Nuke 11

Introduction

Nuke is a node-based compositing package provided by the foundry. The main benefit of Nuke is its non-destructive workflow -- operations are chained together in a tree and you can add/remove operations anywhere in that tree.

In addition that, Nuke comes with both some built-in 3D functionality. You can drop some nodes to set up a simple 3D scene and use the rendered output as part of your comp.

> **NOTE**: The non-commercial version of Nuke has some limitations on what kind of data it can deal with and what type of data it can output…
>
> 1. If you're going to read in a movie file, convert it to an animated gif because no other format seems to work reliably (crashes or non-commercial limitation).

2. If you're going to be writing out a comp, pass it into a Reformat node and resize it such that it's <= 1920x1080.
3. If you're going to use audio, don't. It won't work.
4. You can't write out geometry (WriteGeo node?).

Basic Interface

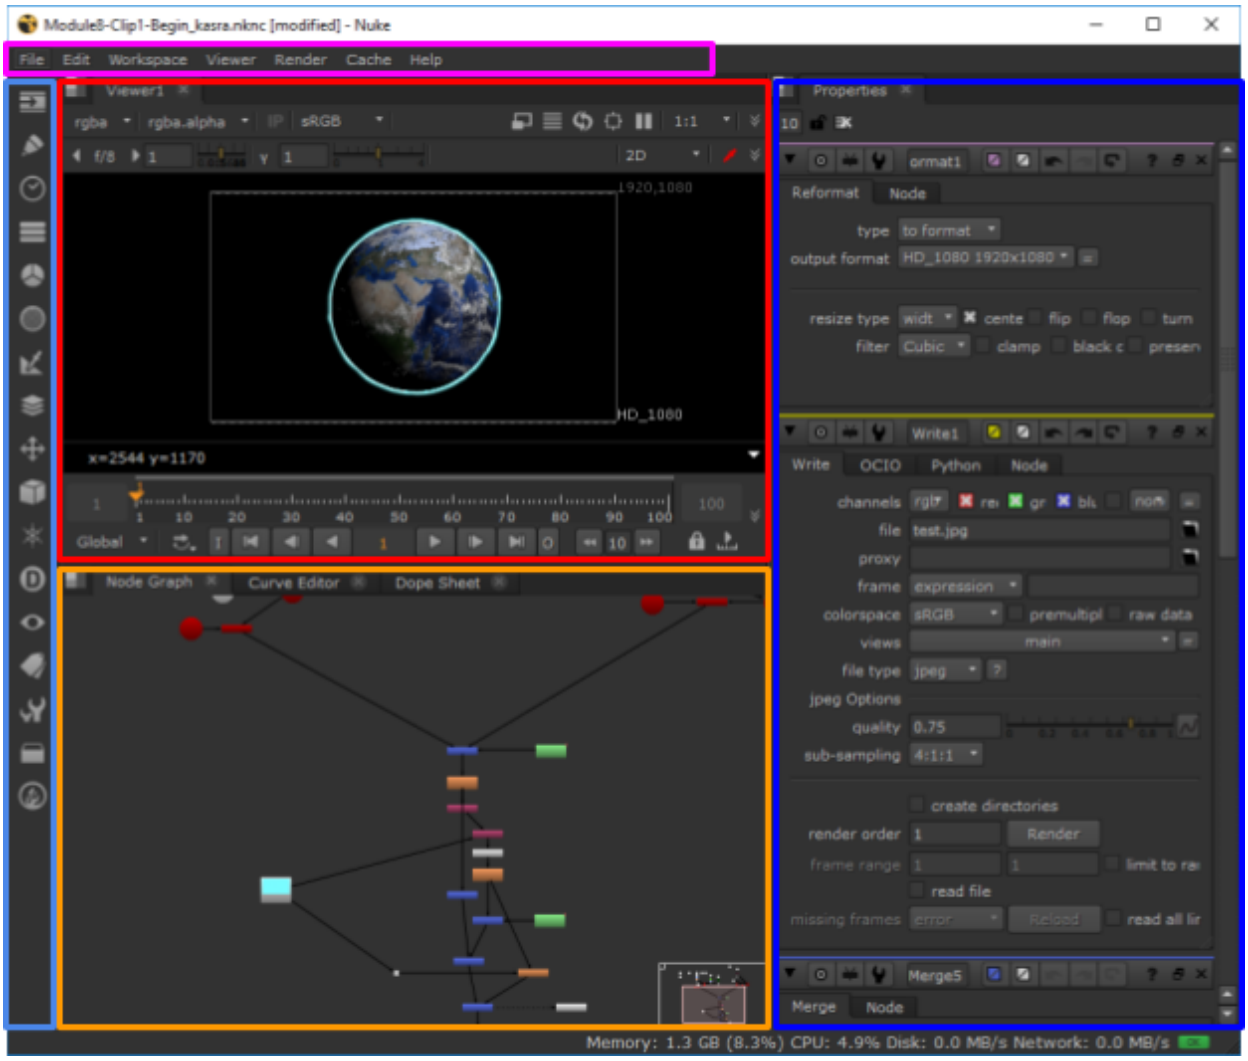

NOTE: If you want to be efficient with Node, you're going to have to make use of keyboard shortcuts. Keyboard shortcuts are context-sensitive: depending on what panel your mouse is hovering over, the keyboard shortcut is going to be for that panel.

NOTE: Just like Houdini, Node provides you with the ability to customize your workspace (your layout of panels). The workspace shown above is the default workspace.

Main Menu

Main menu is always there and provides ways to access basic high-level functionality (e.g. copy/paste/save/open/etc..).

Node Graph

These are where the operations (nodes) for your comp are added and connected together. Note that the dope sheet and curve editor are also available here.

You can add nodes to the node graph via the…
* toolbar (to the left of the node graph)
* tab menu (hit tab while mouse is in node graph)

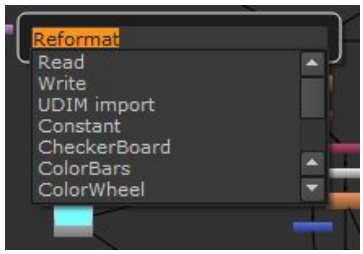

* context menu (right-click in node graph)

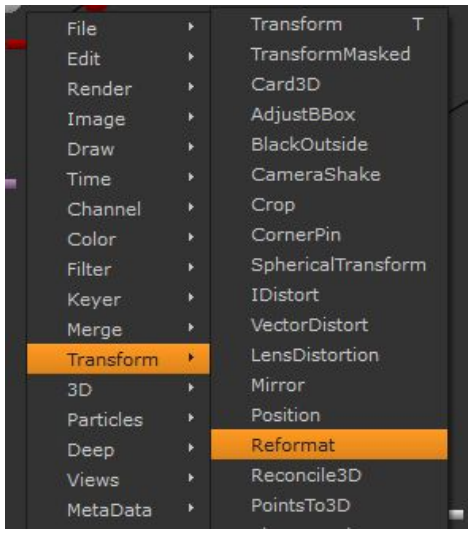

Toolbar

The toolbar provides a way of adding nodes to your node graph.

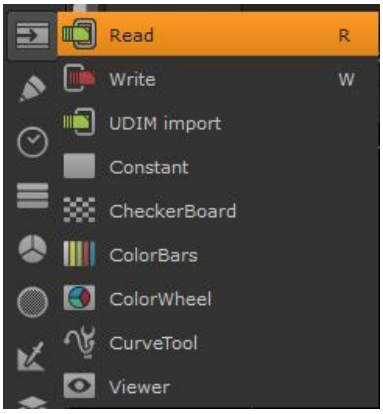

You probably don't want to use this -- better ways of adding nodes are the tab menu and context menu.

Viewer

Viewers are where you can view the output of your comp at various points. Each viewer panel has a corresponding viewer node in the node graph and will display whatever node it's connected to.

You can have multiple inputs to a viewer which you can cycle through or A/B. Also, each viewer can view the 3D output of a scene and do basic 3D manipulations. There'll be more on this later on in the doc.

Properties

Properties display the properties of nodes that you've double-clicked in the node graph.

The properties panel can have the properties for multiple nodes opened at once… That is , if you double-click a node in the node graph, the previously opened properties don't go away. They shift down and the properties for the new node go on top. For example, if you double-click on a lightwrap node and then on a merge node, the properties for the merge node will show up first and then the properties of the lightwrap

node…

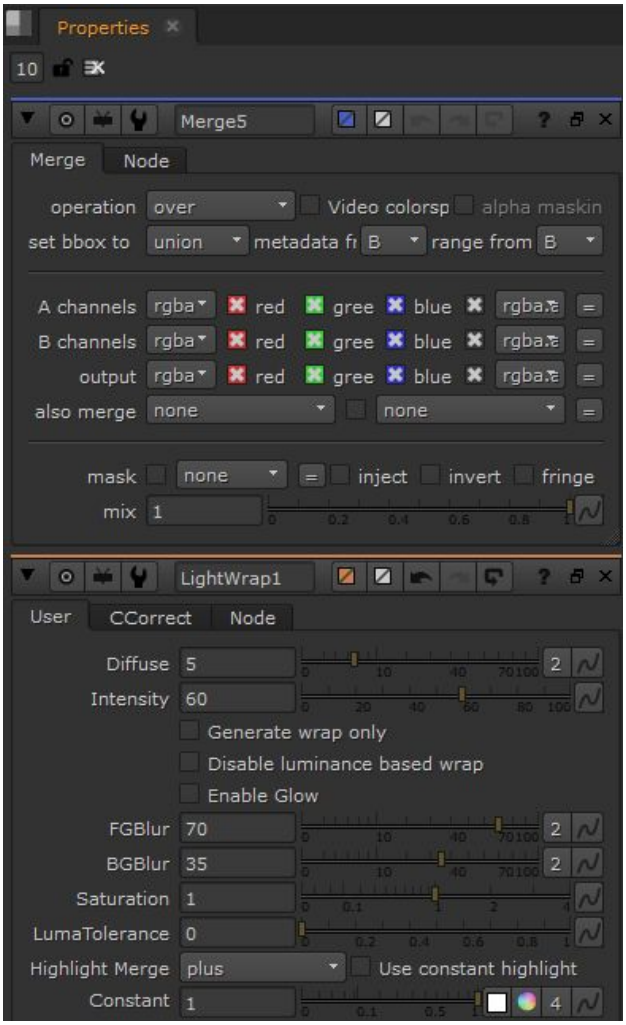

Panels

Like most other packages, …

- panels consist of tabs
- panels can be split either horizontally or vertically
- tabs can be dragged to different panels

In each panel, you'll see a multi-grayscale square on the upper-left hand edge. If you click this square, you'll get a dropdown menu showing you things you can do to the panel…

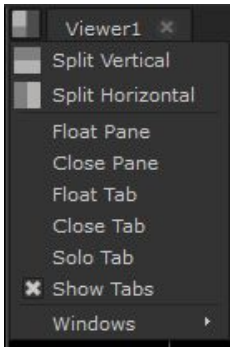

Everything there is pretty much self-explanatory. For example, if you want to split the panel into 2, use one of the split options. Or, if you want to break off the panel into its own window, use the float pane option.

> **NOTE**: Tab menu options (e.g. Float Tab) will operate on the currently active/selected tab within the pane.

Workspaces

Workspaces are a way to organize your panels in Maya (just like Desktops in Houdini or Workspaces in Maya). They're essentially an arrangement of panels and windows that optimize a certain type of workflow.

You can switch between workspaces by going to Workspace in the main menu…

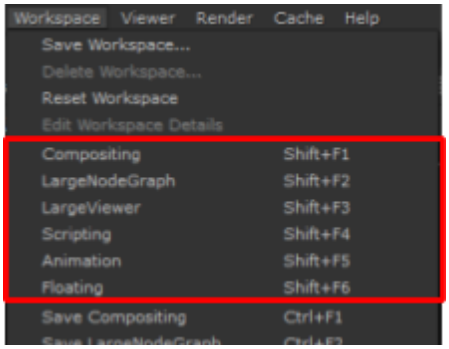

Resetting

If you mess up your workspace and don't know how to get it back to the way it was, you can revert it back to the original way it was set by going to Workspace → Reset Workspace…

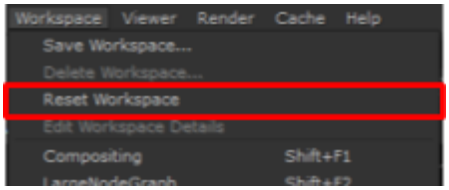

Saving

If you've organized your layout to something that you're comfortable with and you want to save it so you can use it later, you can use Workspace → Save Workspace…

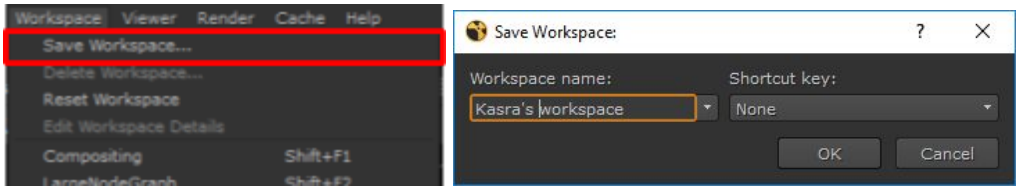

Loading

You can load up that workspace by going to the Workspace → <name>...

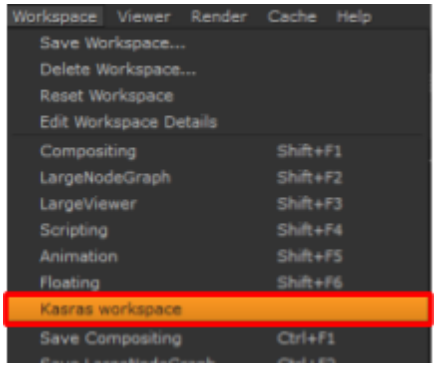

Just like with Houdini… **DO NOT** override the preset workspaces because you won't be able to go back to the original way it was without blowing away all your preferences...

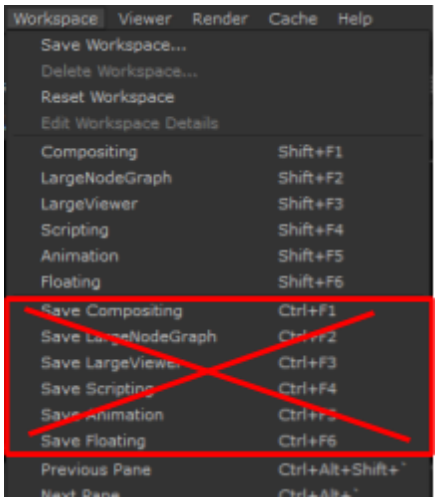

Node Graph Panel

The node graph panel shows you the node graph for your comp -- the operations of your comp and how they fit together…

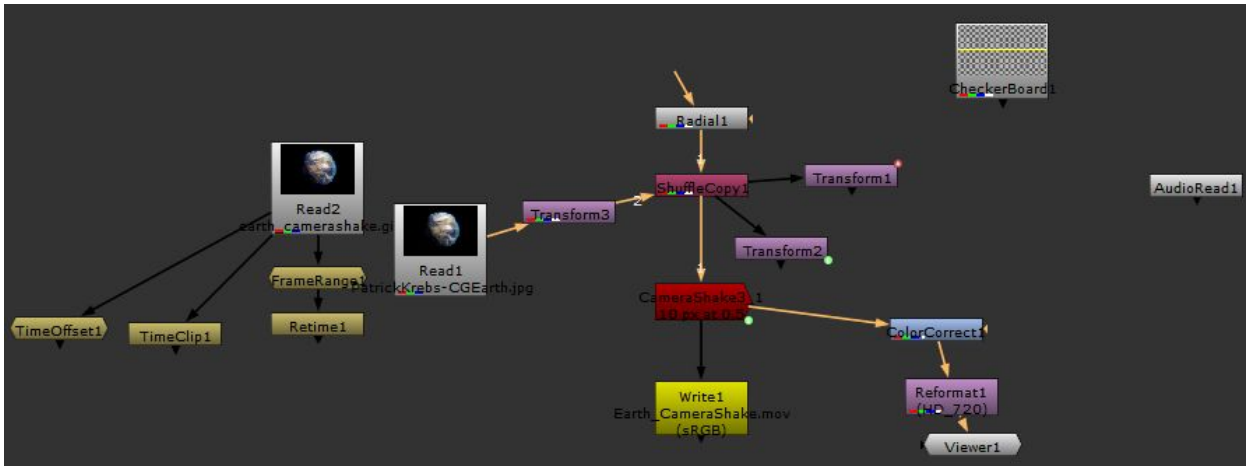

Much like Houdini, if your graph isn't entirely visible a minimap will show up on the lower-right hand side of the node graph…

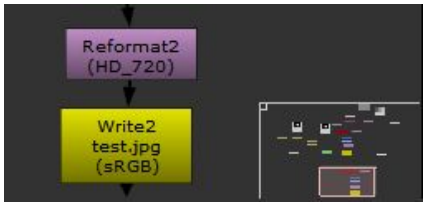

You can…

- click-and-drag inside this minimap to quickly pan around
- resize the minimap by dragging out the top-left corner

Here's are some handy tips for working with the node graph. Remember that your mouse has to be hovering in the node graph for these to work.

How to navigate the node graph...

- F → frame up selected nodes, or everything if no nodes are selected
- Alt+LMB → pan graph by moving mouse
- Alt+MMB → zoom in/out by panning mouse left/right (fine zoom)
- Mouse wheel → zoom in and out (coarse zoom)
- J → jump to a bookmark or backdrop node

  NOTE: MMB click-and-drag can also pan, + and - can also be used for coarse zoom.

Useful keyboard shortcuts for the node graph…

- Tab → add node to graph (tab menu -- need to know the name of the node)
- DEL → delete selected nodes from graph
- D → toggle disable/enable selected nodes
- Alt+D → toggle forcing the selected nodes to show up in the dopesheet
- L → auto-layout selected nodes
- S → brings up project settings in the properties panel
- 0-9 → set current viewer's n^{th} input to selected node (more info in Viewer Panel section)

One common mistake in Nuke is to have the viewer go missing. For a viewer to show up in Nuke, you need a Viewer node in your node graph. If you notice that your viewer has gone missing, it's probably because you deleted your Viewer node. Simply add it back in and connect it back up to the node you want to view. See the Add Nodes subsection below.

Basic Operations

Add Nodes

There are 3 ways to add a node...

1. Hit tab and type out the name of the node you want to add…

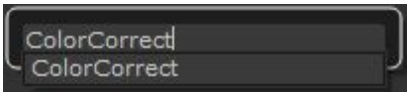

NOTE: Make sure your mouse is hovering over the node graph when you hit tab

2. Right-click in the node graph and use the context menu to add a node…

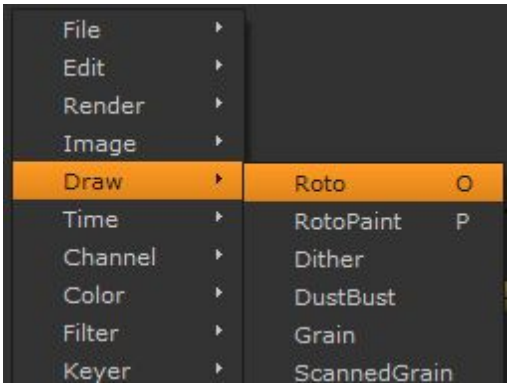

3. Use the toolbar

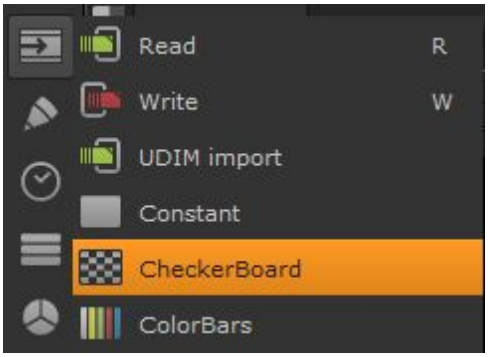

When you add a node, the output of the currently selected node (if you have a node selected in the node graph) will automatically get fed into the newly added node…

Otherwise, the node will be added with nothing connected to/from it…

Remove Nodes

To remove a node, simply select it in the node graph and hit DEL…

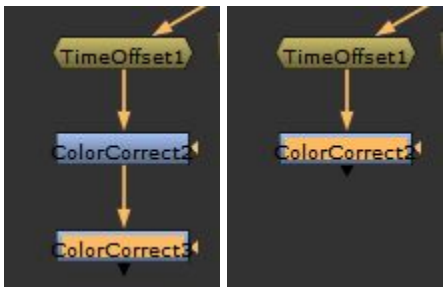

If you delete a node that's sandwiched inbetween 2 other nodes, those 2 other nodes will automatically hook up once you delete the sandwiched node…

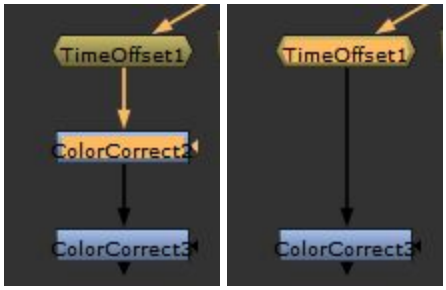

Select Nodes

To select a node, LMB click it…

To select a group of nodes, you can marquee select by LMB click-and-drag…

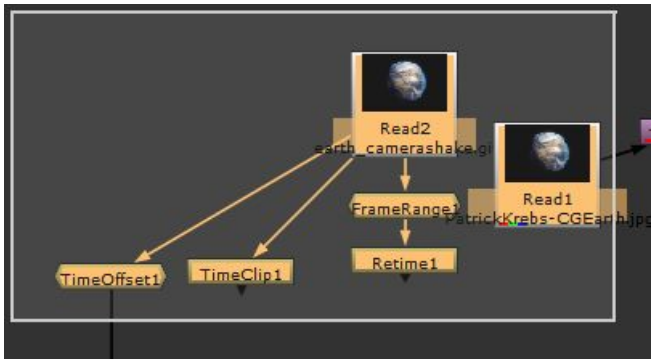

To add/remove from your existing selection, hold Shift while you select nodes…

- If it's already selected, it'll be deselected.
- If it isn't already selected, it'll be added to the selection.

To select everything upstream from the node, hold Ctrl+Shift and LMB click…

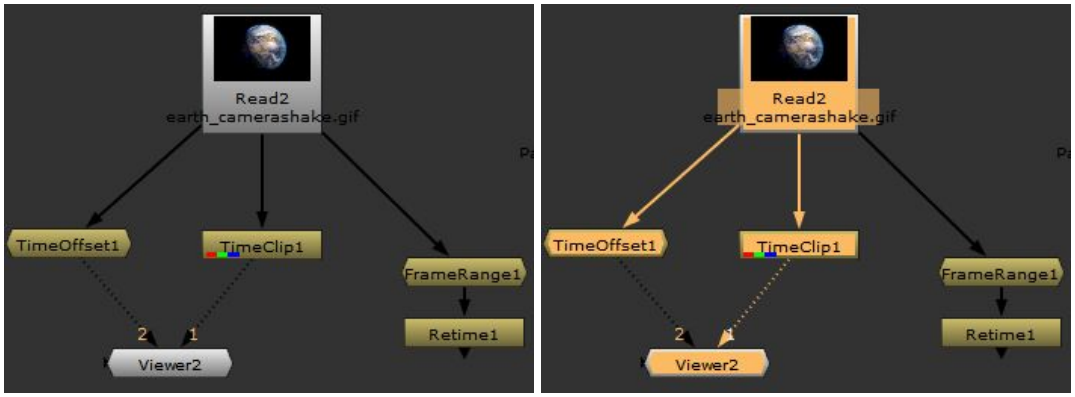

To invert selection (to change your selection such that things that you selected are deselected and everything else is selected), RMB click and goto Edit → Invert Selection…

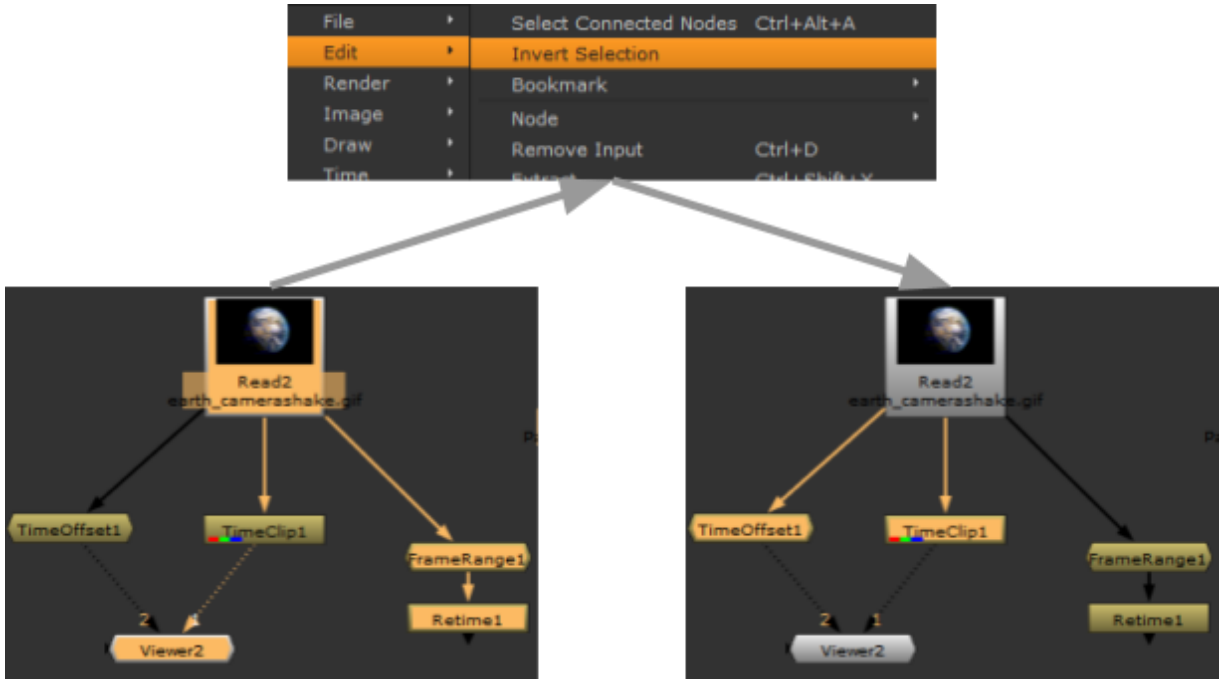

Move Nodes

To move a node, simply select them and LMB click-and-drag…

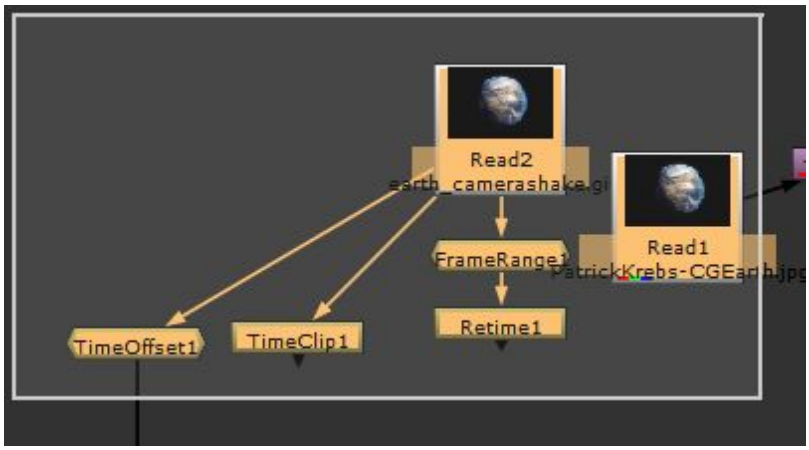

Connect Nodes

Most nodes loose nodes will end up looking like this…

The overall pattern with nodes is...

- inputs show up at the top.
- outputs show up at the bottom.
- special inputs/outputs show up at the sides.

To connect up 2 nodes, you can either…

- LMB click-and-drag the input arrow for some node and connect it to an output…

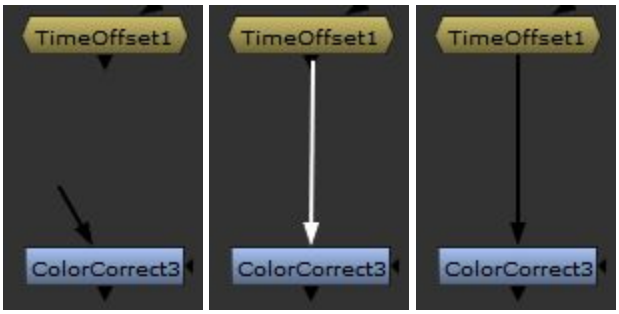

- LMB click-and-drag the output arrow for some node and connect it to the input…

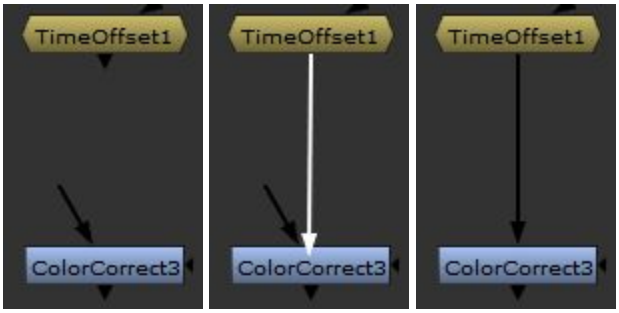

Note that a node can have multiple. For example, a Merge node has 2 inputs…

Another example is a Viewer node, you can drag out the node on the left edge (only shows up if the viewer is already connected to something) and it'll act as a new input…

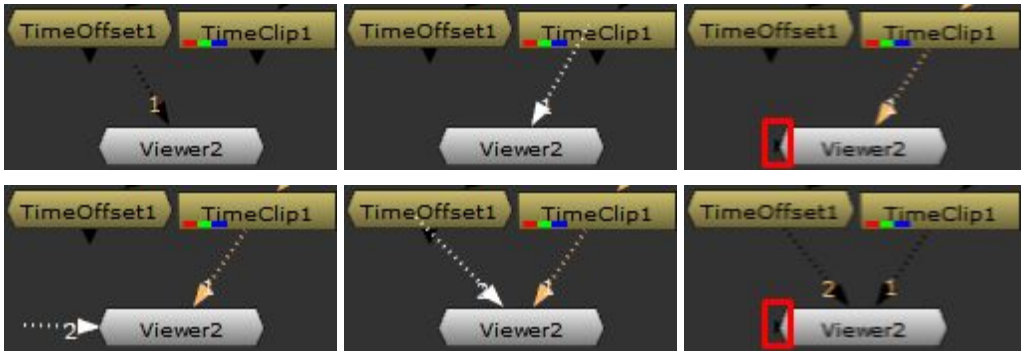

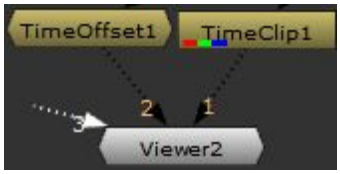

Special inputs are also located on the side. For example, a Merge node can take in a mask via the connector on its right-side…

Open Node Properties

To view a node's properties, simply double-click it. The properties for that node will open up in the properties panel…

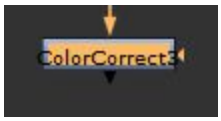

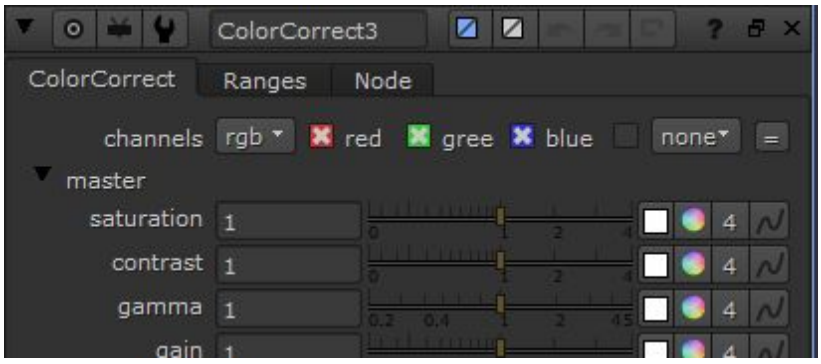

> **NOTE**: See the Properties Panel section for more information on properties and how the properties panel works.

