Shake 4
New Features

Contents

About the New Features in Shake 4

This document covers all of the new features and enhancements introduced in Shake 4.

If you are already familiar with previous versions of Shake, this document lets you focus on the extensive new enhancements and additions.

Note: All of information in this document is duplicated from the *Shake 4 User Manual* included in your Shake box and available in PDF-format from the Shake Help menu.

Changes to the Shake Documentation

The print and onscreen (PDF-format) documentation has been completely reorganized and expanded from prior versions.

User Manual

The user manual has been divided into more chapters to make topics easier to find in the table of contents and in the bookmarks drawer of the onscreen PDF. In addition, the documentation of many prior features has been expanded.

Node Reference Guide

Scripting and variables documentation for each of Shake's functions, previously located in the user manual, has been moved to a new document: the *Shake 4 Node Reference Guide*, accessible from the Shake Help menu.

The Cookbook

The Cookbook—a popular collection of tips and techniques for improving your workflow in Shake—has been relocated from the tutorials book to to the last chapter of the user manual. What's more, content from the Cookbook has been integrated into the rest of the manual, where appropriate, to make these topics more accessible.

Summary of New Features

This section summarizes the new features that appear in this document.

Topics in Chapter 1
- You can open two scripts at once using the Shake graphical user interface.
- You can save and restore favorite views in each Shake interface area.
- You can enable the ability to use relative paths in the File Browser.
- Favorite directories in the File Browser are now saved automatically.
- There are now lock controls for each parameter.
- There is a new "next" button when you create local variables, which makes it easier to create many variables at a time.
- A new console tab lets you view the contents of the console from within Shake.
- New global parameters provide additional ways to control new features.
- An enhanced Node View provides options for visualizing animated nodes, expression links, node concatenation, and noodle color coding and stippling.
- You can now resize the node overview.
- There are four new ways of adding and connecting nodes to a node tree.
- Cloned nodes are now named differently.
- Truelight monitor calibration software is included, which also provide a new Viewer lookup table (VLUT) option for the Viewer.
- The onscreen controls for transform nodes have been updated, with additional hot keys.

Topics in Chapter 2
- Shake now supports the OpenEXR format.
- There is new support for clips and sequences sent from Final Cut Pro 5 to Shake 4.
- New adaptive retiming options in the *FileIn* node use optical flow technology to create smooth motion retiming.
- New remastering options in the *FileIn* node use optical flow technology to deinterlace, change frame rates, and rescale images for high-quality format conversion.
- A new *Cache* node lets you specify key points to be cached in your node tree.
- A new clampMask parameter in the mask subtree of nodes lets you clamp mask image data to a value between 0 and 1.
- The *RotoShape* node has new features, including:
  - The ability to copy and paste shapes between compatible nodes.
  - The ability to cut and paste rotoshape keyframes.
  - Controls to retime shape keyframes
  - New commands for attaching trackers to individual shape control points.
  - The ability to import and export shape data.
- The *SmoothCam* node automatically smooths or locks camera motion in a shot.

- The *AutoAlign* node uses advanced image processing to align up to three input images with one another, automatically.
- The *LensWarp* node provides calibration tools to correct for lens distortion, or to match the lens distortion from one shot in another.

Chapter 3
- Chapter 3 contains complete documentation for the *MultiPlane* node, which provides the ability to do the following:
  - Composite background or foreground elements against a moving background plate using 3D camera tracking data.
  - Arrange multiple layers within a 3D coordinate space for easy simulation of perspective, parallax, and other depth effects.

Interface Enhancements

1

This chapter covers general enhancements throughout the Shake interface.

Opening Two Scripts at Once

Shake is designed to have only one script open at a time. Typically, each script is used to create a single compositing project, with a single frame range and a single node tree. Although Shake supports multiple independent node trees within the same script, all trees share the same duration, defined by the timeRange parameter in the Globals tab.

If necessary, it is now possible to open two scripts simultaneously in two separate interface windows. In this case, what you're really doing is launching two instances of Shake at once. This is primarily useful if you need to copy information from one script to another.

Important: When you open Shake twice, the first instance of Shake is the only one that is able to write to and read from the cache.

Saving Favorite Views

If you find yourself panning back and forth within a particular area to the same regions, you can now create a Favorite View within that area. The Favorite Views apply to the following areas of the Shake interface:

* *Node View:* You can save several views in your node tree where you'll be making frequent adjustments.
* *Viewer:* If you're doing paint work on a zoomed-in image in the Viewer, you can save the position and zoom level of several different regions of the image.
* *Curve Editor:* You can save several different pan, zoom-level, and displayed-curve collections that you need to switch among as you adjust the animation of different nodes in your project.
* *Parameters tab:* You can save the parameters being tweaked, as well as the node being displayed in the Viewer.

Once you've saved one or more Favorite Views in each interface area, you can instantly recall the position, zoom level, and state of that area by recalling the Favorite View that you saved. You can save up to five Favorite Views.

To define a Favorite View:

1 Pan to a position in an area that contains the region you want to save as a Favorite View. If necessary, adjust the zoom level to encompass the area that you want to include.

2 Depending on the area you're adjusting, you can save additional state information particular to that area. Make additional adjustments as necessary so that you can recall the desired project elements:

   • In the Node View, you can save the state of the nodes that are currently loaded into the Viewer and Parameters tabs.
   • In the Viewer, you can save the node that's currently being viewed.
   • In the Curve Editor, you can save the curves that are currently loaded and displayed.
   • In the Parameters tab, you can save the parameters that are being tweaked, as well as the node displayed in the Viewer.

3 To save a Favorite View, move the pointer into that area and do one of the following:

   • Right-click anywhere within the area, then choose Favorite Views > View N > Save from the shortcut menu (where N is one of the five Favorite Views you can save).
   • Press Shift-F1-5, where F1, F2, F3, F4, and F5 correspond to each of the Favorite Views.

Restoring Favorite Views

Once you've defined one or more Favorite Views, you can restore them in one of two ways. Simply restoring the *framing* results in the current contents of that area being panned and zoomed to the saved position. Restoring the *framing and state*, on the other hand, results in the restoration of additional state information that was adjusted in step 2.

To restore the framing of a Favorite View, do one of the following:

   • Right-click in the Viewer, Node View, or Curve Editor, then choose Favorite Views > View N > Restore Framing from the shortcut menu (where N is one of the five Favorite Views you can save).
   • Press F1-5, where F1, F2, F3, F4, and F5 correspond to each of the Favorite Views.

That area is set to the originally saved position and zoom level.

To restore the framing and state of a Favorite View, do one of the following:

• Right-click in the Viewer, Node View, or Curve Editor, then choose Favorite Views > View N > Restore Framing & State from the shortcut menu (where N is one of the five Favorite Views you can save).

• Press Option-F1-5 or Alt-F1-5, where F1, F2, F3, F4, and F5 correspond to each of the Favorite Views.

Depending on the area, the originally saved position and zoom level are recalled, as well as the following state information:

• In the Node View, the node or nodes that were loaded into the Viewer and Parameters tabs when you saved the Favorite View
• In the Viewer, the node that was viewed when you saved the Favorite View
• In the Curve Editor, the curves that were loaded and displayed when you saved the Favorite View
• In the Parameters tab, the parameters that were being tweaked, as well as the node that was displayed, when you saved the Favorite View

Enable/Disable Local File Paths

The Relative Path control, to the left of the Add to Favorites control in the File Browser, gives you the option to enter a relative file path into the File Name field.

Enable local file paths.

Disable local file paths.
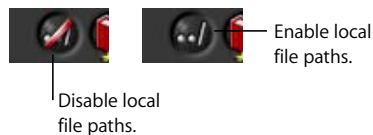
Relative file paths can take one of two forms:

• *./myDirectory/myFile/*
• *../myDirectory/myFile/*

Directories Saved to the Favorites List Are Permanent

If there are one or more directories with content you frequently need to access, you can add them to the Favorites list. The Favorites list is a customizable list of directories that you can add to at any time. As of Shake 4, entries you add to the Favorites list are permanent.

To add an entry to the favorites list:

1 Open the File Browser.

2 Click the Bookmark button.

The currently open directory is added to the Favorites list. All favorite directory paths you add are saved in the *favoritePaths.h* file, located in the *$HOME/nreal/settings/* directory. To remove favorite directories, delete them from this file.

Locking Parameters

Most parameters now have a lock button next to the Autokey button. This control lets you lock that parameter so that it can't be modified.

When you lock a parameter, its value field turns red to indicate that it's locked.
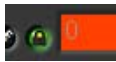
Locked parameters cannot be edited, but if they contain keyframes, an expression, or a link to another parameter, these values continue to animate that parameter.

A Next Button to Create Additional Local Variables

When you create a local variable using the Local Variable Parameters window, you can click OK to create the variable and go back to your project, cancel to close the window without creating a new variable, or next to continue creating new variables.

The Console Tab

The Console tab displays the data that Shake sends to the OS while in operation. It's a display-only tab. Two controls at the top of the Console tab let you change the color of the text, and erase the current contents of the Console. The maximum width of displayed text can be set via the consoleLineLength parameter, in the guiControls subtree of the Globals tab.

