is between a previous and next frame/sample unit such that the next unit equals the seek target. After position is reset using a "seek" operation, the next unit is displayed, which is the seek target. When "play is forward", the next unit is shown, and the position is advanced one unit. When "play is in reverse", the previous unit is shown, and the position is reduced one unit. At the beginning, there is no previous, and at the end, there is no next, but silence is rendered at the end.

17.2.1 Always Show Next Frame

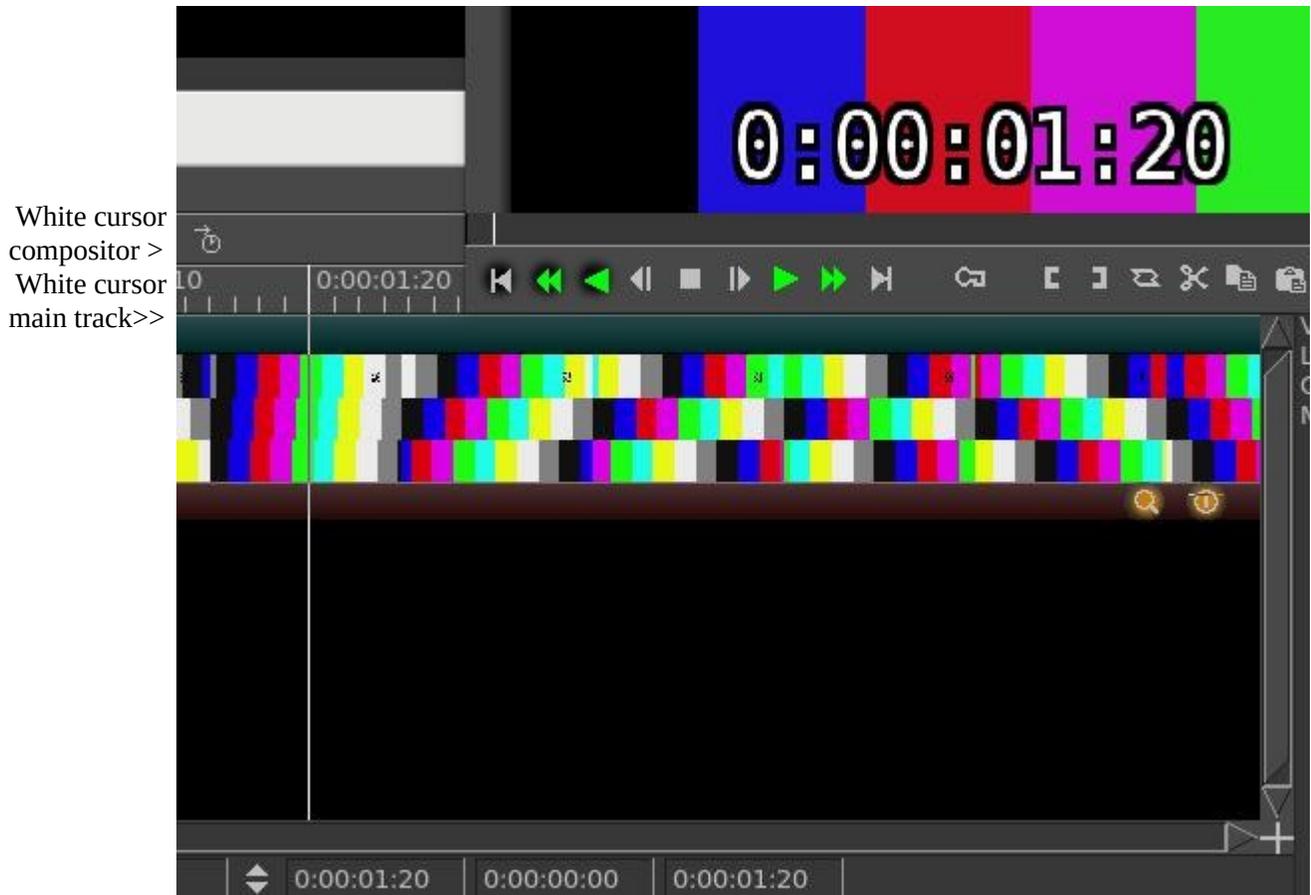
Since some users prefer the insertion pointer to reflect the same as the Compositor a choice is available. For playing forward, there is a preference option which results in what looks like 1 was added to the frame displayed in the Compositor window. To enable this mode, check the box "Always show next frame", and this will be saved to .bcast5. The option checkbox is in the Settings → Preferences-> Appearance tab and when checked, any forward "plays" in the Compositor window show the same frame as you would with a seek. Reverse plays and plays using a selection or In/Out pointers (with Ctrl) work the same as without this preference set. But you will no longer see the odd behavior where if you frame advance forward and then frame advance backward, the displayed frame does not change – instead it will change and look more natural.

A color indicator that shows in the main track canvas timeline and the compositor timeline reminds the user which mode is currently active. The cursor in the compositor turns red for default mode and white for “Always show next frame” mode. The top portion of the insertion cursor in the track canvas mirrors this, with red for default and white otherwise.

Screencast below using the default “playing” method where the frame in the compositor is the one that was just played; in this case play was in the forward direction. Note that the insertion pointer in the main track canvas shows 01:20 but the compositor show 01:19 so you know what you last saw. Also, the cursor/cursor tops in both windows is red.



Screencast below using the “Always show next frame” method where the frame in the compositor is the same one that would have shown with a seek; in this case play was in the forward direction. Note that the insertion pointer in the main track canvas shows 01:20 and the compositor shows 01:20. Also, the cursor/cursor tops in both windows is white.



17.2.2 Examples

(some examples go here)

17.2.3 Seeking Issues

If you have an issue playing a video and not seeing it in the Compositor (just see a black screen), it is most likely due to the media not being designed to be “editable”. It is most likely not damaged. Generally it just does not have keyframes which are needed for seeking which is what is done when you move around the media and start playing in the middle. The media plays just fine in the compositor if you always play from the beginning because then you don’t need keyframes to seek. You can get around this problem if you proxy the media. A good choice to use for the proxy would be “use scalar”, ffmpeg/mp4 and size of ½. The proxied media can then seek and you will see it play in the compositor because keyframes exist.

17.3 Color Space and Color Range Affecting Playback

Playback “single step” and “plugins” cause the render to be in the session color model, while continuous playback with no plugins tries to use the file’s best color model for the display (for speed). This can create a visible effect of a switch in color in the Compositor, usually shown as grayish versus over-bright.

The cause of the issue is that X11 is RGB only and it is used to draw the "refresh" frame. So single step is always drawn in RGB. To make a YUV frame into RGB, a color model transfer function is used. The math equations are based on color_space and color_range. In this case, color_range is the cause of the "gray" offset. The YUV mpeg color range is 16..235 for Y, 16..240 for UV, and the color range used by YUV jpeg is 0..255 for YUV.

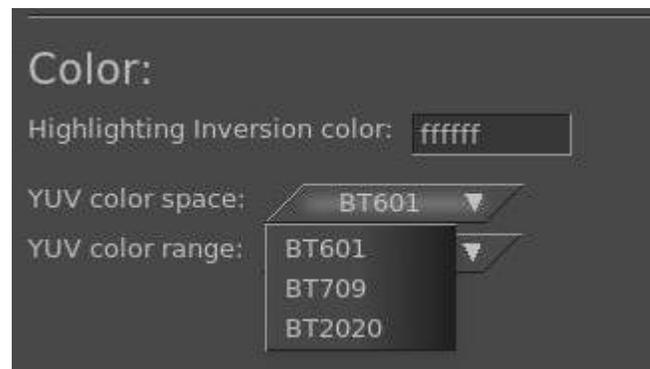
The mpeg YUV color range (16..235) looks sort of like an old TV if it is viewed on a display with jpeg range (0..255). A common expression for the short "mpeg" color range is "compressed" color range. If you are using color compressed data with no display decompression, or using uncompressed data and the display is configured to use compression, the color range will appear "squished" or "stretched", as too gray or too much contrast.

The "use X11 direct when possible" preference with X11 as your video driver, generally means that you value speed over color range. When you use this feature and the color range preference is mismatched, the color switch offsets will still appear. This is a personal choice solely for improved speed.

There is now program code to look for RGB versus YUV color model mismatches. You can override the default setup, which mirrors the original code, via the following.

Settings → Preferences Appearance tab in the lower left hand corner:

YUV color space – default choice is BT601, alternate is BT709 (High Definition), BT2020 (UHD)
YUV color range – default choice is JPEG, alternate is MPEG



17.4 Simple Animation - Filescene using Festival

This functionality was added to cinelerra by the original author to create simple animation. The file type for this animation is “Scene”.

To get started making a simple animated movie copy from the directory: <cin_path>/cinelerra/tests the text2movie and text2movie.xml. You can see what this does via File → Load... text2movie.xml. The file text2movie acts like a normal asset, except changes to it are immediately reflected on the timeline without reloading and the length is infinite. You can just edit the text2movie file to change the script. If the length of the movie increases, drag the right edit handle to extend the edit or use the pulldown Edit->edit length. There is one audio channel created for every character. The frame rate, sample rate, frame size, and camera angles are fixed. To see these values, right click on the asset and look at the “Asset info”.

Currently the functionality that is implemented focuses on dialog between two people. The models are defined in model files saved in Cinelerra's executable directory (for example, /opt/cinelerra/models). The character model and voice is selected separately in the script. The model files have the same name that appears in the script and are usually saved in the directory the script is in, but there is a defined search path, if not. You can create new models for the script without affecting the entire system. These models define the total size of the model along with the images used – the model images are 2D png images because all the animations are baked. Since there is no 3D renderer, no custom movement is supported.

There are currently 2 actions implemented:

- 1) Character2 can cut off character1 if character1's dialog ends in “...”.
- 2) Inserting “[pause]” anywhere causes the character to pause. This is useful for adjusting the timing of dialog.

This is “simple” animation so you can expect speech synthesis not to be that good. And you will have to adjust punctuation and spelling based on the sound. Since the dialog is rendered on-demand, there is a delay when each character starts to speak but you can split dialog into shorter blocks to reduce the delay.

17.5 Textbox Non-standard Character / Unicode Insertion

If you want to enter a special character - like a bullet, an accent grave character, or a mathematical summation symbol – you can use the unicode equivalent in a textbox to do so. In the textbox, keyin Ctrl-Shift-U which puts you into single character unicode mode, then keyin the numerical value for the intended single character followed by the carriage return. For a voluminous list of possible special characters, you can go to unicode-table.com on the internet to choose by highlighting a character to get its numerical equivalence. For example, U+2022 is a bullet. If you make a mistake, you can use the backspace key or if you want to exit unicode-insert-mode, use the ESC key. This feature is especially useful with the Title plugin and for naming Tracks in the main window.

However, it is worth mentioning that some special characters are available via the "compose" key in the current distribution. https://en.wikipedia.org/wiki/Compose_key.

This page intentionally left blank so that chapters begin on a right facing page.

18 - Troubleshooting

You can report potential problems, bugs, and crashes to the Cinelerra-GG website at:

<https://www.cinelerra-gg.org>

Here you can log the problem into the MantisBT bugtracker, or use the forum Q&A for help from other users, or email the problem using the address: cin@lists.cinelerra-gg.org . It is usually more helpful if instead of starting cinelerra from the icon, run from a terminal window so that if there are error messages related to the problem, they can be captured on the screen and passed along. The command to run from a terminal window is: <directory_path of where you installed cinelerra>/bin/cin – for example if installed at directory /mnt0/build5/cinelerra-5.1, you would execute the following command to start the program /mnt0/build5/cinelerra-5.1/bin/cin. The problem you are experiencing may be as simple as a problem with \$HOME/.bcast5 subdirectory so you may want to first delete or rename your current .bcast5 to start with default settings.

18.1 What to Include in Problem Reports

For the best help, if you have a reproducible problem the following list of descriptive items makes it possible to figure out what the problem is and how to fix it. It may be a usage/setup mistake or a real bug which a programmer might like to fix rather quickly. Although it is not always necessary to provide this much information, it eliminates the “try this” or “did you do this?” and a lot of “guessing” rather than heading straight for a solution. Even if you can not provide all of the most helpful information, please report the problem in case other users are experiencing the same issue.