Search for Nodes

> **NOTE**: The hotkey for this is /.

You can automatically select one or more nodes based on the name and bring them into view. To do this, right-click anywhere in the node graph and goto Edit → Search…

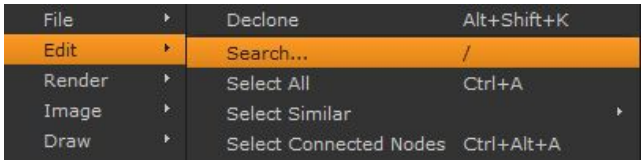

A dialog will pop up asking for a glob-style pattern…

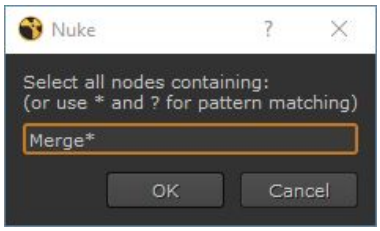

Once you hit OK, it will automatically select all the nodes that match the pattern but only frame up the first one. Remember that to frame up all the selected nodes, hit F.

Layout Nodes

> **NOTE**: The hotkey for this is L.

You can automatically layout nodes by selecting them, then LMB right-clicking anywhere in the node graph and going to Edit → Node → Autoplace…

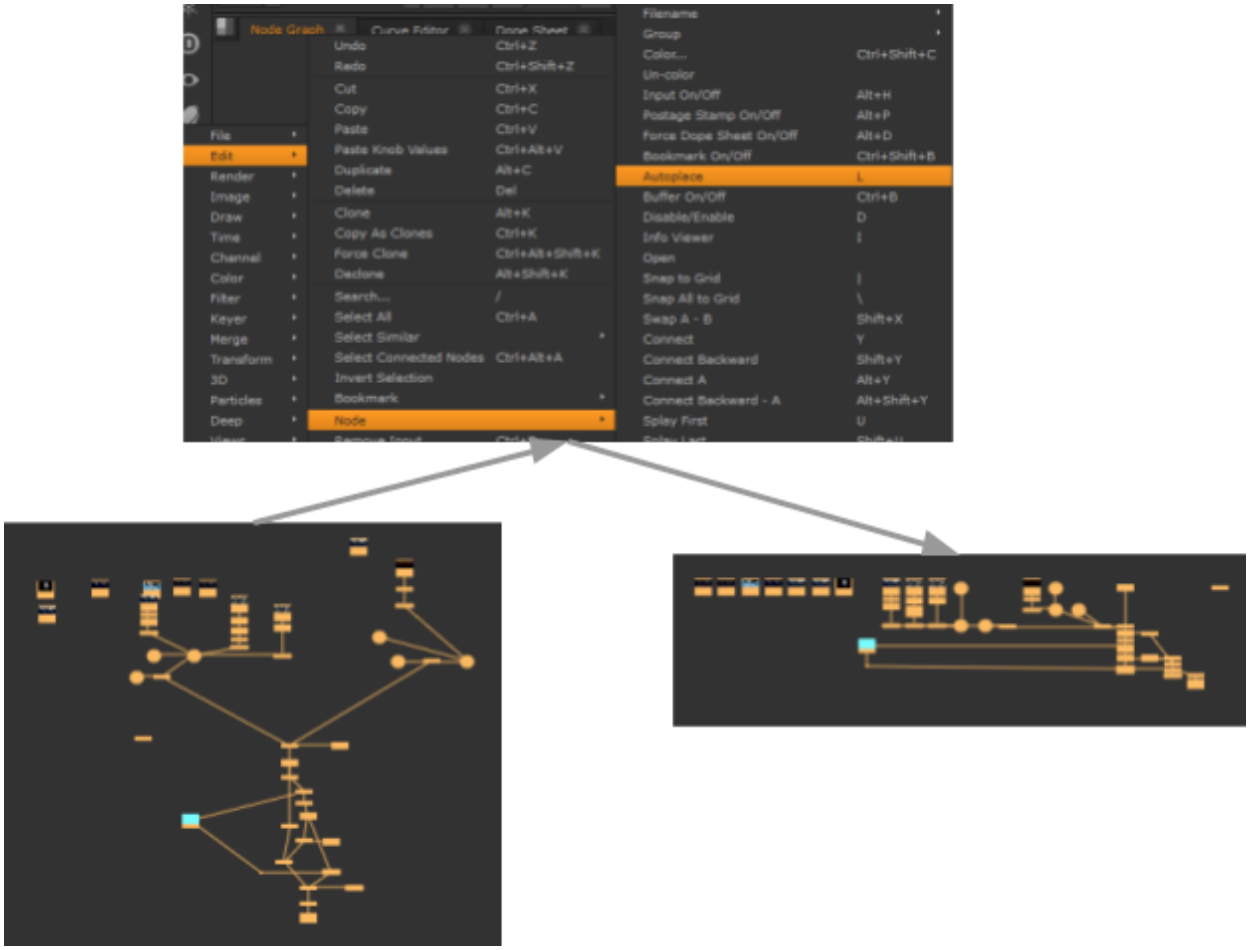

Special Node Types

Dot Nodes

Dot nodes are nodes that allow you break up an edge into multiple pieces, such that your node graph becomes easier to view/navigate. They don't have any functionality as far as the comp is concerned…

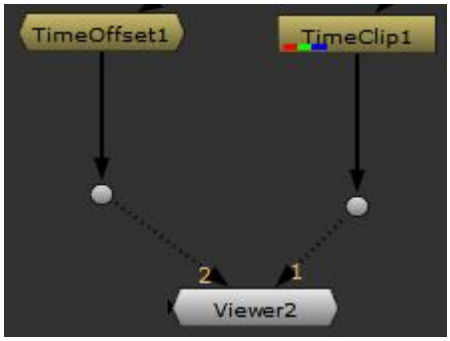

NOTE: Same thing is available in Houdini and Katana.

To split up an edge using a dot node, simply hold Ctrl. Little yellow diamonds will appear on top of edges that can be split. LMB click-and-drag a diamond to create a new dot node...

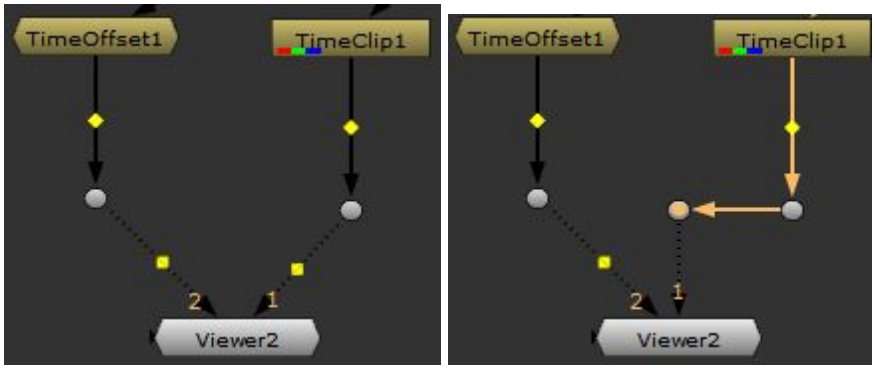

> **NOTE**: You can also drop in a dot node by itself via the Tab menu -- hit tab and type in Dot.

Backdrop Nodes

Backdrop nodes are like Katana backdrop nodes or Houdini's netboxes. To add a backdrop node, hit Tab and type in Backdrop…

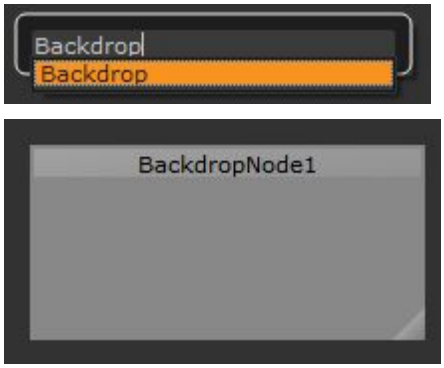

> **NOTE**: There seems to be a bug where the backdrop node doesn't get placed where you hit tab. It drops somewhere random?

Nodes placed inside the backdrop node can be moved all at once by LMB click-and-dragging the titlebar of the backdrop node…

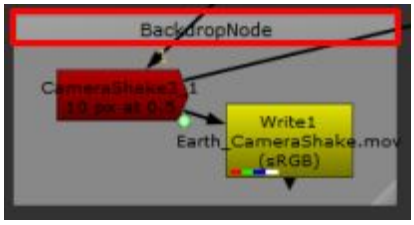

Backdrop nodes are bookmarked by default, so you can quickly get to your backdrop nodes by hitting the J key and typing in the name…

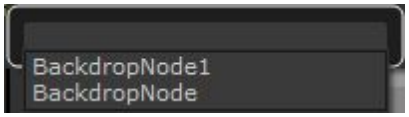

> **NOTE**: Just like with any other node, you can turn off the bookmarking by going to the properties.

You can change the name of the backdrop node by opening up the properties (LMB double-click) and changing it...

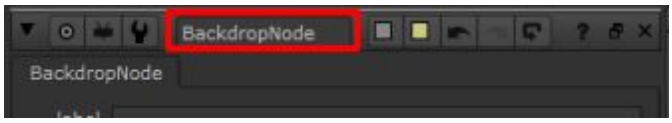

You can resize the backdrop node by dragging the lower-right corner…

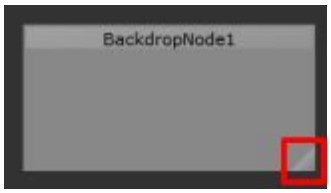

> **NOTE**: Deleting a backdrop node will delete whatever nodes you have inside that backdrop node.

Bookmarks

You can bookmark positions in the node graph based on a node or based on some arbitrary location.

Node

Nodes can be bookmarked by going to the node properties. In the node properties, you can navigate to the Node tab and select the bookmark checkbox…

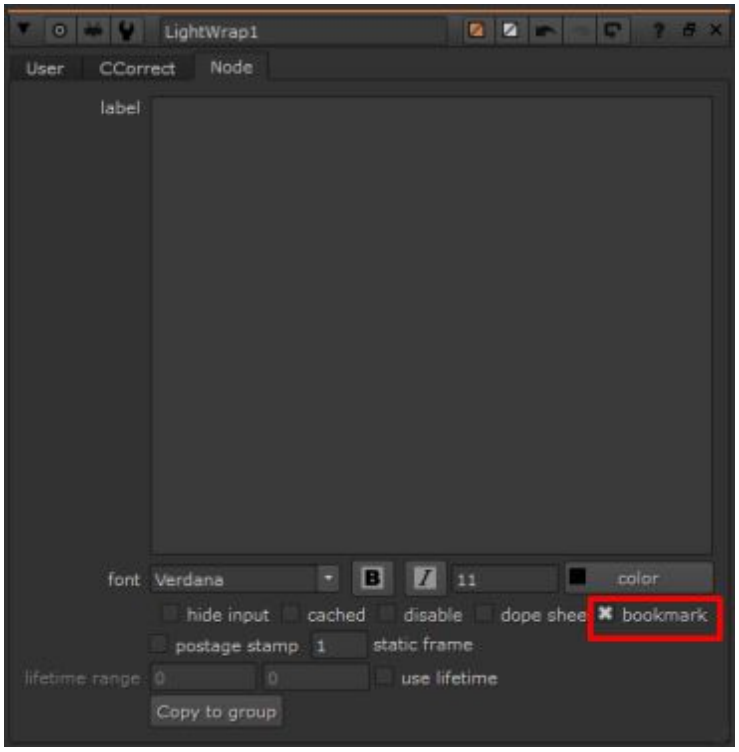

You can jump to any bookmarked node by hitting J and typing in the name of the node (or selecting the node if it's already showing up)…

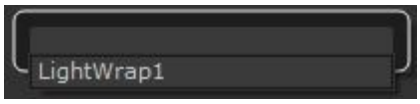

Location

Locations can be bookmarked. When you bookmark a location, you're recording the x, y, zoom level, and the pan level of the location. This is different from a node bookmark -- the J (jump to node) hotkey can't be used to jump to a location bookmark.

To bookmark a location, right-click anywhere in the node graph and goto Edit → Bookmark → Save Location <#>…

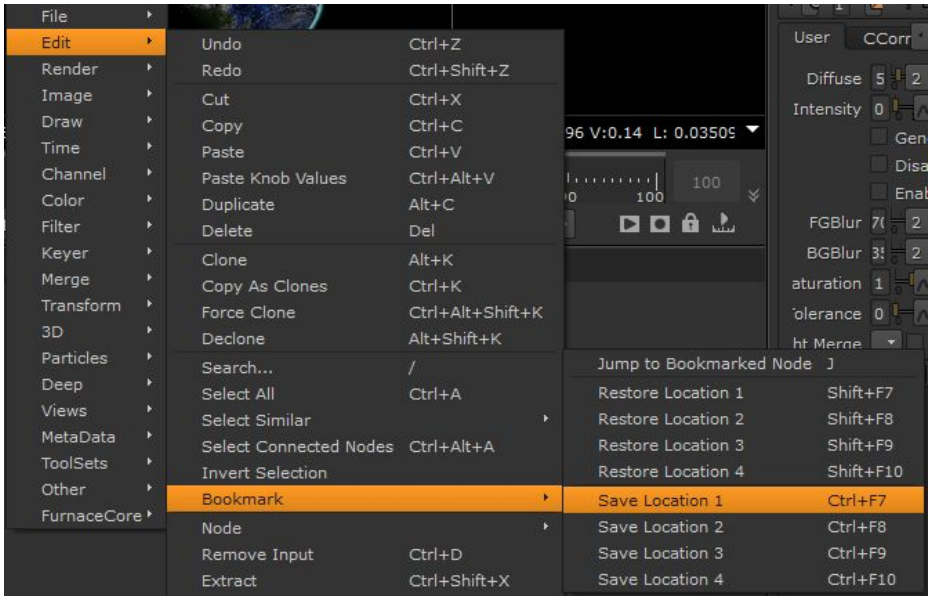

To go back to a bookmarked location, right-click anywhere in the node graph and goto Edit → Bookmark → Restore Location <#>...

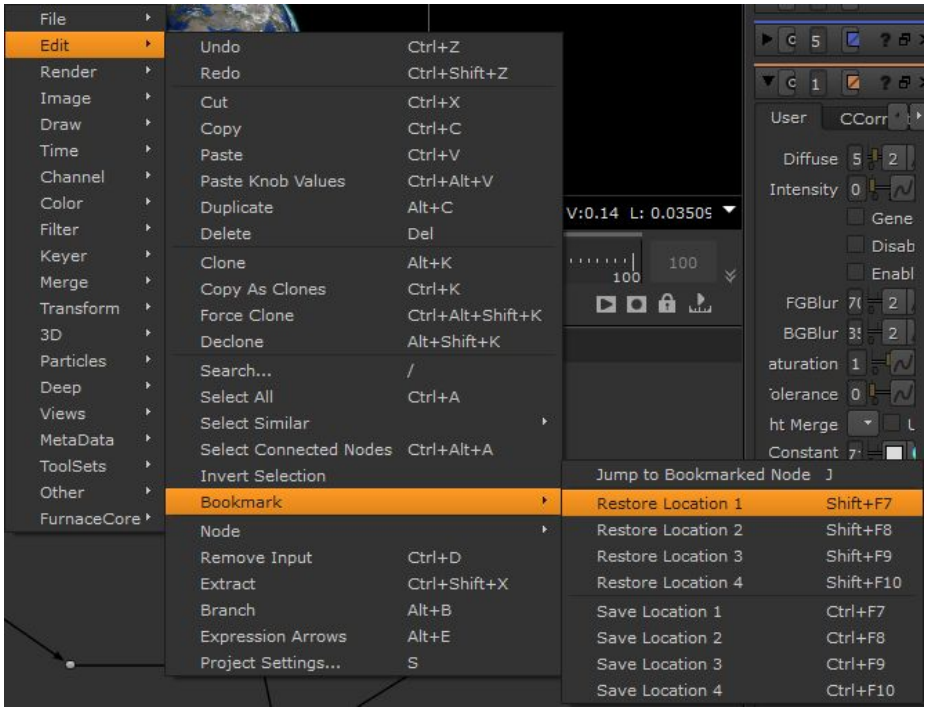

NOTE: Remember that bookmarking a location is different from bookmarking a node. You cannot jump to a location using the J hotkey (jump to bookmark hotkey). You have to use the Edit menu.

Toolsets

You can save a subtree of your comp graph for later use via the toolsets functionality.

Saving

To save a sub-tree as a toolset, begin by selecting all the nodes…

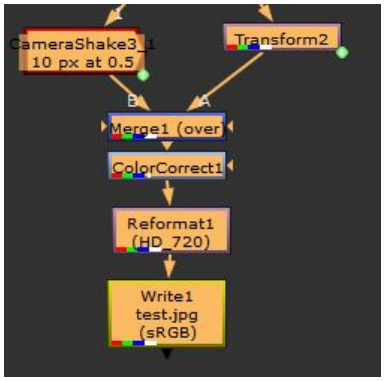

Then, right-click in the node graph and goto Tool Sets → Create…

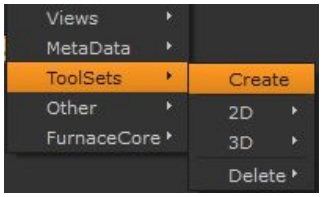

Type in the name of your tool in the dialog that pops up and hit Create…

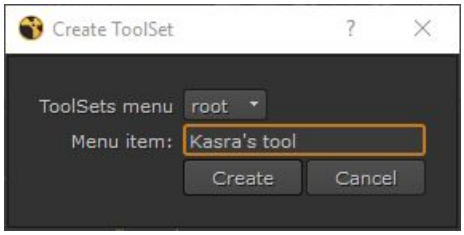

Loading

You can now dump a copy of the nodes that were saved back into the graph. You can do this by going back to the ToolSets menu or by using the Tab menu.

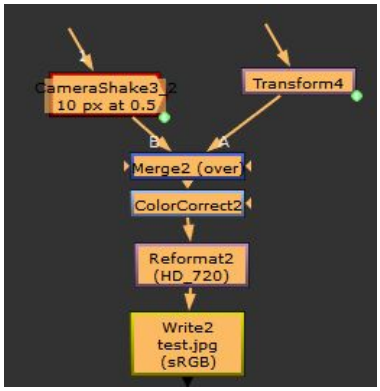

Deleting

If you want to remove the ToolSet, right-click in the node graph and goto Tool Sets → Delete → <name>…

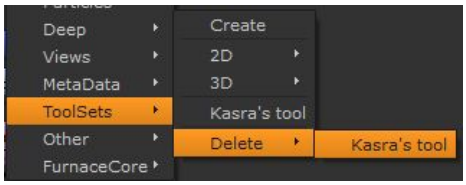

Viewer Panel

The viewer panel allows you to see the comp at a specific point in the node graph...

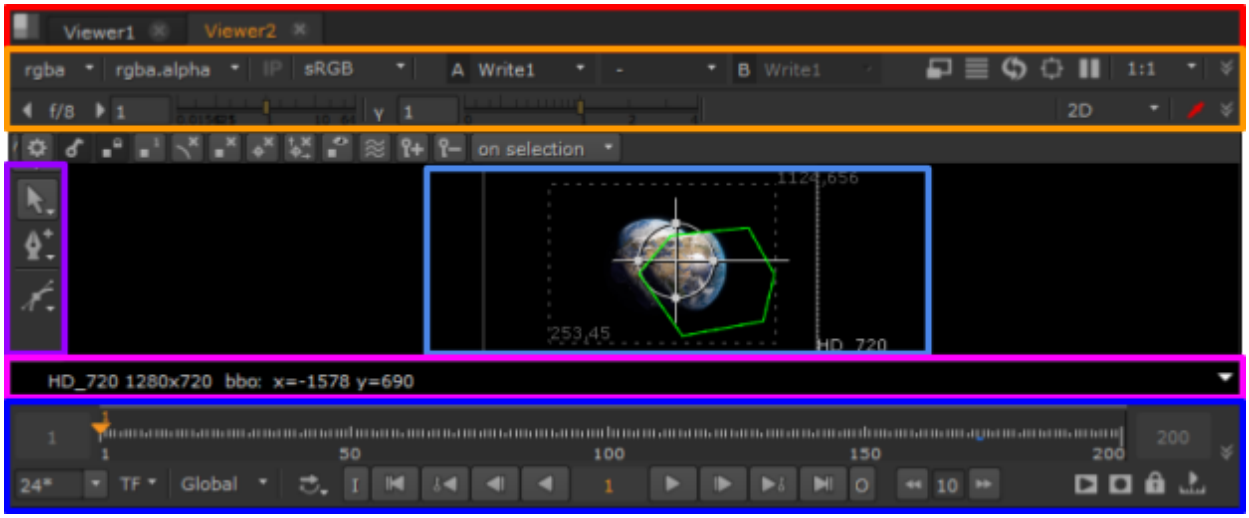

NOTE: Remember that you need to have at least 1 Viewer node in your node graph to get anything to show up here. Connect your viewer node to whatever it is you want to show in the viewer.

Viewers The number of viewers you have available is based on the number of Viewer nodes you have in your node graph. If you have no Viewer nodes

in your graph, you won't have any Viewer tabs in your Viewer panel.

In the example above, there are 2 Viewer nodes: Viewer1 and Viewer2...

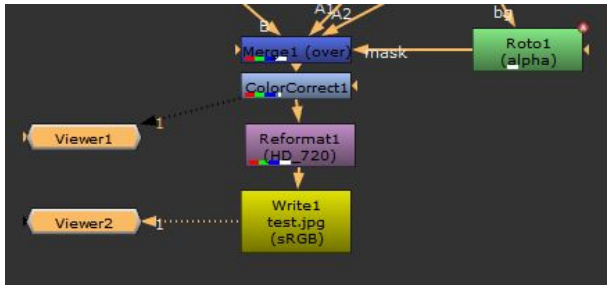

Toolbar

The toolbar gives lets you do a few high-level things. If you have trouble finding some of the items shown here, be aware that toolbar hides items when the window isn't big enough to contain them -- expand your window to view the entire toolbar. Also, you can hover over each item in the toolbar to get a tooltip that describes what it does.

The overall options are…

these are options for color channels and color space...

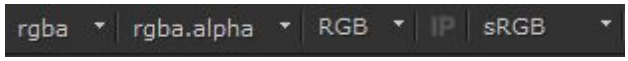

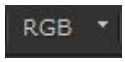

← lets you cycle between RGBA channels (hotkeys are R, G, B, and A)

these are options for A/Bing to outputs (remember that your viewer can have multiple inputs, these inputs will be available in the A and B dropdowns)...

these options control rendering...

→ proxy a lower res output in the viewer?

→ unsure what this does?

→ force render

→ render region (only focus on rendering a certain region of comp)

 → pause (don't render anything if comp tree changes)

 → downscale inputs to the viewer by this ratio (effects )

these are extra options that are collapsed because the toolbar isn't wide enough…

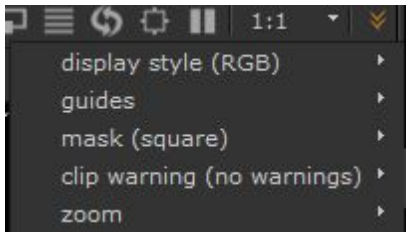

this controls if the viewer is in 2D mode or 3D mode (remember that your comp can have 3D portions that eventually get rendered to an image and gets used in the rest of your comp)...

This toggles the viewer infobar...

Display

The actual display of the viewer. Shows you the output along with…
* any guides that you've set
* tools for any nodes open in the properties panel (even if those nodes aren't part of the comp tree being rendered out to the viewer)

NOTE: The second point is a super important thing to watch out for because it can cause a lot of confusion. For example, if you have the properties of a Roto node open, the mask used will be visible / manipulatable in your viewer. This is true even if that roto node is not part of the comp tree being rendered out.

Viewer Infobar

Provides information about what you're viewing in the viewer. It'll tell you things like size, and if you hover your mouse in it it'll tell you the pixel value of what you're hovering over. Probably other things as well.

Toolbar

Provides tools to update whatever node properties it is that you have open. For example, if you have the Roto node's properties open, you'll get tools to draw or modify the region of the node…

NOTE: Just noticed this for Roto nodes… you need to select the region first with the pointer (1st item) before you can modify it with the pen tool (2nd item) or the slope tool (3rd item).

Scrubber

Animation tools typical with any 3D package…

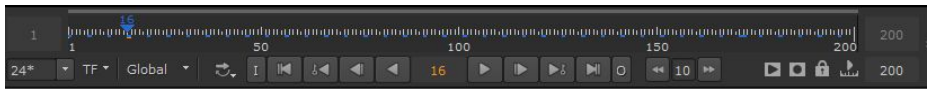

playback controls are pretty straightforward…

the numbers at either end define the input/output range…

setting these will cause playback to be limited within the range specified (by default they're set to the frame range defined in the Project Settings).

first dropdown defines the playback rate…

second dropdown defines points in the scrubber are displayed…

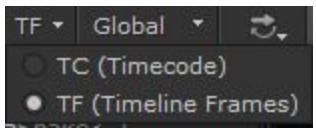

← TF

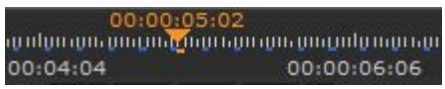

← TC

frames that have keyframes on them will have a blue underline…

if the scrubber position is over a frame with a keyframe, it'll turn blue as well…

as frames are cached, they'll get an orange underline...

buttons on the far left control output and viewer synchronization...

 → flipbook -- fully renders the viewer and displays it

 → capture -- caches out the viewer and displays it, the output here will be whatever is visible in the viewer (including any extra markings such as tools like the region for a Roto node)

 → lock -- syncs the scrubber for this viewer with any other viewer that his the lock enabled as well

How to navigate the viewer panel...

- F → frame up such that the entire image fits into view
- Alt+LMB → pan view by moving mouse
- Alt+MMB → zoom in/out by panning mouse left/right (fine zoom)
- Mouse wheel → zoom in and out (coarse zoom)
- Ctrl+1 → zoom to 100% (native resolution of image)
- Ctrl+2 → zoom to 200%
- ...
- Ctrl+9 → zoom to 900%
- Shift+[→ toggle toolbar (top of viewer) on/off
- Shift+] → toggle animation controls (bottom of viewer) on/off
- +/- → zoom in and out (coarse zoom)

Basic Operations

Adding Viewer

You can add more viewers by dropping extra Viewer nodes into your node graph…

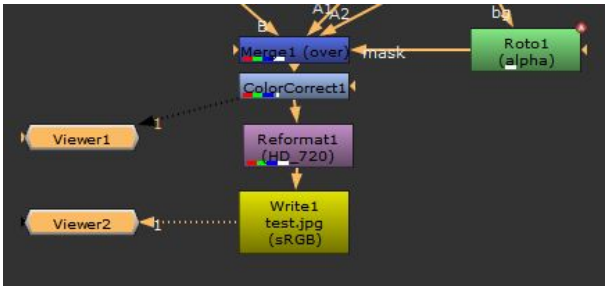

NOTE: If you have no Viewer nodes in your graph, you won't have any Viewer tabs in your Viewer panel.

Connecting Outputs

To view the output of a node, you can connect that node's output into your viewer…

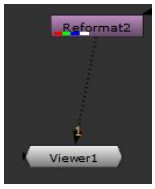

You aren't limited to 1 node. You can have multiple inputs into your Viewer. You can do this by dragging out the arrow on the side of the node and connecting it up to another node…

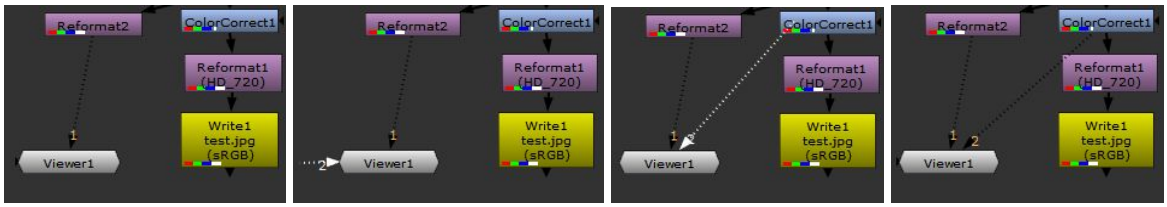

NOTE: Instead of doing this manually, you can just select the node you want to view and hit a number key. For example, if you select a node and hit 5, the 5th input into the viewer will be from that node…

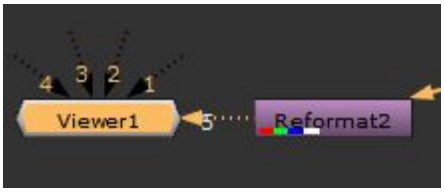

Be aware that these hotkeys only work with the viewer associated with currently open/visible viewer tab. This can get confusing if you have more than 1 viewer node.

Viewing Outputs

You can cycle between which output you're viewing in your viewer by either…

- using the hotkeys 0 to 9 while your mouse is hovered over the viewer to switch.
- using the hotkeys UP and DOWN while your mouse is hovered over the viewer to cycle.

- using the dropdown in the viewer toolbar…

NOTE: Dropdown not visible? It's probably hidden because the window isn't wide enough. Widen the window.