New Global Parameters

The following new parameters have been added to the Globals tab:

consoleLineLength
The maximum line length of information displayed in the Console tab. This defaults to 120 characters.

textureProxy
The proxy level at which texture-rendered images that are used by the *MultiPlane* node's hardware-rendering mode are displayed in the Viewer. This is similar to the interactiveScale setting, in that the proxy level that's set here is used to generate on-the-fly Viewer images.

multiPlaneLocatorScale
Affects all *MultiPlane* nodes within the script. This parameter scales the depth of the virtual space used to distribute the locator points that are displayed in the Viewer (which represent 3D-tracking data clouds that are imported from .ma files). This parameter allows you to expand or compress the relative distance from the camera to the tracked background plate. Adjusting this parameter lets you more easily position layers in space when camera tracking data comes from a subject that's either very far away, or very close. This parameter is for reference only, and has no effect on the data itself. The default multiPlaneLocatorScale value is 50.

noodleColor
Lets you change the base color of noodles in the Node View. Noodles are white by default, but you can use the Color control to change this to anything you like.

noodleColorA, -BW, -BWA, -RGB, -RGBA, -Z, -AZ, -BWZ, -BWAZ, -RGBZ, -RGBAZ
When noodleColorCoding is turned on (in the enhancedNodeView subtree), noodles are color coded and stippled (see "Noodle Color Coding" on page 16), based on the bit depth and number of color channels being propagated by each noodle in your node tree. When turned off, noodles appear using the default NoodleColor.

Different combinations of color channels are represented by different colors, and this group of parameters lets you customize the color used for each representation.

Using the Enhanced Node View

The enhancedNodeView subtree in the Globals tab provides additional display options for the Node View. These options can be turned on all together, or individually, to make it easy to spot image bit-depth at different parts of the tree, animated nodes, node concatenation, and expressions linking one node to another.

Unlike most other guiControls parameters, which are either off or on, these parameters—showTimeDependency, showExpressionLinks, showConcatentationLinks, and noodleColorCoding—have three states. This allows you to toggle all three parameters using the enhancedNodeView command from within the Node View. The three states are:

- *off:* Always off, regardless of whether or not enhancedNodeView is turned on.
- *on:* Always on, regardless of whether or not enhancedNodeView is turned on.
- *enhanced:* On when enhancedNodeView is on, and off when enhancedNodeView is off.

**To toggle enhanced node view off and on:**

1 Before turning on enhanced Node View, make sure the proper subparameters are turned on in the enhancedNodeView subtree of the Globals tab.

2 Move the pointer over the Node View, and do one of the following:

- Right-click, then choose Enhanced Node View from the shortcut menu.
- Press Control-E.
- Open the Globals tab, then click enhancedNodeView.

Enhanced Node View Parameters

There are seven parameters in the enhancedNodeView subtree.

showTimeDependency

This parameter, when turned on, draws a bluish glow around nodes that are animated. This includes nodes that consist of *FileIn* nodes that reference a QuickTime movie or multiple-image sequence, nodes containing keyframed parameters, or nodes utilizing expressions that change a parameter's value over time. In the following screenshot, the alien *FileIn* node is highlighted because it is a multi-frame animation. The *Stabilize1* node is highlighted because it contains motion-tracking keyframes. The *Ramp1* node is not highlighted because it's only a single image, and the *Rotate1* node is not animated.
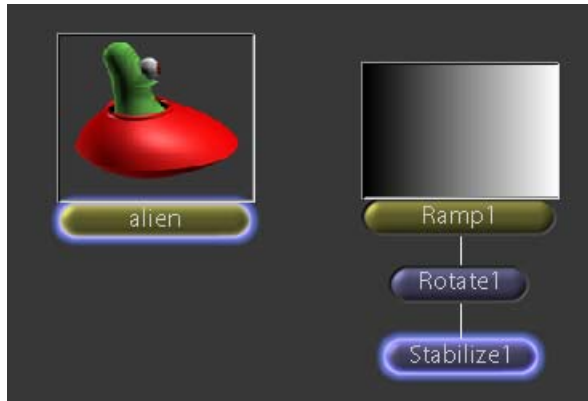
showExpressionLinks

Turning this parameter on draws a light purple line connecting nodes that use expressions to reference other nodes. An arrow pointing from the linked node toward the referenced node indicates their relationship. In the following screenshot, the *Move2D1* node contains expressions that link it to both the *Rotate1* and *Pan1* nodes.
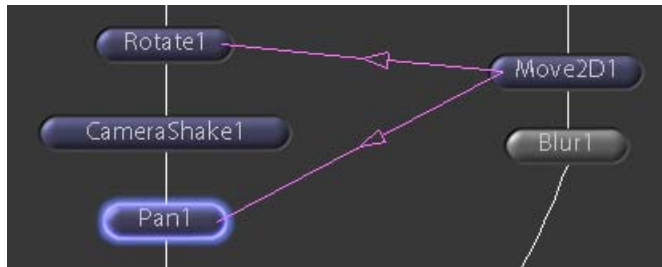
Note: When you clone a node by copying it and then pasting it with the Paste Linked command, the resulting cloned node displays an expression link arrow when showExpressionLinks is turned on.

ShowConcatenationLinks

When this parameter is turned on, a green line connects a series of nodes that concatenate. For example, three transform nodes that have been added to a node tree in sequence so that they concatenate appear linked with a green line connecting the left edge of each node. As a result, nodes that break concatenation are instantly noticeable. In the following screenshot, the *Rotate1*, *CameraShake1*, and *Pan1* nodes are concatenated.
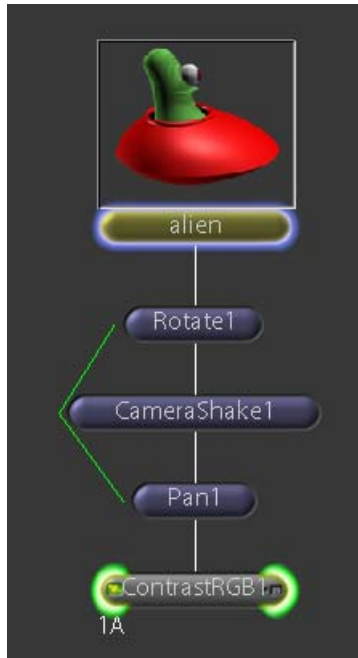
Note: As is often repeated, node concatenation is a very good thing to take advantage of. You are encouraged to group nodes that concatenater whenever possible to improve the performance and visual quality of your output.

Noodle Color Coding

Turning on Shake's noodle color coding parameters provides an additional way to visually distinguish the bit depth and channel information of the image data flowing down your node tree. With color coding turned on, you can see where in the node tree the bit depth is promoted or reduced, which noodles contain RGBA channel data versus BW data, and so forth.

Note: The Node View redraw speed of extremely large scripts may be reduced with noodleColorCoding turned on.

There are two kinds of coding used to identify the image data that noodles represent.

stipple8Bit, stipple16Bit, stipple32Bit
The stipple pattern of a noodle indicates its bit depth. In the following screenshot, three renamed *Bytes* nodes output 8-, 16-, and 32-bit float data respectively. The stippling indicates each bit depth at a glance.
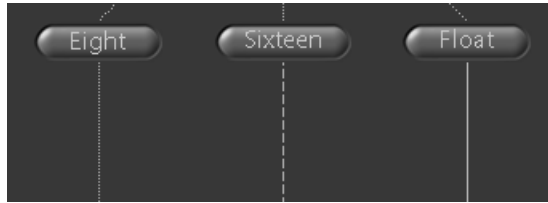
noodleColor Coding
Different combinations of color channels being propagated down the node tree are represented by different colors. In the following screenshot (showing four renamed *Reorder* nodes), the first node is disconnected from any incoming image data, while the next two are respectively propagating BWA and RGB channels. The color channels represented by each noodle are clearly distinguishable (this screenshot is viewable in color via the onscreen help).
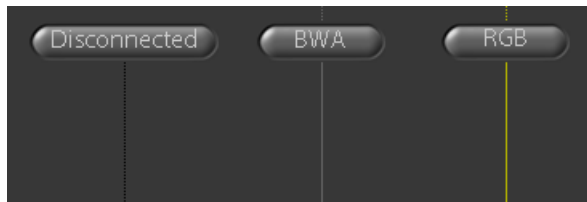
Customizing Noodle Color Coding

In the noodleColors subtree of the colors subtree, there are 12 different parameters for noodle color coding, corresponding to each possible combination of color channels in Shake. Each combination has a color control you can use to set up whichever color scheme makes the most sense to you.