Basically we **need to see what you see** with the input, output, and session file:

- 1) Provide a detailed description of the problem.
- 2) Supply a small representative sample of the original input that exhibits the issue.
- 3) If possible, also provide the rendered output, again using that representative sample.
- 4) Save a session file used with that same sample which will contain a lot of setup parameters; the best method to do this is to use the File pulldown of “Export Project...” with the Copy option. That way all of the files will be in 1 spot and easily loaded onto any other computer.
- 5) To make sure that the same rendering setup is used, it may be necessary to send an additional session file at the definition point just before rendering starts.
- 6) Include the Operating System name and version number and version of Cin that you are running; You can find the date/time “built” in the Settings → Preferences, About tab, bottom left corner.

It is better to upload any files to a datahostfile site as sometimes they can be too large.

How to create a session file: from the menubar pulldown File → Save_as... and choose a filename, for example /tmp/beforerender.xml . Do this after you have the situation setup so can see values in use.

How to create a representative small sample:

- open the input media, setup as you normally would to start the project render;
- select 5-15 seconds of media using the edit mode ptr1 button drag highlighting on the timeline;
- start the render dialog (Shift R), set the filename path, and set “Render range” to “Selection”;

- recommended “Insertion strategy” is “Replace current project”; then press the checkmark OK.
(The result will “load replace” the current project for review.)

18.2 Crash Dumps for Analysis

If you get a SEGV signal crash and can explain what steps you took, a /tmp/cinelerra_<pid>.dmp file is very useful for analysis. You can also use Ctrl-C in the controlling terminal window to force an INTR signal interrupt when you think the program is hung up. It is best to upload the .dmp, input files and other files to datafilehost or similar temporary site as they may be quite large. The .dmp filename looks like /tmp/cinelerra_<pid>.dmp. Log a bug report or email all helpful information on the crash, location of uploaded files, and list of setup steps that illustrate the problem to cin@lists.cinelerra-gg.org .

- 1) Use the latest version of Cinelerra as you want to make sure problem hasn't been fixed.
- 2) Be sure gdb is installed (usually installed but if not, easily downloaded/installed by anyone).
- 3) Run as root (this enables using gdb to create a /tmp/cinelerra*.dmp).
- 4) Be sure Settings->Preferences->Interface->trap SEGV/INTR are checked.
- 5) Invoke the error. This should create /tmp/cinelerra-<pid>.dmp .

For those who can not run as root or can not create the problem while running as root, another method to create a dump is available. Here are the steps to do so – you must first install gdb:

- 1) Temporarily login as root and keyin: `echo 0 > /proc/sys/kernel/yama/ptrace_scope` .
This opens up ptrace, used by gdb (the debugger). You will want to reverse this when done.
- 2) Run cinelerra as an ordinary user from a terminal window using the command line and try to create the problem. If the program crashes, this should produce a file named /tmp/cinelerra*.dmp where * is a number. The dump usually takes about 30 seconds or so to complete after the failure invokes the crash handler so be patient.
- 3) Email the output that shows in the terminal window from where you started up Cin and upload or directly email the /tmp/cinelerra*.dmp file (this file is too large for the Mailing List).
- 4) After you are finished creating/testing the problem, reverse the previous setting while logged in as root by keying in: `echo 1 > /proc/.../ptrace_scope` .

18.3 When Things go Wrong

Sometimes things go wrong and there are some ways to continue your work without much trouble. Below is a list of items to try before abandoning your session.

- 1) If you suddenly get hangs on media or strange looking tracks, you might want to rebuild indexes. You can do this in the Resources window with a right mouse click on the media to get a pulldown with the “rebuild index” option. Better yet, clear them for all of your media with Settings → Preferences, Interface tab, Index Files section, “Delete existing indexes”.
- 2) On an older computer, if you are playing media and it can not keep up, you can turn off “Play every frame” in the Video Out tab of Settings → Preferences, Playback A tab.
- 3) The buffer size can be lowered to 1048 if playback seems choppy or if you have problems with lv2 plugins, or you can increase the buffer size for better flow. This can be changed in Settings → Preferences, Playback A tab, under the Video Out section.
- 4) After saving your current session and exiting cinelerra, you might want to either rename or delete your current \$HOME/.bcast5 directory and start with the default setup. However, all of your

current preferences will be lost.

- 5) Some media has only a single keyframe at the beginning of the file so that moving anywhere on the timeline results in just a black image in the compositor. Cinelerra needs more keyframes to determine position. You can temporarily use Settings → Proxy to ½ size to put in keyframes.
- 6) If the rate at which frames are captured during recording is much lower than the framerate of the source, the video will accumulate in the recording buffers over time and the audio and video will become out of sync. Decrease the number of frames to buffer in the device in preferences->recording so the excess frames are dropped instead of buffered.
- 7) If loading files locks up, usually this is because Cinelerra is building picons for the Resources window. If you load a large number of images, it needs to decompress every single image to build a picon. Go into settings->preferences->interface and disable *Use thumbnails in resource window* to skip this process.
- 8) For an older computer with less CPU power, in Settings → Preferences, Appearance tab, be sure that *Autocolor assets* is disabled; set *View thumbnail size* & *Vicon memory size* & *Vicon color mode* to lower values or turn off *Preview mode* in the Resource window (the right arrow on the right side of the word *Visibility*).

Common Problems

Some messages you may see on a terminal window, may or may not be errors. Here are a few examples.

FFMPEG::open_decoder: some stream times estimated: /your directory/filename

This is not a problem Basically, when you open a file, if a stream has a known duration, there is no message. If the duration is unknown, it is estimated by using the File Size and Bitrate to estimate the duration. It may indicate that the stream number cinelerra uses versus ffmpeg may be counted differently).

AudioALSA::write_buffer err -32(Broken pipe) at sample #

This just indicates that there is something wrong with the audio. Some reasons for this are:

- 1) You simply stopped playing in Cin-gg while the audio is in progress.
- 2) Running on a computer where there is no sound card.
- 3) Incorrect setup of the audio parameters in the Settings->Preferences, Playback tab.
- 4) Your sound system is already in use by another program, like when playing "tunes" outside cinelerra.

Playback does not keep up so you hear the audio ahead of the video.

When this occurs, it most likely means that there is not enough CPU power to keep up. You can go into Settings → Preferences, the Performace tab and uncheck “Play every frame”. Not playing every frame means that as it plays, if it can not keep up with doing all it has to do to change the frame into a vframe and then draw it, the program just notes that it is falling behind so skips frames to catch back up. This does not affect rendering. A different solution is to use Background Rendering as described elsewhere.

Cin will not start and produces error message on the startup window.

There can be various reasons that Cin does not come up. Some of the recent reasons are listed here. Please notify the www.cinelerra-gg.org website when you have any unknown startup issues.

- 1) Cin doesn't come up in Debian with compiz window manager running. Workaround is to use a different window manager or bring up cin first and then compiz. There is also a report that Compiz leads to single frame problems after a certain amount of time in the case where you switch to fullscreen mode and then back out to normal mode – cin stops working and so you will have to restart cin.
- 2) When a library goes from one version to a later version, sometimes a pre-built Cin binary will fail because it was created at a different version than the one the user has on their computer. This seems to happen more frequently on Arch distros because Arch has continuous releases and is always kept up to date. An example of the error message you might see in your startup window would be:

“cin: error while loading shared libraries: libvpx.so.5: cannot open shared object file: No such file”

18.4 Menu Bar Shell Commands

In order to provide some configuration control, the Menu Bar Shell Commands are available for customization purposes. In the main window on the top line containing the File, Edit, ... Window pulldown menus, all the way to the right hand side is the “shell cmds” icon. You will see a small gold-color bordered box with the >_._ inside and if you mouse over it, the tooltip says “shell cmds”. This is a configurable popup which gives you the ability to execute a pre-defined script to do something such as post processing to be done on a repeating basis.

The Shell Cmds popup menu items are configured in:

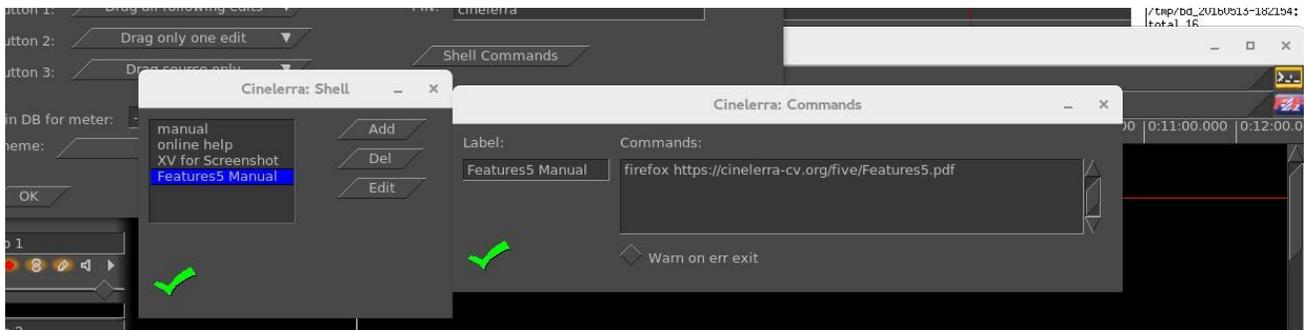
Settings->Preferences->Interface->Shell Commands

This runs a dialog that lets you Add/Del/Edit the items in the main menu popup. Possible options are:

- 1) Add a new menu item/script:
 - click Add
 - in the Label box, change the word "new" to the desired menu label
 - in the Commands box, type the command lines to be included in the script (for example: Label => gimp; Commands => gimp)
 - check-OK in the Commands window
 - check-OK in the Shell window
 - click on OK or Apply in the Preferences window
- 2) Del to delete an existing menu item (for example manual):
 - select to highlight an entry in the Cinelerra: Shell listbox
 - click Del
 - check-OK in the Shell window
 - click on OK or Apply in the Preferences window (the entry is destroyed)
- 3) Edit an existing menu item:
 - click Edit
 - select to highlight an entry in the Cinelerra: Shell listbox

click edit (the edit dialog appears)
set the label and/or commands to the new desired values
check-OK in the Shell window
click on OK or Apply in the Preferences window

Next time you click on the “shell cmds” icon, you will see your new labels.



When a new script is added in cinelerra’s release, it will not be added to your Preferences file automatically unless you create a new one. Or with some help, you can manually correct for this.

There are some Help features currently included in the Shell Cmds menu. Those available are:

- 1) PDF copy of this manual;
- 2) Shortcuts html file for easily looking up a particular shortcut;
- 3) Setting Shell Commands “how to” which explains how to configure your own commands;
- 4) Rendex Mux shell script as described in the Render Farm section.

This page intentionally left blank so that chapters begin on a right facing page.

19 - Performance Tips and Other Tips

Performance and usability of Cinelerra are directly related to the video format being used in conjunction with your computer system hardware - the number of CPUs and its speed, I/O bus speed, and amount of available memory. A basic, less powerful system will be sufficient for users working with audio only or lower resolution video formats. Higher end computers will be needed when playing working with higher resolution formats, like 1080p. Adding effects and several tracks of audio will require more cpu, memory, and various other resources to read, decode and play video. Some help:

- Dual-core and more SMP processors greatly improve Cinelerra speed, making use of threads.
- A large amount of free memory available can help speed up operations by avoiding unnecessary disk swaps and keeping material ready accessible.
- Video editing can be quite I/O intensive. If you are going to produce long pieces in uncompressed or larger resolution formats, you should have plenty of fast access disk space.
- Cinelerra benefits from OpenGL hardware acceleration. A good graphics card is worthwhile to have.
- Multiple monitors really come in handy to increase productivity as you can see more information and in bigger windows so you do not have to keep switching spaces.

Besides the above hardware recommendations, this section covers some tips for performance improvements/speed-ups and tips on some specific issues for older media.