Thumbnailing Outputs

If you want to show multiple outputs in the Viewer at the same time, you can use a ContactSheet node. The ContactSheet node essentially makes a tileset of the nodes feeding into it…

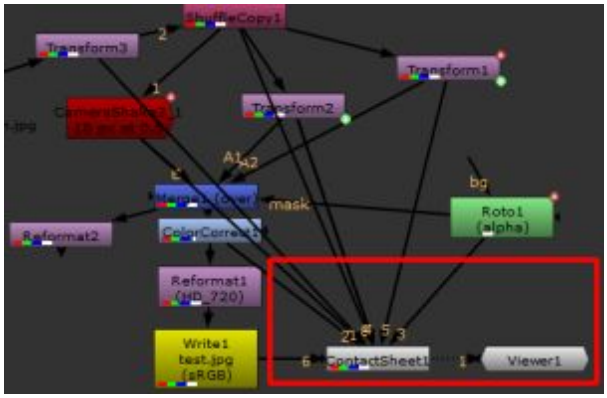

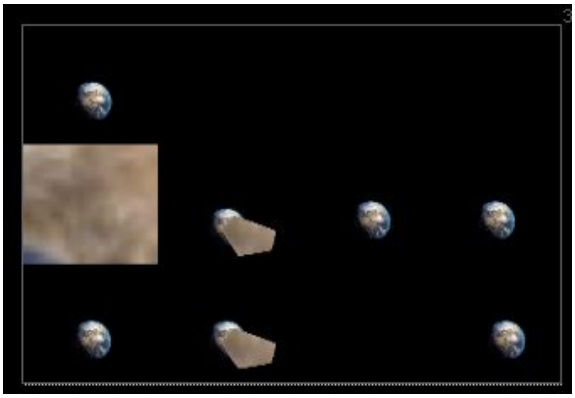

A/Bing Outputs

If you have more than 1 input into your viewer node, you can A/B outputs.

Go to the viewer toolbar and select Wipe for your A/B mode…

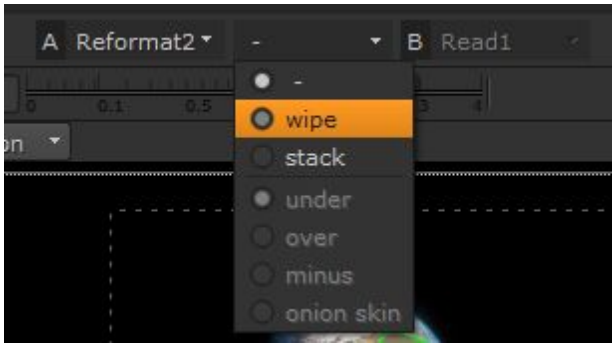

NOTE: Can't find this? It's probably hidden because the window isn't wide enough. Widen the window.

Then select the the A node from the left dropdown and the B node from the right dropdown…

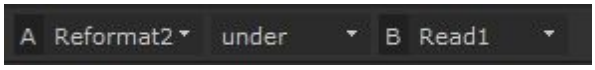

In the center of your viewer, you should have a crosshatch handle that you can drag to let you wipe up/down/left/right + a little circular handle to control the alpha...

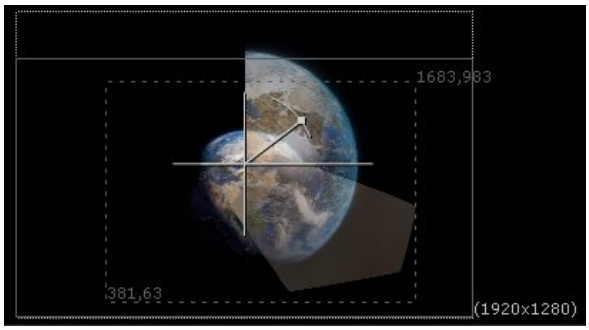

Rendering Output

There are 2 ways to render what you're seeing in your viewer: capture or flipbook. You can find the buttons for these on the lower right-side of the Viewer tab.

> **NOTE**: Can't find the toolbar dropdown? It's probably hidden because the window isn't wide enough. Widen the window.

What's the difference between capture and flipbook?

- Capture () caches out the viewer and displays it. The output here will be whatever is visible in the viewer, including any extra markings such as tools/guides/masks/etc.. (e.g. region for a Roto node or title safe area)...

  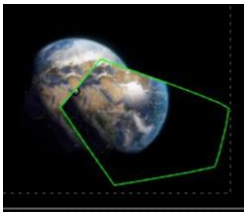

- Flipbook () fully renders the viewer and displays it WITHOUT any markings…

  

RGBA Channels

To view RGBA channels, you can either…

- use the RGB dropdown in the toolbar…

- use hotkeys R, G, B, and A to toggle each channel on/off.

  NOTE: Can't find the dropdown? It's probably hidden because the window isn't wide enough. Widen the window.

2D and 3D Views

Nuke has 3D nodes as well as 2D nodes. The viewer should automatically switch to the correct view type based on what node you're viewing, but you can always manually switch between 2D and a 3D view using the view selection dropdown in the viewer toolbar…

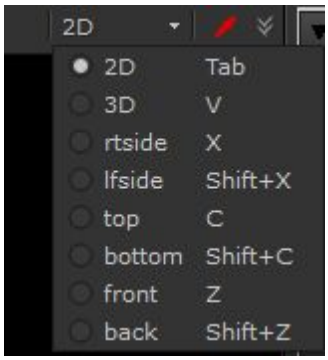

NOTE: Can't find the dropdown? It's probably hidden because the window isn't wide enough. Widen the window.

NOTE: Switching manually makes no sense. If you switch to a 3D view but you're connected to a 2D node, it'll show you the output of the last 3D node you were connected to? Or all 3D leaf nodes? Unsure exactly how the 3D view works…

In 3D view, you can basic camera and selection controls…

- marquee select or single select and move around based on the axis
- Alt+LMB to pan
- Ctrl+LMB to tumble
- Alt+MMB to zoom (fine zoom)
- Mousewheel to zoom (coarse zoom)

2D view...

3D view...

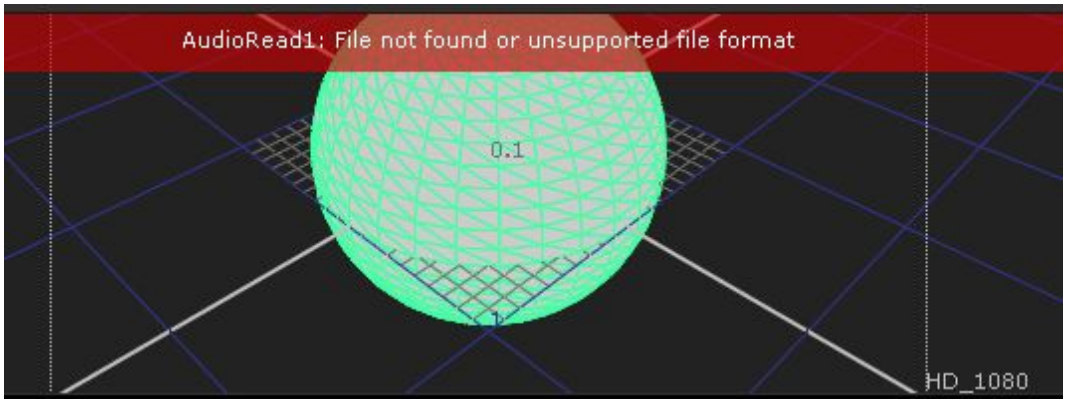

Guides

Guides provide a way to catch problems quickly and/or work within a set of constraints.

Frame Masking

“Frame masking” is when you need to quarantine a part of the shot that you’re working on. The reason why you’d want to do this is because…

- aspect ratio → maybe the client requires that work in a specific aspect ratio.
- title safe → because of the differences between TVs, you need to use the title safe guides when printing text to the screen... generally this happens for end credits, lower thirds/chyrons, or title screens.
- action safe → most TVs extend the image past the viewable area, possibly by a lot… the action safe guide marks what part of the image that will be guaranteed viewable on all TVs… you generally want to keep all “action” in this safety zone.

To show the different frame masks (guides), use the safezone dropdown button  in the viewer toolbar…

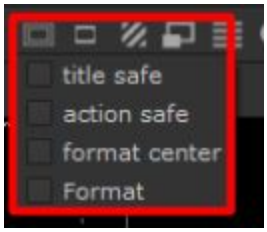

NOTE: Can't find the dropdown button? It's probably hidden because the window isn't wide enough. Widen the window.

The item…

- title safe → described above

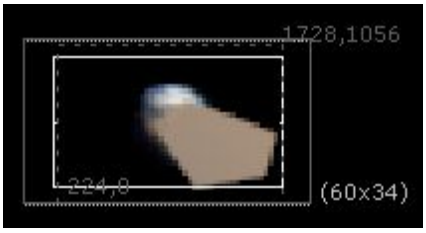

- action safe → describes above

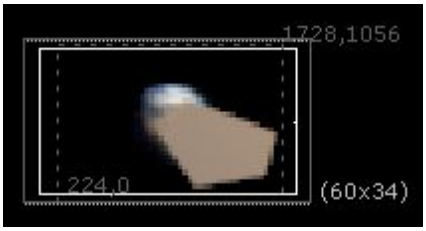

- format center → puts a little crosshair in the center of whatever your native format is (e.g. HD_1080 1920x1080)

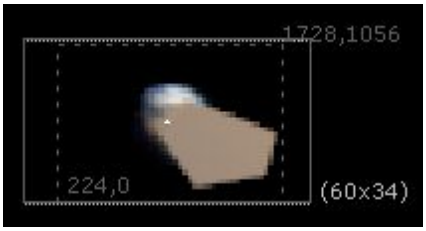

- Format → puts red guides showing you the upper/lower region of whatever your native format is (e.g. HD_1080 1920x1080)

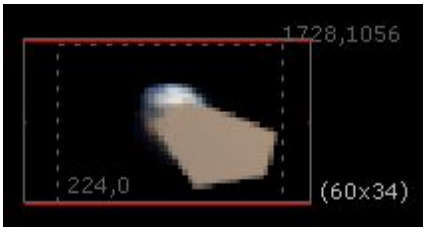

Aspect Ratios

You can set aspect ratio (guides), using the aspect ratio dropdown button  in the viewer toolbar…

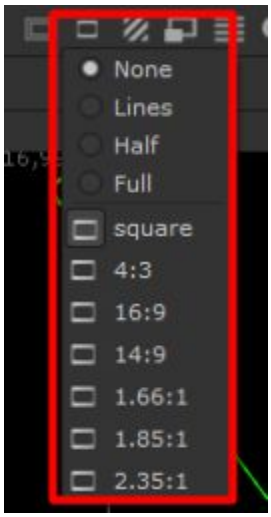

> **NOTE**: Can't find the dropdown button? It's probably hidden because the window isn't wide enough. Widen the window.

The first part of the dropdown defines how you want the aspect ratio guides to show up…

- None → nothing visible

- Lines → white lines

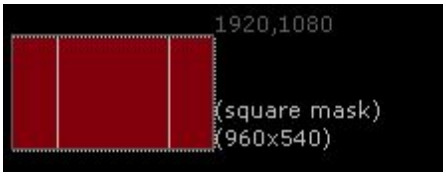

- Half → dimmed bars (half intensity of original pixels)

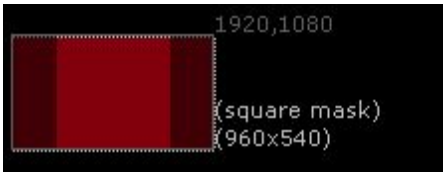

- Full → full black bars

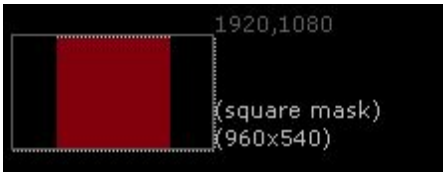

The second part of the dropdown defines the aspect ratio. This should be self-explanatory.

Overexposure

The clip warning / zebra striping button can be used to find areas of the image that have become over exposed. Over exposed images will cause older TVs to bleed color out of the intended boundaries.

To find overexposed portions of your image, first turn on the clip warning button in the viewer toolbar…

Now, if you have parts of your image that are overexposed, that part of the image will show up as stripes…

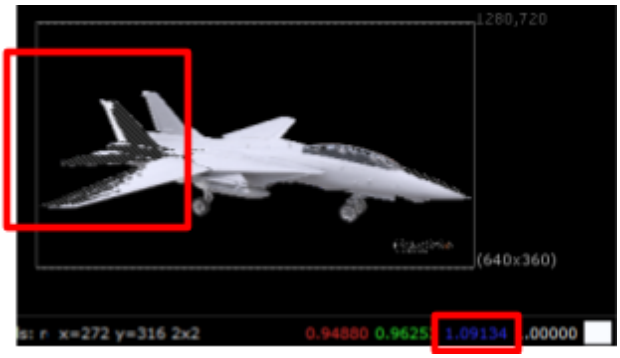

Computation Speed

Just like with 3D renders, your comp may be a computationally intensive process. If things are taking too long to render out in your viewer, you can try some of the methods in the subsections below.

Region of Interest

If things take too long to render out in your viewer, you may want to limit your viewport to rendering a small part of your comp. This is what the Region of Interest (ROI) tool does.

To use the ROI tool, first, select it in the viewer toolbar…

> **NOTE**: Can’t find this button? It’s probably hidden because the window isn't wide enough. Widen the window.

A little box will show up in your viewer…

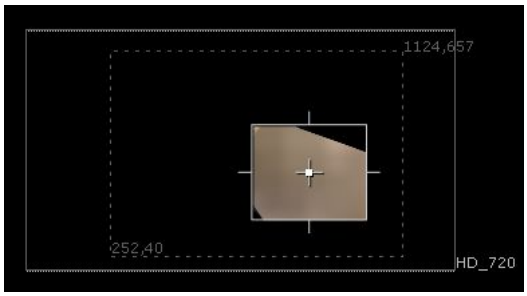

This is the region that gets rendered. You can drag it around and resize it however you see fit.

Proxy Modes

A proxy is a version of your comp, but set to a much lower resolution. As such, your comp tree should take much less time to compute.

To set the proxy mode, select the downres ratio in the viewer toolbar…

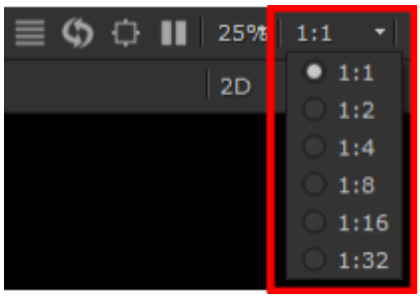

> **NOTE**: Can’t find this dropdown? It’s probably hidden because the window isn't wide enough. Widen the window.

So for example, if we tell the viewer to downres to 1:32, we’re going to get a very pixelated picture in the viewer…

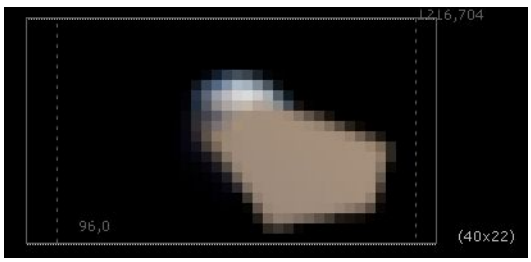

NOTE: Remember that this only affects the viewer in which it's assigned to. If you're piping to a Write node, you can downsample using a Reformat node… although I don't know if this has the same effect as reducing computation time across the entire chain of operations? Does downsampling the viewer even do this? Maybe the proper way to do it would be to add a Reformat node right after whatever it is you're reading in, and disable that node when you want to write it out to production???

Node Caching

Each node's properties has a Node tab. In that tab, there's a checkbox called cache…

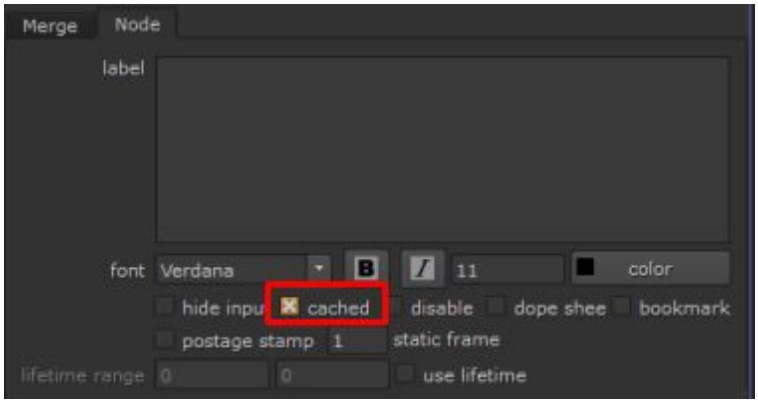

If selected, the node's contents are forcefully cached in memory such that if any node downstream changes, it will use the cached contents of this node instead of going further up the tree and recomputing everything.

Properties Panel

The properties panel allows you to see and manipulate the properties of nodes…

NOTE: Remember that to view the properties of a node, you need to double-click it in the node graph.

At the very top of the properties panel are the tools/options for the properties panel. Immediately following that are the actual properties for the nodes you've opened up.

Options and Tools

These 3 items are the options/tools for the properties panel…

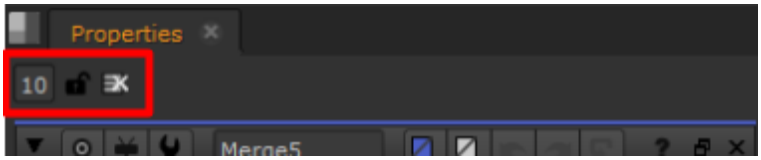

→ Textbox that defines max number of nodes the properties panel can have open. In this example, it's set to 10. If you try to open up the properties for more than 10 nodes, the oldest node open will get closed.


→ Toggle that locks the node properties (prevents opening new node properties). If this is on, double-clicking on nodes in the node graph to open up their properties will open then in a new floating window.

 → Closes all node properties in the properties panel.

Node Properties

Each node type has a unique set of properties. For example, the lightwrap node...

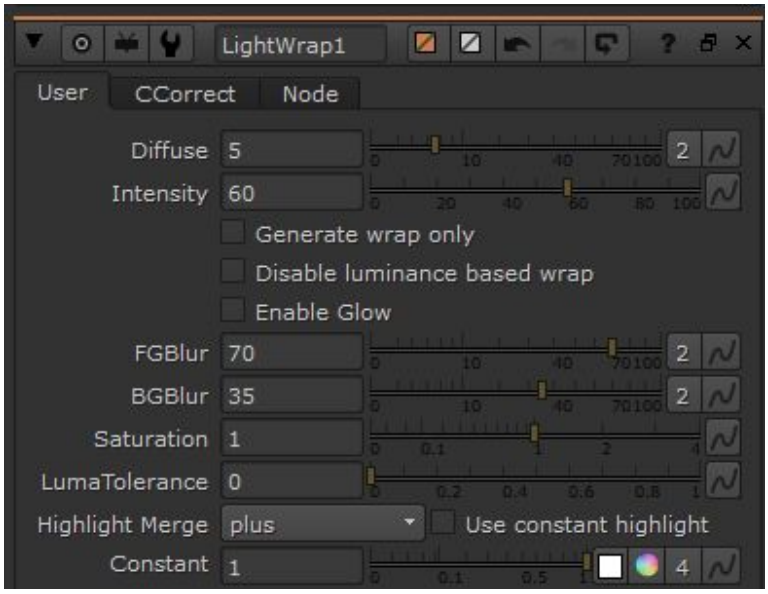

Each node opened in the properties panel has a toolbar followed by one or more tabs that group organize the properties for that node.

> **NOTE**: Be aware that when the properties of a node are open in the properties panel, any viewer manipulators for that node will show up in your viewer(s). For example, if you have a Transform node's properties open, the manipulator for that Transform node will show up in your viewers. This happens even if that transform node isn't connected to a viewer.

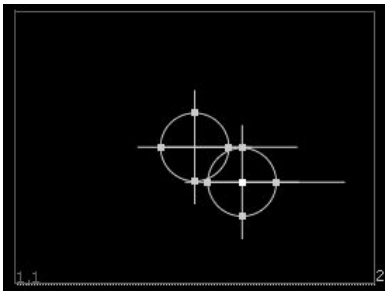

Toolbar

The header/toolbar provides quick access to a set of useful UI operations…

 → Collapses the properties for that node…

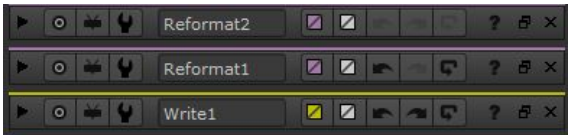

 → Centers the node in the node graph (but doesn't select it).

 → Select one of the inputs into this node in the node graph…

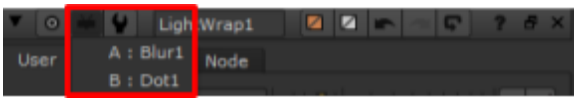

NOTE: Once selected, it'll be selected in the node graph but it won't be brought into view. To bring it into view, hover your mouse over the node graph and hit F (frame up).

 → Presets button. You can use this to save the properties state of your node and load it back up again (presets can be applied to all nodes of the same type)...

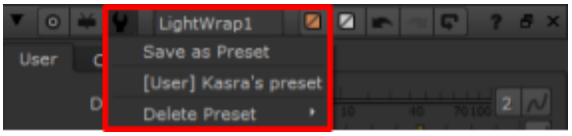

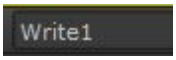 → Name of the node -- changing this will also change the name in the node graph.

 → Color of the node -- changing this will also change the color in the node graph.

 → Color of the tool in the viewer -- for nodes that have a tool that you can use in the viewer, this is the color that'll get used. For example, setting the mask in a roto node can be done in the viewer. When you draw the mask, this is the color that's used for

the outline (set to  in this example)…

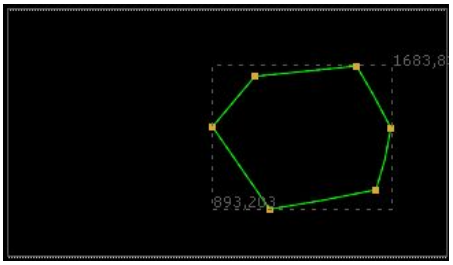

 → Undo last property change.

 → Revert last undo.

 → Revert the properties to the state they were in BEFORE the properties were opened.

Tabs

The tabs following the toolbar are how a node type organizes its properties...

The tabs for a node are unique to that type of node, so there isn't really much to cover here. But, almost all nodes seem to have a Node tab (covered in the Common Properties section).

Setting

Absolute

If the property is numeric, chances are you'll be provided with a slider…

The slider will have some range set to it, but you're not limited to that range. You can always click into the textbox and set the property to a value that's outside of the range of the slider…

When you do this, note that the slider goes to the edge and then changes to an arrow. This indicates that the value is out of the slider's range. If you click and drag the arrow back, it'll change the value to wherever you place it on the slider.

There are a couple of ways to update values other than using the slider / typing it in...

1. You can increment/decrement a numeric value by using Ctrl+UP and Ctrl+DOWN when you're in the text box. Depending on where the carrot is when you do this, the granularity of the increment/decrement changes.

For example, putting the carrot at the tens place and hitting Ctrl+UP twice will increment the value by 20…

2. You can increase/decrease a numeric value by clicking in the textbox, the click-and-dragging MMB left or right…

Unlike keyboard increments, the placement of the carrot doesn't change anything. However, if you hold Shift while you're dragging the increments/decrements will be more coarse.

Expressions

You can enter expressions into a textbox. For example, instead of typing in 960 you can type 1920/2. Once you press enter it'll evaluate and change it to 960.

Links

Properties for a node can be bound/linked to properties of another node by copying/pasting links.

Open the properties on node you want to bind from, go to the property you want to link, right-click it and choose Copy → Copy Links…

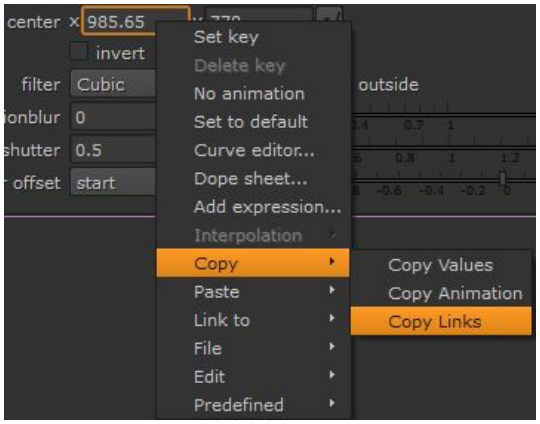

Then, open the properties for the node you want to bind to, go to the property you want to link, right-click it and choose Paste → Paste Absolute…

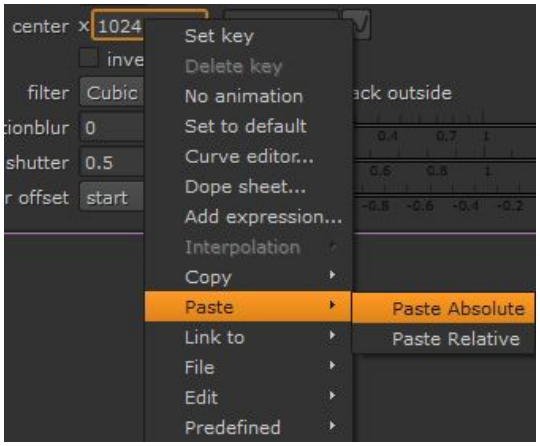

Once you do this, the destination property will go dim blue and a little green line will link to the node you've bound to…

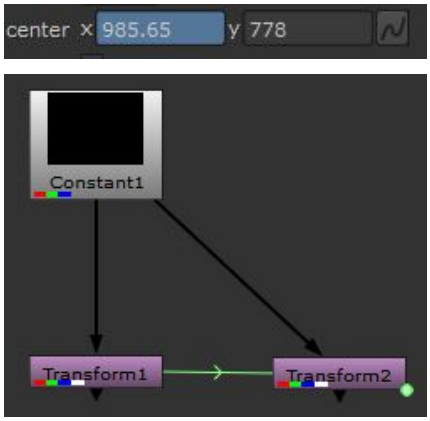

> **NOTE**: This is setting an expression. If you want to see what the expression is, right-click on the property and choose Edit Expression…

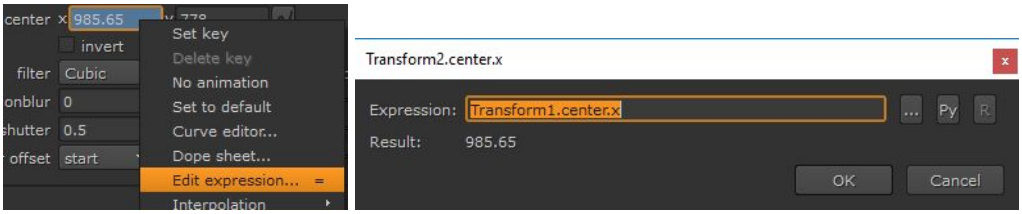

Expanding

If you see a button with a number in it (), it means that there are actually 4 different inputs for this property that are being collapsed down into 1 -- all 4 inputs are using the value specified.

You can customize all 4 inputs by clicking the button to expand them out…

> **NOTE**: You can click the button again to shrink it back down a single value, but be aware that this will make all values get set to the first value.

Color Picking

If the property you're trying to change is for a color, you can use either a color wheel ( ) or a color picker ( )…

The color wheel ( ) will expand the property to give you finer color selection controls…

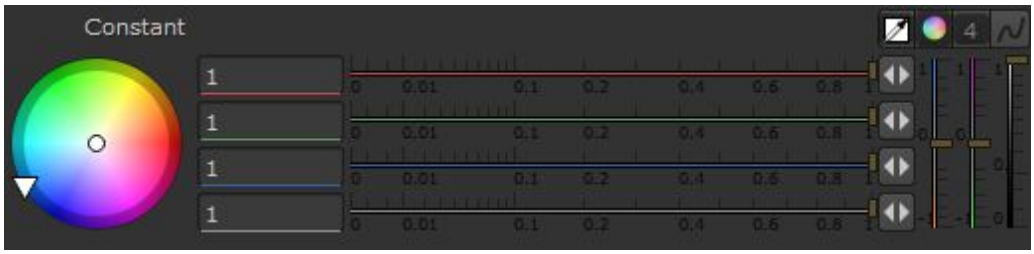

The color picker ( ), when enabled, will allow you to go to a viewer panel and Ctrl+LMB to copy that color into the property.

Keyframing

Properties that have a little curve button ( ) are keyframe-able. If you click the button, you'll get options for setting keyframes…

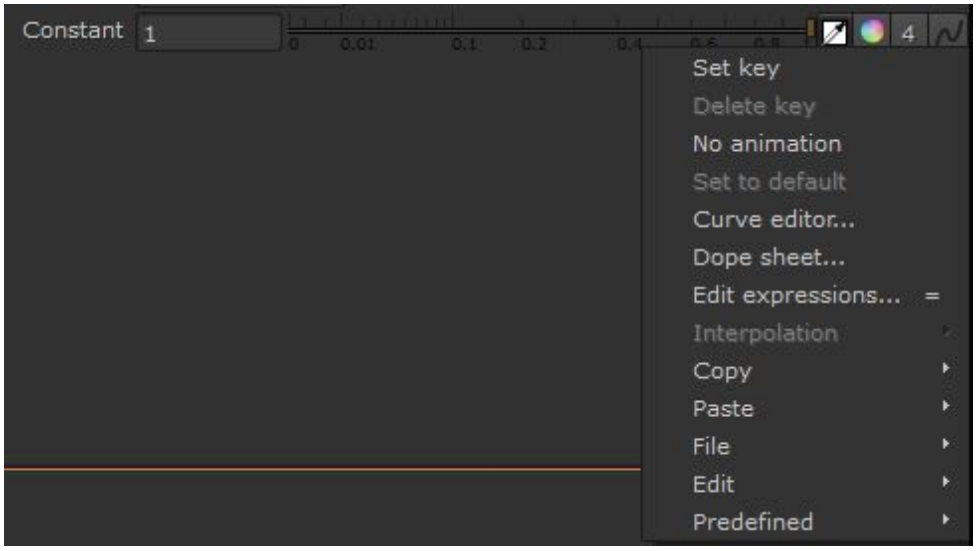

When a property has a keyframe set on the current frame, it'll have a bright blue background…

When a property is keyed but the current frame is not a keyframe, it'll have a dim blue background…

NOTE: Unlike other some other 3D packages, Node has auto-keying enabled. That means that if your property is keyed and you change it, it'll automatically add a keyframe at that frame.

NOTE: It looks like pretty much all properties are keyframe-able, even ones that are non-numeric. You can right-click on the input (e.g. the checkbox or the text input or whatever) and you'll pretty much get the same menu…

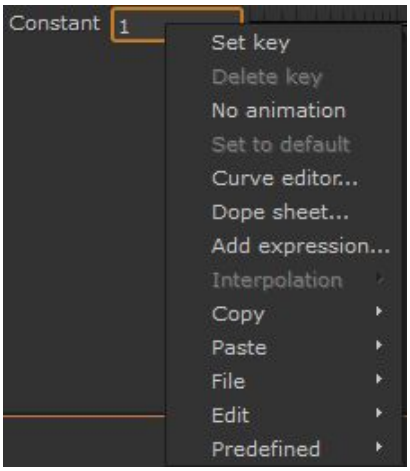

The most important options here are…

Set key → Add/update a keyframe at the current position (wherever the scrubber for the viewer/dope sheet/curve editor is at).