The top parameter, noodleColor, lets you change the base color of noodles in the Node View. Noodles are white by default, but you can use the color control to change this to anything you like.

noodleColors				
noodleColor	⬜	1	1	1
noodleColorA	⬜	1	1	1
noodleColorBW	⬛	0.15	0.15	0.15
noodleColorBWA	⬛	0.4	0.4	0.4
noodleColorRGB	🟨	0.7	0.7	0
noodleColorRGBA	⬜	0.7	0.7	0.7
noodleColorZ	🟦	0	0	0.4
noodleColorAZ	🟦	0	0	0.4
noodleColorBWZ	🟦	0	0	0.4
noodleColorBWAZ	🟦	0	0	0.4
noodleColorRGBZ	🟦	0.2	0.2	0.8
noodleColorRGBAZ	🟦	0.4	0.4	1
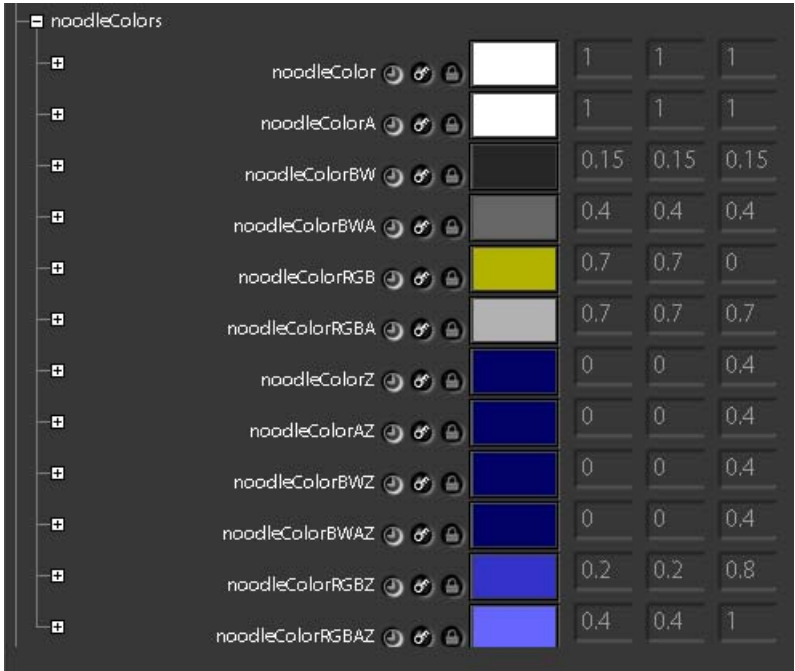
Customizing Noodle Stippling

You can customize the stipple patterns of noodles in the enhancedNodeView subtree of the Globals tab. You can also customize the colors used to identify noodle bit depth in the noodleColors subtree of the colors subtree of the Globals tab.

In the enhancedNodeView subtree, the stipple8Bit, stipple16Bit, and stipple32Bit parameters each have five different stipple patterns you can choose from, for maximum clarity.

Creating Custom Stipple Patterns

Different stipple patterns can be set in a .h preference file. Each stipple pattern is defined by a four-byte hex number that, when converted to binary, provides the pattern of the line drawn for each bit depth—each 1 corresponds to a dot, and each 0 corresponds to blank space.

For example, 0xFFFFFFFF is the hex equivalent of 1111111111, which creates a solid line. 0xF0F0F0 is the hex equivalent of 1111000011110000, which creates a dashed line.

The default settings are:

gui.nodeView.stipple8Bit = 0x33333333;
gui.nodeView.stipple16Bit = 0x0FFF0FFF;
gui.nodeView.stipple32Bit = 0xFFFFFFFF;

Resizing the Node Overview

The overview can be resized, making it easier to see.

To resize the node overview:

▪ Drag the upper-right corner, the top, or the right of the overview.
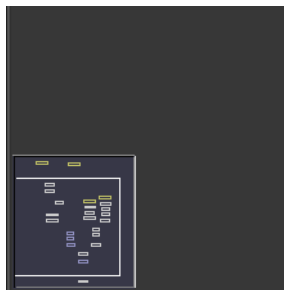
Default size
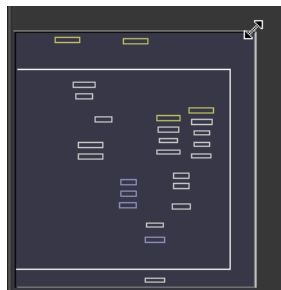
After resizing

New Ways of Adding and Connecting Nodes

There are several new ways of working in the Node View.

To add one type of node to multiple nodes in the Node View:

1 Select two or more nodes in the Node View.

2 In the Tool tabs, right-click the node you want to add, then choose Insert Multiple from the shortcut menu.

The new node is inserted after each selected node.

To connect one node to another by Shift-clicking:

1 Select the first node you want to connect.
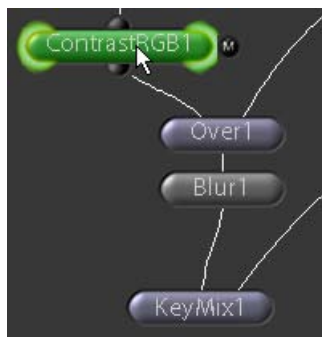
2 Move the pointer over the second node you want to connect so that the knot is visible, then Shift-click the knot to connect both nodes together.
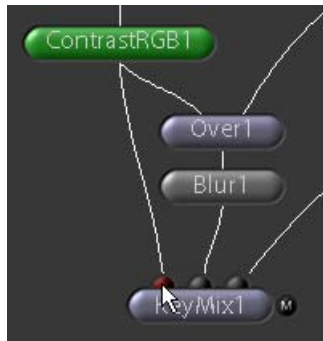
You can also use this technique to connect a group of nodes to a single multi-input node, such as the *MultiLayer, MultiPlane,* or *Select* node.

To connect several nodes to a multi-input node at once:

1 Select all of the nodes you want to connect to the multi-input node.

2 Shift-click the plus sign input of the multi-input node.
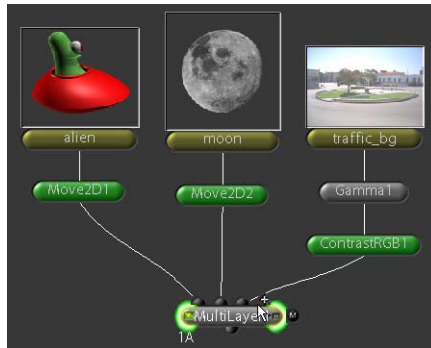
All selected nodes are connected.

To insert a disconnected node between two existing nodes:

1 Drag a node directly over a noodle so that both its knots appear highlighted. You can select which input of a multi-input node to insert by dragging the node to the left or right, so that the desired input knot is highlighted.
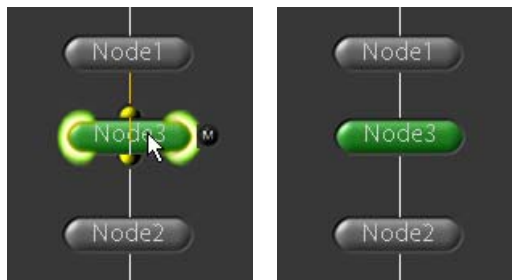
2 When you release the mouse button, the node is automatically inserted.

Cloned Node Naming

Cloned nodes are automatically named *OriginalNode_CloneN*, where *N* is the number of clones that have been made. You can clone a node as many times as you want, and all cloned nodes are linked to the original node. In the following screenshot, the *Gamma1* node has been cloned twice, in order to apply an identical gamma correction to the other two images in the tree.
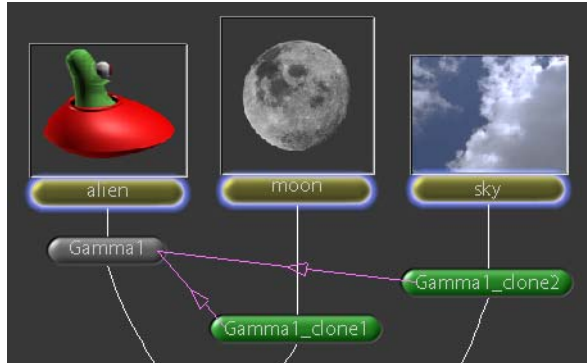
The links between cloned nodes can be viewed by turning on ShowExpressionLInks in the enhancedNodeView subtree of the Parameters tab. (and then turning on enhancedNodeView).

Monitor Calibration With Truelight

Shake includes Truelight, a color management system that lets you simulate on your Shake preview monitor, as closely as possible, the image that will eventually be printed to film or displayed on a high definition monitor or projector. This simulation is based on calibration data acquired from a variety of film stocks, film recorders, monitors, digital projectors, and film projectors, and on calibration profiles that you can generate yourself.

There are three parts to the Truelight tools:

• *TLCalibrate*, in the Other tab, is a utility node that you can use to accurately calibrate your monitor's color characteristics. This node allows you to create a calibration profile by eyeballing adjustments to a series of ten images using the controls within this node. Once you've made your adjustments, you can save this profile for future use.

 Note: This node also allows you to use calibration profiles generated by a Truelight Monitor probe.

- The calibration profiles generated using *TLCalibrate* can then be used by the Truelight VLUT (Viewer lookup table) control in the Viewer shelf.

The Truelight VLUT lets you previsualize the image in the Viewer as it will look after being output from a film recorder, or displayed by a high definition monitor or projector. You can use the Load Viewer Lookup command to make adjustments to the Truelight VLUT parameters in the Parameters tab, choosing which device profile you want to emulate.

- A second node in the Color tab, *Truelight*, performs the same function as the Truelight VLUT, except that it can be added to the node tree. The Truelight node has parameters that are identical to the Truelight VLUT that let you specify the device profile, current display profile, and color space for the preview. Additional controls let you fine-tune the preview.

Important: The Truelight functions are intended for previsualization only. They are not intended for use as color correction operations.

For more information on using the Truelight plugins, see the Truelight documentation, located in the Documentation folder on the Shake installation disk.

Updated Onscreen Controls for Transform Nodes
The pan, scale, and rotate onscreen controls have been updated, and have additional keyboard shortcuts for transforming an image without having to position the pointer directly over it.