19.1 Optimized Playback of Large Format Video – X11 Direct

Normally, when Cinelerra 5 reads a video frame, it is copied into a "Vframe". This frame may also need other actions performed on it, such as a color model change. In addition, ffmpeg and libzmpeg "can_scale_input". So the read can be color transformed and scaled just by asking the library to do that. This means that if the compositor is in a "good" state with no zoom, the VFrame read can be done in the fastest render color model, and already scaled to the correct size for the compositor. In reality, this is not what you need for editing, but quite often the "virtualconsole" is not used because the render is media only – that is "just data". If no data transforms are needed and the input scaling can be done, the vrender program detects this, and tells the codec to transmit the data in a compatible way to the compositor canvas. This is the X11 "direct" data path.

With the X11 video driver choice these large format files, such as 4K, will playback faster than either X11-XV or X11-OpenGL. However, you still have the option to turn off the X11 direct data path if you use Settings → Preferences, tab Playback A, set video driver to X11 and uncheck "use direct X11 render if possible".

19.2 Proxy Settings

Working with videos that have large image geometry can greatly impede editing. Instead you can substitute "proxies" which will create smaller video image files from the original file that can then be edited more quickly. When you are done working on this smaller scale, you will need to bring up the Proxy settings menu again, and change the Scale factor back up to Original size so that all of your edits/work take affect on that original higher quality video on the timeline. You can not nest clips while in a proxied state; you will get the error "Nesting not allowed when proxy scale != 1".

To use this feature, select Settings → Proxy settings... (shortcut of Alt-r) and change the Scale factor from Original size to your downsized choice. You can choose Ffmpeg as the File Format and a choice

of various codecs associated with that. A good choice is the default of “mpeg” which can usually be quite fast. In addition, to modify values for that codec, click on the wrench. When you have completed your choices, just click OK, and then the video tracks will be rendered. This may take some time, but previous proxy renders will be reused. The proxy videos will be added to your assets in a separate Proxy folder, and the video track edits will use the “proxies”. Proxy downsizing renders all loaded tracks, but only work on the 1st video layer of multi-layer media. Rendered proxy media is saved in the same directory as the original media. As usual, you can delete proxy files from the project or disk in the Resources window if you no longer want to retain them.

There are two ways that the proxy files can be used, with or without input scaling. When the proxy is done without rescaling, the Mask, Camera and Projector automations are re-scaled accordingly. In this case, the entire project will be re-sized so that the session is in the resized geometry. Not all plugins are useful when the project is rescaled, because the keyframe data must be in the original geometry. In this case, you can use the rescaler, by enabling “Use scaler (FFMPEG only)”. It has the advantage that the project size does not change, the proxy media is down-scaled as usual and up-scaled on read-in, which means the project editing is done in full scale. Since decoding is done on smaller video, there is a time savings, but all rendering is done full scale. The main reason for using “scaler” is that it does not change the image coordinate data, so that automation and plugin parameters will be in the original project geometry. This is not as fast as the first option, but is a performance gain, and may be needed if you are using plugins that need coordinate data such as the Titler. As noted, it only works on ffmpeg video formats.

In the upper right hand corner of the main window, there is a toggle button to easily switch back and forth when you have a proxied file on the timeline. The icon is to the left of the FF icon. It will have the letter “P” as the icon for Proxy or if “Using Scaler”, the letter “S”. This is especially useful when editing and you need to see a better image temporarily.

Screencast here shows the Use scaler checked so you can still use plugins and the original project size is kept. The Scale factor pull-down gives you available size options. Note the new media dimensions shown (partially covered). If the size is an odd number, 1 is added to make the dimensions both even #'s.

A convenient “Beep on done” checkbox is included so that you can work on other tasks until there is an audible notify of completion.

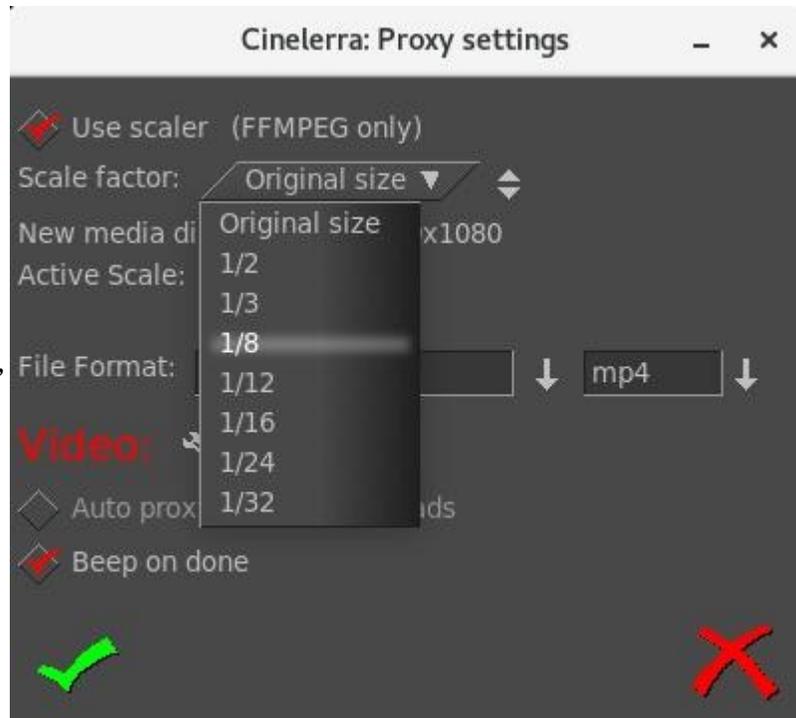


Figure 126: Proxy settings dialog

Checking the “Auto proxy/scale media loads” results in any additional media loads to be automatically proxy scaled. However, single frame media, such as PNG or JPEG “stills”, can not be scaled to “stream” media. If this type of media exists, you should “use scaler”.

If you get error messages when creating proxies, check the Video wrench settings. These usually default to values that are expected to work correctly for the “File Format” and codec you selected but they can be changed and may result in errors.

More specific information on which plugins need to use scaler follows. If the keyframe data uses coordinate data that is absolute, then the scalar should be used. If the data is normalized (like always 0-100%) then the proxy can be done without the scaler. The session geometry format, shown in Settings → Format as width x height, is changed if the scaler is not used to cause all of the data to be in the reduced format. If this affects the plugin operation, then the scalar should be used. Examples of plugins that need the scaler are: Title, AutoScale, Scale, ScaleRatio, and Translate. Most others are safe to use without scaling.

19.3 Some Settings Parameter Values

“Cache” in Settings → Preferences, Performance tab is used to store images on the timeline. One 1080 frame uses about 10 MB. The default setting is 256 and this is enough for testing and running. However, why not use more memory if it is available. To experiment for testing a good number tuned to the way you use your computer, set the cache to 0, bring up cin5, load a typical media file, play it and run “top” on another terminal window to see how much memory is being used. In the “top” display, look at “free” memory. Whatever your computer is not using, is a good number to use for

cache. If you start other programs, or change the design of the session so that it uses a lot of frame storage, you may need to re-experiment over time.

Also in the Performance tab is “Seconds to preroll renders” configurable setting. At this time it is unknown as to its purpose.

For system “swap” 1 GB seems to be more than sufficient. If the amount of memory being used by the program is “close”, then swap might save you but often if swapping becomes necessary, it presents more problems and you end up killing cinelerra anyway.

HDV 1080i editing using proxy files

Working with high definition video requires a lot of processing power. Therefore, you may want to do your editing on low resolution files instead, and use HDV material only for the final rendering.

19.4 Tips Improving Smaller Computers Use

A list of items to check for smaller computers that will help to use less cpu/memory/resources follows:

For large media files, use *proxy* to do your main editing.

In Settings → Preferences, Appearance tab, uncheck *Use thumbnails in resource window*.

In Settings → Preferences, Appearance tab, uncheck *Autocolor assets*.

Speed-up certain time-consuming FFmpeg plugins through use of a carefully selected OPTS file.

For large media files, in Settings → Preferences, Playback A, Video Driver set *use direct X11 render if possible*.

For the Video Driver in Settings → Preferences, Playback A, if using a good graphics card, choose *X11-OpenGL*.

If you have multiple cpus or multiple computers, even if slow, take advantage of using *Render Farm*.

When editing, *background rendering* causes temporary output to be rendered constantly while the timeline is being modified. The temporary output is displayed during playback whenever possible so it does not have to be recalculated – very useful for transitions and previewing effects that are slow.

In Settings → Preferences, Playback A, uncheck *Play every frame* which means frames will be skipped as playback of the video falls behind.

Adjust *Cache size* in Settings → Preferences, Performance, to not exhaust the memory and yet still provide decent playback.

19.5 General Crash Handling Tips

This section is a handy guide for describing various kinds of software computer system failures. Only some of these various lockups or crashes can be dealt with. Hopefully, it will help to have some hints to know what kind of failure it is, or to save your work or to avoid future problems. For most of this, your name must be root, although you can certainly try to see if it works for you when not root.

System lockups:

When the system locks up, it is usually a system problem. Normally an application program cannot hose the system. It is a major goal of system design to prevent an application (app) from failing a system interface. This does not mean an app can't mess things up.

Cinelerra crash:

This is covered in Section 6 – Crash Dumps for Analysis. Just a reminder that you must be root and by providing a crash dump and as much other information as possible, you will be helping the developer to analyze the problem and fix it so that it can be avoided in the future.

X Server crash:

Keyboard does not respond, screen is frozen, caps lock may operate LED light. Sometimes using ctrl-alt-F1 ... ctrl-alt-F7 (etc.) will allow you to regain control of a VT console. You can use this to login and check logs: eg. /var/log/Xorg.0.log, dmesg, journalctl ... etc. If you have another computer, make sure a terminal server is configured (for example: rsh, ssh, or telnet), then remote login via this other computer and check the logs. Most important is to immediately note the current software state, and the very last thing that preceded the crash, i.e. last button click, last keystroke, ... or whatever.

Kernel crash:

The machine goes completely dead. The keyboard caps lock LED will probably be flashing. Most likely the only way to see anything after the kernel crashes is to use a serial port console log and usually kdb, the kernel debugger, and special cabling. This requires a lot of setup, and is normally reserved for experts. Login from another computer will not work. Pinging the ip address will not respond since the network stack is part of the kernel. There are some virtual machine setups that will let you debug a “guest” kernel, but this also requires a lot of setup, and affects which kernel is currently under test. The kdb route is preferable.

Keyboard grabs, Server grabs, and Deadlocks:

A grab is an X-server state where all events are forced to just one window event stream. This forces the user to respond to the dialog. Things seems to be working, but no keypresses do anything useful. The system clock and other programs will still be working. The network will work for remote logins. Grabs can be canceled if the /etc/X11/xorg.conf X config contains special setup as shown below:

```
Section "ServerFlags"
    Option "HandleSpecialKeys" "Always"
    Option "AllowDeactivateGrabs" "True"
    Option "AllowClosedownGrabs" "True"
EndSection
```

```
Section "InputDevice"
```

```
Identifier "Keyboard"
Driver "evdev"
...
Option "XkbOptions" "terminate:ctrl_alt_bksp"
Option "XkbOptions" "grab:break_actions"
EndSection
```

or to \$HOME/.xinitrc, add:

```
# xkb terminate/grab actions disabled in xorg.conf, use:
setxkbmap -option "grab:break_actions"
setxkbmap -option "terminate:ctrl_alt_bksp"
ctrl-alt-bksp = terminate the X-server, may restart automatically
```

If your system is running in runlevel 5 (you might be at runlevel 3 for development), it also allows the user to cancel a keyboard grab using:

```
ctrl-alt-keypad-* = ungrab
ctrl-alt-keypad-/ = ungrab, kill grab owner process
```

Modal forms (always on top, and usually ptr/kbd grab) dialog boxes can hose a system by putting a form over another form holding a grab. This means the form that needs input may never get any, and the effect is a deadlock. Usually you will have to restart X (ctrl-alt-bksp).