Delete key → Remove the keyframe at the current position (wherever the scrubber for the viewer/dope sheet/curve editor is at).

No animation → Remove all keyframes for this property.

Edit expressions → If animated, this will be set to “curve”. “curve” means that you are defining the keyframes and transitions between them…

You can change this to an expression (e.g. “frame/5”) such that the animation for this property is defined by an algorithm instead.

Common Properties

The Node tab contains a set of properties that configure how the node looks in the node graph as well as controls some common backend options…

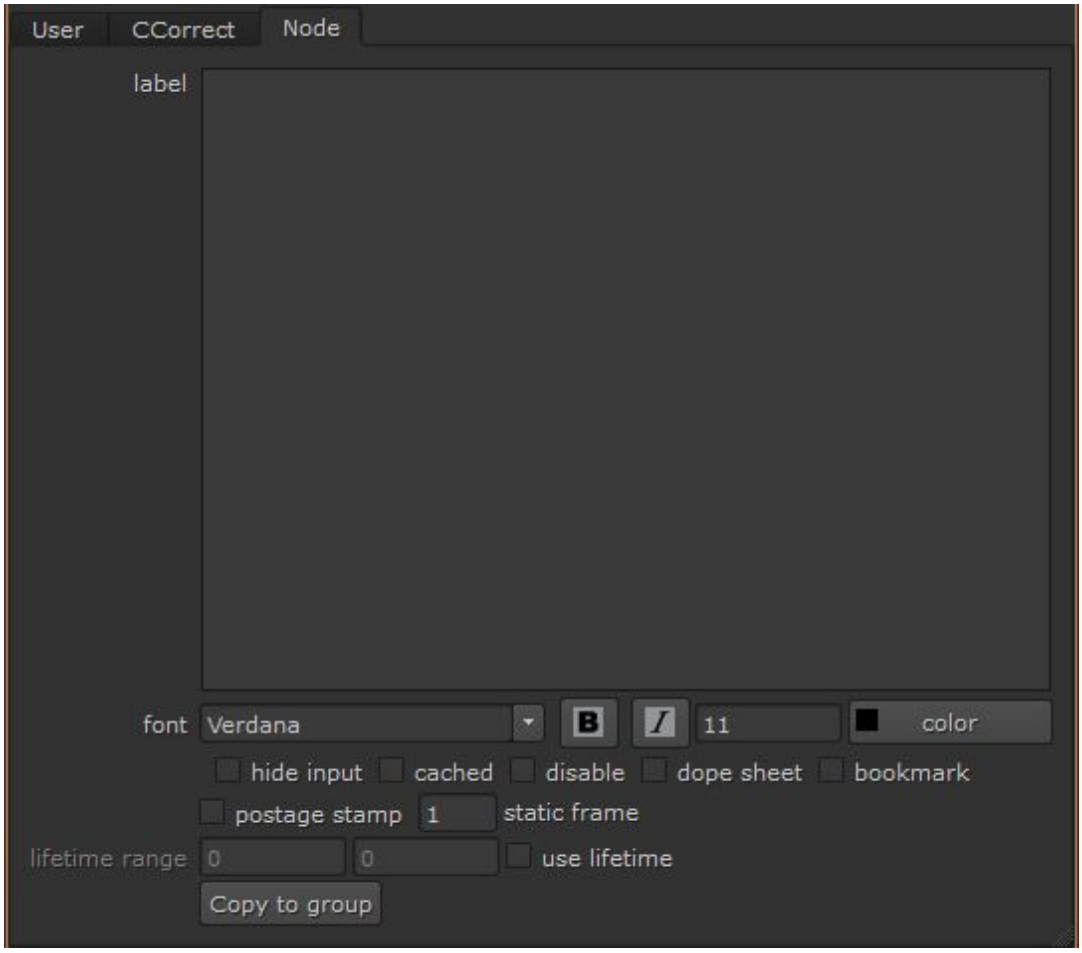

UI Properties

Label → Note displayed on the node in the node graph.

Font → Font used for the node in the node graph.

Hide Input → If the node isn't selected, the inputs into that node are hidden in the node graph.

Dope Sheet → For keyframes to show up in the dopesheet, you need to have the node properties open. If this option is turned on then the keyframes for this node will always show up in the dopesheet (even if the node's properties aren't open).

Bookmark → If selected, the node will be bookmarked such that you can go into the node graph, hit J, and it'll show up as one of the nodes you can quickly navigate to…

Postage Stamp → If selected, the node will save the frame specified in the textbox directly to the right of this option and display it as a thumbnail on the node (in the

node graph)…

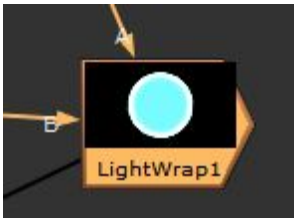

Backend Properties

Cached → Caches the output of this node such that if something changes downstream, the time to recompute the tree will be shorter.

Disable → Disables the node in the node graph -- like the pass through flag in Houdini.

Use Lifetime → If selected, the lifetime range defines the frames for which this node is active. This is like Disable except that it it only disables the node if it's outside the specified range.

Preset Properties

Presets are properties that have been saved for some specific node type. These preset properties can be loaded up again for other nodes of the same type.

Saving

To save your current node's properties as a preset, click the presets button in the toolbar () and choose "Save as Preset"...

A dialog box will pop up asking you what you want to name your preset. Fill in a name and hit Create…

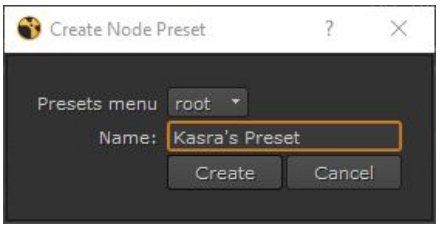

Loading

You can apply your presets to other nodes of the same type by clicking the presets button in the toolbar ( ) and choosing the preset to apply…

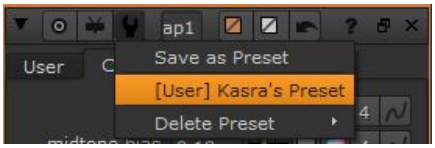

> **NOTE**: When you do this, it overrides ALL properties, not just the ones that were changed from the default when the presets were created.

Deleting

You can delete a preset by clicking the presets button in the toolbar ( ) and going to Delete Preset → <preset name>…

Dope Sheet Panel

The dope sheet panel is exactly like the dope sheet in almost every other 3D tool. There's almost nothing new here...

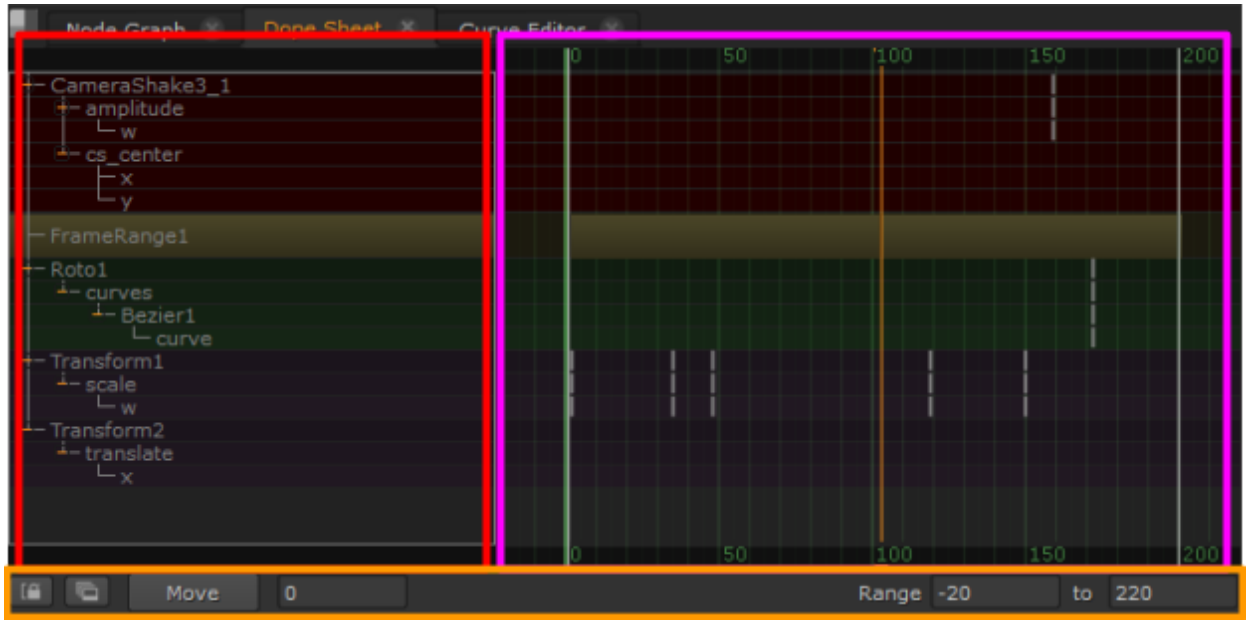

NOTE: Remember that for keyed properties to show up in the dope sheet, the node which the property sits in has to either…
* have that node's properties open
* have that node's dopesheet property turned on (under the Node tab)...

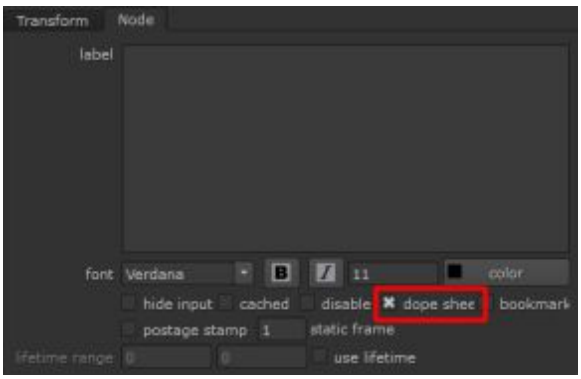

hotkey to force dopesheet property on is Alt+D (mouse must be hovered in node graph)

Hierarchy View The nodes and properties for which keyframes exist.

Keyframes The keyframes and scrubber...
* small white ticks are keyframes
* long white lines (at 1 and 200) define the start/end frame for shot
* long orange line is the scrubber (synced with scrubber in viewports and curve editor).

NOTE: Want to change the start/end for the shot? Hover your mouse over the node graph and hit S to open the project settings in the

properties pane.

Time and Read nodes are displayed different…
* time nodes (e.g. TimeOffset or TimeClip) are yellow blocks

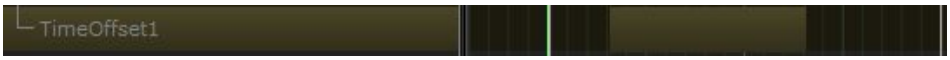

* read nodes (e.g. Read) are gray blocks

NOTE: The hierarchy of stuff in the dopesheet seems to context sensitive to the current Viewer. That is, the contents of the dopesheet change depending on if/where a Viewer node is connected. Here's the graph…

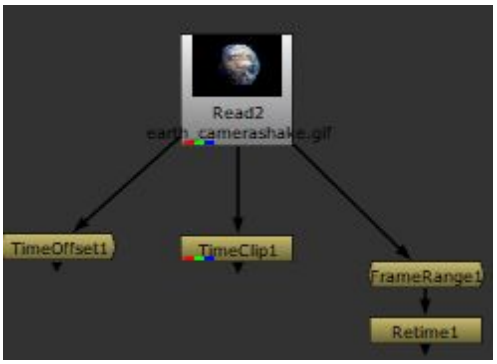

Here's what it looks like normally (no Viewer connected)...

Here's what it looks like with a Viewer connected to the TimeOffset…

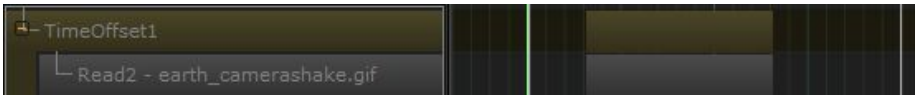

Here's what it looks like with a Viewer connected to the TimeClip…

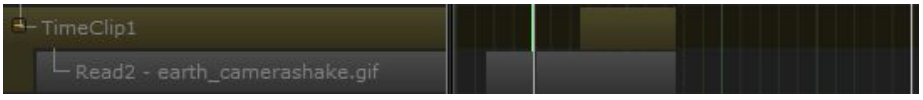

Here's what it looks like if you move the TimeClip such that it's after the TimeOffset, then hook the output up to a Viewer...

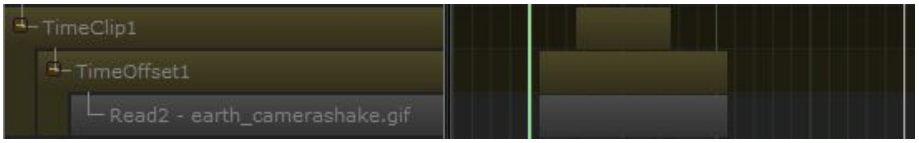

Settings Bar Dope sheet settings/options.

* first toggle () will keep the frame ranges synced between the dope sheet and curve editor.

* second toggle () will make it so that all Read nodes show up in the dope sheet.

* move button (

) will move keyframes you've selected in the dope sheet by the offset specified in the textbox. Note that this doesn't have to be a whole number.

* frame range fields (

will change the frame range that the dope sheet shows.

How to navigate the dope sheet...

- F → frame up selected keyframes, or everything if no keyframes are selected
- Alt+LMB → pan graph by moving mouse
- Alt+MMB → zoom in/out by panning mouse left/right (fine zoom)
- Mouse wheel → zoom in and out (coarse zoom)

Basic Keyframe Operations

Remember that keyframes are represented in the dopesheet as small white ticks…

Select Keyframes

You can select keyframes in the dope sheet just like you do in the node graph. That is, you can either select them individually or you can use a marquee select....

If you're using marquee select, you can hold Shift when selecting to add to and remove from your current selection.

NOTE: This won't work with individual selections, it'll only work with marquee selection.

Move Keyframes

Once selected, you can move keyframes by either…

- LMB click-and-dragging to a new location

- using the Move button in the settings bar at the bottom

NOTE: By default, if you use the mouse to move you're going to be moving in whole number increments. But, if you hold Shift while you move it will fractionally increment…

Scale Keyframes

If you've selected keyframes in multiple different frames, you can expand/contract your selection by dragging the edges of the box surrounding your selection. The cursor will turn into a horizontal resize cursor (⇔) when you reach the edges…

Basic Clip/Time Operations

Remember that clips and timing nodes are represented in the dope sheet as solid gray/yellow blocks…

Move Clips/Time Nodes

NOTE: Remember clips and time nodes show up as solid blocks (clips as gray blocks and time nodes as solid blocks).

When you load up either in the dope sheet, you can move it around just like you would keyframes: click-and-drag it to some new position…

However, you probably shouldn't do this directly in the dope sheet. Instead, you should connect your clip to a TimeOffset node and do any adjustments to the offset in there.

> **NOTE**: The main reason you want to do this is because it's non-destructive. Messing with Read nodes directly is going to cause havoc downstream for other people that use your comp.

Trim Clips/Time Nodes

> **NOTE**: Remember clips and time nodes show up as solid blocks (clips as gray blocks and time nodes as solid blocks).

When you load up either in the dope sheet, you can trim it just like you would when you resize keyframes: click-and-drag the edge to some new position…

> **NOTE**: Look at the frame's frame range property to figure out what you want it to do for the frames that are being trimmed off…

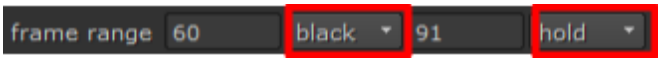

However, you probably shouldn't do this directly in the dope sheet. Instead, you should connect your clip to a TimeClip node and do any adjustments to the values in there.

> **NOTE**: The main reason you want to do this is because it's non-destructive. Messing with Read nodes directly is going to cause havoc downstream for other people that use your comp.

If you trim a clip like this, you can hover over the bottom of the clip and you'll get a horizontal-resize type icon. This is the slip icon, and you can click-and-drag to slip which frame the clip starts on vs which frame it ends on (it'll still be clipped by the same amount)…

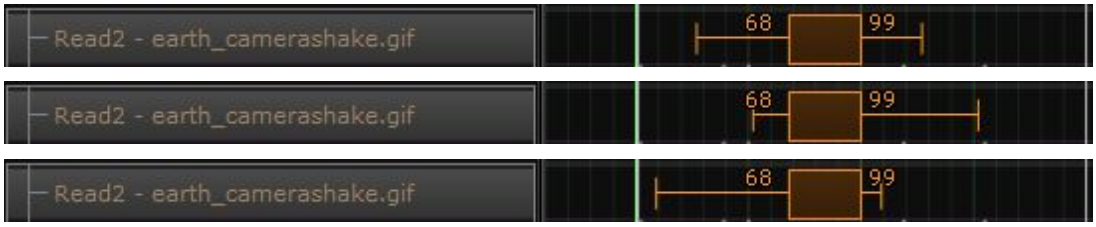

Show Audio

If you have audio loaded into your comp (via a Read/AudioRead node), you sometimes want to get that audio showing the dope sheet. To do this, right-click anywhere in the dope sheet and goto View → Audio → Draw Style → Behind OR Below.

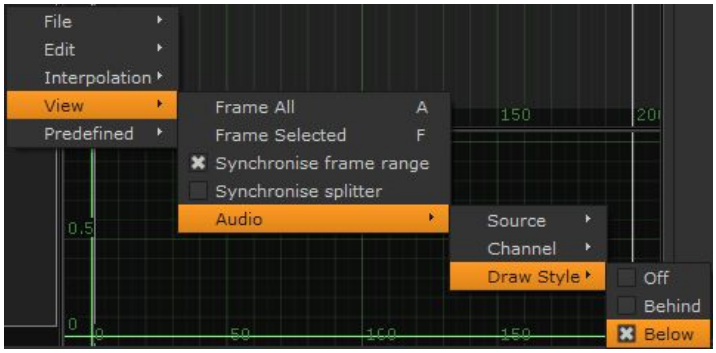

If you choose…

- below, a bottom panel will open and show the waveform.
- behind, the waveform will show behind your dopesheet.

  NOTE: In certain cases, the frame ranges will go out of sync with the waveform pane if you choose Below. If this happens, turn off the audio view and turn it back on.

  NOTE: Audio doesn't seem to work in non-commerical mode. Here's how it should look if you choose Below (from the lesson)...

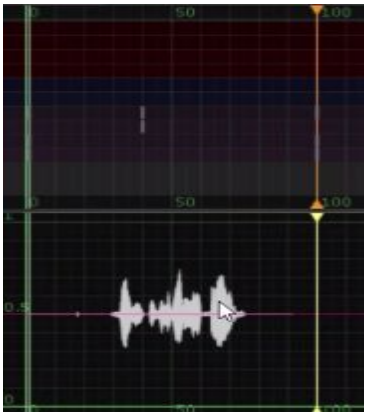

Duplicate Entries

In certain cases, the same node may show up more than once in your dope sheet. For example, imagine the following node graph…

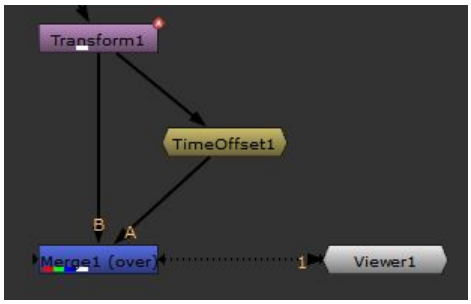

The transform has keyframes set. The output from the Transform is going down 2 paths and then being merged again. If you look at the dopesheet for these nodes you'll see that the Transform1 node shows up twice…

These 2 entries in the dopesheet are referencing the same node. If you move the keyframe for one, it'll automatically move the the keyframe for the other…

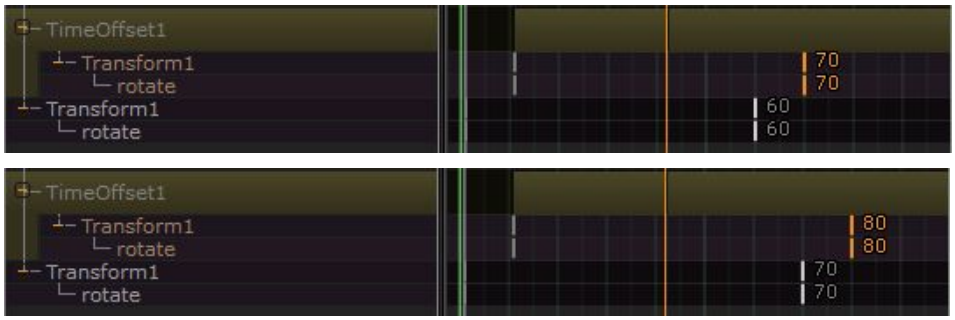

Curve Editor Panel

The curve editor panel is exactly like the curve editor in almost every other 3D tool. There's almost nothing new here...

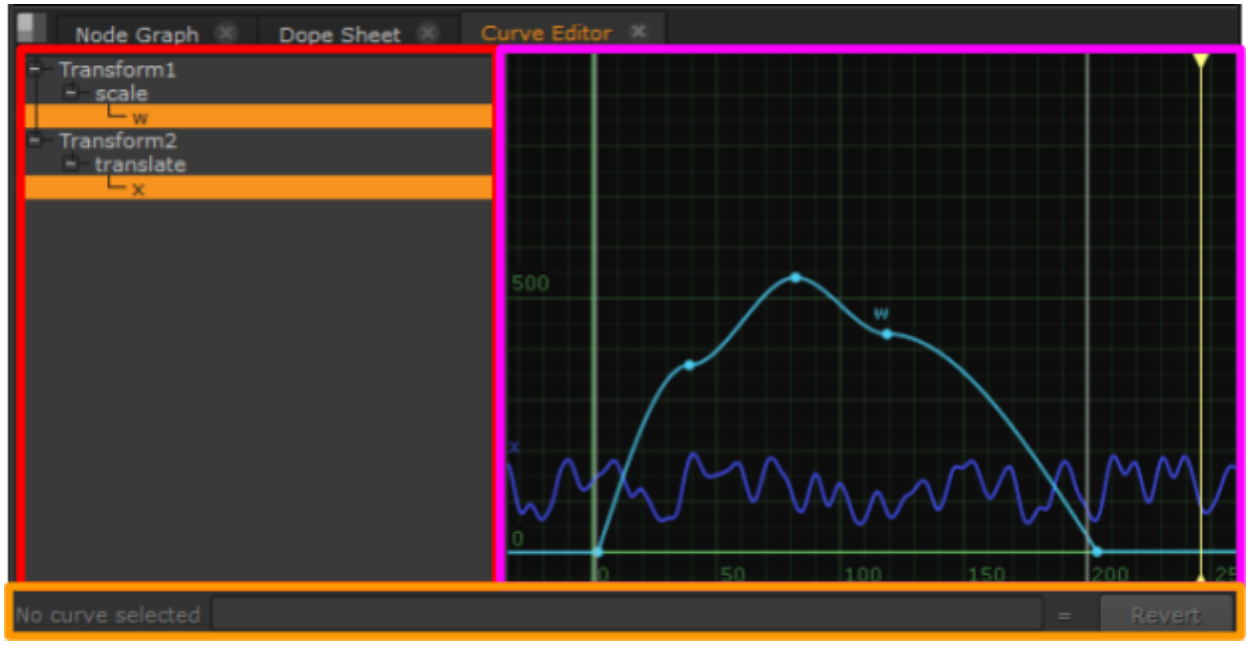

NOTE: Remember that for animated properties to show up in the curve editor, the node which the property sits in has to have that node's properties open. Unlike with the dope

sheet, there is no node property to force a node to show up in the curve editor.

NOTE: Remember that you can sync the frame range showing in the curve editor with the dopesheet. You have to click the lock icon in the bottom of the dope sheet to get this to happen (see the dope sheet section for more info). There is no lock icon in the curve editor.

Hierarchy View The properties for which keyframes/animations exist.

Curves The curves that show how the properties changes over time.

Be aware that…
1. for a curve to show up, the property must be selected in hierarchy view. In the example snapshot, we have 2 items selected. You can also select a parent item to show all curves under that item.

2. curves based on expressions cannot be modified visually in the curve editor. Note that the cyan curve (w) has 3 dots -- these 3 dots are the keyframes set on the curve that you can change/drag around. The blue curve is a curve that's based off of an expression/algorithm -- it has no dots because there are no actual keyframes.

Expression The expression bar shows you the expression for the curve. You have to physically click on the actual curve line to get the expression to show up here.

* curves based on keyframes will just show "curve"...

* curves based on expressions will show the acutal expression…

How to navigate the dope sheet...

- F → frame up selected keyframes, or everything if no keyframes/curves are selected
- A → frame everything, regardless of what's selected
- Alt+LMB → pan graph by moving mouse
- Alt+MMB → zoom in/out by panning mouse left/right (fine zoom)
- Mouse wheel → zoom in and out (coarse zoom)
- Ctrl+Alt+LMB → add a keyframe to the selected curve (you must physically click on the actual curve beforehand to select it)

Basic Keyframe Operations

Remember that keyframes are represented in the curve editor as dots on a curve…

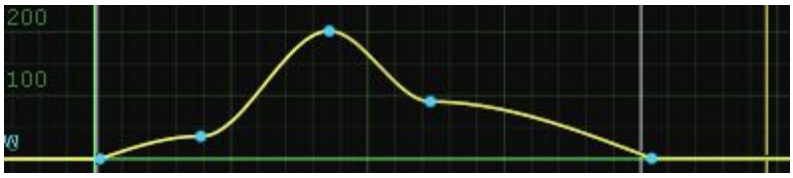

Select Keyframes

You can select keyframes in the curve editor just like you do in the node graph. That is, you can either select them individually or you can use a marquee select....

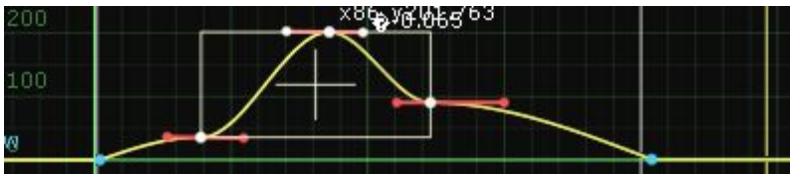

If you're using marquee select, you can hold Shift when selecting to add to and remove from your current selection.

> **NOTE**: Unlike the dope sheet, this WILL invert both individual selections and marquee selections.

Move Keyframes

Once selected, you can move keyframes by either…

- LMB click-and-dragging to a new location
- Right-clicking and going to Edit → Move

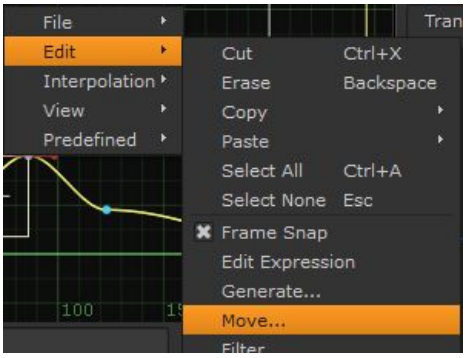

> **NOTE**: By default, if you use the mouse to move you're going to be moving in whole number increments. But, if you hold Shift while you move it will fractionally increment…

Scale Keyframes

If you've selected keyframes in multiple different frames, you can expand/contract your selection by dragging the edges of the box surrounding your selection. The cursor will turn into a horizontal resize cursor (⇔) when you reach the edges…

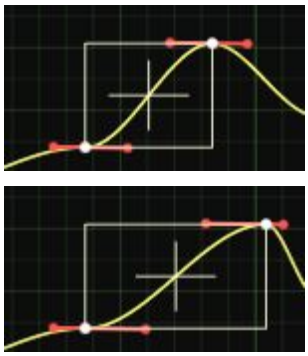

Add Keyframes

There are multiple ways to add keyframes...

Single Keyframe

To add individual keyframes, first select the physical curve (the actual curve line). Then, Ctrl+Alt+LMB click will add a keyframe at the point you clicked…

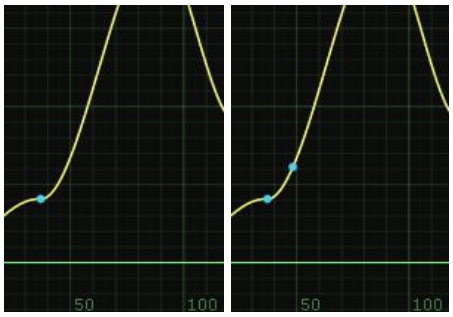

Copy Keyframes

To copy keyframes, select them with Ctrl+C, then move the scrubber to where you want to paste them, then hit Ctrl+V…

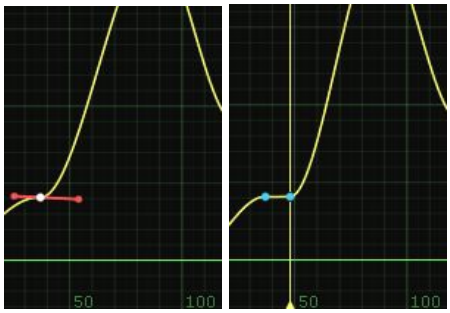

Generate Keyframes

To generate a group of keyframes, first show the curves that you want to add keyframes to -- this operation will add keyframes to all visible curves. Then, right-click and choose Edit → Generate. A dialog box will show up. Fill it in with the appropriate values (e.g. generates keyframes starting at frame 20 and ending at frame 100, in 30 frame increments)...

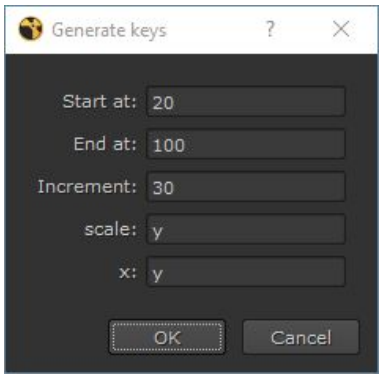

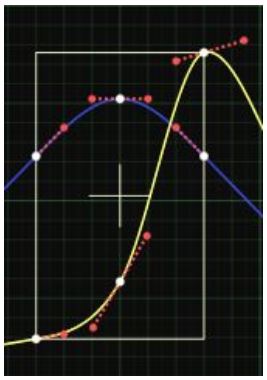

Draw Keyframes

To draw keyframes, first select the physical curve (the actual curve line). Then, hold Alt+Ctrl+Shift+LMB click-and-drag to draw out the curve…

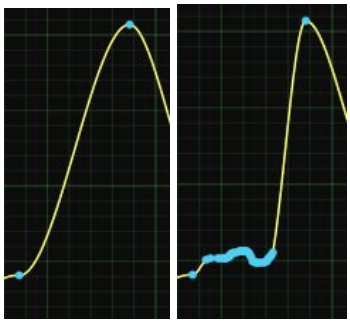

Change Tangents

You can change the tangent for a keyframe by selecting it and moving the red handles…

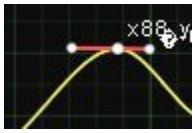

You can also change the type of tangent used by the keyframe by right-clicking and selection one of the options under Interpolation…

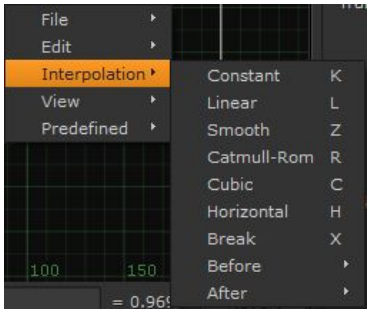

All the standard tangent options are available.