Pan

The pan controls provide two ways to move the image within the Viewer. Drag anywhere within the image bounding box to reposition the image freely. Drag the vertical arrowhead to restrict position changes to the yPan parameter. Drag the horizontal arrowhead to restrict position changes to the xPan parameter.
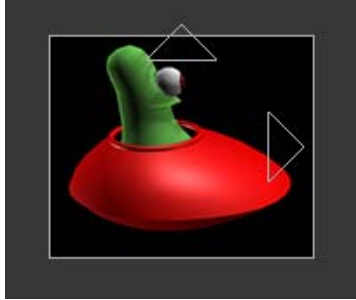
You can also press Q or P while dragging anywhere within the Viewer to pan an image without having to position the pointer directly over it.

Scale

Drag any of the corner controls to scale the entire image up or down, affecting the xScale and yScale parameters. Drag the top or bottom controls to scale the image vertically, affecting only the yScale parameter. Drag the left or right controls to scale the image horizontally, affecting only the xScale parameter.

Drag the center control to move the point about which scaling is performed, affecting the xCenter and yCenter parameters.
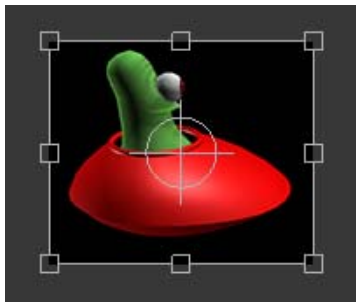
There is an additional method you can use to scale images in the Viewer.

To scale an image in the Viewer without using the scale handles:
1 Select an image.
2 With the pointer positioned over the Viewer, press E or I.

3 When the dimension pointer appears, drag in the direction you want to scale the image.

The colors in the dimension pointer correspond to the pan controls.
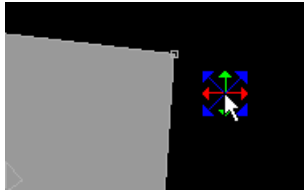
When you drag, the dimension pointer turns into an axis arrow, indicating the dimension in which you are scaling the image -- horizontal, vertical or diagonal.
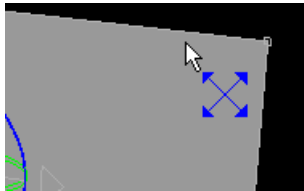
Rotate

Drag the blue rotate control to rotate the image about the center point, affecting the angle parameter. Drag the white center control to move the center point itself, affecting the xCenter and yCenter parameters.
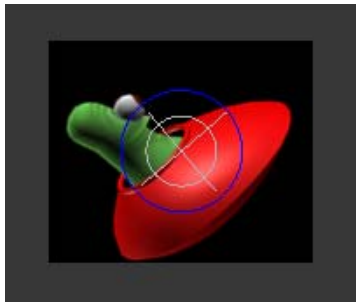
To rotate an image without positioning the pointer directly over it, press W or O, then drag in the Viewer. The dimension pointer appears, allowing you to to rotate the image in any direction.

Move2D

The *Move2D* node combines the onscreen transform controls for the *Pan, Scale*, and *Rotate* nodes into a single, all-purpose transformation node.
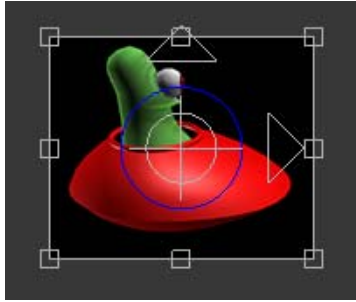
After an image is rotated with the *Move2D* node, the horizontal and vertical panning controls (arrowheads) lock movement to the new orientation of the image.
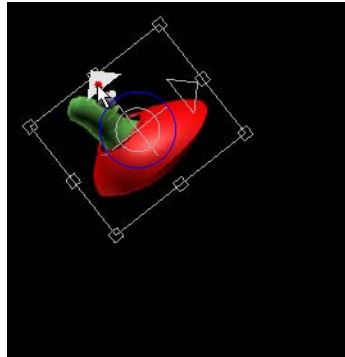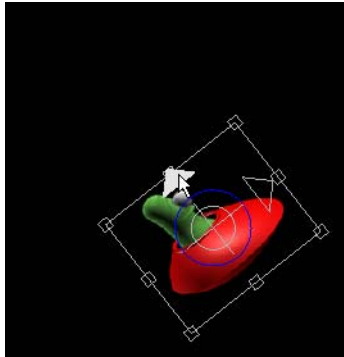
Move3D

Similar to the *Move2D* node, the *Move3D* node adds three colored dimensional angle controls to control xAngle (red), yAngle (green), and zAngle (blue) parameters. This lets you simulate 3D transformations with 2D images.
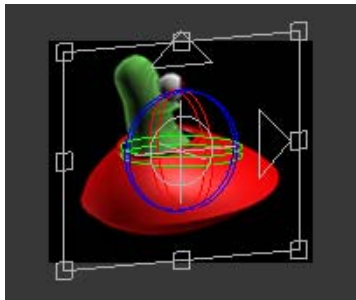
As with the *Move2D* node, when an image is panned with the *Move3D* node, the horizontal and vertical panning controls lock movement to the new orientation of the image, instead of the absolute orientation of the Viewer frame.

There are keyboard shortcuts (identical to those used within the *MultiPlane* node) that you can use to manipulate images without having to position the pointer directly over them.

Note: Unlike similar controls available for layers connected to the MultiPlane node, the constrained pan controls do not move the image forward or back if the yAngle is tilted. To manipulate images within a more fully featured 3D environment, connect them to a MultiPlane node, and use the available controls.

Major Features and New Nodes

2

This chapter covers major new features, and new nodes that have been added to Shake 4.

Support For the OpenEXR Format

OpenEXR (.exr) is an extremely flexible, cross-platform file format developed and maintained by Industrial Light & Magic. Key features of the OpenEXR format include support for the efficient storage of high dynamic-range image data using the 16-bit float "half" format, and support for auxiliary image data channels. OpenEXR 16-bit float and 32-bit float data channels can be read directly into Shake's RGBAZ data channels. In addition, OpenEXR 32-bit unsigned integer channel data can be read into Shake's Z data channel, although Shake's image processing nodes cannot process this data.

Note: 32-bit unsigned integer channel data will only be useful to custom plug-ins with built-in logic capable of processing the data within the Z channel.

A major feature of the OpenEXR format is its ability to support an extremely wide dynamic range. Thanks to its floating-point support, a contrast range of up to 30 f-stops can be supported with no loss of precision. Color resolution in 16-bit float ("half") files is 1024 steps per f-stop.

Another advantage of the OpenEXR format is support for any number of additional auxillary data channels, in addition to the standard RGBAZ channels. For example, additional channels can be written to store luminance, surface normal direction channels, velocity channels, and even individual lighting passes written from a 3D rendering application.

Shake Support for Auxiliary OpenEXR Data Channels

Internally, Shake only supports the processing of RGBAZ channels down the processing tree. However, the *FileIn* node provides channel remapping options in the Image tab of a *FileIn* node's parameters. Each channel that Shake supports has a corresponding pop-up menu. Each menu presents a list of every compatible channel within the referenced OpenEXR file. Color channels that correspond to Shake's supported channels are mapped by default.

FileIn channel pop-up menus for two files with different image channels

To remap any image channel:

1 Load a *FileIn* node that references an OpenEXR file into the Preferences tab.

2 Choose a new channel to map to from any color channel pop-up menu.

Channel remapping has the following restrictions:

• Any 16-bit or 32-bit float channel can be remapped within the RGBAZ channels.
• 32-bit integer channels can only be remapped to the Z channel.

   Note: If you want to access multiple 32-bit integer channels within a script, duplicate a number of *FileIn* nodes equal to the number of channels you want to access, then remap each 32-bit integer channel to a Z channel of one of the duplicate *FileIn* nodes.

Support for Data Compression

The OpenEXR format supports several codecs, with options for either lossless or lossy compression. Compression ratios range from 2:1 to 3:1.

Note: By default, *FileOut* nodes set to output OpenEXR images default to the Piz codec.

The following codec information appears courtesy of Industrial Light & Magic:

• *none:* No compression is applied.
• *RLE:* (Lossless) Differences between horizontally adjacent pixels are run-length encoded. This method is fast, and works well for images with large flat areas. But for photographic images, the compressed file size is usually between 60 and 75 percent of the uncompressed size.

- *ZIP:* (Lossless) Differences between horizontally adjacent pixels are compressed using the open source zlib library. ZIP decompression is faster than PIZ decompression, but ZIP compression is significantly slower. Photographic images tend to shrink to between 45 and 55 percent of their uncompressed size.
- *PXR 24:* (Lossy) After reducing 32-bit floating-point data to 24 bits by rounding, differences between horizontally adjacent pixels are compressed with zlib, similar to ZIP. PXR24 compression preserves image channels of type HALF and UINT exactly, but the relative error of FLOAT data increases to about 3°—10-5.This compression method works well for depth buffers and similar images, where the possible range of values is very large, but where full 32-bit floating-point accuracy is not necessary. Rounding improves compression significantly by eliminating the pixels' eightleast significant bits, which tend to be very noisy, and difficult to compress.
- *Piz:* (Lossless): This is the default compression method used by Shake. A wavelet transform is applied to the pixel data, and the result is Huffman-encoded.This scheme tends to provide the best compression ratio for typical film images. Files are compressed and decompressed at roughly the same speed. For photographic images with film grain, the files are reduced to between 35 and 55 percent of their uncompressed size.