Window Manager issues:

The "desktop" window manager can intercept and modify all kinds of user input. Mostly, this is a good thing, but can be a nuisance. If user keypresses can be programmed to trigger actions, then they may be useful to send KILL or INTR to an app that seems to be holding X's attention. For example:

```
killall -INTR cinelerra, killall -9 cinelerra,
killall X, but you must run as root to be able to do this
```

The ALT and META keys may be intercepted by the window manager, and this can cause unexpected interface operation.

19.6 Usage Tips for Specific Operations

Encoding into Dolby Pro Logic

Dolby pro logic is an easy way to output 6 channel audio from a 2-channel soundcard with degraded but useful results. Rudimentary Dolby pro logic encoding can be achieved with usage of some effects.

- 1) First, create the front left and right channels. Create 2 audio tracks, each carrying either the left or right channel. Pan the left channel to the left and the right channel to the right with pan.
- 2) Now create the rear left and right channels. Create another 2 audio tracks as above - the left channel panned left and the right channel panned right. Then apply invert audio to both new channels and the signals will come out of the rear speakers.
- 3) Next, create the center channel by creating a single audio track with monaural audio from a different source. Center it with the pan control and the signal will come out of the center speaker.

- If a copy of the signal in the back speakers is desired in any single front speaker, the signal in the back speakers must be delayed by at least 0.05 seconds and a single new track should be created. Pan the new track to orient the signal in the front speakers.
- If the same signal is desired in all the speakers except the center speaker, delay the back speakers by 0.5 seconds and delay either the front left or front right by 0.2 seconds.
- If you want to hear something from the subwoofer, create a new track, select a range, drop a synthesizer effect, and set the frequency below 60 Hz. The subwoofer merely plays anything below 60Hz or so.

Other tricks you can perform to separate the speakers are parametric equalization to play only selected ranges of frequencies through different speakers and lowpass filtering to play signals through the subwoofer.

Improving Analog TV Quality

The picture quality on analog TV is not always good but you can do things in Cinelerra to make it look more like it did in the studio.

First, when capturing the video, capture it in the highest resolution possible. For Europeans it is 720x576 and for North Americans it is 720x480. Do not bother adjusting the brightness or contrast in the recording monitor, although maxing out the color is useful. Capture it using MJPEG or uncompressed Component Video if possible; if not then capture it using JPEG or RGB as a last resort. Now on the timeline use Settings->Format to set a YUV colorspace. Drop a Downsample effect on the footage. Set it as follows:

```

Horizontal:      2
Horizontal offset: 0
Vertical:        2
Vertical offset: 0
    red
    x green
    x blue
    alpha
  
```

Use the camera tool to shift the picture up or down a line to remove the most color interference from the image. If you have vertical blanking information or crawls which constantly change in each frame, block them out with the Mask tool. This improves compression ratios. More invasive quality improvement techniques involve deinterlacing.

Remove Interlacing

Interlacing is done on older video sources, such as camcorders. Playing this video results in jagged images on a computer monitor, but with Cinelerra you can use deinterlacing effects to solve this.

Line Doubling: done by the Deinterlace effect when set to Odd lines or Even lines. When applied to a track it reduces the vertical resolution by 1/2 and gives you progressive frames with stairstepping. This is only useful when followed by a scale effect which reduces the image to half its size.

Line averaging: the Deinterlace effect when set to Average even lines or Average odd lines does exactly what line doubling does except instead of making straight copies of the lines it makes averages of the

lines. This is actually useful for all scaling. There is an option for adaptive line averaging which selects which lines to line average and which lines to leave interlaced based on the difference between the lines. It does not work.

Inverse Telecine: this is the most effective deinterlacing tool when the footage is an NTSC TV broadcast of a film. It is described in the Plugin Effects section.

Time base correction: the first three tools either destroy footage irreversibly or do not work at times. Time base correction is last because it is the perfect deinterlacing tool. It leaves the footage intact. It does not reduce resolution, perceptually at least. It does not cause jittery timing.

Frames to Fields effect: converts each frame to two frames, so it must be used on a timeline whose project frame rate is twice the footage's frame rate. In the first frame it puts a line-averaged copy of the even lines. In the second frame it puts a line-averaged copy of the odd lines. When played back at full framerates it gives the illusion of progressive video with no loss of detail. Best of all, this effect can be reversed with the Fields to frames effect. That one combines two frames of footage back into the one original interlaced frame of half the framerate. Be aware that frames to fields inputs frames at half the framerate as the project. Effects before frames to fields process at the reduced framerate. Unfortunately, the output of Frames to Fields can not be compressed as efficiently as the original because it introduces vertical twitter and a super high framerate. Interlaced 29.97 fps footage can be made to look like film by applying Frames to Fields and then reducing the project frame rate of the resulting 59.94 fps footage to 23.97 fps. This produces no timing jitter and the occasional odd field gives the illusion of more detail than there would be if you just line averaged the original.

HDTV exceptions:

1920x1080 HDTV is encoded in a special way. If it is a broadcast of original HDTV film, an inverse telecine works fine. If it is a rebroadcast of a 720x480 source, you need to use a time base and line doubling algorithm to deinterlace it.

Making video look like film

With an older camcorder video which has low quality video, you can improve the results by turning it into progressive 24 fps output; at least as close as possible. You should only do this for low quality video.

- 1) Set project framerate to twice the video framerate.
- 2) Apply a Sharpen effect. Set it to sharpness: 25, no interlacing, and horizontal only.
- 3) Drop a Frame to Fields effect on the same track. Set Average Empty Rows to on and play through the video a few times to figure out which field is first. If the wrong field is first, the motion is shaky. Secondly, any editing in the doubled frame rate may now mess up the field order. It is not yet clear which is the easiest way to support warnings for field glitches but for now you need to go back to the normal framerate to do editing or play test to make sure the fields are right.
- 4) Render just the video to the highest quality file possible.
- 5) Import the video back to a new track. Set the project framerate to 24. The new track should now look more like a file with sharper images than the original footage.

This entire procedure could be implemented in one non-realtime effect, but the biggest problem with that is you will most often want to keep the field based output and the 24 fps output for posterity. A non-realtime effect would require all that processing just for the 24 fps copy.

Clearing out haze

You probably photograph a lot of haze and not as much blue sky. Horizon shots also usually need more depth. This is what the gradient effect is for.

Drop the gradient effect on hazy tracks and set the following parameters:

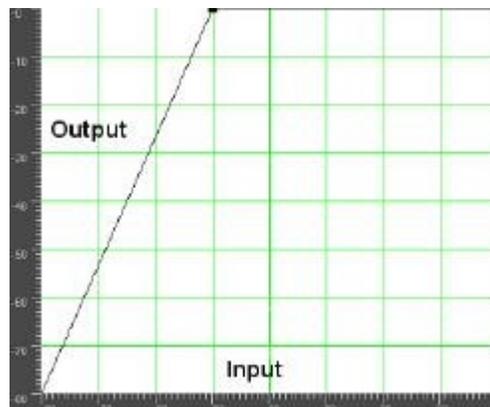
Angle: 0
Inner radius: 0
Outer radius: 40
Inner color: blue 100% alpha
Outer color: blue 0% alpha

It is important to set the 0% alpha color to blue even though it is 0% alpha. The color of the outer alpha is still interpolated with the inner color. This is a generally applicable setting for the gradient. Some scenes may work better with orange or brown for an evening feel.

Making a ringtone for a cell phone

- 1) Go to File->Load files... and load a sound file with Insertion strategy: Replace current project.
- 2) Go to Settings->Format change Channels to 1 and Samplerate to 16000 or 22050.
- 3) Highlight a region of the timeline to use for the ringtone. To improve sound quality on the cell phone, you need the maximum amplitude in as many parts of the sound as possible.
- 4) Right click on track Audio 1 and select Attach effect... Highlight the Compressor effect and hit Attach in the attachment popup.
- 5) Make sure the insertion point or highlighted area is in the region with the Compressor effect.
- 6) Right click on track Audio 2 and select Attach effect...
- 7) Highlight Audio 1: Compressor and hit Attach.
- 8) Click the Audio1 Compressor's magnifying glass to bring up the compressor GUI.
- 9) Set the following parameters:

Reaction secs: -0.1
Decay secs: 0.1
Trigger Type: Total
Trigger: 0
Smooth only: No



- 10) Click Clear to clear the graph. Click anywhere in the grid area and drag a new point to 0 Output and -50 Input. The graph should look like the above.
- 11) Go to File->Render. Specify the name of an mp3 file to output to. Set the file format to MPEG Audio. Click the wrench for Audio and set Layer to III and Kbits per second to 24 or 32. Check Render audio tracks and uncheck Render video tracks. Hit OK to render the file.

The resulting '.mp3' file must be uploaded to a web server. Then, the phone's web browser must download the '.mp3' file directly from the URL. There also may be a size limit on the file.

Time stretching audio

It may appear that time stretching audio is a matter of selecting a region of the audio tracks, enabling recording for the desired tracks, going to Audio->Render Effect, and applying TimeStretch. In actuality there are 3 audio effects for time stretching: Time Stretch, Resample, and Asset info dialog.

Time Stretch applies a fast Fourier transform to try to change the duration without changing the pitch, but this introduces windowing artifacts to the audio. It is only useful for large changes in time because obvious changes in duration make windowing artifacts less obtrusive.

For smaller changes in duration, in the range of 5%, Resample should be used. This changes the pitch of the audio but small enough changes are not noticeable. Resample does not introduce any windowing artifacts, so this is most useful for slight duration changes where the listener is not supposed to know what is going on.

Another way to change duration slightly is to go to the Resources window, highlight the media folder, right click on an audio file, click on Info. Adjust the sample rate in the Info dialog to adjust the duration. This method also requires left clicking on the right boundary of the audio tracks and dragging left or right to correspond to the length changes.

Panning and zooming still images

Cinelerra's powerful keyframe features allow you to use pan and zoom effects on still pictures.

- 1) Load and create a clip from a still image. Make the clip 10 seconds long.
- 2) Activate the automatic generation of keyframes.
- 3) Using the transport controls, go to the beginning of the clip.
- 4) Using the compositing camera control set the clip's initial position.
- 5) Using the transport controls, move forward a couple of seconds on the clip.
- 6) Dragging on the compositing camera move the camera center to a new position further along.
- 7) Now, rewind to the beginning of the clip and play it.

You can see that the camera smoothly flows from keyframe point to next keyframe point, as Cinelerra automatically adjusts the camera movement in straight lines from point to point.

20 - Language Translations

There are several “po” files for various languages to make cinelerra more usable for non-English countries. A program, xlat.C, assists in providing several variations of text files that can be used in order to allow anyone to help make meaningful translations.

To build a new cin.po file, do the following after moving to cinelerra's directory:

```
./po/xlat.sh > /tmp/cin.po
```

To use the xlat.C program, first compile it with “c++ xlat.C”. You can see the usage help here:

```
# ./a.out
list csv ./a.out csv < data.csv > data.po
list po ./a.out po < data.po > data.csv
list po ./a.out dups < data.po
list po ./a.out nodups < data.po
get strings ./a.out key < xgettext.po
gen xlation ./a.out xlat xgettext.po xlat.csv
gen xlation ./a.out xlat - text,xlat ... < xgettext.po
```

This program has 6 commands where the desired command is the first parameter to a.out.:

- 1) “csv” = comma separated value file; to recreate a po from a csv file.
- 2) “po” = convert po to csv data; for example, what you need to convert ru.po to ru.csv.

You can even open the resulting *.csv in ooffice and update the key->value replacements (fields separated by "," only and check quoted fields as text during import). These results can be "Saved As" a csv file, and then used by xlat.C to reformulate a “po”.

- 3) “dups” = list only key/value items where either the key=value or value=””.
- 4) “nodups” = list only key/value items where key!=value and value!=””.
- 5) “key” = list cin.po key set.
- 6) “xlat” = overlay translation. This is the most important use and is described next.