Basic Curve Manipulation Operations

These are “predefined” operations that change curve’s animation. When you apply a predefined operation to a curve, the original curve will show up as a dotted line behind the the new predefined curve...

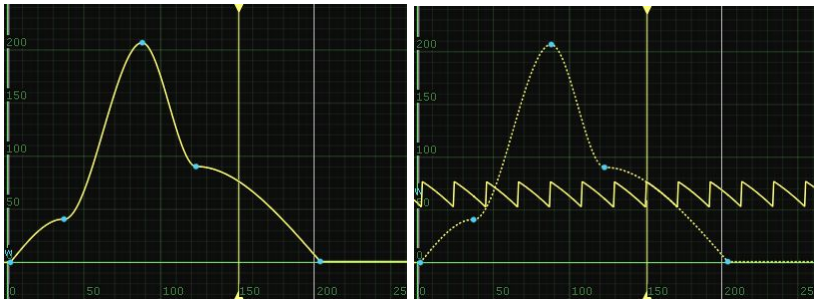

You can still manipulate the original curve (note the keyframes are still there) -- the predefined curve will update to match.

Loop Curve

To loop a part of the curve, select it then right-click anywhere in the curve editor and goto Predefined → Loop…

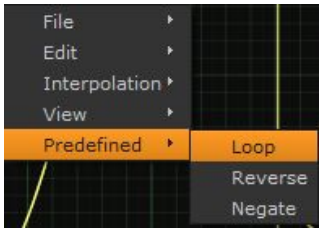

A dialog box will show up asking you which portion of the curve you want to loop…

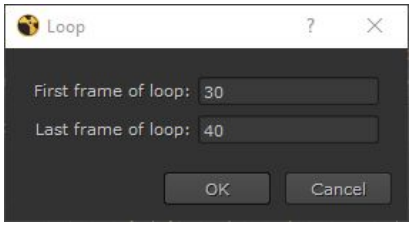

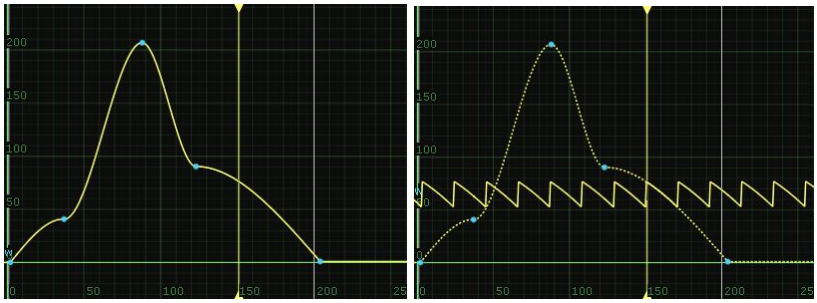

Reverse Curve (Flip over Y-axis)

To flip a curve over the Y-axis, select it then right-click anywhere in the curve editor and goto Predefined → Reverse…

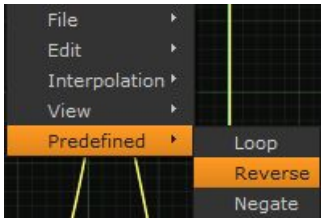

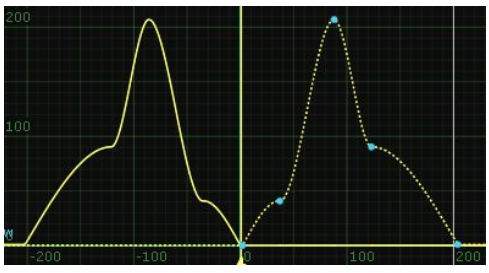

This may look like a useless operation because the curve is being reversed over the y-axis -- it's in the playback frame range for the reversed curve is in the negatives (e.g. -200 to -1).

However, if you select all the keyframes on the original curve shift them left, it'll shift the reversed curve right.

Select all keyframes, then move them into the negative space. The reversed curve should start moving into positive space as you move the original curve into negative space…

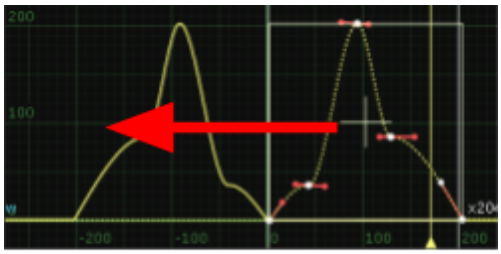

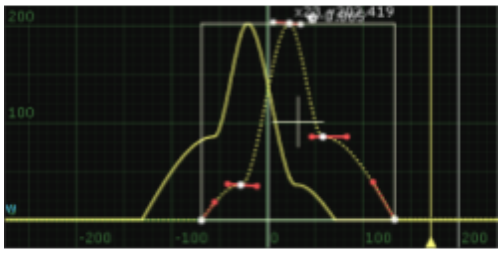

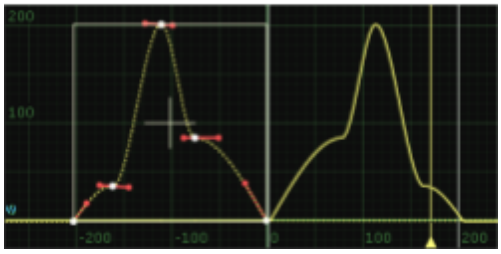

NOTE: If you notice that you can't shift past a certain point, it's probably because you don't have ALL keyframes selected. Make sure you have ALL keyframes selected. Ctrl+A if you're unsure.

Negate Curve (Flip over X-axis)

To flip a curve over the X-axis, select it then right-click anywhere in the curve editor and goto Predefined → Negate…

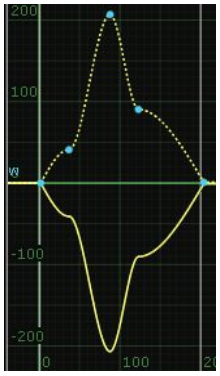

Select the relevant keyframes and move them around to adjust the negated curve however you want...

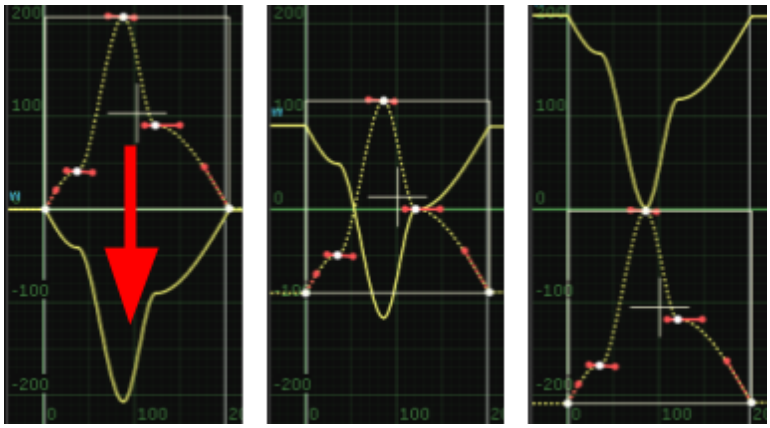

Curve Expressions

Remember that you have the ability to animate via expressions instead of keyframes.

At the bottom of the curve editor, you'll see the expression bar…

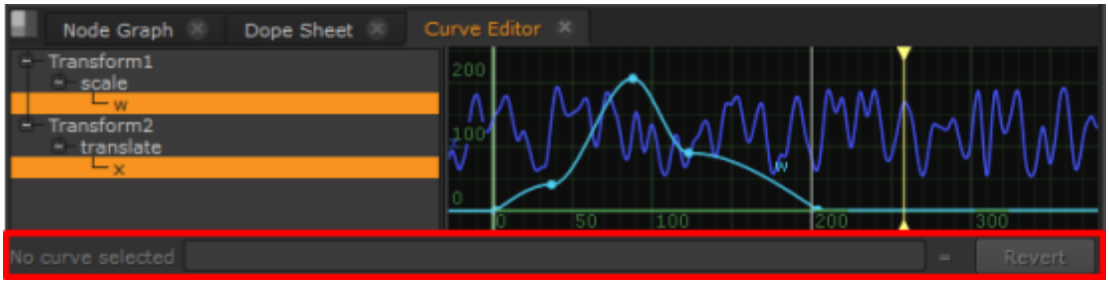

The expression bar shows you the expression for the curve you have physically selected...

- curves based on keyframes will just show "curve"

  

- curves based on expressions will show the actual expression

  

NOTE: “Physically selected” means physically clicking on the line for the curve.

You can change the expressions to fit whatever you need. For example, if you have keyframe based animation that you want to lower the amplitude/intensity of, you can change “curve” to “curve/2”...

NOTE: Notice that, just like with negating/looping/reversing a curve, the original curve shows up as a dotted line and the new less intense curve shows up as a solid line. The keyframes are still on the original curve. Updating those keyframes will also update the new curve.

Show Audio

If you have audio loaded into your comp (via a Read/AudioRead node), you sometimes want to get that audio showing the curve editor. To do this, right-click anywhere in the curve editor and goto View → Audio → Draw Style → Behind OR Below.

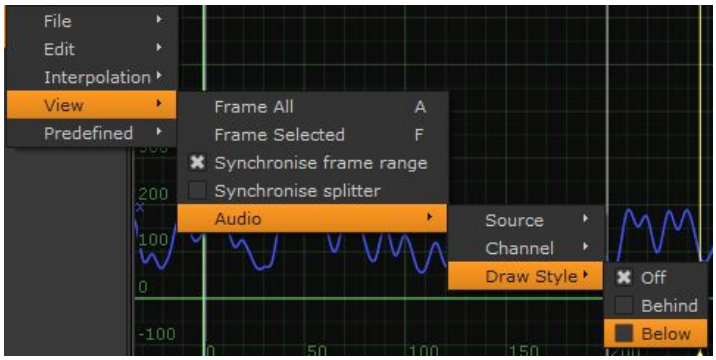

If you choose…

- below, a bottom panel will open and show the waveform.
- behind, the waveform will show behind your dopesheet.

NOTE: In certain cases, the frame ranges will go out of sync with the waveform pane if you choose Below. If this happens, turn off the audio view and turn it back on.

NOTE: Audio doesn't seem to work in non-commerical mode. Here's how it should look if you choose Below (from the lesson)...

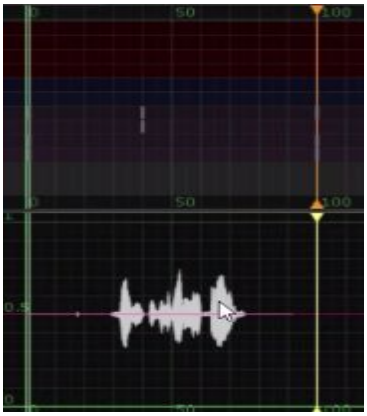

Common Nodes

Input/Output

Read → read image/image sequence/audio/movie from disk
ReadAudio → like a Read node but only for audio? (doesn't work in non-commercial)
ReadGeo → like a Read node but for geometry?
Write → write comp out to disk
WriteGeo → like a Write node by for geometry? (doesn't work in non-commercial)
Viewer → creates a viewer and shows the output of whatever node it's connected to

Animation Timing

TimeOffset → shift an input by a certain amount of time
TimeClip → clip an input to fit a certain range of time
Retime → change playback speed

Graphics

2D

Blur → apply a blur
EdgeBlur → apply a blur to the edges (alpha required?)
Roto → apply a mask to the operation
ColorCorrect → adjust color settings (e.g. adjust gamma)
Resample → resize your output

3D

Sphere → a sphere
Cube → a cube
Light → a light (point, directional, or spot -- can be configured)
Camera → a camera
Scene → “scene graph” node where all your 3D nodes eventually pipe into
ScanlineRender → renders your 3D input as a 2D image

Nuke 11 Channels and Layers

Introduction

2D image processing in Nuke revolves around the concept of layers and channels.

Channels store floating-point numbers that represent the data that Nuke is processing. This is typically normal color information (e.g. RGBA channels), but can also be data such as depth or velocity at a certain point in the image.

Layers are used to organize channels into logical groups. For example…
- color data (RGBA) would go into a layer of its own.
- specular/diffuse reflection data would go into a layer of its own.
- subsurface scattering would go into a layer of its own.
- depth data would go into a layer of its own.
- etc..

There is no limit to how many channels a layer can hold, but Nuke's nodes seem to limit you to handling 4 channels per layer.

Node Badges

Nodes in Nuke will display a badges at the bottom to tell the user which channels that node is outputting out. For example...

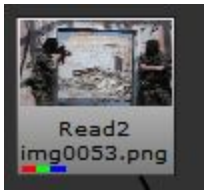

In the example above, there's a red, green, and blue badge on the lower-left. This indicates that the node is outputting a red, green, and blue channel. If it were outputting an alpha channel, it would also contain a white badge. For example…

If it were outputting non-color channels (e.g. depth), the node would also have a bunch of other badges following the RGBA badges. For example...

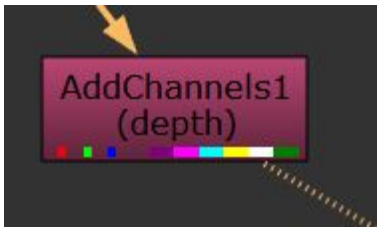

> **NOTE**: I have no idea what all badges past RGBA mean or why some of them are only half-filled. The lesson said that dark green is the inclusion of non-color data channels but I don't think this is true? The node in the above example adds depth, motion, disparity, and a custom layer called fancylayer with 8 channels.

Channel Manipulation Nodes

Nuke provides a set of nodes for manipulating channels/layers in your comp…

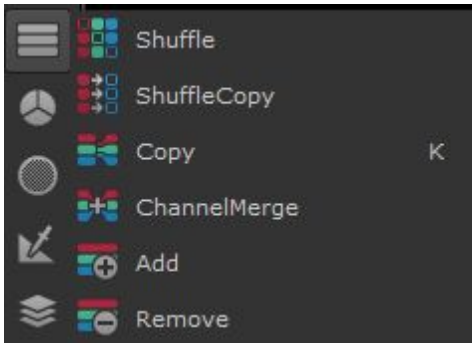

To add any of above nodes, you can use any of 3 basic methods described in the main Nuke document…

- Tab menu → hit tab in the graph and type in the name of the node…

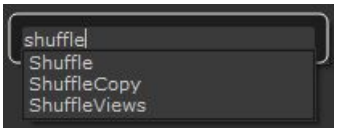

- Context menu → right-click in the graph and goto Channels…

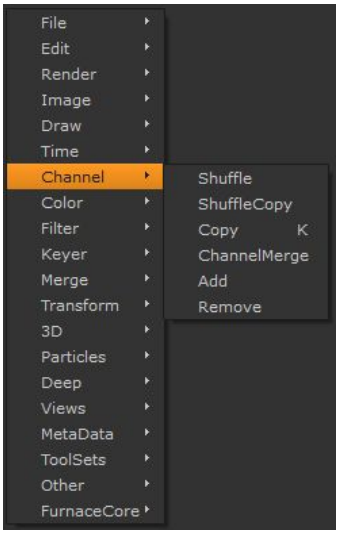

- Toolbar → click the stack…

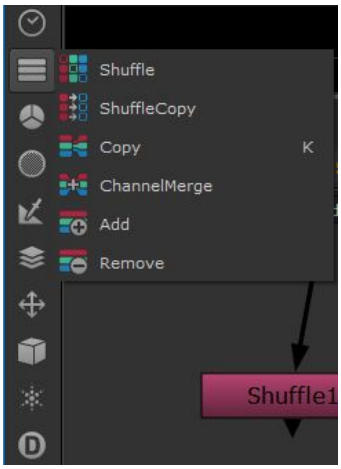

If you already have a node selected when you add a node, it'll get added to the output of the selected node. If the selected node already has an output, it'll get added between the nodes.

Add Channel

Add Channel node adds layers/channels.

To add a new channel, goto the properties panel and select the channel from the drop-down + adjust the color slider to whatever you want the initial value to be. The new channel will get filled with that value…

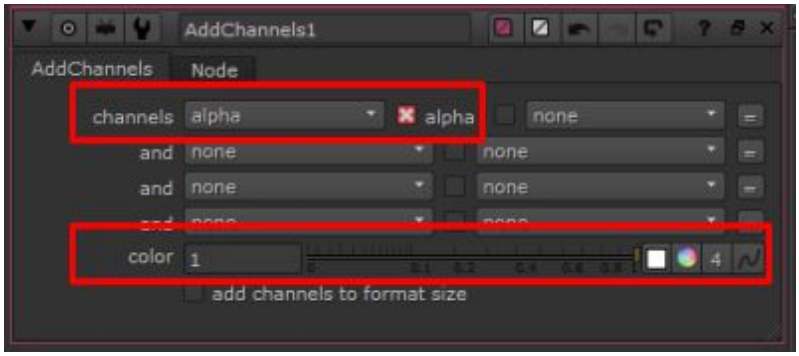

NOTE: If what you want to add isn't listed, there's a add option in the dropdown that'll let you define it. Or, you can go under the other layers submenu and choose from there. The lesson said that if it's under other layers, the layer isn't present (if it existed it would have been listed in the main dropdown?)...

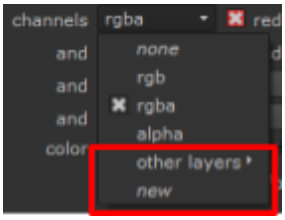

You can add multiple layers/channels, but the design of the properties panel is slightly confusing. Each channel to add is a row…

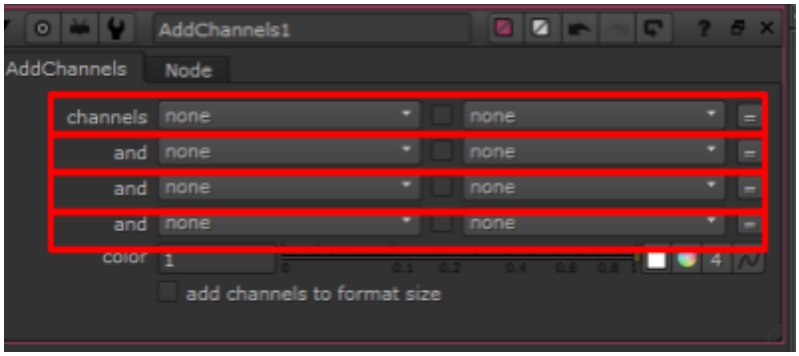

The first dropdown is where you select actually select the layer you want to add. When you select the layer, the first 3 channels of that layer show up as 3 checkboxes, and the 4th checkbox is to enable the 2nd drop-down where you can select a 4th channel from the layer. For example…

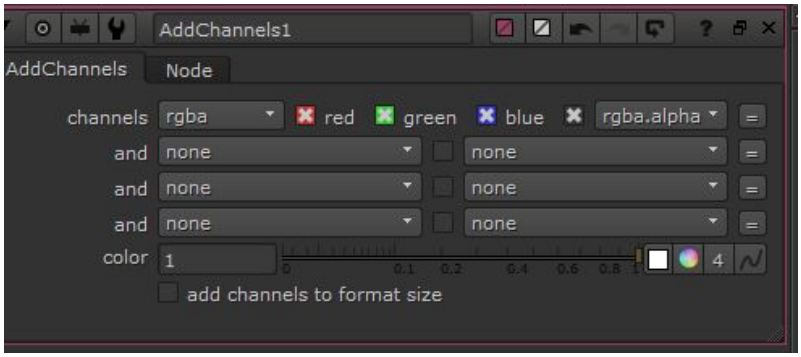

In the above example, we're adding in all channels in the rgba layer (red, green, blue alpha). If we deselect red, it won't add in the red channel. If we deselect the last checkbox (before the dropdown) it won't add in the alpha channel.

> **NOTE**: If we had a layer that had > 4 layers, we can select the 4th channel to add for the layer from this dropdown (I think). What happens if we want to add all the channels for the layer? My guess is that we would have to add the layer multiple times…

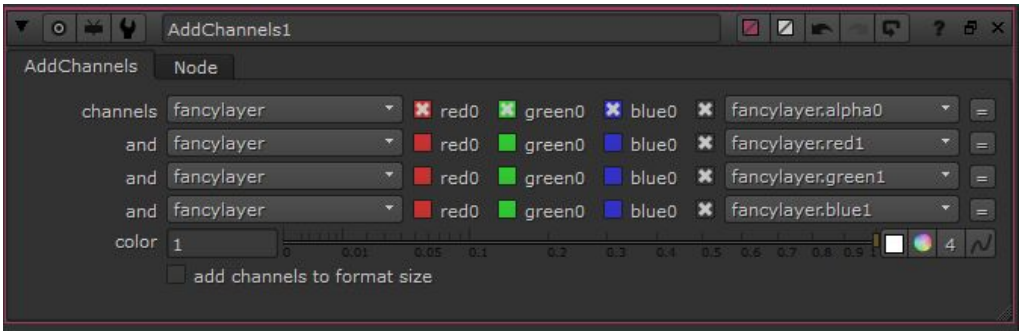

Remove Channel

Remove Channel node removes layers/channels.

This node's properties are very similar to the Add Channel node's properties, except that you're removing channels/layers instead of adding them…

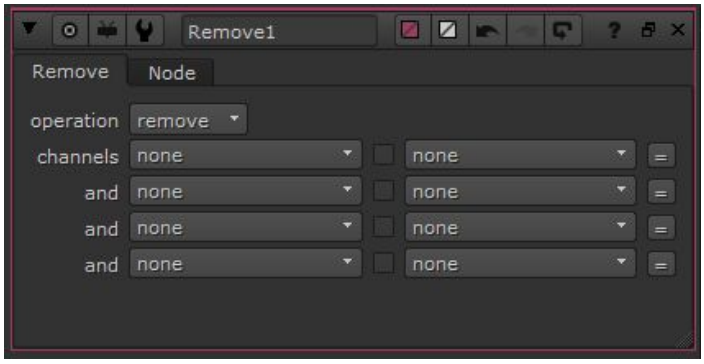

NOTE: To see how the layer/channel selection works, read the Add Channel section.

The one exception is the operation dropdown located at the top…

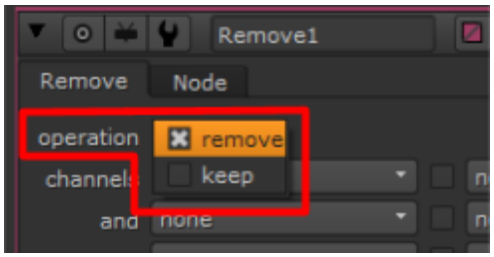

Choosing …

- remove → will delete the channels you've selected.
- keep → will delete all channels EXCEPT FOR the ones you've selected.

Copy

Copy node lets you copy layers/channels from one input to another.

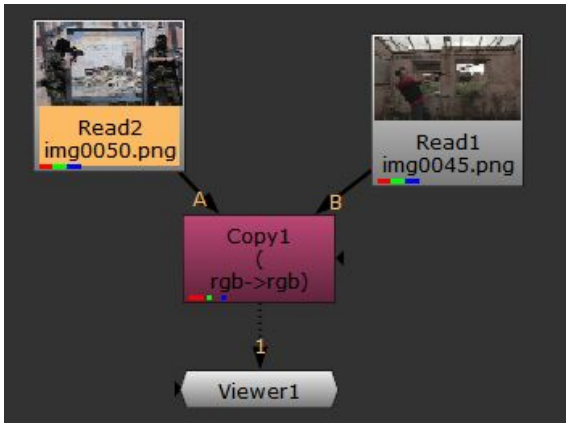

To copy a channel, goto the properties panel and select the channel to copy from in A to the channel to copy to in B. You can do this for up to 4 channels….

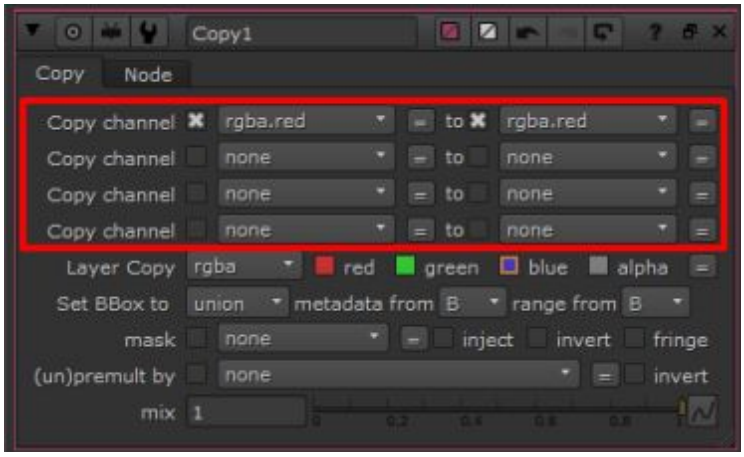

NOTE: Note that you can choose the destination channel as well. That means you can map a channel from a A to a different channel from B (e.g. A.rgba.red to B.rgba.blue).

Or, if you want, you can select the layer to copy and use the checkboxes provided to select which layers to copy…

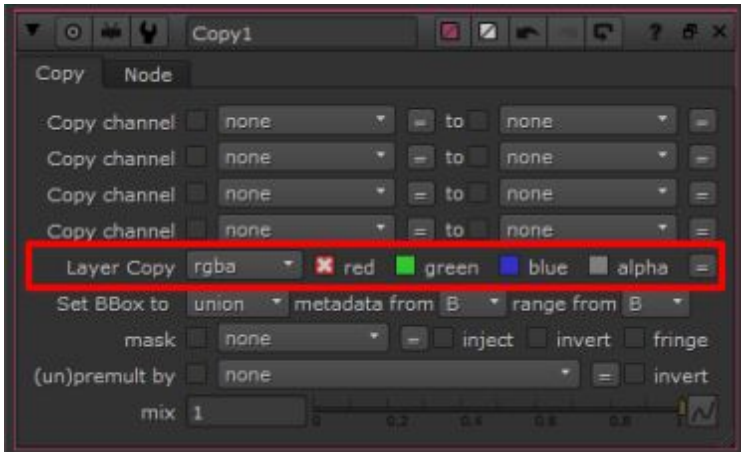

Shuffle Copy

Shuffle Copy is like a mix of the Add node and Copy node, but with a convoluted properties interface that doesn't really make sense.

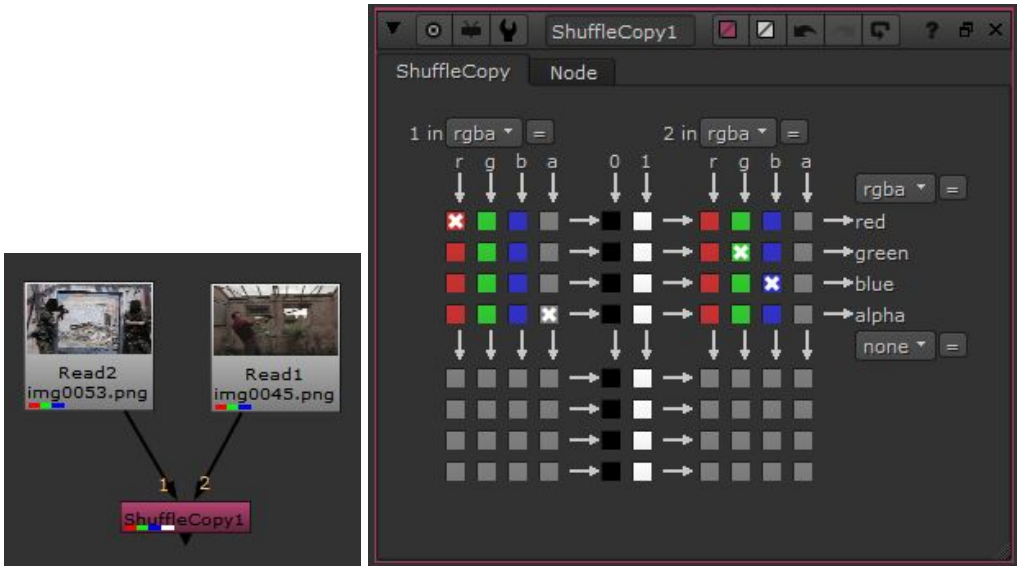

The columns represent inputs and the rows represent outputs. So for example, if I wanted my final image to have the R channel from input 1 and a B and G channel from input 2, it would look like this…

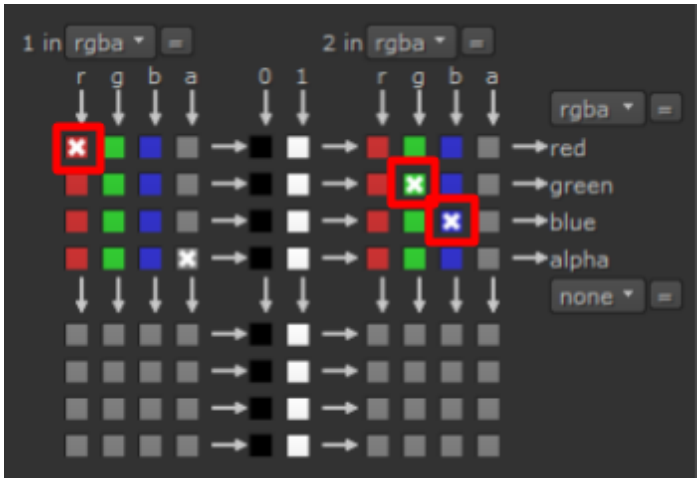

You can also force an output to be set to all 0s or 1s by selecting the appropriate constant 0 or 1. In the example below, we're forcing the G channel to 0…

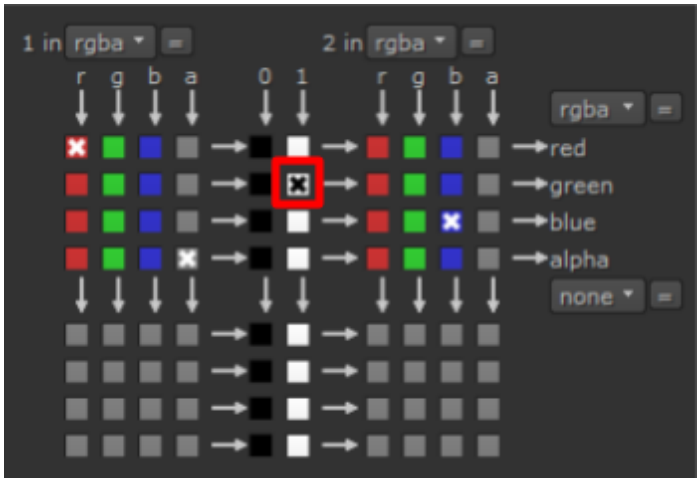

You can choose 2 layers to output to. For example, if you wanted output the R channel from input 1 to the depth channel…

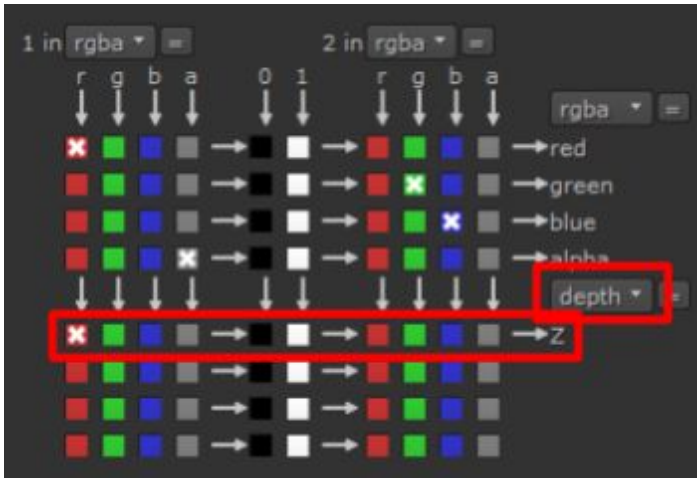

> **NOTE**: It looks like this entire interface breaks done if you select a input or output layer with > 4 inputs. A lot of the functionality here is a duplicate of Add node and Copy node -- maybe use those instead of this.