OpenEXR Proxy Handling
Shake can read both tiled and scanline OpenEXR images. Scanline files contain a single image at a set resolution, but tiled files hold several versions of the same image at a variety of resolutions, for use as proxies in supporting applications.

Shake's proxy mechanism does not take advantage of tiled images. As a result, Shake defaults to reading in the highest available tiled resolution.

For More Information
More information about the OpenEXR format can be found at http://www.openexr.com.

Using Shake With Final Cut Pro 5
A new command in Final Cut Pro 5, *Send to Shake,* provides an automated way to move media back and forth between both applications. Using the Send to Shake command in Final Cut Pro 5 exports one or more selected clips into a Shake script, opening it immediately in Shake while Final Cut Pro 5 is running. When you do this, a placeholder is created in the originating Final Cut Pro 5 project file that automatically corresponds to the media that will be output from Shake.

Note: Each exported clip from Final Cut Pro 5 is brought into the Shake script using individual *FileIn* nodes. This is true even if two or more clips originate from the same master clip in the original Final Cut Pro 5 project.

For example, you can use Final Cut Pro 5 to superimpose a group of clips that you want to turn into a single composite using Shake. Final Cut Pro 5 makes it easy to set the In and Out points of each clip, and how they overlap. You can then send the media to Shake along with each shot's edit decision information, freeing you from having to reconstruct the media arrangement within Shake.

You can also move an entire sequence of clips into a Shake script. For example, you might do this to add operations to each individual clip in that scene to perform color correction, or keying.

Once you're finished in Shake, you can render the *FileIn* node that was automatically created when you used the *Send* command from Final Cut Pro 5, and easily relink the resulting media in the original Final Cut Pro 5 project.

How Sent Clips Are Arranged in Shake

Regardless of how you move Final Cut Pro 5 clips into Shake, how they're assembled in the newly created Shake script depends on whether they were sequentially arranged within a single video track, or vertically superimposed using several video tracks. Imported Final Cut Pro 5 clips are arranged within the node tree using *Select* and *MultiLayer* nodes:

- Clips edited sequentially on the same video track in Final Cut Pro 5 are connected to a single *Select* node when exported to Shake. The *Select* node switches between clips at their In and Out points, reflecting the editing decisions made on the track in Final Cut Pro. If the clips were originally superimposed across multiple video tracks, each video track that contains a clip results in a corresponding *Select* node being created in the Shake script. All clips that were edited into the same video track are connected to the same *Select* node.

  Note: The actual edit points for each *FileIn* node attached to the *Select* node are stored within the branch parameter. The data stored within this parameter is *not* intended to be editable; any attempt to do so will disrupt the edit points of the affected nodes.

- All the *Select* nodes are connected to a single *MultiLayer* node, which determines which clips are in the foreground of the composition, and which are in the background. Their arrangement reflects the arrangement of video tracks in the original Final Cut Pro sequence.

For example, if you used the Send to Shake command on the following three sequentially edited clips:
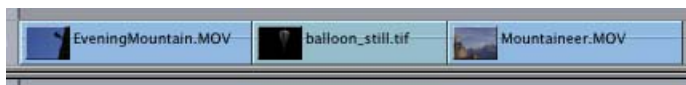
Sequentially edited clips in Final Cut Pro

The result would be the following Shake script, with one *Select* node and one *MultiLayer* node.
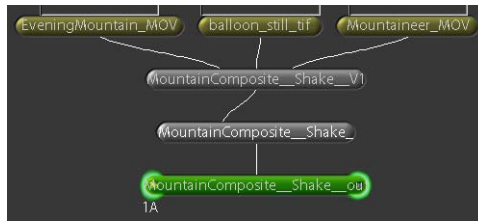
Resulting arrangement in Shake

If you used the Send to Shake command on the following superimposed clips:

Sequentially edited clips in Final Cut Pro 5

The result would be the following Shake script, with three *Select* nodes, and one *MultiLayer* node:
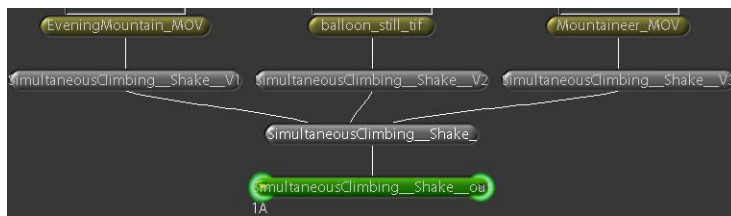
Resulting arrangement in Shake

While it is possible to slide footage within edits by adjusting the placement of imported clips in Shake's Time View, you are better off making these adjustments in Final Cut Pro and re-sending the media to Shake. Shake's Time View makes it difficult to determine whether there is sufficient underlying footage to prevent gaps in the sequence.

Warning: Audio clips and tracks from the original QuickTime files are not imported into Shake. Any timing changes you make in Shake will result in the adjusted clips going out of sync with the audio in the originating Final Cut Pro 5 project file.

Unsupported Media and Effects

Since QuickTime is the file format used for all media exchange between Final Cut Pro 5 and Shake, the following media and settings are not imported into Shake from Final Cut Pro 5:

- QuickTime audio tracks
- Standalone audio files
- Still image files
- Generators
- Composite modes
- Transformations (referred to in Final Cut Pro 5 as motion effects)
- Filters

Sending Clips From Final Cut Pro 5

If you want to send one or more selected clips (or a single sequence), from Final Cut Pro 5 to Shake, you should use the Send to Shake command in Final Cut Pro 5.

To send one or more clips from Final Cut Pro 5 to Shake:

1 Clean up your project timeline, so that you are able to select only the clips you intend to send.

2 Do one of the following:

- Select one or more clips in the Timeline or Browser.
- Select a sequence in the Browser.

3 Do one of the following:

- Choose File > Send > Send to Shake.
- Control-click (or right-click) the selected clips or sequence, then choose Send > To Shake from the shortcut menu.

4   When the Send to Shake dialog appears, select the appropriate options:
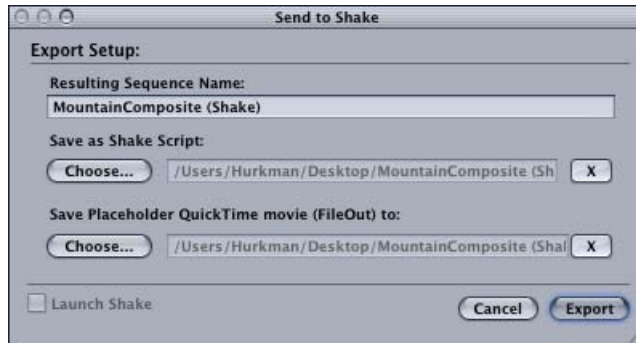
- *Resulting Sequence Name:* Type a name for the sequence that you've selected, prior to sending the media to Shake.
- *Save as Shake Script:* Type a name for the Shake script to be created, then click Choose to pick a location on disk to save it to.
- *Save Placeholder QuickTime movie (FileOut) to:* Type a name for the placeholder QuickTime movie that will correspond to the *FileOut* node in the newly created Shake script, then click Choose to pick a location on disk to save it to.

5   Check the Launch Shake box if you want to automatically open the newly created Shake script and start working on it.

6   Click Export.

When you click Export, four things happen:

- A duplicate sequence appears in your Final Cut Pro 5 project, containing duplicates of the selected media.
- A Shake project is created on disk.
- A placeholder QuickTime file is created on disk.
- The placeholder QuickTime file appears in a new video track that is created as the topmost track in your sequence (the original media remains where it was).
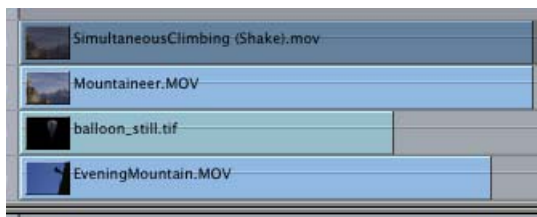
The placeholder QuickTime clip in your Final Cut Pro 5 project corresponds to the media that will eventually be rendered out of Shake—specifically, from the *FileOut* node appearing at the end of the generated Shake script.

The TimeRange of Scripts Generated From Final Cut Pro 5
The timeRange Global parameter in the Shake script that's created by the Send to Shake command is automatically set with the appropriate range of frames for the media it references.

Important: Clicking the Auto button to update the timeRange is not recommended. This can result in many more frames being referenced than expected, depending on the total duration of the source media files that are referenced.

Sending Media Back to Final Cut Pro 5
When you're finished working in the Shake script that was generated from Final Cut Pro 5, all you have to do is render the originally created *FileOut* node. The newly rendered media file takes the place of the original placeholder QuickTime file, ready for use by the original Final Cut Pro 5 project.

When you reopen the originating Final Cut Pro 5 project file containing the original placeholder QuickTime file, you'll need to use the Reconnect Media command to relink the clip in your Timeline to the media that was rendered out of Shake.

New Adaptive Method for Retiming
The retiming options in the FileIn node now accomodate high-quality adaptive interpolation of in-between frames. This option is available from the retimeMode pop-up menu.

Adaptive

This option in the retimeMode pop-up menu uses advanced image analysis to generate new in-between frames, creating seamless slow and fast-motion effects. Selecting Adaptive reveals additional parameters.