The xlat command line parameters specify a new cin.po template, usually created with xlat.sh, and a list of key/value files which are used to build a mapping for the desired translation. The mapping files are added to the mapping in the order they appear on the command line, and any existing key is replaced with the newest definition; so typically the newest key/value data is last in the command parameters. Once the mapping is built, the first parameter, the new po template, is scanned and the keys it contains are used to find the latest mapping in the key/value files. The new value replaces the existing value in the template. For example, to overlay a new map onto an existing po:

```
c++ xlat.C
./a.out po < xx.po > /tmp/xx.csv
./a.out po < new.po > /tmp/new.csv
./a.out xlat xx.po /tmp/xx.csv /tmp/new.csv > /tmp/new.po
```

The first run preserves the existing mapping of xx.po, the second creates new mappings from new.po, and the third merges the original and new mappings to create a po with new included/overriding xx.po.

Of course, you can also use the freely-available “poedit” program to provide translations!

Additionally: some words and abbreviations can lead to ambiguous language translations. Therefore, the usage of C_ and D_ in the program code was added to represent Contextual and Definitional exceptions to the usual _ and N_ . You will see the following:

C_(“msgstr”) is translated to D_(“qual#msgstr”) by xlat.sh,
and invokes gettext with msgid = “qual#msgstr”.

When no 'po' translation is supplied, the qual# is removed, and only the default 'msgstr' text is displayed. If a 'po' translation is defined for the current locale, then the translated msgid = 'qual#msgstr' is used to access the translated msgstr.

The default MSGQUAL is the basename of the C source file. For the file src_file.C, the default MSGQUAL is

```
#define MSGQUAL “src_file”
```

It is used to define the qualifier needed to transform:

```
C_(“str”) to D_(“src_file#str”)
```

The resulting xlat.sh'd source is scanned by xgettext to create the initial cin.po. In other words:

_(s) does normal international text translation as always -

The msgid line is: msgid “s”

C_(s) makes it appear as if you used D_(“src_file#” s) -

The msgid line is: msgid “src_file#s”

If it does not translate, the default msgtext is “s”, not “src_file#s”

D_(qual#s) tries to look up _(qual#s) -

This msgid line is: msgid “qual#s”,

If it does not translate, the default msgtext is “s”, not “qual#s”

Summary:

use _(s) as usual in almost all cases

use C_(s) if this source file has specialized meaning for s

use D_(qual#s) if this source file has several meanings for s

21 - Licenses

Cinelerra-GG codebase is licensed GPLv2+

GREYcstoration plugin is licensed CeCILL v2.0

Google SHA1 implementation is licensed BSD 3-clause

Cinfinity icons are licensed Creative Commons CC-BY 4.0

Neophyte theme is licensed Creative Commons CC-BY 4.0

OpenCV is licensed 3-clause BSD License

freeverb components and the Tapeworm font are in the public Domain

License: GPLv2+ and CeCILL and BSD and CC-BY and Public Domain

GPLv2+ GNU General Public License version 2

Cinelerra-GG codebase is licensed GPLv2+ .

See: <https://www.gnu.org/licenses/old-licenses/gpl-2.0-standalone.html>

Creative Commons Attribution 4.0 International License

Cinfinity icons (c) by "sam" and Neophyte theme are licensed under a Creative Commons Attribution 4.0 International License.

see: <<http://creativecommons.org/licenses/by/4.0/>>.

License Agreement for OpenCV

License Agreement

For Open Source Computer Vision Library
(3-clause BSD License)

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the names of the copyright holders nor the names of the contributors may be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided by the copyright holders and contributors “as is” and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall copyright holders or contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

This page intentionally left blank so that chapters begin on a right facing page.

Appendix A – QuickStarts

A.1 Cinelerra-GG Quick Start Guide

Cinelerra is a software program NLE, Non-Linear Editor, that provides a way to edit, record, and play audio or video media on Linux. It can also be used to retouch photos, watch TV, and create DVDs.

- [1 – Install the software](#)
- [2 - Start Cinelerra program](#)
- [3 - Load media](#)
- [4 – Choose Output Format](#)
- [5 - View and Listen](#)
- [6 - Edit/Compose](#)
- [7 - Backup your work](#)
- [8 - Create your new media](#)
- [9 - Play your new media](#)

1 - Install the Software

On the internet, click on the Download page at: <https://cinelerra-gg.org/downloads/>

Here you will see several Operating System distro packages that are already built for you to download. Click on your preference and read the specific instructions for usage.

Packages

Select your distro, you can find the [documentation here](#).



ubuntu



debian



arch



opensuse

...

However, if you want to get going as quickly as possible, just do this so that everything is in 1 place:

Download your Operating System's tar file from <https://cinelerra-gg.org/download/tars/> to /tmp .

Keyin: `cd /name-of-directory-where-you-want-the-software` (for example, `cd /software`)

Keyin: `mkdir cin`

Keyin: `cd cin`

Keyin: `tar -xJf /tmp/cinelerra-5.1-*.txz` (if you put the tar in /tmp AND replace * with full name)

2 - Start Cinelerra program

Depending on how you installed the software, you can log in as root or as a user if you used a package.

Keyin: /your-software-directory-path/bin/cin

Or if you installed using the pkg method, click on the Cin icon.

You will now see 4 separate windows appear. The top 2 windows from left to right are the *Viewer* which is most useful for previewing clips and media and the *Compositor* which displays the current working frame at the timeline position. The bottom 2 windows are the *Cinelerra Program*, also called the timeline, which is where the real work gets done and the *Resources* window showing a selection of media or effects.

Figure 127: Viewer Window

Compositor Window

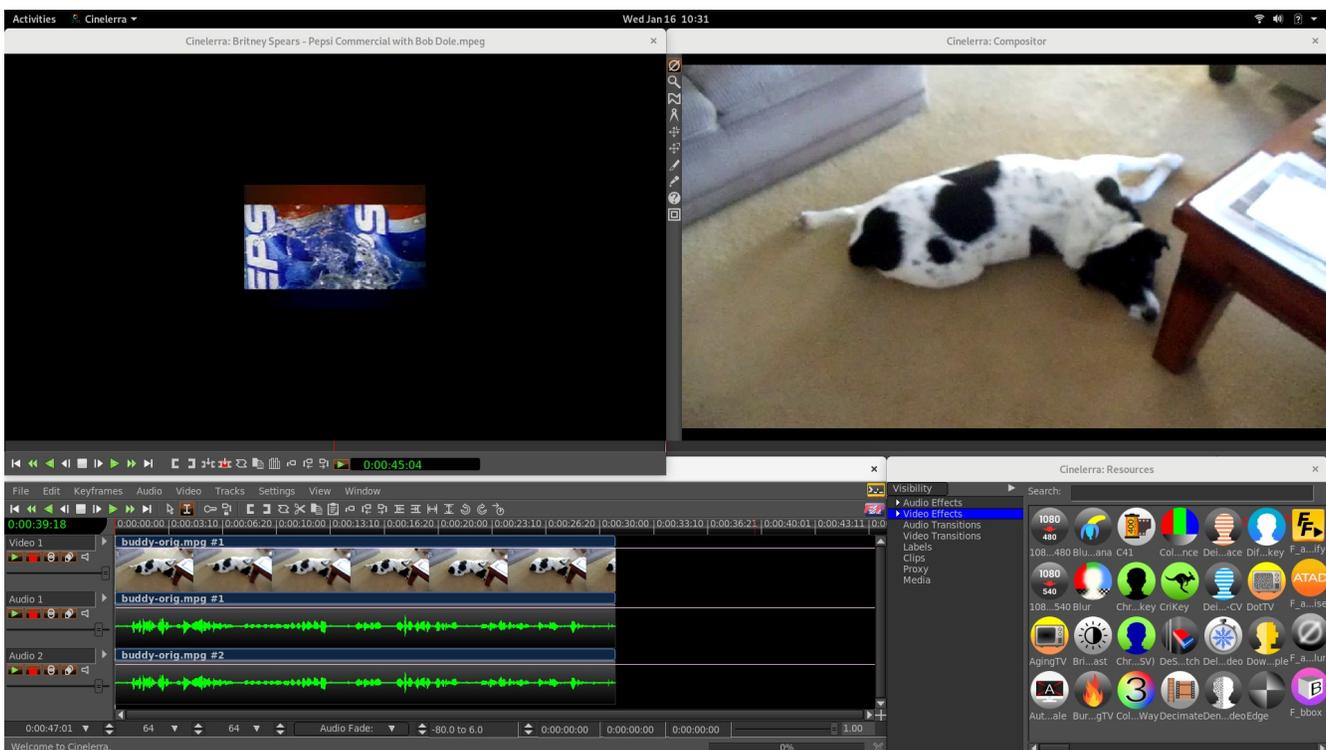


Figure 128: Main / Program / Timeline

Resources Window

Any of these windows can be resized to better suit your needs. Note that if your system's native language is not English, some of the words you see on the screen will be correctly translated for you, others will be in english, and some will have not very good translations.

It is important to know that Cinelerra does not directly change your media. It writes all changes to what is called the EDL, Edit Decision List. This way you original media remains completely intact.

3 - Load Media

On the main timeline program window are many pulldowns, the first of which is “File”.

- 1) Click on File for a list of available options and note that in the right hand column are shortcuts for many of the options that will come in handy if you use Cin often.
- 2) Next click on the second one down - “Load files” - which brings up the Load menu.

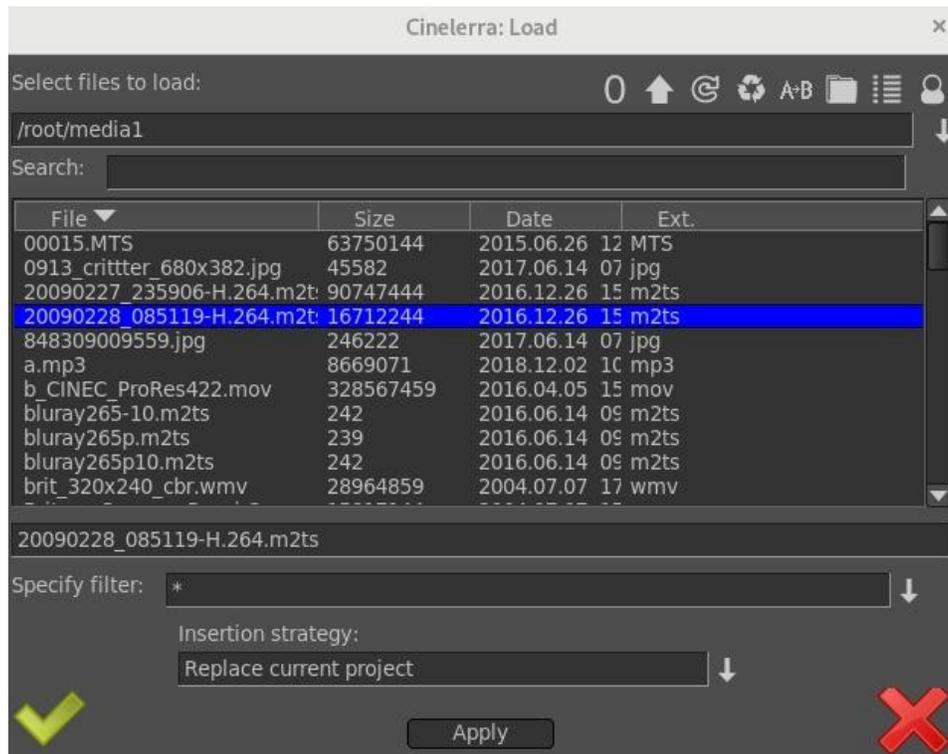


Figure 129: Load media window - note the icons top right for more options

- 3) Below “Select files to load” on the top left side is a textbox and if you look all the way to the right side of the textbox, there is a down arrow which you use to navigate your file system. Highlight the desired file system and you will see that directory name appear in the textbox and the files below.
- 4) Scroll to the media file you would like to work on and highlight that file. When you do, you will see that filename also appear in the textbox below the listing of files. You could have directly keyed in that file in that textbox instead.
- 5) On the bottom of the Load menu, is a box called “Insertion strategy”. For getting started the default of “Replace current project” is sufficient. But you can click on the down arrow to see what is available for future use.
- 6) Now click on the green-colored checkmark on the bottom left-hand side to actually load the file and see it appear on the timeline in Cinelerra’s main window and a single frame in the compositor. The first track will most likely be video thumbnails and the next tracks may be audio waveforms.
- 7) Press the space bar in the main Program window and your video will start playing and press the space bar again to stop the play. While playing, you should see the video in the Compositor window in the upper right hand side of your screen and if you have your audio hooked up, you will hear the sound. To get back to the beginning of the video, hit the home key on your keyboard.