Shuffle

Shuffle is like the Shuffle Copy node, except that it only takes in 1 input.

What makes Shuffle even more confusing than Shuffle Copy is that the properties are exactly the same as Shuffle Copy. That is, it makes it seem like there are 2 inputs when there's only actually 1...

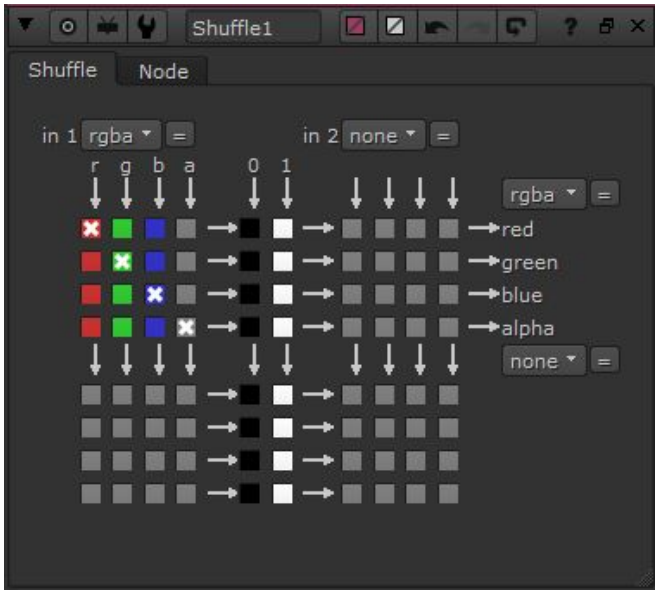

My understanding is that this works like the Shuffle Copy node, but both of the input layers come from the single input (even though the second input layer dropdown is labeled as "in 2").

See the Shuffle Copy node for more information on how to use the properties panel.

> **NOTE**: It looks like this entire interface breaks done if you select a input or output layer with > 4 inputs. A lot of the functionality here is a duplicate of Add node and Copy node -- maybe use those instead of this.

Channel Merge

Channel Merge is like a Merge node except that it…

1. only operates on channels.

2. provides a simpler set of operations.

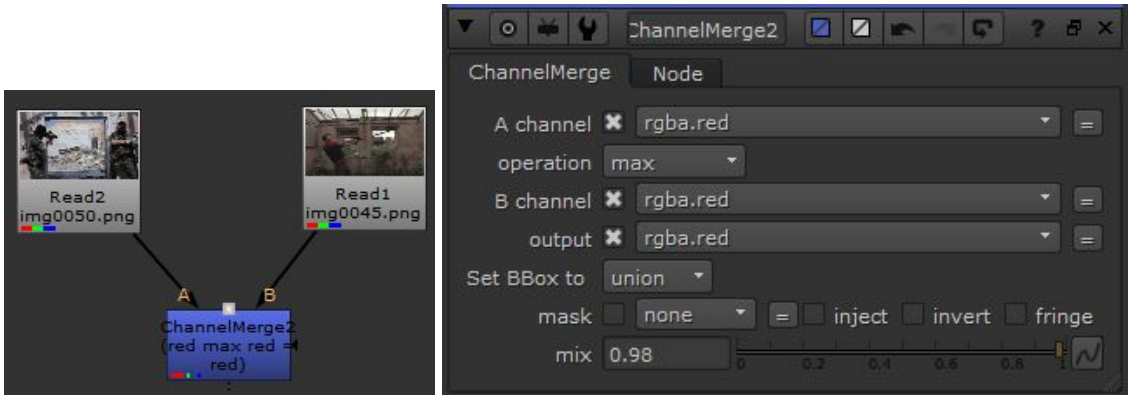

The most important parts of the properties panel..

- A channel → channel to use from the A input.
- B channel → channel to use from the B input.
- operation → math operation to perform.
- output → channel to output the result to.

Below is a copy of the Foundry docs that talk about what each operation boils down to (math-wise). What they do conceptually do can be found here: https://help.thefoundry.co.uk/nuke/8.0/content/user_guide/merging/merge_operations.html.

- **absminus** abs(A-B) - how much the pixels differ.
- **b if not a** A?A:B - shows A wherever A exists; otherwise shows B.
- divide A/B, 0 if A<0 and B<0 - divides the values but stops two negative values from becoming a positive number.
- **from** (B-A) - subtracts A from B.
- **in** Ab - only shows the areas of image A that overlap with the alpha of B.
- **max** (max(A,B)) - selects the lighter of the two colors as the resulting color. Only areas darker than B are replaced, while areas lighter than B do not change.
- **min** (min(A,B)) - selects the darker of the two colors as the resulting color. Any parts that are lighter than B are substituted. Any parts of the image that are darker than B don't change.
- **minus** (A-B) - subtracts B from A.
- **multiply** (AB, A if A<0 and B<0) - multiplies A by B. The result is always darker. Blending with black gives black and with white returns the color unchanged.
- **out** A(1-b) - only shows the areas of image A that do not overlap with the alpha of B.
- **plus** A+B - the sum of the two colors. Increases brightness to lighten A and reflect B.
- **stencil** B(1-a) - this is the reverse of the out operation. Only shows the areas of image B that do not overlap with the alpha of A.
- **union** A+B-AB - shows both image A and B.
- **xor** A+B-2AB - shows both image A and B where the images do not overlap.

Channel Helper Nodes

Layer Contact Sheet

Layer Contact Sheet is like a normal Contact Sheet node, except that it displays thumbnails of all the layers in the image.

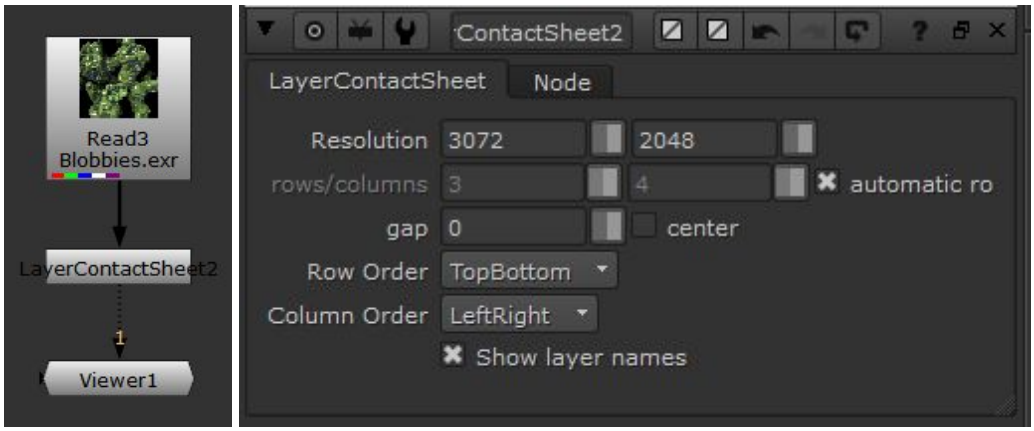

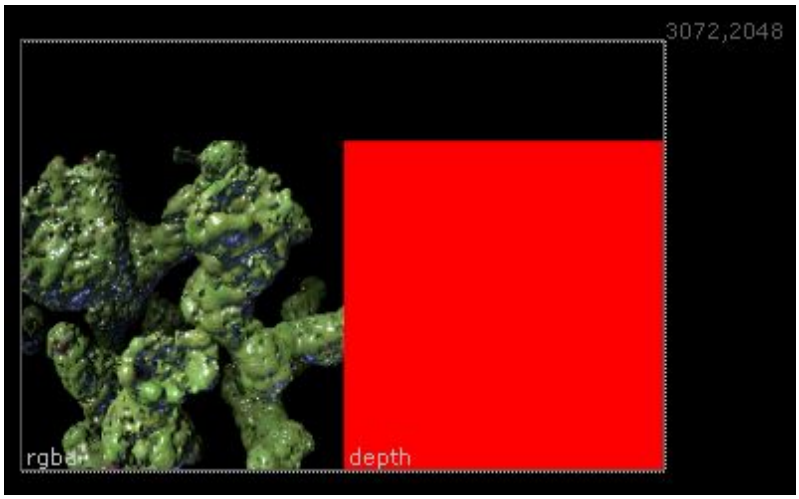

> **NOTE**: The property show layer names is checked on. This shows the name of the layer on the thumbnails being output.

The example has 2 layers: rgba and depth. The rgba layer which contains the image pixels and the depth buffer contains the z-buffer. The z-buffer for this image is set to all 1s. The layer contact sheet doesn't know how to display that, so it just translates depth.z to as the color red.

For non-color channels, they automatically get translated to RGBA when being displayed by layer contact sheet… The

- 1st channel of the layer gets output as red.
- 2nd channel of the layer get output as green.
- 3rd channel of the layer gets output to blue.
- 4th channel of the layer gets output to alpha.

  NOTE: What happens if layer has more than 4 channels? I don't know.

Nuke 11 Tracking

Introduction

Tracking refers to the concept of locking onto one or more features in an image sequence and following them as that image sequence progresses. Typical use-cases for tracking include…

- camera stabilization
- replacing a moving region in a scene (e.g. a screen for a phone or watch)
- mapping out the movement of an object in the scene

Features

The term feature refers to any static / semi-static point in the shot. The user typically starts out by picking these points and setting a few parameters. The tracking algorithm will take those points and parameters and try it's best to follow the point as the image sequence progresses.

> **NOTE**: Even with algorithmic tracking, the user may need to tweak the track manually. The track may get offset or lost due to quick movements that cause the portion being tracked to blur. These frames must be manually tracked by the artist.

What makes a good feature for algorithmic tracking? The feature should …

- be identifiable and static throughout the frames
- be well lit and remain so throughout the frames

- not drastically change between frames
- not enter a heavy shadow between frames
- have a high contrast difference between its foreground and background

In the following example, each of the cats 2 eyes are being as a feature. Each eye is easily identifiable and doesn't drastically change, and there's a good contrast between the iris (cream coloured) and pupils (black)...

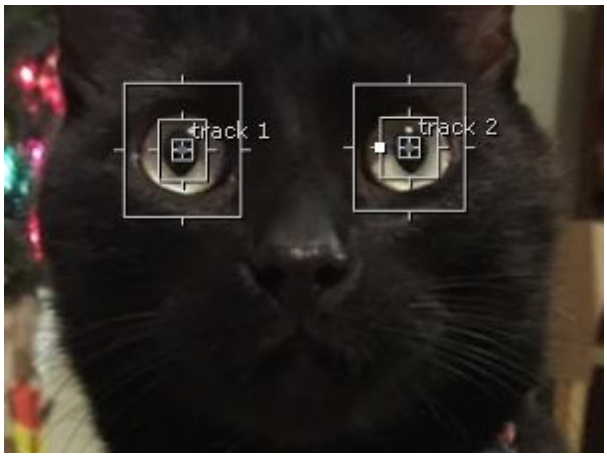

Transformations

The number of features defines what type of transformations can be tracked. If you're tracking only 1 feature, you're only going to get 2D position (translation) information.

If you're tracking 2 (or more) features, you can get 2D scale and rotation information as well. What essentially happens is that the 2 features form the endpoints of a line. The line endpoints are compared to the line endpoints from a previous/reference frame to determine how much rotation/scale to add (the line can twist or grow or shrink).

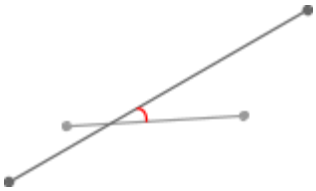

You can also use multiple tracks to get better information. For example, you can use multiple tracks for just translation. The result of the translation will be the average movement between the them.

> **NOTE**: Nuke recommends a 3-point track if you're planning on dealing with rotation and scale as well? But 2-point tracks will work as well.

Tracker Node

In Nuke, tracking is done using the Tracker node. Image data is fed into the Tracker node and the artist defines the features to track in the viewer. These tracked features can then be linked to by other nodes for use down the line (the tracker node also provides some common tracking functionality directly within the node).

Adding Node

To add a Tracker node, you can use any of 3 basic methods described in the main Nuke document…

- Tab menu → hit tab in the graph and type in Tracker…

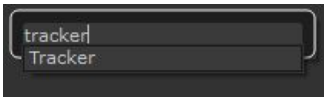

- Context menu → right-click in the graph and goto Transform → Tracker…

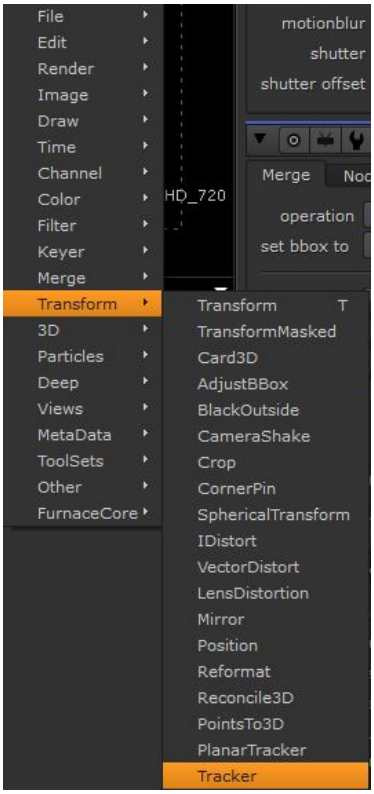

- Toolbar → click the cross-hatched arrows and select Tracker…

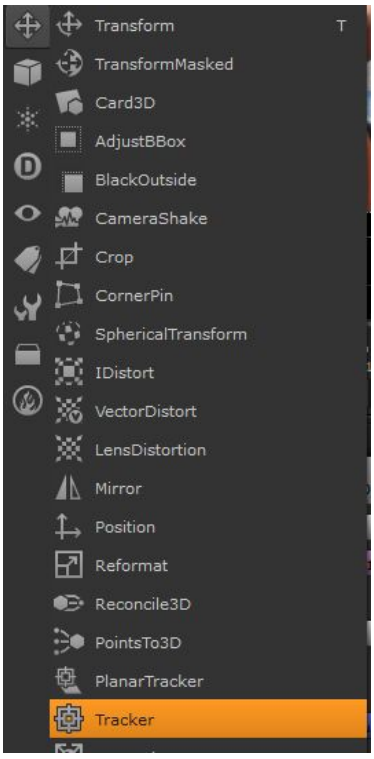

If you already have a node selected when you add the Tracker node, it'll get added to the output of the selected node. If the selected node already has an output, it'll get added between the nodes.

Tracking Features

To add track a feature in a tracker node, open up the properties, go to the Tracker tab, and hit the Add Track button…

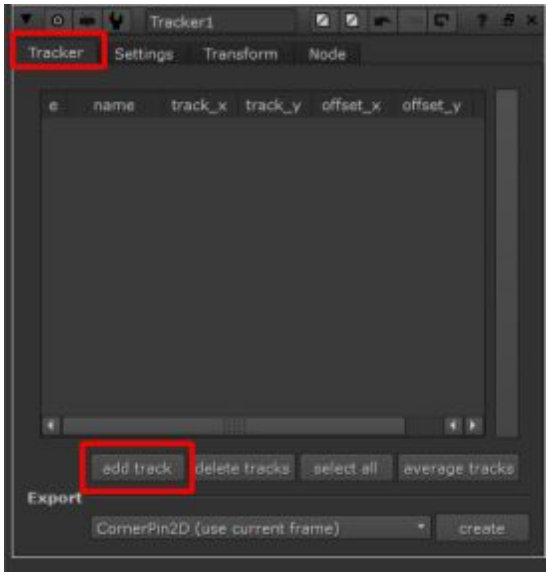

> **NOTE**: Forgot how to open the propertie for a node? Double-click it in the node graph. Remember that you can have the properties for multiple nodes open at once. If you have the properties for 2+ Track nodes open at once, things will get confusing because the tracking markers for all nodes will show up in the viewer. It's best to close down the properties of any other tracker nodes.

Once you've added the tracker, the marker for that tracker will show up in the middle of your viewer and in your node's properties...

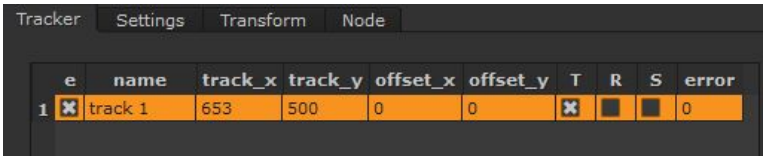

The marker is what defines, identifies, and follows the feature...

- anchor → part of the feature to fixate on
- pattern box → area encompassing the feature (what to search for)
- search box → area to search for where feature has moved to when moving frames

Scrub to your reference frame and move the marker by its anchor to the feature you want to track. Once placed, you can resize the pattern box by dragging it to encompass the feature you're tracking. Same thing with the search box -- you'll want it large enough so that as the feature moves between frames, it's still in the search box area.

> **NOTE**: You can ONLY MOVE THE MARKER USING THE ANCHOR. Mouse manipulation on the pattern/search box just resizes those boxes around the anchor.

Creating Tracks

There are 2 types of tracking: auto tracking and keyframe tracking.

Auto Tracking

> **NOTE**: After reading this section, check out the improving tracks section. It contains some extra notes on how to get better auto tracking.

In many simple cases, you can get away with auto tracking. Auto tracking is where Nuke will use the parameters you defined (search box / pattern box / etc..) to try to automatically identify and follow the feature between frames.

To do this, make sure that the marker is correctly placed/sized on the feature, then go to the toolbar in the viewer and select either the track to end or track to start buttons…

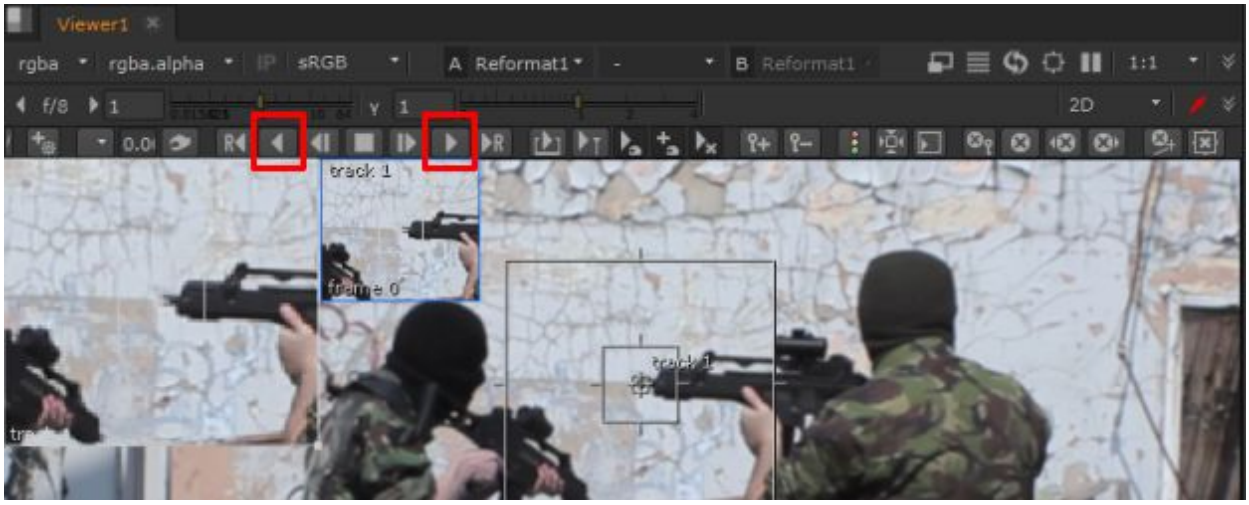

> **NOTE**: Don't see this toolbar in the viewer? Make sure the properties for your tracker node are open and the track is selected in the properties.

If you click the …

- forward button, it'll track the feature while moving frames forward.
- backward button, it'll track the feature while moving frames backward.

Once one of these buttons gets clicked, the autotrack will start tracking until it either hits the last frame or gets to a point where it can no longer identify the feature.

> **NOTE**: If you set your marker at a frame > 0, you'll likely want to use both the forward and backward buttons. Start from the frame you set the marker at and hit forward button, then go back to that frame and hit the backward button.

> **NOTE**: If you click the R buttons to the side of the forward/backward buttons, it'll let you specify a maximum number of frame to autotrack instead of tracking until the end.

In many cases, autotrack will continue tracking but track the wrong thing because…

- the object got obscured
- the search box wasn't big enough
- the pattern didn't stand out enough
- etc…

In the following example, the gun barrel is tracked properly into the 2nd frame, but by the 3rd frame it gets obscured by the soldier's body. The autotrack tries its best to identify the feature, but since the feature isn't there anymore it locks onto what it thinks is the next closest thing…

In the tracker toolbar, there's a button that looks like a traffic light. If you enable it, the points of your track will turn red/yellow/green depending on how bad/good the identification of the feature was. If we turn it on for the above example, we see that the 3rd frame is a solid red, which means that it's very unsure about the track…

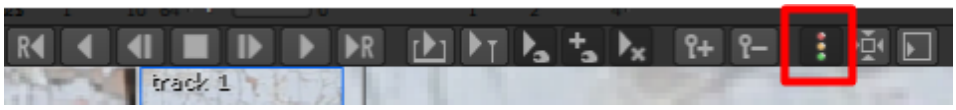

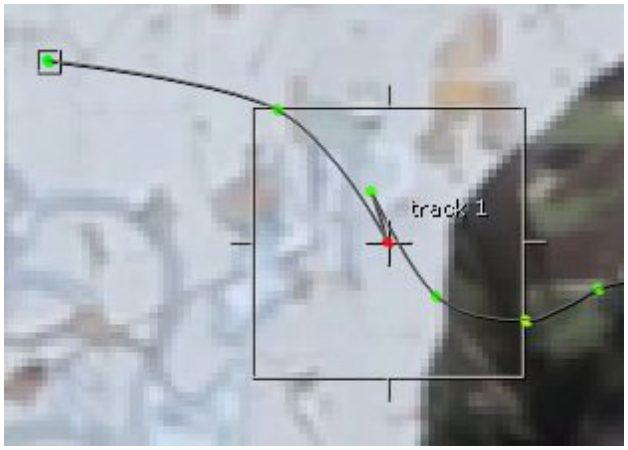

To fix this, we need to use keyframe tracking (discussed below), but essentially everything past the 2nd frame is a bad and needs to be discarded. To do this, navigate to the last good frame then go to the toolbar and hit clear forward button…

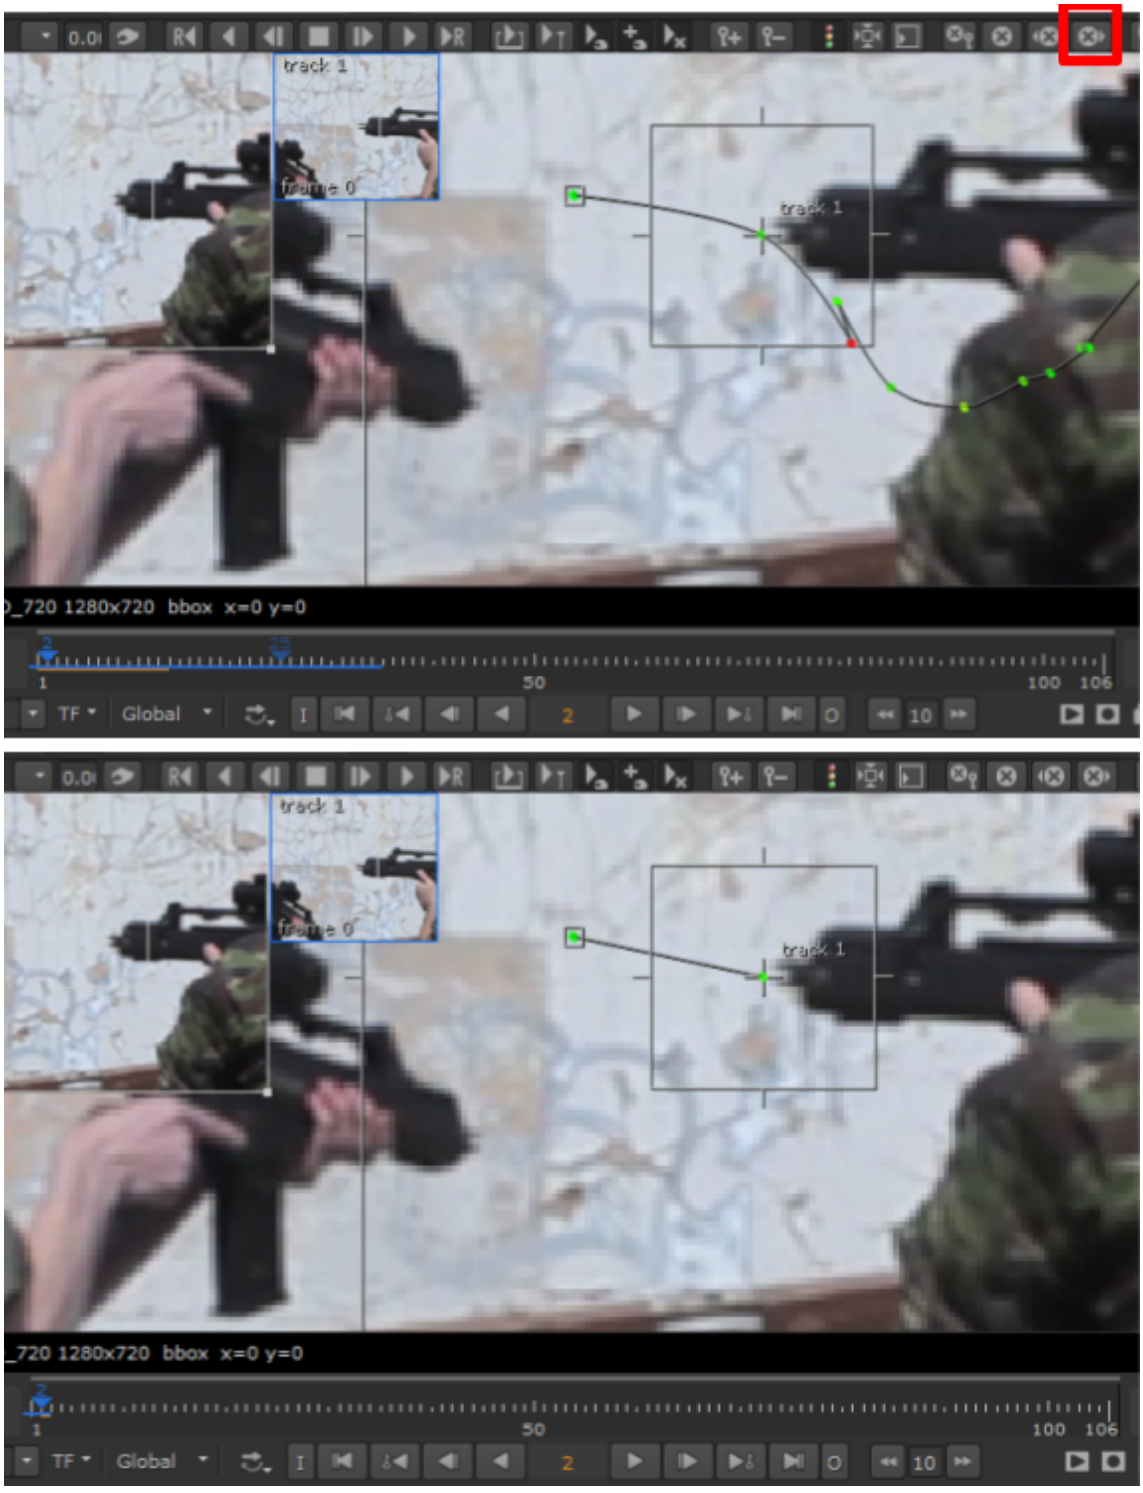

NOTE: The button to the right is the clear backwards button. Use this if you're autotracking backwards instead of forwards.

Keyframe Tracking

Keyframe tracking is where you manually track the feature between frames. It's a bit tedious to do if you're doing it for every frame, but the typical workflow is to track using auto track until it screws up, use keyframe tracking to get it back on the correct path, then go back to using autotrack.