Fast versus Best Settings for Adaptive Retiming

When setting up an adaptive timing operation, you might be tempted to simply choose Best across the board for every parameter. This would probably be a mistake—producing dramatically longer render times in exchange for a potentially undetectable increase in quality, especially at higher resolutions.

That said, every clip's requirements are different. One group of settings is unlikely to produce equal results for shots with widely different subjects and movement. You're encouraged to do some limited tests prior to committing yourself to a particular group of retiming settings. As you experiment with different settings, be sure you always compare output from the Fast settings to that from the Better or Best settings, to make sure that it's worth committing yourself to the more computationally intensive Best settings.

- *Motion:* Two options determine the trade-off between image quality and processing time—Fast and Best. Fast makes motion interpolations using a mesh warp, and works well in most situations. Best, the most processor-intensive mode, uses motion vectors to track each pixel from frame to frame, interpolating new frames.
- *DeInterlacing:* Similar to the Motion parameter, two options let you determine the trade-off between image quality and processing time—Fast and Good. These settings represent two levels of mesh warping used to interpolate data from both fields of each frame. Here are some tips for using the DeInterlacing parameter:
  - Good is actually a better setting for frames without any moving subjects.
  - If you're working with standard definition video, there is no difference between using Fast or Good.
  - The DeInterlacing operation can also be used to improve clips that were poorly deinterlaced in other applications prior to importing into Shake.

- *AlwaysInterpolate:* With AlwaysInterpolate turned off, the final result of a retiming operation is a mix of original, unprocessed frames, and interpolated frames. For example, setting the speed to 0.5 to slow an image sequence down by 50 percent results in an alternating series of original frames and interpolated frames. In cases where there is a significant visual difference (softening, for example) between the unprocessed and reprocessed frames resulting from a retiming operation, the image sequence may appear to flicker. Turning on AlwaysInterpolate forces Shake to process every single frame in a retimed sequence, resulting in a series of frames with consistent visual quality.

If the Motion parameter is set to Best, three additional parameters become available:

- *BackwardFlow:* Turning on BackwardFlow evaluates the flow of images in both directions in order to generate interpolated in-between frames. This mode is usually visually superior, but significantly slower.
- *FlowSmoothness:* Higher or lower values may improve the quality of interpolated frames, depending on the shape of the subject in the image.
  - Use *low* values to improve the quality of subjects in the frame that change shape, like a person or animal that's running or jumping.
  - Use *high* values to improve the quality of static objects that don't change shape, such as trees, buildings, or cars.
- *FlowPrecision:* This is the last parameter you should adjust, after obtaining as much quality as possible with all of the above settings. Increasing the value of this parameter increases the overall precision of the Adaptive retiming operation by increasing the resolution at which the optical flow is estimated. A value of 0 is fine for most situations.

Remastering Media Options in the FileIn Node

The Convert option of the reTiming parameter provides a method for converting media from one format to another using advanced image processing to deinterlace, rescale and retime the incoming media. For example, if you have a high definition image sequence that you want to convert into a standard definition image sequence, or a PAL clip that you need to change to NTSC, the Convert option provides the tools to do so.

Choosing Convert reveals a series of parameters within the *FileIn* node that allow you to change the frame rate, resize the output resolution, anti-alias and sharpen the resulting image, and deinterlace the media being referenced by that *FileIn*. These options provide the highest-quality means of resizing and deinterlacing available in Shake, with results that are superior to the transform nodes that are available from the Tool tabs. These options are only available within the *FileIn* node.

You can use these options to convert individual shots that you're compositing within Shake, or you can read in an edited sequence from an application like Final Cut Pro 5 for format conversion using Shake.

Important: If you're converting a clip from a video frame rate to that of film with the intention of adding 3:2 pulldown back to the video (to achieve a film look for video), render the 24 fps conversion first. Add 3:2 pulldown to the shot in another operation by processing it in a second script, or by adding 3:2 pulldown with a command-line render.

Fast Vs. Best in Remastering Parameters
When setting up a remastering operation, you might be tempted to simply choose Best across the board for every parameter. This would probably be a mistake—producing dramatically longer render times in exchange for a visually undetectable increase in quality, especially at higher resolutions.

That said, every clip's requirements are different. One group of settings is unlikely to produce equal results for shots with widely different exposures, grain, and camera movement. You're encouraged to do some limited tests prior to committing yourself to a particular group of mastering settings. As you experiment with different settings, be sure you always compare the output from the Fast settings to that of the Better or Best settings, to make sure that it's worth committing yourself to the more computationally intensive Best settings.

Automatic Scene Detection for Multiple Shots
If you're reading in a sequence of pre-assembled shots, the remastering operators in Shake use automatic scene detection to eliminate artifacts at the frame boundaries between shots. This edge detection works well for cuts and dissolves. but other types of transitions may produce unwanted artifacts.

Convert Parameters

The Convert mode has the following parameters:
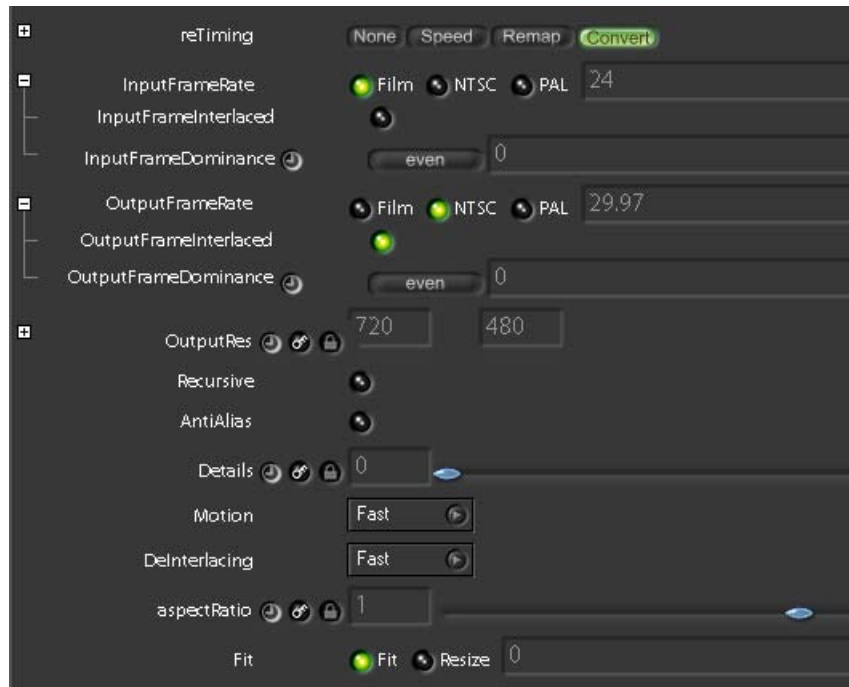
InputFrameRate
Specify the original frame rate of the input media here. This parameter is also a subtree with two additional subparameters.

InputFrameInterlaced
If the input media is video, enable this parameter if it's interlaced.

InputFrameDominance
If the input media is interlaced, specify the field dominance here.

OutputFrameRate
Specify the output frame rate here for format conversion. This parameter is a subtree with two additional subparameters.

OutputFrameInterlaced
If you want interlaced video output, turn this parameter on. Leaving it off results in Shake outputting progressive-scan (non-interlaced) media.

OutputFrameDominance
If OutputFrameInterlaced is turned on, specify the field dominance of the output image here.

OutputRes
Two fields where you enter the horizontal and vertical output resolution you want to convert the media to, to scale it up or down. Scaling an image sequence using the OutputRes parameter of the Convert options results in higher-quality output than using Shake's *Transform* nodes.

Recursive
Turning this parameter on enables a different resizing method, which can be sharper when enlarging some kinds of images. Try it on one or more representative frames to see if it helps.

Note: The recursive setting may also enhance unwanted noise, depending on the image.

AntiAlias
Turning this parameter on improves the quality of conversions when you're scaling media up. For example, when converting standard definition video to high definition, turning on AntiAlias smooths out jagged edges that might appear in the image.

Details
A built-in sharpening control that lets you add detail back to an image being enlarged. Unlike other sharpening operations, the details setting is able to distinguish between noise and feature details, and generally doesn't increase unwanted grain. Increasing this parameter may introduce jagged edges, however, which can be eliminated by turning on the AntiAlias parameter.

Motion
Two options determine the trade-off between image quality and processing time. Fast makes motion interpolations using a mesh warp, and is generally the only setting necessary for purposes of remastering.

Note: Best, the most processor-intensive mode, uses motion vectors to track each pixel from frame to frame to interpolate new frames, and should only be necessary when doing retiming. For more information on the parameters available for the Best setting, see page 37.

DeInterlacing
Similarly to the Motion parameter, two options let you determine the trade-off between image quality and processing time—Fast and Good. These settings represent two levels of mesh warping used to interpolate data from both fields of each frame.

AspectRatio

This parameter is a multiplier that allows you to convert pixels of one aspect ratio into another aspect ratio—for example, from NTSC to PAL, or from high definition (square) to NTSC. The default value of 1 makes no change to the output. The following table contains common conversion values:

Operation	Conversion Value
Square to NTSC (4:3)	0.9140, or 1 if Fit is set to Resize
Square to PAL (4:3)	1.1111, or 1 if Fit is set to Resize
NTSC (4:3) to Square (4:3)	1.1111, or 1 if Fit is set to Resize
NTSC (4:3) to PAL (4:3)	1, set Fit to Resize
PAL (4:3) to Square	.9375, or 1 if Fit is set to Resize
PAL (4:3) to NTSC (4:3)	1, set Fit to Resize

Note: In some conversion cases, AspectRatio may be left at 1 if the Fit parameter is set to Resize, which rescales the image horizontally and vertically to match the new OutputRes.