4 – Choose Output Format

You can skip this step if you want the format of your output to be the same as your input. However, to create output media that is widely viewable on many platforms, to include phones and television, you should set your format accordingly.

- 1) On the main timeline, use the Settings pulldown (about the 7th pulldown from the left side top) and click on “Format” which is the first option in that list.
- 2) A “Set Format” menu will appear that shows what the current format is for your loaded media in an Audio and a Video tab. In the United States, the Video Frame rate is usually expected to be 29.970 and usually the Color model is only changed if you have a personal preference.
- 3) The “Canvas size” is probably the only thing you will want to change here in order to get to the most commonly viewable settings. On the right hand side of the Width parameter is a down arrow. Left click the down arrow to see your options.
- 4) Highlight 1280x720 HD for a good common option.
- 5) Click OK to have this option take effect. When you do, the Compositor window may change to fit this option and may look wrong sized.

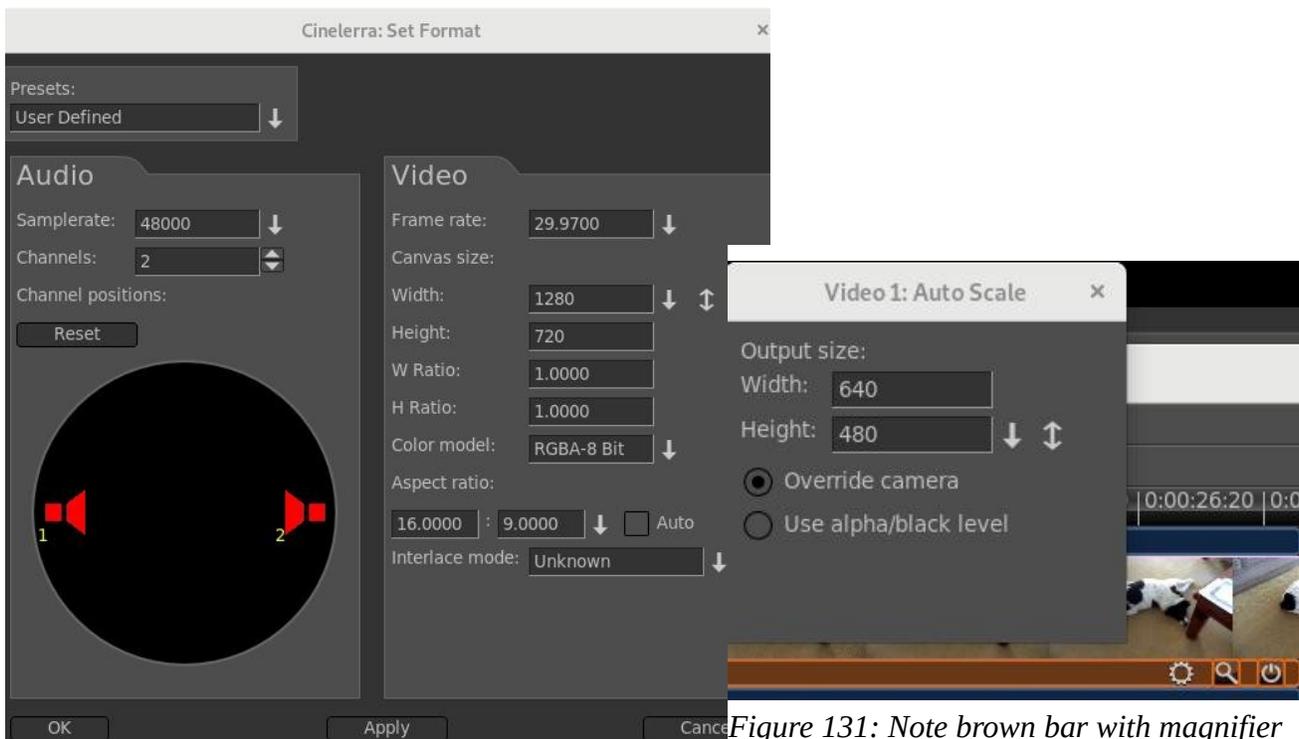


Figure 131: Note brown bar with magnifier

Figure 130: Format menu to change settings

- 6) If the video now looks too small or too large in the Compositor, you will want to “autoscale” it to look correct when the new media is created. To do this, mouse over to the Resources window in the lower right hand corner and under the word Visibility, highlight “Video Effects” to see some plugins.
- 7) Mouse over the “Auto Scale” icon, left click to highlight the words underneath the icon, and mouse drag the icon to the timeline video track. When you see a white colored outline show on that track, drop the Auto Scale icon there and you will see that the video may now automatically scale to a

new value. Click on the magnifying glass icon on the brown colored line beneath the main timeline video which opens a new window. In that window, again use the down arrow to choose 1280x720 HD, then dismiss this window.

8) If not needed, to remove the Auto Scale plugin, right mouse on the brown line and choose Detach.

5 - View and Listen

1) On the second line, below the pulldowns, are transport buttons to move back and forth on the timeline and play forward or reverse, fast or slow, or a single frame. When you mouse over one of these buttons, a yellow-colored tooltip appears to tell you its function along with a key shortcut inside of parenthesis. When you left click the mouse on the transport button it starts the play and click again to stop it. As you use these buttons, watch the Compositor to watch your video.

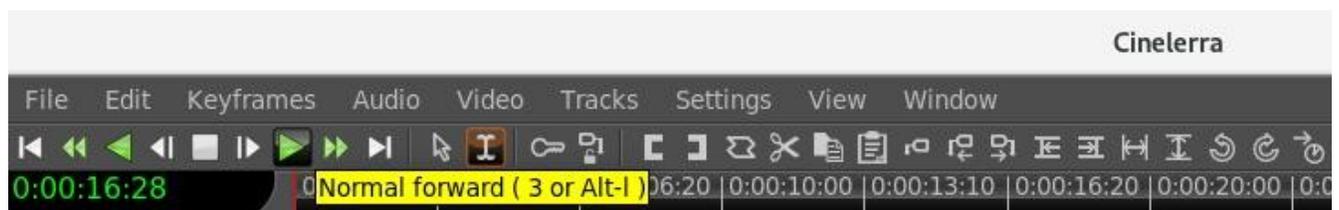


Figure 132: Menu pulldowns at the top with Transport buttons below. Note the yellow tooltips too.

2) On the timeline, you only see thumbnails and not every single picture/waveform. You may want to use your keyboard's down arrow to expand the thumbnails and the up arrow to unexpand them – on United States keyboard, the arrow keys are generally together on the lower right hand side of the keyboard, a little to the right of the space bar. This is a more cpu intensive operation and for very large video can be time-consuming.

6 - Edit/Compose

There may be sections of your media that you want to delete, or audio that is hard to hear and needs to be enhanced, or there is a need for a descriptive title that you want to add. Here are a few basics. But first be sure that you are in "cut and paste" mode (this is the default) by checking to verify that you see a gold color around the "I" i-beam mode icon as in the figure above. If the arrow to the left is gold instead, you are in "drag and drop" mode so switch to "cut and paste" by clicking on the "I" instead.

- 1) You should look at the **"Edit" pulldown** - 2nd from the upper left on the main timeline to see the most common options to use. The first option in the list is "Undo" followed by a terse comment of the last operation that you performed that can be undone.
- 2) To **delete a section** of video/audio is described next. Various ways to do that are available but the easiest is to move your mouse and left click at the beginning of the section you want to delete on the timeline and while holding down the left mouse button, drag to the end of the section to be deleted. When you do this, a white colored highlighted section becomes visible. Use the edit pulldown and choose the split/cut option to cut out the highlighted area (note the shortcut of x). Remember if you cut the wrong thing out you can always use the Edit pulldown to Undo that.
- 3) To **add a transition** where there is deleted section which may make your video look disjointed, do the following. Go back to the Resources window in the bottom right hand corner. Change to "Video Transitions" by highlighting that underneath the word "Visibility". Highlight a transition like

BandSlide with the left button mouse click, hold down and drag to the video track and when you see a white-colored box around the area that you deleted above, drop the icon. Right mouse click the icon on the track to vary some parameters like length.

4) To **insert another clip** from a different video, first you have to load the other video on another track. Go to File pulldown again and choose the Load files option. Type in a directory at the top again and erase any specific file that you may have chosen previously in the bottom 2 textboxes. It is very important to now change the Insertion strategy to “Append in new tracks” or you will write over your current work. But if you make this mistake, you can use the Edit pulldown and Undo that!

a) Once the new video is on the track below your current work, you want to work with only this new track, so disarm your other tracks by looking to the left of each track’s timeline and click the 2nd button beneath the track name, for example Video 1 or Audio 1. The track name textbox will turn red to remind you that the track has been disarmed. The boxed area is called the patchbay.

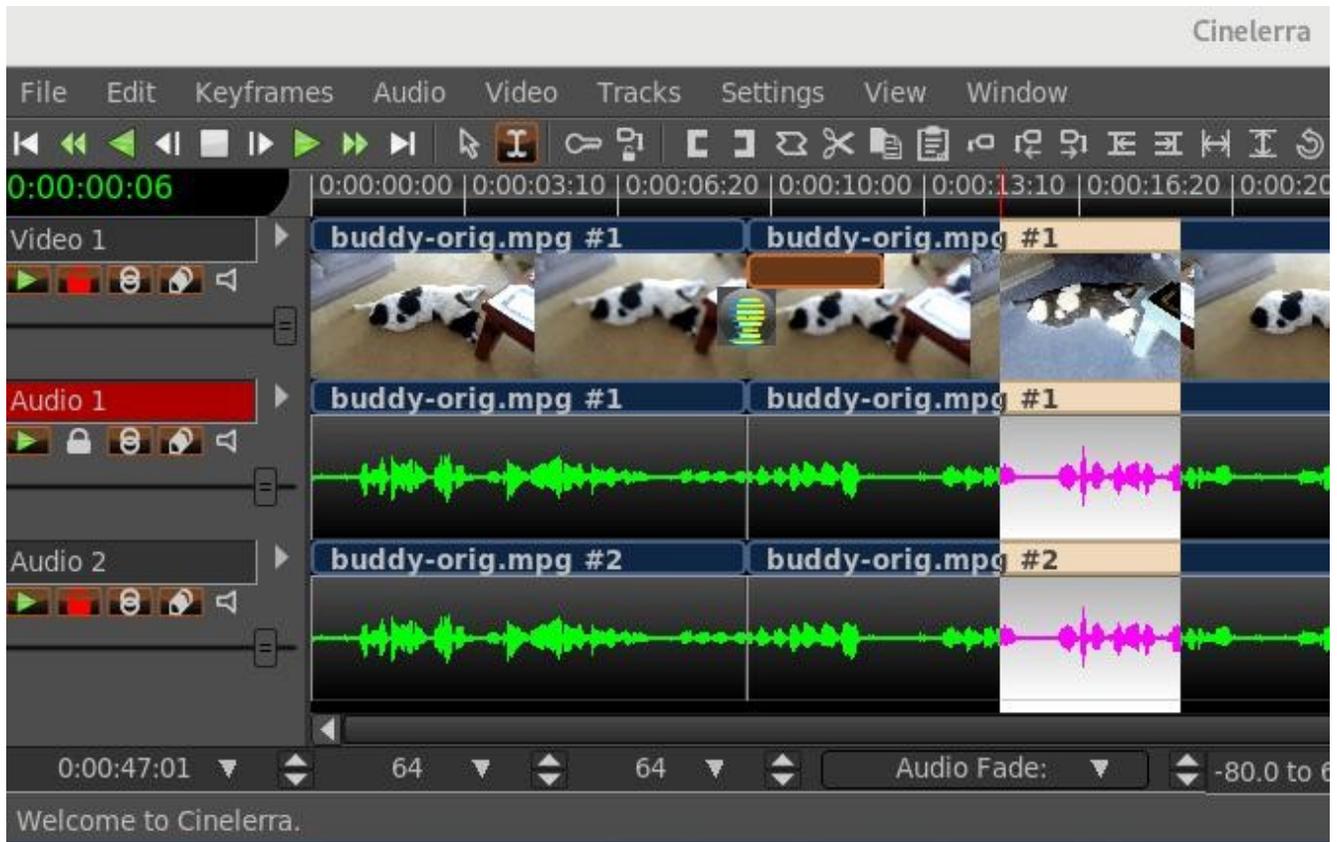


Figure 133: Audio 1 is disarmed. BandSlide transition above. A highlighted section.