To set a tracking keyframe, navigate to the frame and just move the tracking marker. As soon as you release the mouse button, a tracking keyframe gets set…

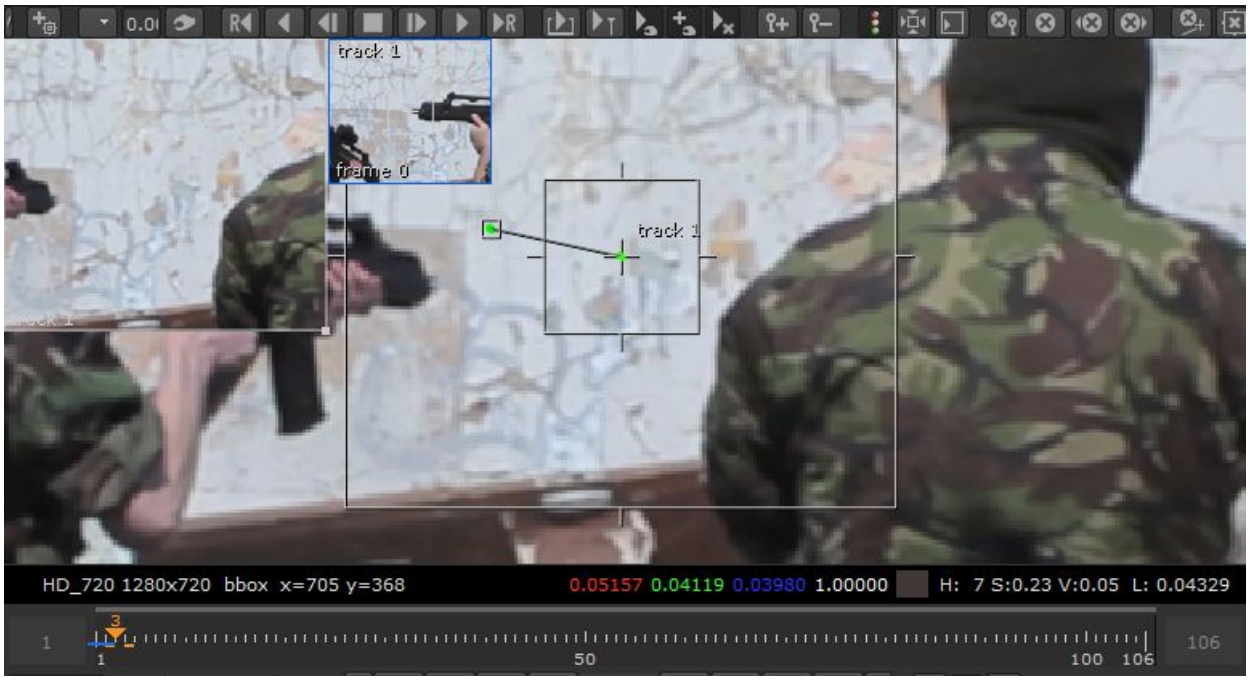

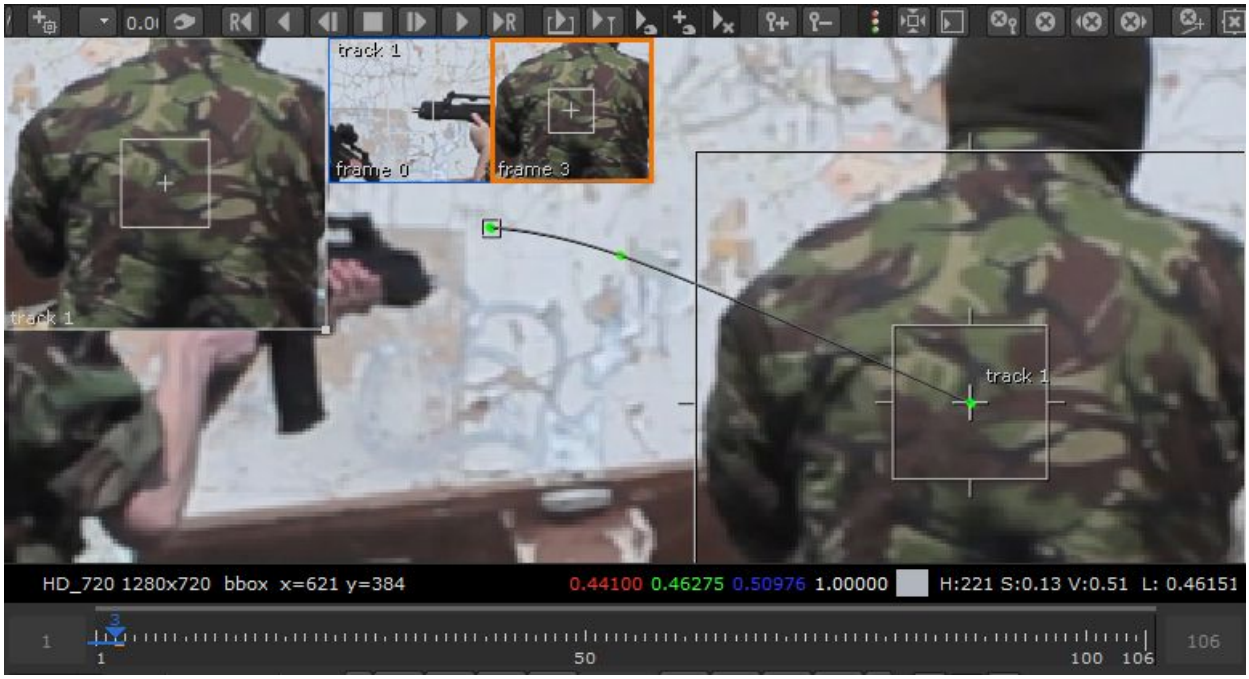

NOTE: Notice how a new thumbnail shows up at the top of the viewer when you manually set a keyframe marker. That's how you know it's a keyframe.

Another way to set a tracking keyframe is to have autotrack calculate where it thinks the feature will be in the next frame, and then manually adjust it. You can do this by using either the track next or the track prev toolbar buttons…

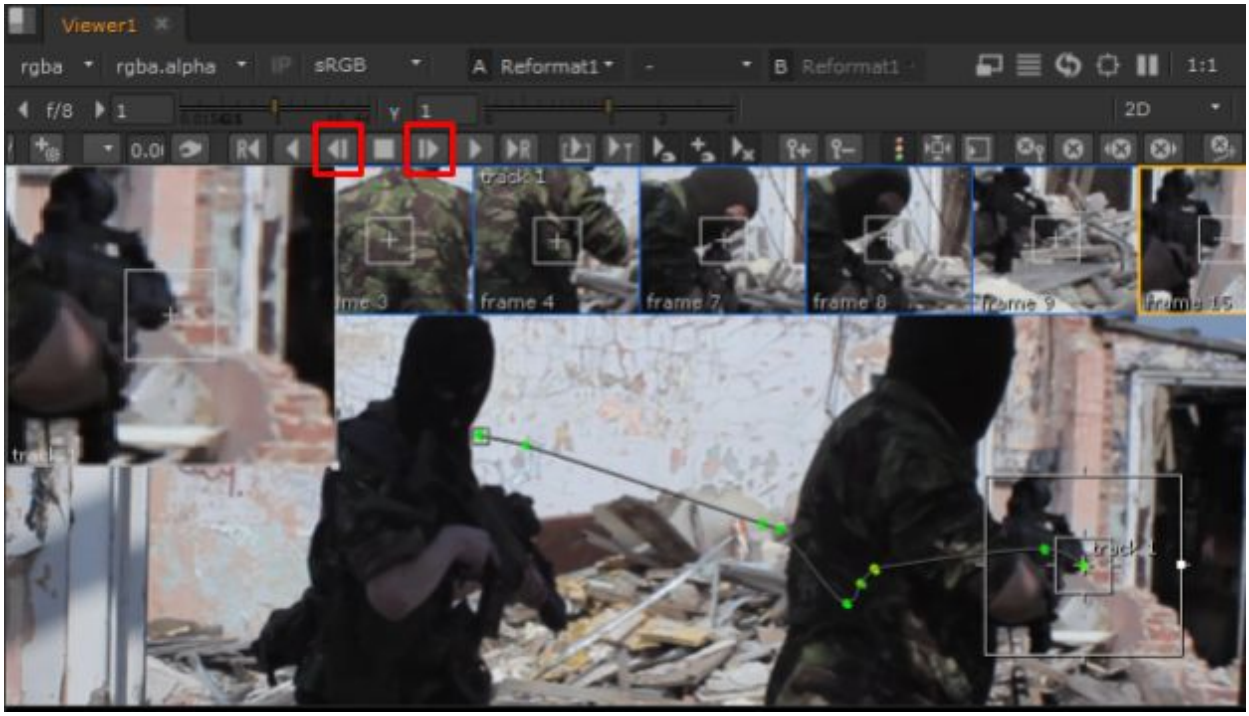

NOTE: Don't see this toolbar in the viewer? Make sure the properties for your tracker node are open and the track is selected in the properties.

NOTE: Notice all the thumbnails up top -- these are for each frame that's using keyframe tracking. Even though we're using the track next/track prev buttons, as soon as we move the marker the track for that frame will get converted from an auto track to keyframe track.

Offsetting Tracks

If the feature you're tracking gets obscured or goes off-screen, you can temporarily offset your marker onto another feature until it comes back into view. The feature you swap to will act as an offset for the main feature. That means that the movement of this temporary feature will be controlling the original track -- the track point will get offset by the temporary feature's movement.

For example, imagine that we're tracking the following black wall stain…

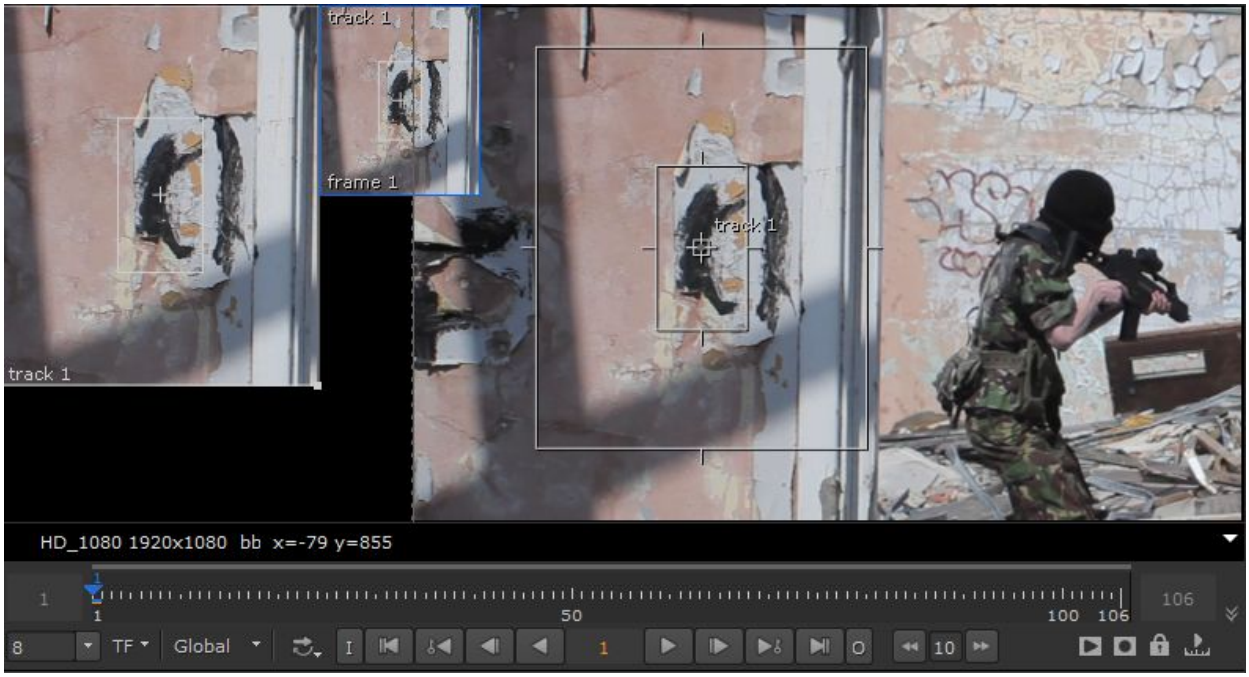

By the 29th frame, this feature will be obscured by the soldier and won't be trackable anymore…

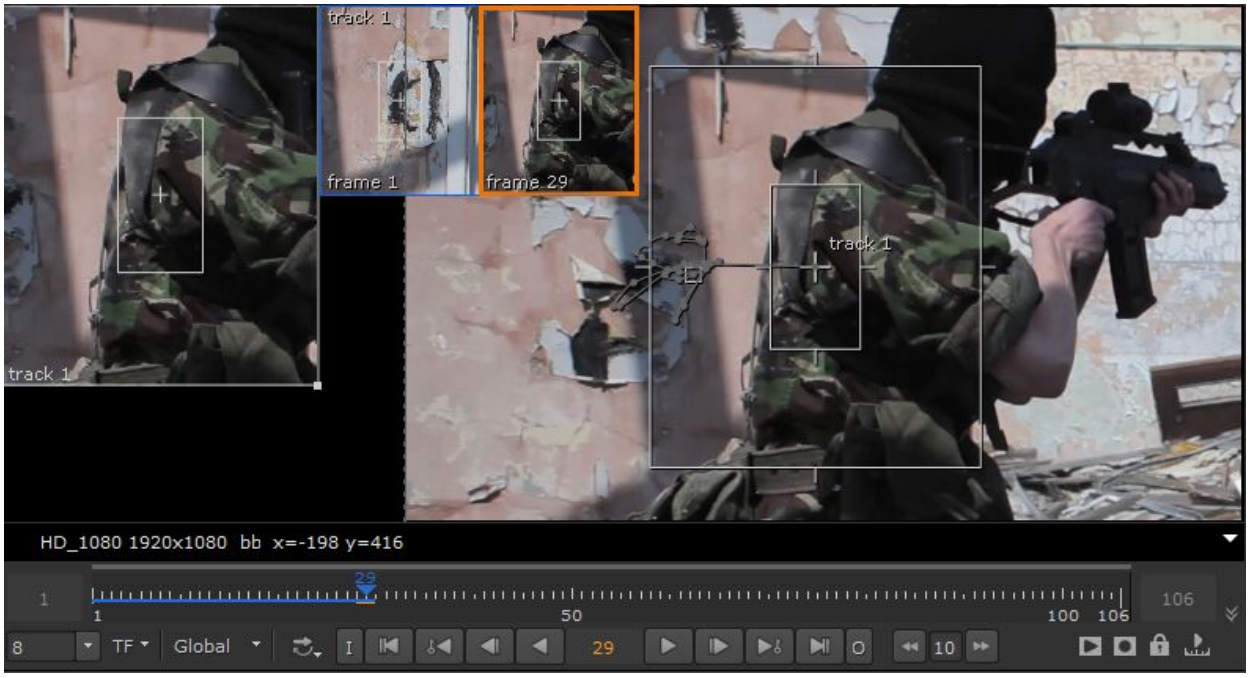

If this happens, we can temporarily move our marker to another feature (preferably on the same rigid body) to continue tracking. We do this on the frame IMMEDIATELY BEFORE the frame where the feature gets obscured by Ctrl+LMB click-dragging the marker's anchor to the new feature. That new feature will then be used to offset the original point being tracked.

In the example above, we'll Ctrl+LMB click-drag the anchor on frame 28 (just before our original feature gets obscured) to the blotches just to the left of our original blotch…

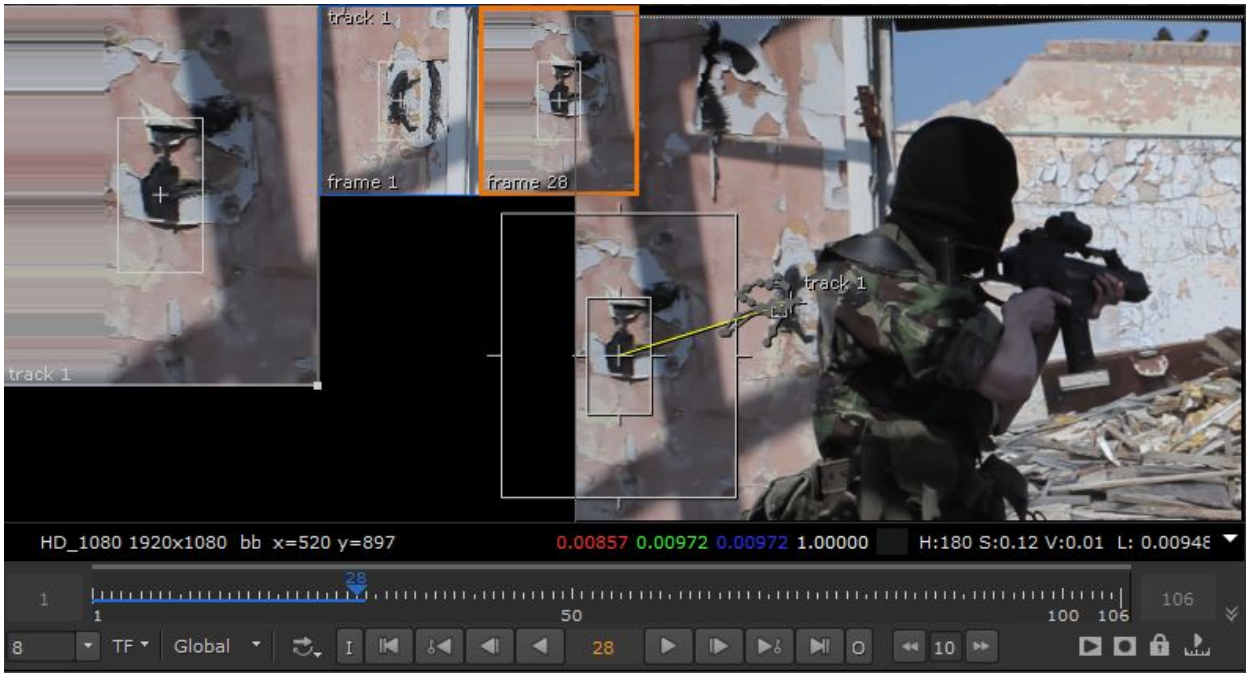

NOTE: The yellow line in the screenshot above denotes that this new feature's track will be used to offset the original track position moving forward. As this new feature moves, the movement offsets will be applied to the original track.

NOTE: Don't forget to hit the clear forward toolbar button when you go to the frame immediately before where the obscuring happens. Remember that it clears the track of everything after the frame you're on -- that data is no longer good because the feature was obstructed. See previous sections to find out which button is the clear forward/backward (search this doc using Ctrl+F).

We can move this temporary anchor as many times as we need to before putting it back onto the original feature being tracked. Just follow the same process: on the frame IMMEDIATELY BEFORE it gets obscured, Ctrl+LMB click-drag the marker's anchor to a new feature and continue the autotrack.

Once the feature is visible again, you can click the clear offset button in the toolbar and move the marker's anchor back to the original feature. If you didn't mess anything up, when you hit the clear offset button the track should come back onto (or very close) the original feature...

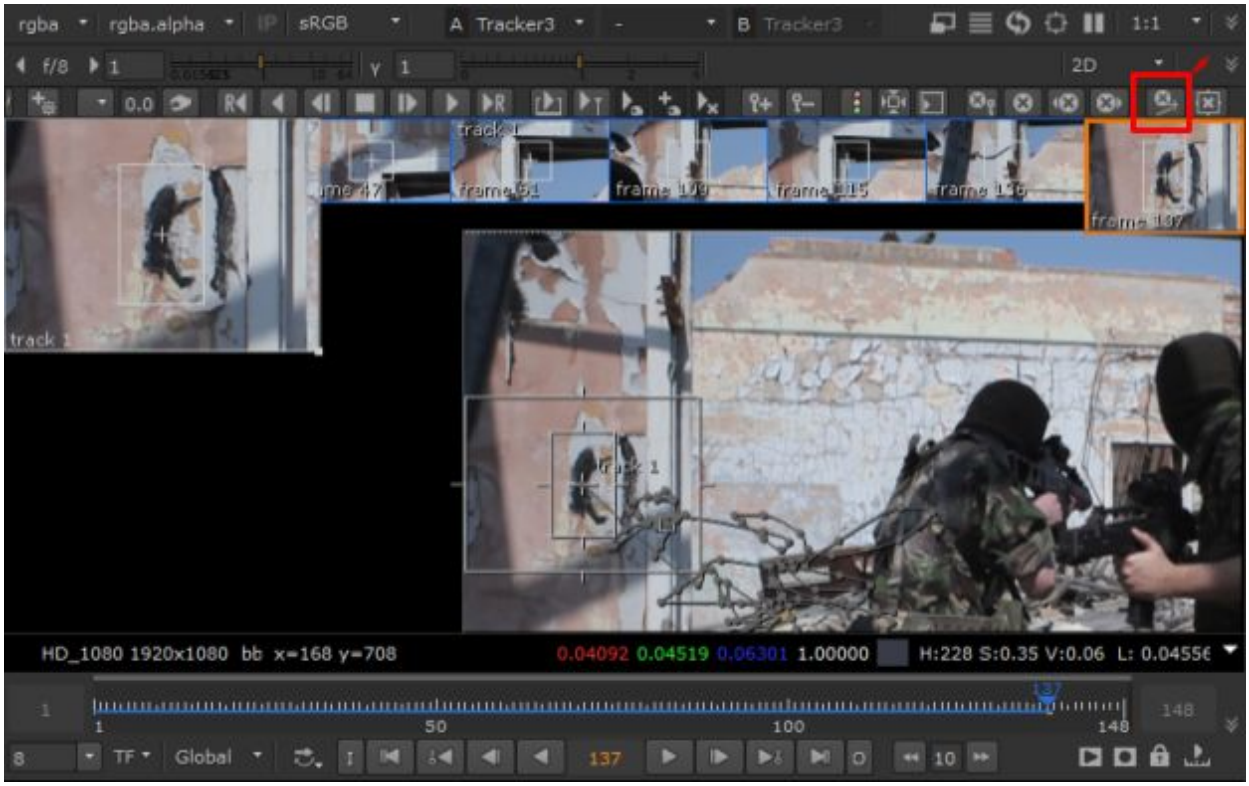

Rotation/Scale Tracks

As already mentioned in the introduction, tracking only 1 feature will only give you translation (position) information. Tracking more than 1 feature will allow you to track rotation and scale as well.

To define which track gets used for what, you can set the T R S checkboxes on each track…

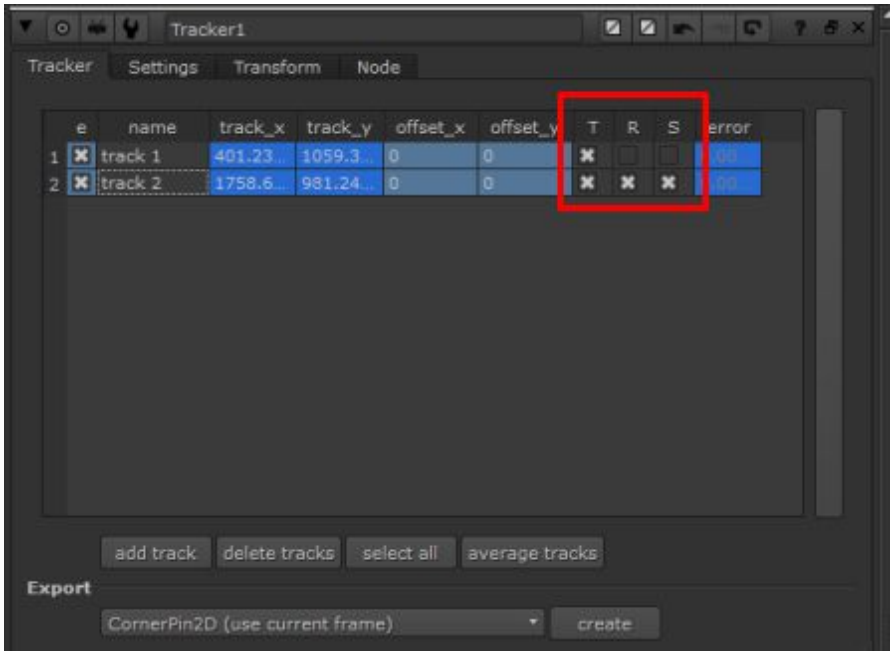

The meanings of T, R, and S should be obvious…

- T → track used for calculating translation (position)
- R → track used for calculating rotation
- S → track used for calculating scale

In the above example, multiple tracks are assigned to T. All that means is that translation (position) from both tracks are combined using averaging.

Also, notice that R and S is only set on the second track. For rotation and scale information you need to have at least 2 tracks. If you have…

- only 1 track, setting R/S on it does nothing.
- only 2 tracks, setting R/S on at least one of the tracks will give you rotation/scale.
- more than 2 tracks, setting R/S on multiple tracks will give you rotation/scale by averaging the results from each track.

NOTE: I'm not 100% sure on any of this. None of it is clearly documented, but when you try it out and go to the transformation tab, it looks like the transformation crosshairs support what's stated above.

Improving Tracks

There are some settings you can tweak to improve your feature tracking.

Pre-track filters

For autotracking, you may be able to get better results by fiddling with the options under the settings panel of the tracker node's properties…

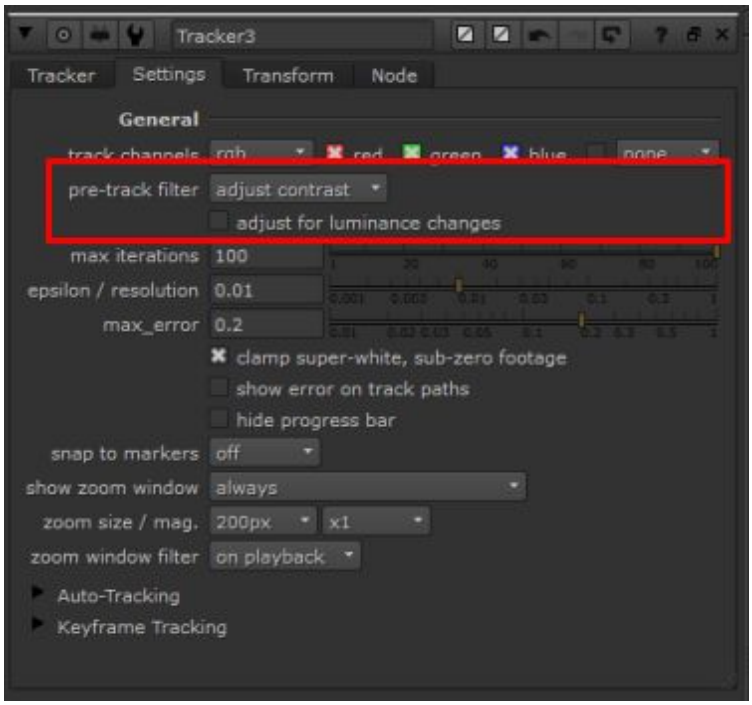

The main options here are...

- pre-track filter → This is a image filter applied to the area where the tracking is happening. The goal is to make it easier for the autotracking algorithm to follow the feature. The filter can be one of 3 values:
    - none -- no filter
    - adjust contrast -- modifies the contrast such that it becomes easier to track (this is the default). Remember that the lesson said that features that highly contrast are good for autotracking (see intro section for more info).
    - median -- attempts to denoise such that it becomes easier to track. This is only useful if your image has noise.
- adjust for luminance changes → Remember that the lesson said that features that enter in and out of shadows may not be good for autotracking (see intro section for more info). Apparently this option is to help with that. From the Nuke docs…

When enabled, Tracker does some extra pre-filtering to compensate for changes in brightness. This option slows the tracking process and can reduce the accuracy of tracks, so only enable this control if there are known changes in brightness.

Smoothing

Smoothing can be used when you have jittery tracks.

According to the lesson, some types of shots are known for having jittery/noisey tracks. The example that was given was a shot taken with a dolly. Even though the dolly is keeping the camera fixed and moving it in a straight line, it often travels over a floor that isn't a perfect plane. Even something as benign as having the dolly travel over carpet fibers causes slight jitters in the shot.

Nuke gives you the option to smooth out the transformation (rotation/scale/translate) of the track by averaging out the n frames together. To do this, update smooth values under the transformation panel of the tracker node's properties…

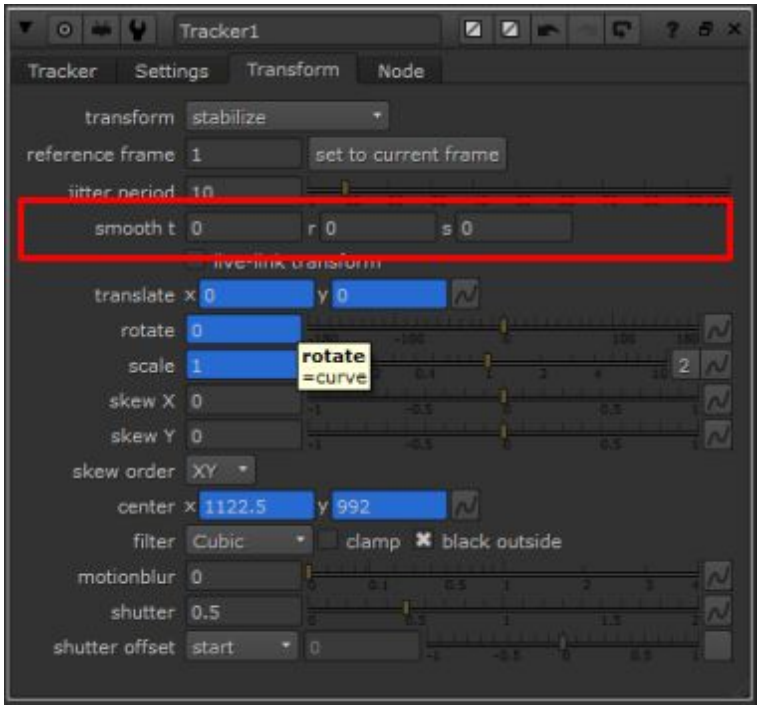

Transform Output

Once you've set up the tracks for your tracker node, you can choose what the tracker node does with that information by going to the Transformation tab…

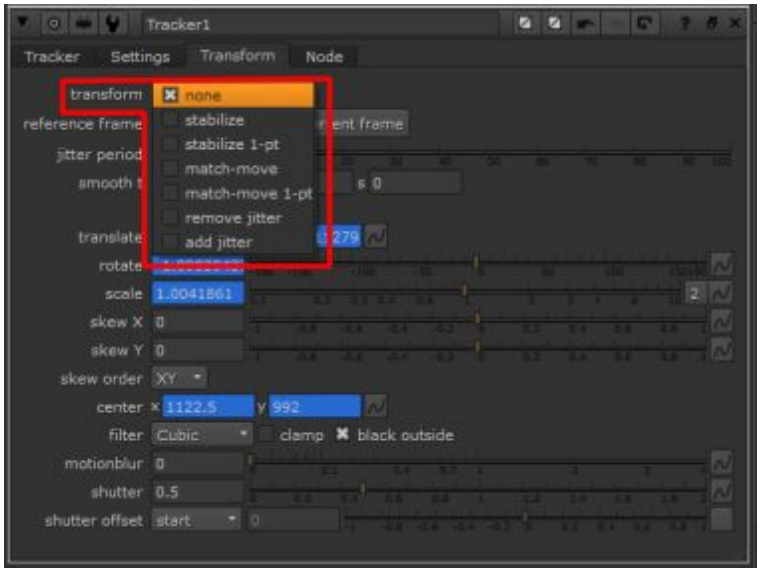

The operation performed is selected under the transform dropdown...

- none → does nothing.
- stabilize → transforms such that the points stay inplace (camera stabilization).
- match-move → transforms based on the movement (so another image can follow along).

> **NOTE**: I don't know about add/remove jitter. Use the 1-pt variants if you aren't tracking translation and scale information (you only have 1 tracker).

Another option to be aware of here is reference frame (the next one after transform). According to the Nuke docs, this is the reference frame from which the calculations for the track are done? So the rotation/scale is always calculated from this point?

> When transform is set to stabilize or match-move, sets the frame in the input sequence to use as the reference or identity frame.

Linked Output

If you want to use the final translation/rotation/scale output of your tracker node in some other node, you can link it. The transform tab in the tracker node properties will contain the computed values (they'll be highlighted in blue meaning they're animated/keyframed fields)...