Fit

The Fit parameter determines how an image fits into the new frame that's determined by the OutputRes parameter if the aspect ratio of the *FileIn* image is different than that of the final resolution set by the OutputRes parameter. There are two options:
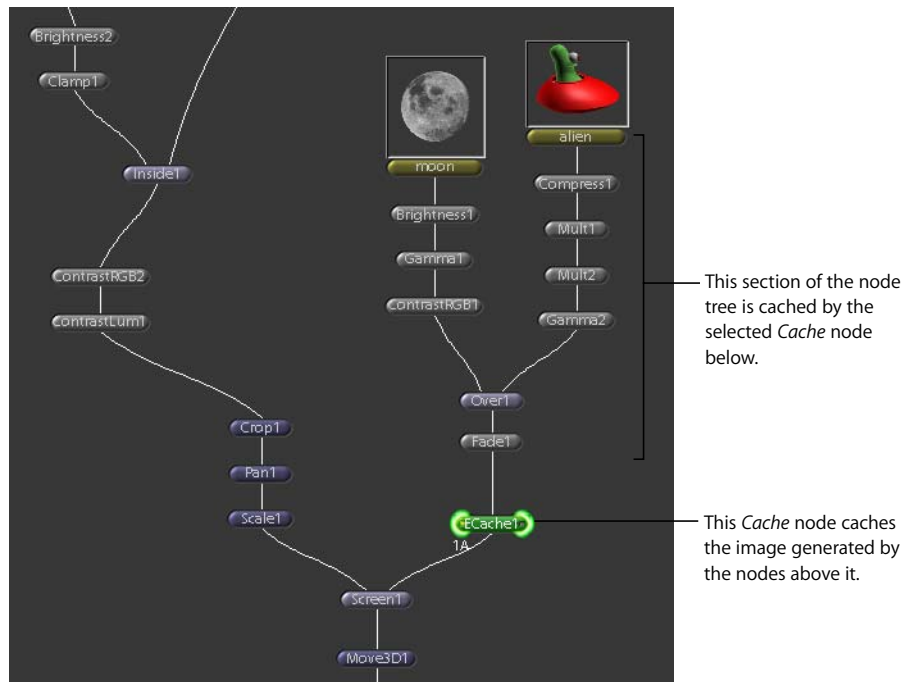
- *Fit:* Enlarges the image until either the vertical or horizontal resolution of the image matches that of the outputRes, depending on which is greater. This option maintains the original aspect ratio of the image, enlarging it as much as possible and leaving black to fill in the resulting gaps at the sides, or the top and bottom.
- *Resize:* Rescales the vertical and horizontal resolution of the image so that it fits within the entire frame of the OutputRes. The image is stretched as necessary to fit the new resolution.

The Cache Node

The *Cache* node lets you tell Shake to cache image data at specific points in the node tree. This gives you explicit control over which parts of the node tree require rendering while you work. For example, if there is a processor-intensive branch of your node tree that you're no longer working on, you can insert a *Cache* node in between the last node of that branch and the section of the node tree in which you're now working. Afterwards, the currently displayed frame is immediately cached. If you want to cache a range of frames in order to pre-render that branch of the node tree, you can use the Render Cache Nodes command. All cached image data is stored within the same cache, in memory or on disk.

Note: *Cache* nodes cache image data at the currently set proxy resolution.

From that point on, Shake references the cached image data within that node, instead of constantly reprocessing the nodes that precede it, unless a change is made to one of the preceding nodes.

This section of the node tree is cached by the selected *Cache* node below.

This *Cache* node caches the image generated by the nodes above it.

Important: Using a *Cache* node crops the image to the current image size, eliminating data from the Infinite Workspace from that point in the node tree on.

When the Cache Becomes Full
Cache nodes that can't be cached appear red in the Node View.

This node is able to cache. **This node is unable to cache.**

There are two possible situations when a *Cache* node won't be able to actually cache:

The input image size is larger than the maximum allowable cache file size.
You can easily tell if this is the case by opening the indicated *Cache* node into the Parameters tab, then checking to see if the value of the imageSize (an input image's Bit-depth * its Width * its Height) is larger than the value of the imageSizeLimit. If this is the case, you need to either increase the value assigned to the diskCache.cacheMaxFileSize global plug, or change the size of the incoming image.

The total cache memory limit has been exceeded.
The second possibility is that the amount of memory needed by all the *Cache* nodes in your script exceeds the memory assigned to the cache by the diskCache.cacheMemory global plug. In this case, no additional *Cache* nodes may be cached without increasing the diskCache.cacheMemory global plug.

Caching and Updating Frames

The *Cache* node updates whenever the playhead moves, caching additional frames if necessary because of changes that have been made in the preceding nodes. If necessary, you can also render one or more *Cache* nodes and cache a range of frames in advance using the Render Cache Nodes command.

If you later make changes to one or more nodes in a section of the node tree that's been cached, the affected cached frames are discarded, and can be re-cached.

To use the Cache node:

1 Insert a *Cache* node after the last node of a section of the node tree that you want to cache.

2 Load the *Cache* node's parameters into the Parameters tab.

3 Select an option from the forceCache parameter. The disk+memory option is the default forceCache setting, and is almost always the preferred setting to use.

4 If you want to immediately cache that section of the node tree for a specified duration, choose Render > Render Cache Nodes.

5 The Cache Render Parameters window appears, which automatically updates the timeRange to the Global timeRange.
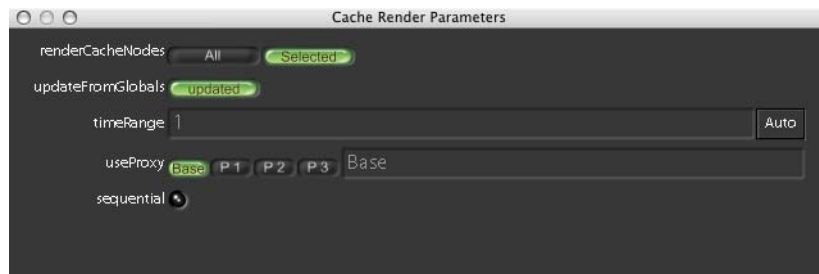
6 Click Render to render the *Cache* node. A Flipbook appears, allowing you to view the progress of the render, and play through the cached image sequence.

Parameters in the Cache Render Parameters Window
The Cache Render Parameters window has the following parameters:

renderCacheNodes
If you have multiple *Cache* nodes in the node tree, you can select one or more of these and render them simultaneously by setting renderCacheNodes to Selected in the Cache Render Parameters window. Or you can render all *Cache* nodes in the node tree by setting renderCacheNodes to All.

timeRange
If necessary, you can change the timeRange to cache a different frame duration.

useProxy
Images are cached using your script's current proxy setting. You can manually override the proxy setting in the Cache Render Parameters window, but those cache files won't actually be used by Shake until you change the script's proxy setting to match. This gives you the option to render multiple sets of cache images to match each proxy setting you plan on using.

sequential
Turning sequential on causes each *Cache* node to process the tree above it for each frame *before* allowing the next *Cache* node to process. When sequential is turned off, all *Cache* nodes are rendered simultaneously. This is more efficient in cases where *Cache* nodes share upstream nodes/trees. However, if there are too many processes running at the same time they will compete for CPU and memory resources, which may cause the overall processing time to increase.

Parameters in the Cache Node
The *Cache* node has the following parameters:
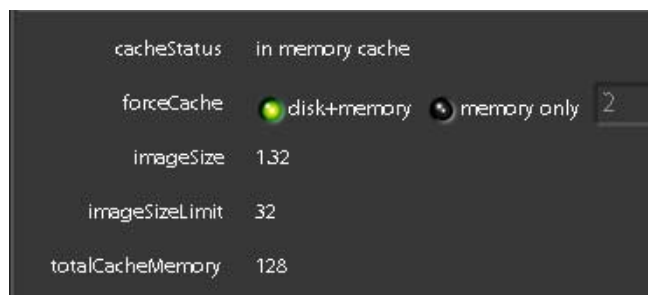
cacheStatus

This is a display-only parameter that shows whether the input image has been cached or not.

- *not cached:* Nothing has been written to the cache.
- *in disk cache:* Input image data has been moved to the disk cache. This is a result of the memory cache becoming full, or cache images having been saved after exiting a previous Shake session.
- *in memory cache:* The input image data has been written to the memory cache.
- *in transient memory cache:* The input image data has been written to the transient memory cache.

forceCache

This parameter lets you set how cached image data is stored when you update the cache with the Render Cache Node command. The selected forceCache behavior bypasses the Global cacheMode caching behavior. There are two options:

- *disk+memory:* The input image is written to the memory cache whenever the cache node is updated, and then transferred to the disk cache when the memory cache is full. All frames in the memory cache are moved to disk when Shake quits. In most cases this is the preferred behavior.
- *memory only:* Moves the input image into the memory cache every time the *Cache* node is updated, but never writes cache data to disk.