- b) Move to the area you want to make a clip of on your newly loaded track, hold down the left mouse button and drag the area to be made into a clip which will turn the color white. Remember, you disarmed the other tracks so only this track is relevant at this time. On the second line of the main window to the right of the transport buttons, are action buttons and as you mouse over them a yellow colored tooltip explains its purpose. Find the one that says “To clip” which is on the right hand side of the right bracket symbol.
 - c) Click on “To clip” and a small window comes up which you can comment in but it is not necessary so just click on the green checkmark and now you will have a clip.
 - d) Disarm that new track and re-arm your original tracks so you can go back to working on them.
 - e) Move your cursor to the spot in your original video where you want to insert the clip. Make a “Split” with the “Split | Cut” option.
 - f) Go to the Resources window and under the word Visibility, highlight Clips so you can see your recently created clip in the box to the right. Highlight that clip and drag it to where you did the blade cut and drop it in.
- 5) To **add a Title** or any wording you will use the Title plugin. In the Resources window, under the word “Visibility”, highlight Video Effects. In the box to the right, many plugin icons appear. Scroll to the right using the scroll bar at the bottom of the Resources window to locate Title. Highlight the Title icon and drag/drop to your video track. By now your video track may be in sections as you deleted, added blade cuts, and inserts so where you drop the Title icon will be surrounded with a white colored box. It will take effect in that entire area so you may want to highlight a section as usual with left mouse click on the timeline and drag to the end of the desired area.

- a) Right click on the brown colored bar that appeared below your video track to get to options and then left click on Show to get the Title window to appear.
- b) Now for the fun part. First type in some words in the bottom large text box just to see what it does. There are so many variable parameters here and they are a lot of fun to play around with.
- c) You can dismiss the Title window when finished BUT be sure to leave the brown colored Title bar on the track. And if you enabled the “drag” feature, you should disable it so you do not forget.
- d) Right mouse click on the bottom Text box to see many more interesting parameters.

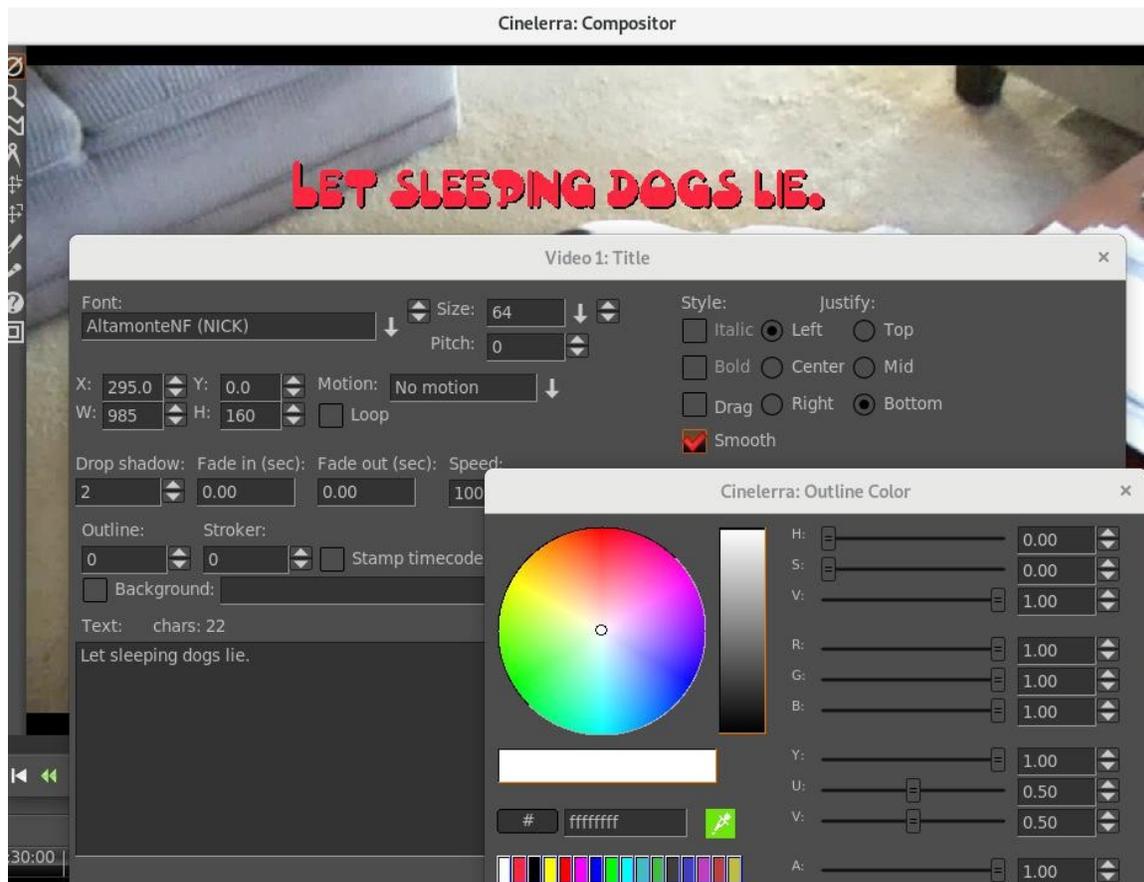


Figure 134: Compositor + Title menu for setting parameters + the Color Picker

7 - Back up your work

At this time, or even earlier if you think you might make a mistake or if you are concerned about computer crashes, you should save your work. Use the File pulldown, and you can use “Save as” to designate a directory and filename. Then click the green checkmark. You are saving the EDL which is the set of changes that you have made – this file is separate from your original media.

8 - Create your new media

- 1) Once again in the main Program window, click on the File pulldown and highlight/click the Render option which is about the 9^h option down from the top of the list. A Cinelerra Render menu will appear.

- 2) First keyin the first textbox the file to render to under “Select a file to render to”
- 3) For the File Format, click on the down arrow and select FFMPEG (because this is the most commonly used format; later you may want to experiment with others). To the right of that box, click on the down arrow and highlight mp4 – again because this is common. When you click on mp4, notice that if there is an extension to your filename in the “file to render to” above, it may change it to mp4 and if there is none, it will add “.mp4” because that is what is expected.
- 4) Make sure there is a red colored checkmark next to the words Audio and Video right below if you have/want both audio and video. To the left of that checkmark box, is a symbol that looks like a wrench. Click on this for Audio just to see the default Preset options which are just fine so dismiss the menu. Then click on the wrench for Video and check Pixels by using the down arrow to the right to be yuv420p – this is most commonly usable option. And click on the green checkmark.
- 5) Check the Insertion Strategy in the Render Menu window. You might want to change that to a different strategy than the default of “Append in new tracks”. If not, then when the Render is done, your new video will automatically be loaded in another set of tracks below your work tracks. Click on the green checkmark in the lower left corner to start the render.
- 6) As the render is running, you will see the video play by in the Compositor. Rendering is usually slow, especially with plugins added.

9 - Play your new media

The file you created in the Render step should now be playable. You can test this in cinelerra most easily by going to the Resource window in the lower right corner, clicking on the Media folder, and dragging and dropping the last video to the Viewer window. There is a separate set of transport buttons on the bottom on that screen to use for playing.

A.2 YouTube with Cinelerra + Basic Video Creation in 8 Steps

To create a youtube or dailymotion video, you can easily follow the steps below. You will have to learn a lot more about Cinelerra to take full advantage of its capabilities and make some really special videos, but this is just to get a start and to see the possibilities.

- 1) Start Cinelerra; usually you can do this by clicking on Cinelerra icon or keyin {cin_path}/bin/cin .
- 2) In the “Program” window on the lower left side of your screen, left mouse click the “File” pulldown.
- 3) You will see “Load files” as the second choice so left mouse click this and find your video file to load, highlight it, and check the green checkmark in the lower left hand corner to get it loaded.
- 4) Edit your video in the Program window using the basic commands of:
 - play and then stop using the space bar
 - move the mouse and then left click to move the insertion (location) pointer
 - cut a section out by holding down the left mouse and drag, then keyin “x” to cut or “c” to copy
 - paste a copy or cut section by moving the insertion pointer, then keyin “v”
- 5) Add a title by highlighting the “Video Effects” in the right hand side “Resources” window; then highlighting the Title icon and dragging it to the “Program” window video track and dropping.
- 6) Click on the middle icon button (looks like a magnifying glass) on the brown-colored Title bar to bring up the Title window bottom text box and keyin a title.
- 7) Use the “File” pulldown to select “Render” to create the desired video. In the “Render” window just next to the empty box to the right of the “ffmpeg” file format, click on the down arrow shown there to see the choices and pick “youtube”. Then move back up to keyin the path and filename to render to. It will pick all of the defaults automatically for you so then just click on the green checkmark to have it start. There is a progress bar in the main window, very bottom of the right hand side.
- 8) Keyin “q” in the main window to get out of Cinelerra and yes or no to save your edit session.

Youtube will allow the upload of the resulting rendered file as named. However, Dailymotion requires that the file be named with an acceptable extension so you must rename the output file to have the extension of .webm instead of .youtube.

There are currently 6 specific variations within the ffmpeg (file format) / youtube (file type) for different video options. You see these when you click on the wrench to the right of the word Video & then the Compression down arrow in the Video Preset window. The first 3 are based on Webm/Vp9 (credit Frederic Roenitz) and contain basic comments of usage and where to find more information.

The first 3 below, plus any of the VP9 files under the file type of “webm” are the recommended options to use because they are freely usable in any circumstance.

sd.youtube	Standard Definition use with default audio/Opus stereo.youtube
hd.youtube	High Definition “ “
uhd.youtube	Ultra High Definition “ “

Alternatives based on h264 and for non-commercial use are listed below. For Dailymotion, these must be renamed to have a different extension of .mp4 instead of .youtube before uploading.

sd_h264.youtube	Standard Definition – must change to audio stereo_with_h264.youtube
hd_h264.youtube	High Definition - “ “
uhd_u264.youtube	Ultra High Definition - “ “

This page intentionally left blank so that chapters begin on a right facing page.

Appendix B - Developer's Section

B.1 How Builds Really Work and Even More Options

This section describes how builds really work if you want to know more about making changes or understanding the process probably only for a developer or system administrator.

Builds occur in 4 basic steps:

- 1) unpack/patch source code
- 2) configure build
- 3) make build targets
- 4) installation

So, an example of what happens in 4 steps for a single-user build would be as follows:

- 1) unpack/patch source code
git clone --depth 1 "git:/.../target" cinelerra5
./autogen.sh
- 2) configure build
./configure --with-single-user
- 3) make build targets
make 2>&1 | tee log
- 4) installation
make install

A lot of things can be tweaked to change the results. Mostly these changes are parameters to the configure step, which can change important build related items, like the application name, or where and what the target system directories should be. This makes it possible to have several versions at the same time on the same computer if needed. To see what it is that the makefiles use to build cinelerra, look at the resulting toplevel global_config file which is created by the ./configure step.