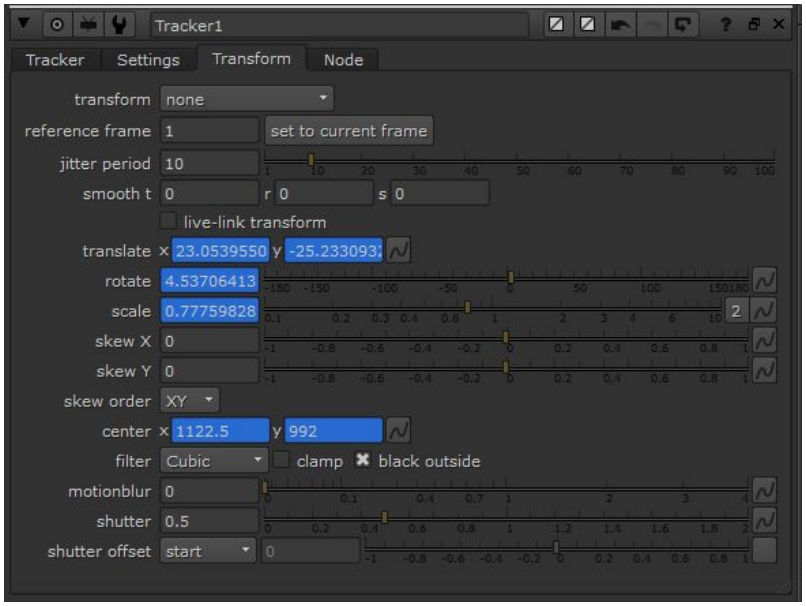

In addition, you can use transformation values from the individual trackers inside the tracker node. You can do this by going to the property which you want to set, clicking the squiggly line button next to it, and choosing the tracker node and tracker to link to…

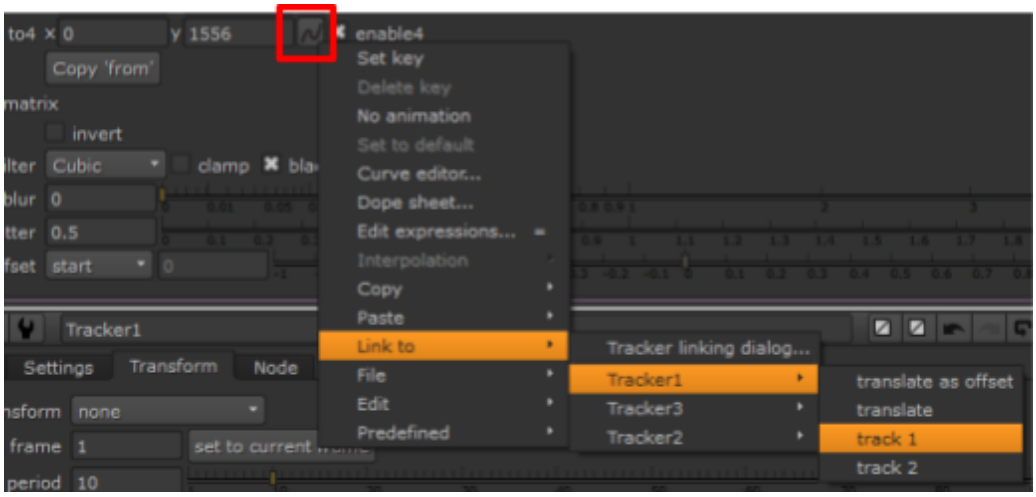

A common use-case for linking to individual trackers in a tracker node is when you want to superimpose an image on a screen. You can have a tracker node with 4 trackers, and link the position of each tracker to the corners of a CornerPin2D node.

NOTE: For more information on linking, see the Linking section of the main document.

Nuke 11 Rotoscoping

Introduction

Rotoscoping is a fancy term for creating masks on image sequences, such that you can apply some set of operations on that mask only. For example, color correct only a certain object in the image sequence or copy a certain object from the image sequence into another image sequence.

The term rotoscoping seems to be an esoteric VFX-specific term. It has some historical context in that it was originally a technique used in the early 1900s to do similar work (tracing/masking). From Wikipedia…

> Rotoscoping is an animation technique used by animators to trace over motion picture footage, frame by frame, when realistic action is required. Originally, photographed live-action movie images were projected onto a glass panel and re-drawn by an animator. This projection equipment is referred to as a rotoscope, developed by Polish-American animator Max Fleischer. Although this device was eventually replaced by computers, the process is still referred to as rotoscoping.

Other terms closely related to rotoscoping include…

- matte (photography) - combine 2 or more image elements, usually a foreground and background, into a single image. For example, a picture of an actor and a image of a

landscape. The term matte is often used to refer just to the background that will end up in the final image.

- composite (vfx?) -- compositing is the combining of several image elements into a single image / sequence of images, often with the goal of making the final output look as if its a single cohesive scene.
- chroma key (vfx?)
- green screen (vfx?)

Roto Node

In Nuke, rotoscoping is done using the Roto node. Image data is fed into the Roto node and the artist defines the masks/traces by manually drawing it out in the Viewer. These masks/traces are then output from the node as a channel (e.g. alpha channel, red channel, etc…) and can be used further down the node graph or dumped out for some other tool to use.

Adding Node

To add a Roto node, you can use any of 3 basic methods described in the main Nuke document…

- Tab menu → hit tab in the graph and type in Roto…

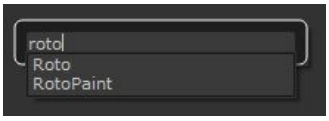

- Context menu → right-click in the graph and goto Draw → Roto…

- Toolbar → click the pencil and select Roto…

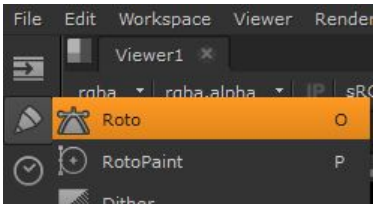

If you already have a node selected when you add the Roto node, it'll get added to the output of the selected node. If the selected node already has an output, it'll get added between the nodes.

Tracing Shapes

You can trace/draw the shapes for a Roto node directly in the viewer, but you need to make sure that you first have that Roto node's properties open.

> **NOTE**: To open up the properties for a node, double-click on the node.

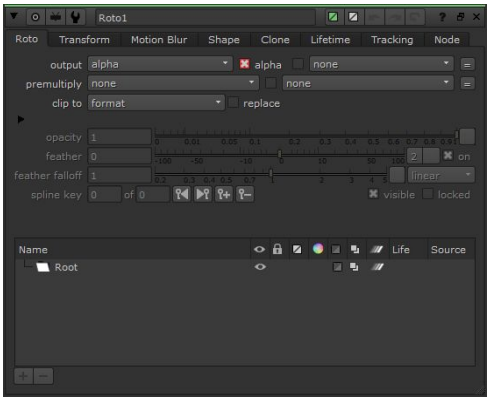

> **NOTE**: Remember that you can have the properties for multiple nodes open at once. If you have the properties for 2+ Roto nodes open at once, the one that's currently selected will be the one you're drawing on.

Once open, you'll see 3 buttons in a vertical toolbar hugging the left of the viewer you have open…

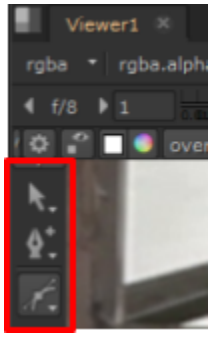

NOTE: Remember that the Viewer doesn't have to be the Viewer that the Roto connects to. So long as the properties of the Roto node are open, you can draw in any open Viewer.

The top button is for selecting parts of an existing shapes, the middle button is for manipulating points on existing shapes, the third button is for drawing new shapes…

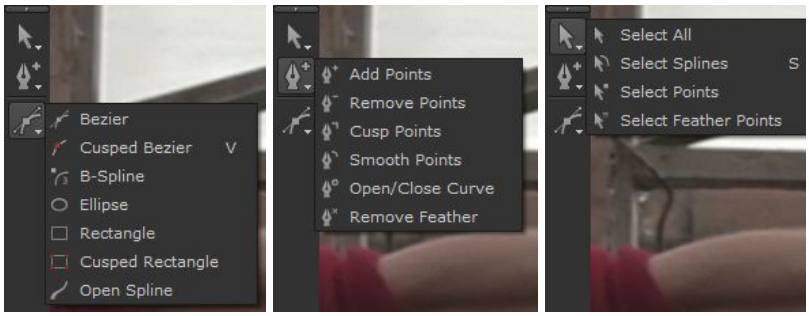

NOTE: Be careful when clicking the 3 main buttons. If a button is already selected but you click it anyways, it'll switch to the next item in it's list. For example, if I have the first button already selected and I click on it again, it'll cycle to the next item in it's list…

If you want the dropdowns to show up, you need to click-and-hold until it pops up.

Creating Shapes

To draw/trace shapes on a Roto node, you need to have the node's properties open and the node selected.

In the Viewer, make sure you have the Bezier tool selected...

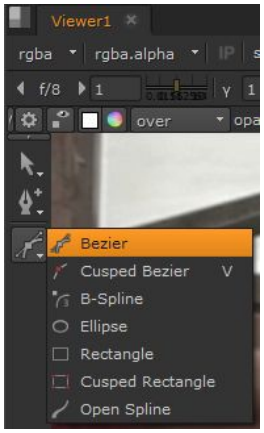

NOTE: You can play around with other tools in this list as well. They all draw stuff but Bezier is the easiest to work with. Cusped Bezier is like Bezier except that it won't give you bezier handles as you draw (discussed further below).

You can then start clicking in the viewer to start drawing the shape. If you…

- LMB click → it'll drop a point without bezier handles

Don't worry if you forget to add in bezier handles. You can always add handles for a point after you're done tracing (discussed further on in this doc).

- LMB click+drag → it'll drop a point where you clicked and the drag defines the length of the bezier handles.

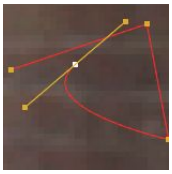

The bezier handles here are symmetric. Once you've dropped the point, you can do Ctrl+LMB on a handle to change it without changing the other handle. You can also do this after the fact once you're done tracing (discussed further on in this doc).

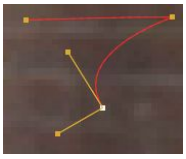

NOTE: If you aren't in drawing mode and you accidentally hit Ctrl+LMB, it'll lock the color picker onto a particular pixel. That pixel will get highlighted in red. If this happens by accident, Ctrl+RMB will remove it.

NOTE: The shape you draw must be a closed -- it can't have any gaps/openings. If you don't close it yourself, Nuke will close it for you as soon as you switch to another tool.

Once your shape is completed, you'll see it in the list of shapes in that Roto node's properties. A Roto node can hold on to multiple shapes. The final output of the Roto node will combine all these shapes together (so don't worry if there's overlapping happening between your shapes).

The Node properties like you do things such as hide shapes, lock shapes, etc...

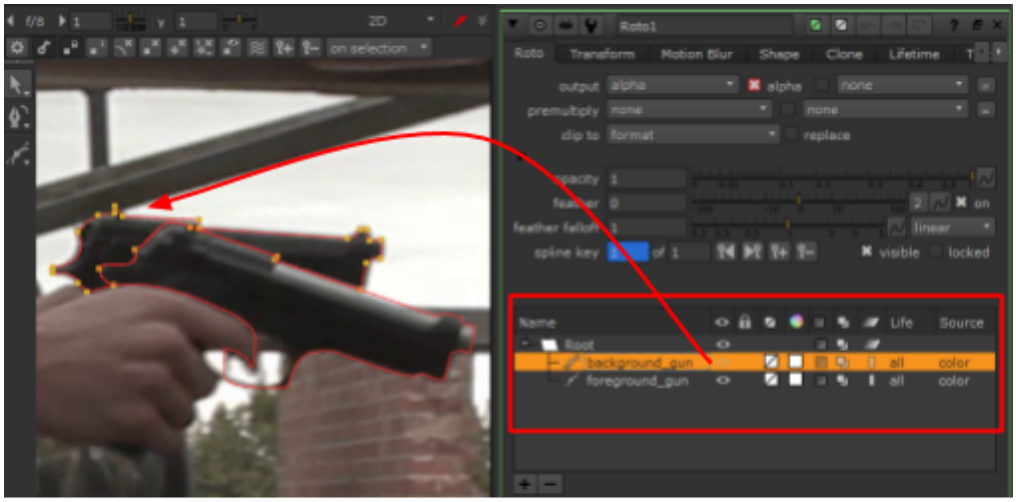

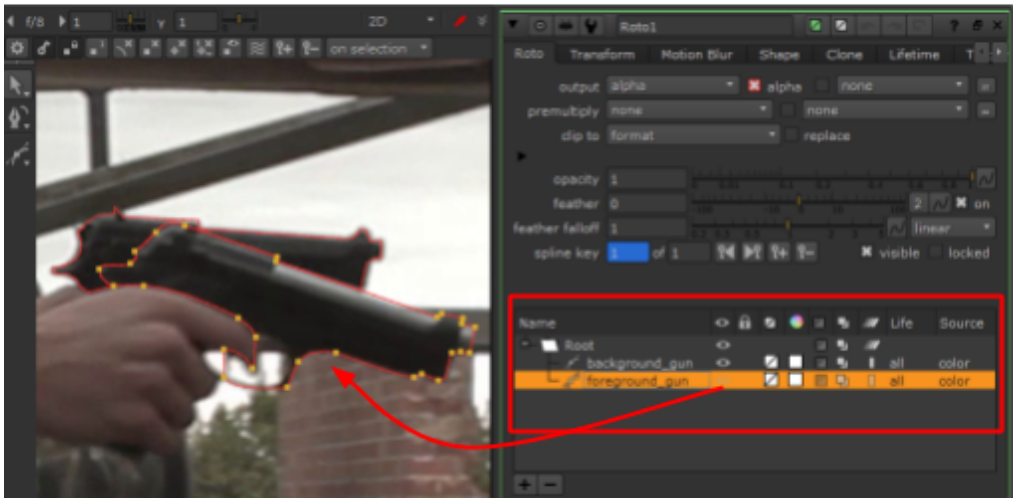

NOTE: To rename a shape, double-click on the name in the properties panel.

Modifying Points

> **NOTE**: Make sure you have the Roto node's properties open and the shape you want to edit selected. You'll know the shape is selected when you see the points highlights in the viewer / when the shape is selected in the Roto node's properties.

You can move points via the 1st toolbar button in the Viewer.

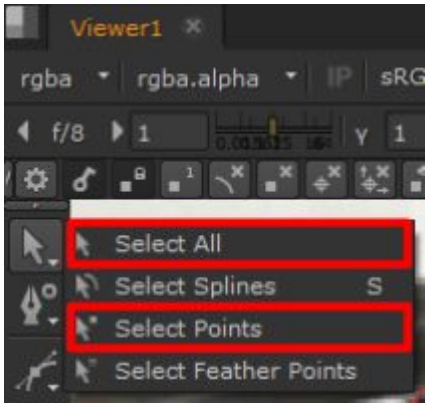

- Select All → select everything, including points
- Select Points → select just points

> **NOTE**: Typical people use Select All for everything. Rarely ever will you be tweaking points only -- what's more likely to happen is that you'll tweak a point, then maybe tweak its feathers or it's bezier handles, and then go back and tweak that point again, etc.. etc.. It's a mixed workflow.
>
> The problem with Select All is that Nuke uses proximity-based selection. So, if 2 selectable things are very close to each other, when you click it'll always select one over the other (even if you're very careful with your clicks). The work around to this (aside from not using Select All) is to zoom in until the 2 selectable things are separated enough to make distinctly selectable.

You can manipulate many points at once by first selecting the shape and then either marquee selecting over points or Shift+LMB clicking points...

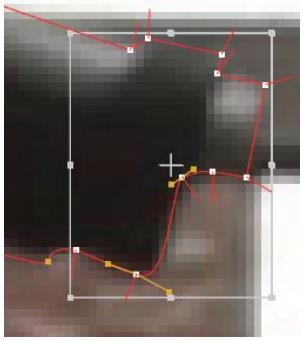

You can add and remove points from a shape points via the 2nd toolbar button in the Viewer.

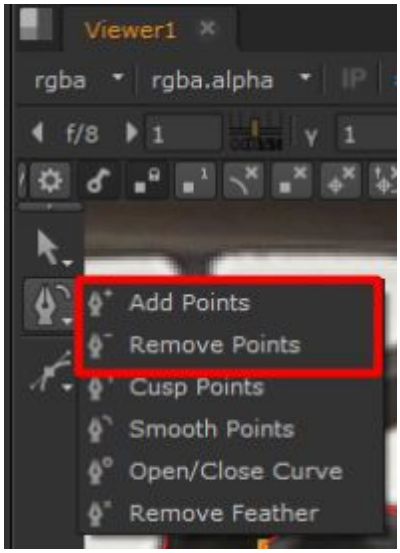

The items listed here are self-explanatory...

- Add Points → adds points to where you click
- Remove Points → removes points you click

Modifying Bezier Handles

> **NOTE**: Make sure you have the Roto node's properties open and the shape you want to edit selected. You'll know the shape is selected when you see the points highlights in the viewer / when the shape is selected in the Roto node's properties.

You can change around the bezier handles for a point by selecting it via the 1st toolbar button in the Viewer.

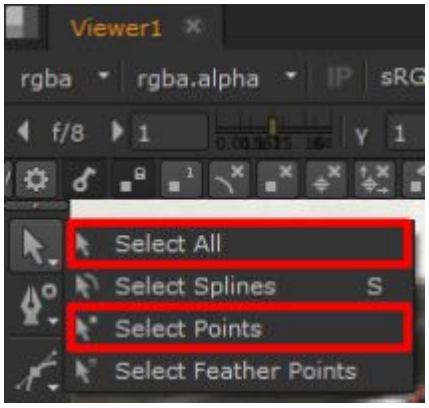

Once selected, if the point actually has has bezier handles, you'll see 2 other lines+points protruding from that point…

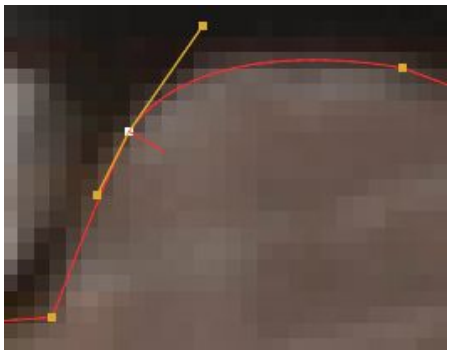

If you…

- LMB click-and-drag a handle → it'll symmetrically move around the handles, meaning that a change in one handle will make a proportionally opposite change in the other handle.
- CTRL + LMB click-and-drag a handle → it'll asymmetrically move around the handle, meaning that the change to the handle won't affect the other handle.

You can add and remove bezier handles from a point via the 2nd toolbar button in the Viewer.

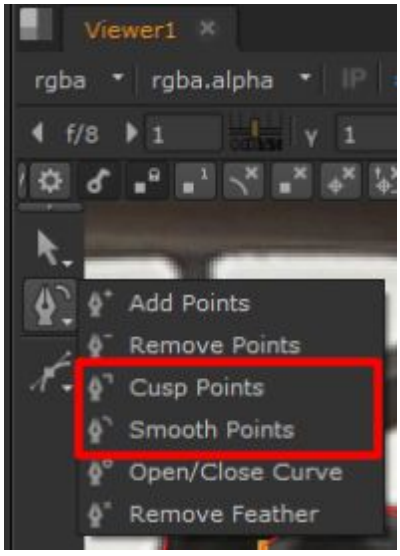

The items listed here are self-explanatory...

- Cusp Points → removes the bezier handles on points you click
- Smooth Points → add bezier handles on points you click (or smooth if handles already exist)

Feathering Shapes

Feathering is used when what you're tracing has a blur (e.g. motion blur). You trace over the object as normal and provide a feather for the region where the blur is occuring. The feathered region will less intensely apply whatever operation the output of the Roto feeds into.

There are 2 ways to apply feathering. The first is to use the global feathering option found in the Roto node's properties.

Shape Feathering

Select the curve you're operating on and then change feathering sliders…

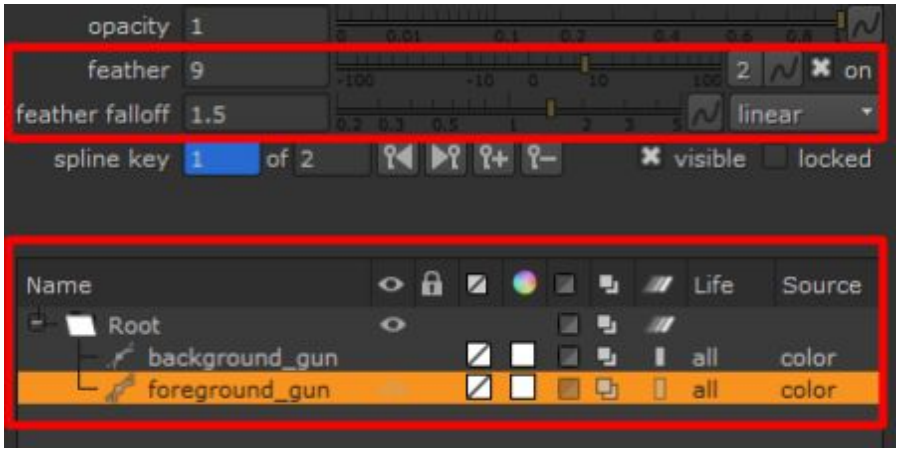

The output won't be immediately visible on the trace, but you can see it if you view the alpha channel…

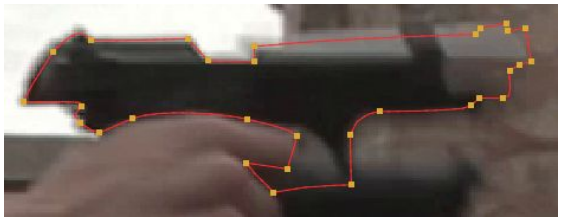
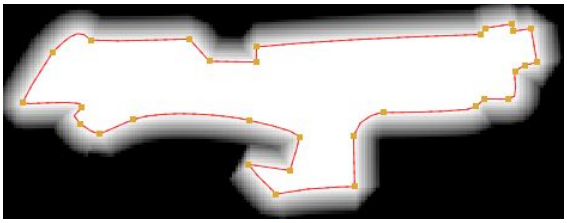

Point Feathering

The second way to feather shapes is to apply the feather directly on the points. When you select a point, you should see a little red line protruding from that point. If you click-and-drag that line, it'll create a feathered region. That region will be denoted by a dotted line...

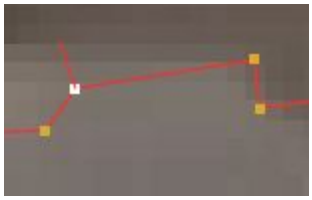
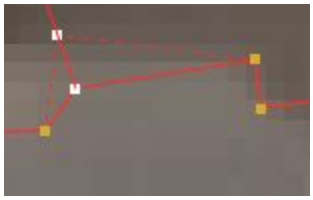

> **NOTE**: If you can't see the protruding line, it may be that the line is parallel to the lines connecting neighbouring points. Either way, you can Ctrl+LMB drag to create the feathered region -- it does exactly the same thing as pulling out the line.

The points on this new dotted line will be a copy of the original points. If the original point had bezier handles, so will the point on the dotted line (and the handles will be pointing in the same direction and have the same length). You can manipulate the points independently from the original, but they'll always be tied to the original. You'll see a line going through the dotted line point and the original point when you select either…

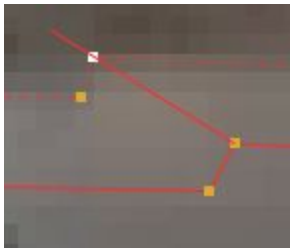

The output won't be immediately visible on the trace, but you can see it if you view the alpha channel…

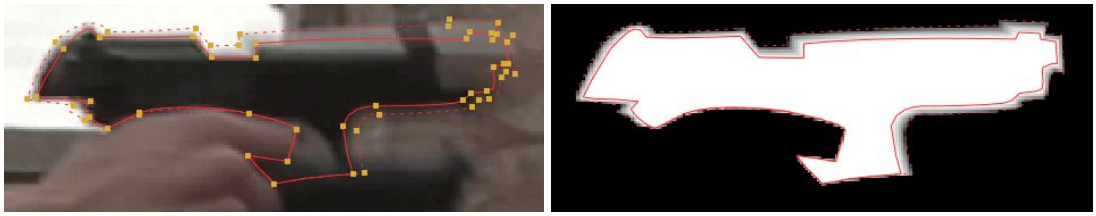

Animating Shapes

If you're tracing over an image sequence, you can animate the points on your shapes using normal keyframe animation.

To keyframe a shape…

1. select the shape to keyframe
2. move the frame where the keyframe should be set
3. manipulate the points on that frame

As soon as you change something on a shape, a keyframe automatically gets set. You can then change stuff however you see fit using the dope sheet or curve editor.

> **NOTE**: I haven't tested this but I don't think you can add/remove points on a keyframe (outside of feathering points).

Values are keyed only on the selected shape. If your Roto node has multiple shapes, the scrubber will only show keyframes for the currently selected curve. If Root is selected, it'll show the keyframes for all shapes.

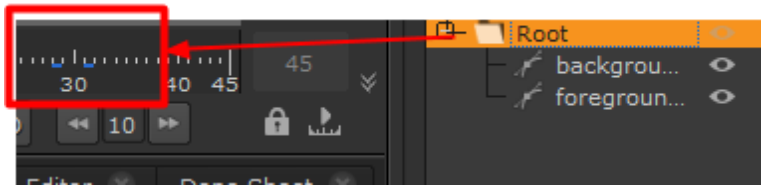

Note the number of blue ticks in the above examples. They're the keyframes. The background_gun shape only has 1 keyframe, while the foreground_gun shape has 2 keyframes. Depending on what curve is selected, the keyframes shown in the scrubber/dope sheet/curve editor will be different.

Copy/Paste Shapes

You can copy/move shapes to other Roto nodes by RMB-clicking the shape in the properties panel and selecting Copy/Cut, then going to another Roto node and pasting it.

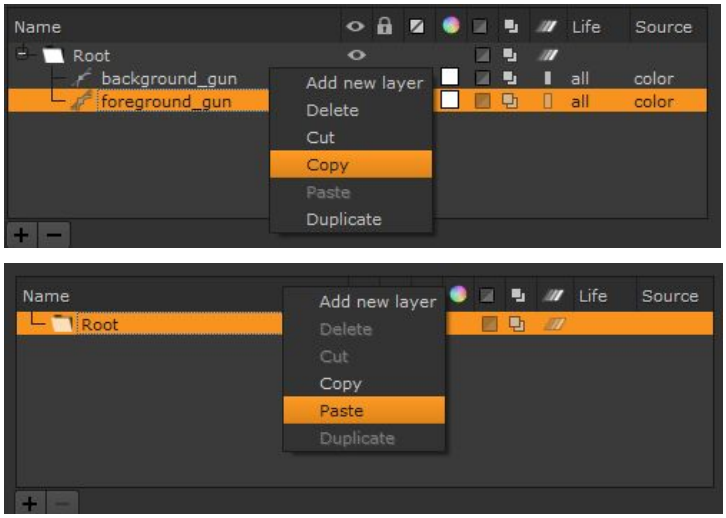

You can also drag-and-drop shapes between the properties panel of different Roto nodes.

Best Practices

Some important points to be aware of when making your shapes…

1. start from frame where object being traced is most stable -- the most stable/easiest to trace still of the object is the one you should start with.
2. add points around joints -- if the object bends in a certain spot, that spot should probably have a point.
3. keep the number of points to a minimum -- this helps maintain "point consistency".
4. make points symmetrical -- if there's a point, there should probably be another point across from it.

   NOTE: The lessons didn't really explain why these are important. These seems to apply only to image sequences -- individual images probably don't matter as much?

Channel Output

By default, the shapes you traced/drew will show up in the alpha channel (combined with the original input), so it won't be visible. You can control which channel the Roto node dumps its output to via the properties…

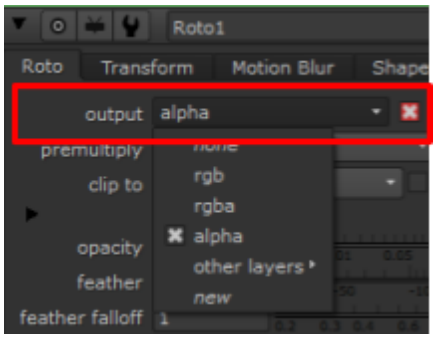

The output is typically left as alpha. If you want to have the output go to another channel, you can use a Copy node to change it…

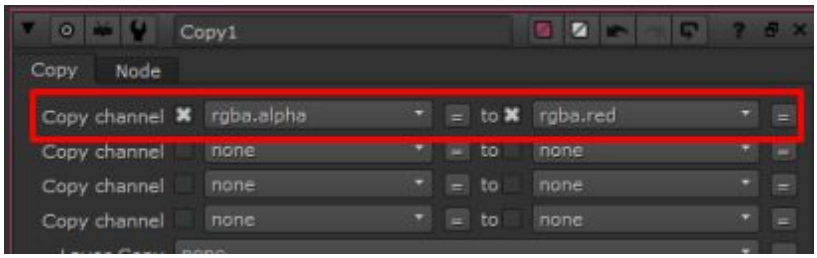

> **NOTE**: This is super useful if you want to see the mask you're creating as well as the original image -- just each the output of the copy node go to another Viewer node.

Cache Output

In many cases, the user wants to cache the traces/drawings out to disk to an image sequence (maybe they want to feed the masks downstream to some other application). This is done by feeding the Roto node's output to a Write node....

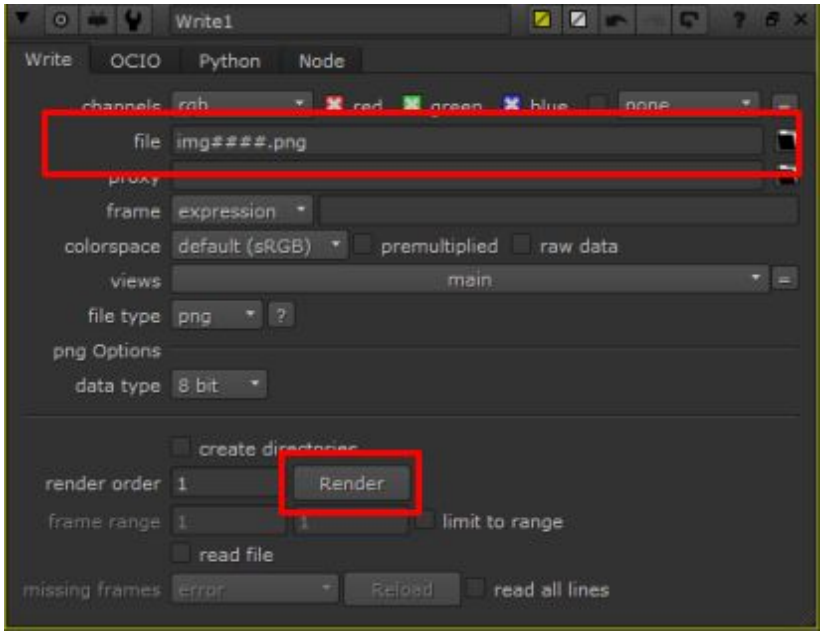

NOTE: Have an image sequence? Make sure your Write node is outputting as an image sequence by using # signs in the filename… e.g. img####.png.