Building cinelerra requires many thirdparty libraries, and it is recommended that you use the static build version included in the git repo. Some of them are patched, and fix significant bugs. It is important to note that because system installation historically has been with as many shared objects as possible, the defaults are that any system library detected during configuration setup will be used, when the package is built "--without-single-user", which is the default build. To build with static thirdparty libraries for system install to the system /usr area, use:

```
.configure --enable-static-build --prefix=/usr
```

Sometimes, additional package parameters and variables are needed during thirdparty builds. These optional values occur before and after the "configure" and "make" commands during a build. A presentation of the format of the package qualified variable names and how they appear in the build procedure are:

```
pkg.cfg_vars      prepended to configure
```

pkg.cfg_params	appended to configure
pkg.mak_vars	prepended to make
pkg.mak_params	appended to make
pkg.cflags	added as CFLAGS+=\$(cflags) to pkg.vars
pkg.cppflags	added as CPPFLAGS+=\$(cppflags) to pkg.vars

These steps are done for EACH of the packages in the thirdparty build:

```
<pkg.cfg_vars> ./configure <pkg.cfg_params>
<pkg.mak_vars> make <pkg.mak_params>
```

The thirdparty Makefile has a set of default vars and params used to build each of the needed thirdparty packages, but you can specify new or overriding values for these Makefile substitutions. These thirdparty build config changes are specified in the toplevel file: cin_config. By using this file, you can save the configuration changes made for the current build to use the next time you do a new build. For example, to add an include file path to the giflib build, add this line to cin_config:

```
giflib.cflags := -I/usr/local/include/giflib5
```

To have a param/var change apply to all thirdparty builds, use:

```
CFG_VARS, CFG_PARAMS, MAK_VARS, MAK_PARAMS
```

CFLAGS, CXXFLAGS and LDFLAGS are forwarded to the thirdparty build environment via:

```
CFLAGS=-ggdb ./configure --with-single-user
```

However, there is no guarantee that the thirdparty build will honor the environmental flags.

Finally, there are build controls, which enable/disable and set build features.

A few of the more useful ./configure --parameters are:

```
--with-jobs=n      where n=number of build jobs; defaults to 1.5*cpus+2.
--enable-static-build  build all 3rd party libs; defaults to yes if single-user, else no.
--with-single-user    build installs to <build_path>/bin; no system installation.
```

The ./configure command builds “global_config”. The global_config is read by the thirdparty/Makefile to create the recipes and definitions used by the thirdparty build.

There are a lot of different options. Thirdparty library build control is available in the configure step of the build. Thirdparty libraries are built on a demand basis. So if you use:

```
--enable-libname=auto # the static-build enable, or the lack of a system library causes a build
--enable-libname=yes  # this forces the thirdparty build
--enable-libname=no   # this forces no thirdparty build
```

Ffmpeg is a “strongly connected component” in the build linkage and widely influences the cinelerra library demands. It is possible to make small additions to the ffmpeg configuration step using the environment variable FFMPEG_EXTRA_CFG. For example, to eliminate the use of libvdpau (an nvidia support library) in the ffmpeg configuration step after you have determined that it is causing an error, use:

```
make clean
autogen.sh
export FFMPEG_EXTRA_CFG="--disable-vdpau"
./configure ...
```

B.2 Thirdparty Parallel Build

The Makefile in the thirdparty build directory employs a set of macros used to create a build rule set of thirdparty library build dependencies. The standard build sequence of [source, config, build] is used to prepare thirdparty products as static libraries. Build package dependency can be specified in the Makefile std-build macro call. These Makefile macro calls define the rules used for each thirdparty build. The expanded rule definitions may be viewed by using:

```
make -C thirdparty rules
```

Individual package libraries can be rebuilt, via:

```
make -C thirdparty <pkg>-clean; make -C thirdparty <pkg>
```

The rule targets create the set of thirdparty packages which are built from local source archive copies of thirdparty source code and patches, if needed. The build rule set of dependencies allows for compiling multiple thirdparty programs simultaneously using maximum computer resources. This parallel build speeds up the process considerably. For example, these are full static build timings on the production build machine (full build includes building all thirdparty programs as well as all of cinelerra):

```
1 cpu = 61 mins
12 cpus = 7.5 mins
24 cpus = 2 mins
```

B.3 Find Lock Problems with Booby Trap

A Booby Trap is used in CinGG for setting a trap to catch lock problems that might have been missed. It will trap boobies only if compile by adding “--with-booby” on the configuration command line. This is the default if you compile using ./bld.sh from the GIT repository. It should not interfere with normal execution.

If you have the time and inclination, enable --with-booby and send any trap output that you find. Maybe you will catch some boobies and if you do, send a snapshot of any boobies you find.

There are 2 potential traps:

- 1) If you try to unlock a lock when it is not locked;
- 2) or if you execute a drawing operation without holding the window lock.

The trap prints the following in the controlling terminal window:

```
BOOBY!  
<backtrace>
```

An example backtrace is below along with a hint below on how to analyze:

```
-----  
/home/cin5/bin/./cin(_Z5boobyv+0x3f) [0x557069fa9b2f]  
/home/cin5/bin/./cin(_ZN13BC_WindowBase9draw_lineEiiiiP9BC_Pixmap+0x3b)0x557069fb9a9b]  
/home/cin5/bin/./cin(_ZN10BC_ListBox11draw_borderEi+0x73)[0x557069f7dc73]  
/home/cin5/bin/./cin(+0x9707fb) [0x557069f7e7fb]  
/home/cin5/bin/./cin(_ZN10BC_ListBox16center_selectionEv+0x4e)[0x557069f7f2ae]  
/home/cin5/bin/plugins/video/sketcher.plugin(_ZN17SketcherCurveList6updateEi+0x1a0)  
    [0x7f1b8002a4c0]  
/home/cin5/bin/plugins/video/  
sketcher.plugin(_ZN18SketcherCurveColor17handle_done_eventEi+0x76)  
    [0x7f1b8002a5f6]  
/home/cin5/bin/./cin(_ZN15BC_DialogThread3runEv+0xd8)[0x557069f6fb78]  
/home/cin5/bin/./cin(_ZN6Thread10entrypointEPv+0x45)[0x557069fc5995]  
/usr/lib/libpthread.so.0(+0x7a9d) [0x7f1b91b4ea9d]  
/usr/lib/libc.so.6(clone+0x43) [0x7f1b90accb23]  
-----
```

To see which routine is reporting the booby keyin: `c++filt`

And then the 2nd line in the backtrace above: `_ZN13BC_WindowBase9draw_lineEiiiiP9BC_Pixmap`

It comes back with the routine as: `BC_WindowBase::draw_line(int, int, int, int, BC_Pixmap*)`

B.4 Valgrind Support Level

Valgrind is a memory mis-management detector. It shows you memory leaks, deallocation errors, mismanaged threads, rogue reads/writes, etc. Cinelerra-GG memory management is designed to work with Valgrind detection methods. This assists in developing reliable code. Use of Valgrind points out problems so that they can be fixed. For example, when this version of Cinelerra shuts down, it deallocates memory instead of just stopping, thus making memory leak detection possible.

The best way to compile and run valgrind is to run the developer static build. This takes 2 steps and you must already have gdb and valgrind installed:

- 1) The standard static build:
`cd /path/cinelerra-5.1`
`make clean`
`./bld.sh`
- 2) run the incremental rebuild for debug objs
`CFLAGS=-ggdb make -j8 rebuild_all`

Now your cinelerra obj has all of the debug stuff. Next run valgrind as root for the most useful results:

```
cd /path/cinelerra-5.1/cinelerra
valgrind --log-file=/tmp/log --leak-check=full ./ci
```

This runs cinelerra under the control of valgrind, and produces a log file in /tmp which will list information about any leaks, usually clearly identifiable. But it runs very slowly, and is basically single threaded, which means that race conditions may be impossible to catch... like one thread deletes memory that another thread is currently using. But overall it is a big help and if you test any new features email the log output. A lot of effort when writing the code was put into trying to be sure that all of the object constructors have matching destructors so that the leaks can be identified. There are already several libraries that create predictable memory leaks and valgrind does a good job for most of these.

It is impossible to test everything with valgrind because some things are just too big and slow for a practical test. Occasionally you can find a leak or an illegal memory access. There are several false alarms that are difficult to avoid "Conditional jump" messages, and "unhandled DW_OP_", but anything with the word "illegal" in the message is important. Memory leaks that originate in cinelerra are good to find and fix, but are usually not deadly.

B.5 CFLAGS has -Wall

When compiling Cinelerra-GG Infinity a CFLAGS option used is "Wall" where the "W" represents warnings and "all" means all. This causes the compile to check for simple mistakes that can be detected automatically and issue warnings when the code is questionable. It can also detect situations where the compiler will generate incorrect code, like type-punned pointer. By turning on this flag, when new code is vetted for predictable mistakes, the code can be corrected before becoming manifested in the application.

This page intentionally left blank so that chapters begin on a right facing page.

C.2 Image Sequence Creation

Example script to create a jpeglist sequence file is next:

```
#!/bin/bash
out="$1"
dir=`dirname "$out"`
shift
geom=`jpegtopnm "$1" | head -2 | tail -1`
w=`(set - $geom; echo $1)`
h=`(set - $geom; echo $2)`
exec > $out
echo "JPEGLIST"
echo "# First line is always JPEGLIST"
echo "# Frame rate:"
echo "29.970030"
echo "# Width:"
echo "$w"
echo "# Height:"
echo "$h"
echo "# List of image files follows"
while [ $# -gt 0 ]; do
  if [ x`dirname "$1"` = x"$dir" ]; then f=./`basename "$1"`; else f="$1"; fi
  echo "$f"
  shift
done
```

Example usage of this script follows:

```
jpeglist.sh outfile infiles*.jpg
```

C.3 Webm / Vp9 Usage and Example File (credit Frederic Roenitz)

There are some common VP9 rendering options files that support creation of video for YouTube, Dailymotion, and other online video services. Webm / VP9 is a media file format which is free to use under the BSD license and is open-source; thus there are no licensing issues to be concerned about. The Webm container is based on Matroska for video and Opus for audio.

Youtube easy startup steps are documented in section 41.7. These same steps have been verified to work for creating Dailymotion videos – however, the created files must be renamed before uploading to change the youtube extension to webm instead for Dailymotion.

Below is one of the VP9 rendering options file with documentation for specifics:

```
webm libvpx-vp9
```

```
# 20171114-2203
# from https://developers.google.com/media/vp9/settings/vod/
# 1280x720 (24, 25 or 30 frames per second)
#
#
# Bitrate (bit rate)
#
# VP9 supports several different bitrate modes:
# mode
# Constant Quantizer (Q) Allows you to specify a fixed quantizer value; bitrate will vary
# Constrained Quality (CQ) Allows you to set a maximum quality level. Quality may vary within
bitrate parameters
# Variable Bitrate (VBR) Balances quality and bitrate over time within constraints on bitrate
# Constant Bitrate (CBR) Attempts to keep the bitrate fairly constant while quality varies
#
# CQ mode is recommended for file-based video (as opposed to live
# streaming). The following FFMpeg command-line parameters are used
# for CQ mode:
#
# FFMpeg
# -b:v <arg> Sets target bitrate (e.g. 500k)
# -minrate <arg> Sets minimum bitrate.
# -maxrate <arg> Sets maximum bitrate.
# -crf <arg> Sets maximum quality level. Valid values are 0-63, lower numbers are higher quality.
#
# Note: Bitrate is specified in kbps, or kilobits per second. In video
# compression a kilobit is generally assumed to be 1000 bits (not
# 1024).
#
# Note: Other codecs in FFMpeg accept the -crf parameter but may
# interpret the value differently. If you are using -crf with other
# codecs you will likely use different values for VP9.
bitrate=1024k
minrate=512k
maxrate=1485k
crf=32

# Tiling splits the video into rectangular regions, which allows
# multi-threading for encoding and decoding. The number of tiles is
# always a power of two. 0=1 tile, 1=2, 2=4, 3=8, 4=16, 5=32.
tile-columns=2

# modified from https://trac.ffmpeg.org/wiki/EncodingForStreamingSites
# To use a 2 second GOP (Group of Pictures), simply multiply your output
# frame rate * 2. For example, if your input is -framerate 30, then use
# -g 60.
g=240

# number of threads to use during encoding.
```

threads=8

May be set to good, best, or realtime
quality=good

This parameter has different meanings depending upon whether quality
is set to good or realtime. Speed settings 0-4 apply for VoD in good
and best, with 0 being the highest quality and 4 being the
lowest. Realtime valid values are 5-8; lower numbers mean higher
quality
speed=4