Internal Cache Parameter Display

The *Cache* node also displays the following parameters:

imageSize

The size (in megabytes) of the input image. The imageSize is determined by the following formula:

```
Bit-depth * Image Width * Image Height
```

imageSizeLimit

The maximum allowable size of the input image, in megabytes. This is set with the diskCache.cacheMaxFileSize global plug. The default value is 32 MB.

Note: The imageSizeLimit display lets you easily spot situations where a *Cache* node is not rendering because the imageSize is greater than the imageSizeLimit.

totalCacheMemory

The total RAM (in megabytes) available for caching to memory. This is set with the diskCache.cacheMemory global plug. The default value is 128 MB.

New Parameter in the Mask Subtree of Nodes

A new parameter has been added to help you manage mask input.

clampMask
Turning this parameter on clamps mask image data to a value between 0 and 1.

New Features for RotoShape Nodes

RotoShape nodes, and other nodes that use shapes, have additional features.

Copying and Pasting Shapes Between Nodes

There are several nodes that use shapes besides the *RotoShape* node. These include:

- *LensWarp*
- *Morpher*
- *Warper*

Shapes can be copied and pasted between all of these nodes, so that a shape drawn in one can be used in any other. Animated shapes are copied along with all of their keyframes.

Note: You cannot copy shapes from *RotoShape* nodes that were created in Shake 3.5 or earlier.

To copy a single shape:
- Right-click the transform control, outline, or any point of the shape you want to copy and do one of the following:
  - Choose Copy Shape from the shortcut menu.
  - Press Control-C.

To copy all visible shapes:
- Right-click anywhere in the Viewer, then choose Copy Visible Shapes from the shortcut menu.

  Note: Since this command only copies visible shapes, you can turn the visibility off for any shapes you don't want to copy.

To paste one or more shapes into a compatible node, do one of the following:
- Right-click anywhere within the Viewer, then choose Paste Shapes.
- Move the pointer into the Viewer, and press Control-V.

Cutting and Pasting RotoShape Keyframes

You can copy and paste rotoshape keyframes from one frame of the Time Bar to another. Whenever you copy a keyframe, you copy the entire state of that shape at that frame.

To copy a keyframe:

1 Move the playhead in the Time Bar to the frame where you want to copy the current state of the shape.

2 Right-click the transform control of the desired shape, then choose Copy Keyframe of Shape from the shortcut menu.

Note: You can copy the state of a shape at any frame, even if there is no keyframe there. Simply position the playhead anywhere within the Time Bar and use the Copy Keyframe command. That data can be pasted at any other frame as a keyframe.

To paste a keyframe:

1 Move the playhead in the Time Bar to the frame where you want to paste the copied keyframe.

2 Right-click the transform control of the desired shape, then choose Paste Keyframe of Shape from the shortcut menu.

Shape Timing

Three parameters within the timing subtree of the *RotoShape* parameters allow you to modify when a rotoshape starts and ends. An additional retimeShapes control lets you retime all keyframes that have been applied to that *RotoShape* node, speeding up or slowing down the animation that affects the shapes within.

timeShift

Offsets the entire rotoshape, along with any keyframes that have been applied to it. This parameter corresponds to the position of that rotoshape in the Time View.

inPoint

Moves the in point of the rotoshape, allowing you to change where that rotoshape begins. This parameter corresponds to the in point of the rotoshape in the Time View.

outPoint

Moves the out point of the rotoshape, allowing you to change where that RotoShape ends. This parameter corresponds to the out point of the rotoshape in the Time View.

Retiming RotoShape Animation

The retimeShapes button, within the timing subtree of the *RotoShape* Parameters tab, lets you retime all of the keyframes that are applied to that rotoshape.

Using this command, you can compress the keyframes that are animating a rotoshape, speeding up the changes taking place from keyframe to keyframe, or expand them, slowing the animation down.

When you click retimeShapes, the Node Retime window appears. The retimeShapes command has two modes, Speed and Remap, which affect a *RotoShape* node's keyframes similarly to the Speed and Remap options found within the Timing tab of a *FileIn* node.

Speed Adjustments

The default Operation, Speed, lets you compress or expand all of the keyframes within the *RotoShape* node by a fixed multiplier.

Suppose you have the following group of keyframes:

Using the default Amount value of 2.0 and clicking apply contracts the keyframes proportionally—resulting in the following distribution:


invert

Turning on the Invert button expands the keyframes by the Amount value, instead of contracting them. This has an identical effect to setting Amount to a decimal value between 0 and 1.

Remap Adjustments

Setting Operation to Remap provides a way for you to use curve expressions to retime the current keyframe distribution. This lets you apply a retiming curve from a *FileIn* node to the keyframes of a shape that you've already animated to rotoscope that image.

Note: If you want to create a curve specifically to use with the Node Retime command, you can create a local variable within the *RotoShape* Parameters tab, load it into the Curve Editor, and create a curve expression which you can then copy and paste from the local variable into the Retime Expr field.
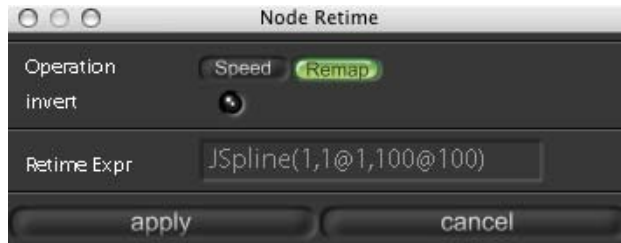
Attaching Trackers to Individual Control Points
You can attach separate trackers to each of a shape's individual control points. You can attach as many trackers to as many separate control points as you like.

To attach an existing track to a single control point:
1 Select a shape control point in the Viewer.
2 Right-click the selected points, then choose a tracker from the "Attach tracker to selected CV's in current shape" shortcut submenu.

   The selected tracker now animates the position of that control point. The offset between the position of the track and the original position of the shape control point is maintained as the point is animated along the path of the tracker.

   You can also create a new *Tracker* node that is referenced by a specific control point.

To create a new Tracker node attached to one or more control points:
1 Select one or more points in the Viewer.
2 Right-click one of the selected points, then choose "Create tracker for selected points" from the shortcut menu.
3 Attach the new tracker to the image you want to track, and use the tracker's controls to track the desired feature.
4 The control point you selected in step 1 automatically inherits the tracking data.

To remove a tracker from one or more control points:
1 Select one or more shape control points in the Viewer.
2 Right-click one of the selected points, then choose "Remove tracker reference on selected points" from the shortcut menu.

Importing and Exporting Shape Data

Two controls let you import and export shape data between Shake and external applications. These controls are located in the Viewer shelf when editing *RotoShape*, *Warper*, or *Morpher* nodes.

To export shape data:

- Click the Import Shape Data button in the Viewer shelf, choose a name and destination for the export file in the File Browser, then click OK.

To import shape data:

- Click the Export Shape Data button in the Viewer shelf, choose a compatible shape data file using the File Browser, then click OK.

To support this new feature, a new shape data file format has been defined—named SSF (Shake Shape File). This format standardizes the manner in which shape data is stored by Shake for external use. Before shape data can be imported from an external application, it must first be converted into the SSF format. For more information on the SSF format, see the *Shake 4 SDK Manual.*

The SmoothCam Node

This transform node differs from the other tracking nodes in that it doesn't focus the track on small groups of pixels. Instead, it evaluates the entire image at once, using advanced motion analysis techniques to extract transformation data.

Once this information is derived, this node has two modes. It can smooth the shot, eliminating unwanted jitter while maintaining the general motion of the camera, or it can lock the shot, stabilizing the subject. This node can affect translation, rotation, zoom, and perspective, making it more flexible for certain operations than the other tracking nodes.

The *SmoothCam* node is primarily useful for removing unwanted trembling from less than stable crane or jib arm moves, eliminating teetering from handheld walking shots, or reducing vibrations in automotive shots. Secondarily, the *SmoothCam* node can be used to stabilize shots that may be difficult to lock using the *Stabilize* node.

As useful as the *SmoothCam* node is, be aware that motion blur that is present in the image will remain, even though the subject in the shot is successfully smoothed or locked. This may or may not affect your approach to the composite.

Important: Interlaced images *must* be deinterlaced prior to use with the *SmoothCam* node.

Masking Important Features

The *SmoothCam* node has two inputs. The first one is for the input image to be processed. The second input is for an optional matte with which you can isolate a subject or area that you want the *SmoothCam* node to ignore while performing its analysis.

When creating a matte to use with the *SmoothCam* node, white areas are ignored, and black areas are analyzed.

Using the SmoothCam Node

The *SmoothCam* node works similarly to the *AutoAlign* node, in that the input image must be analyzed first in order to derive the data that is used to smooth or lock the image. This data is stored within the Shake script itself, so that each clip needs to be analyzed only once.

To analyze the media using the SmoothCam node:

1 Attach a *SmoothCam* node to the image you want to process.

2 Load the *SmoothCam* node's parameters into the Parameters tab.

3 If necessary, set the alalysisRange to the frame range you want to process.

   By default, the trackRange is set to the number of frames possessed by the media of the *FileIn* node to which it is attached.

4 Choose an analysisQuality mode. Start with Normal, which provides excellent results in most cases.

   Note: If, at the end of this process, the result is not what you'd hoped, then you should try again with the analysisQuality mode set to high. Be aware this will significantly increase the time it takes to perform this analysis.

5 Click the "analyze" button.

   The frames within the frame range are analyzed, and this data is stored within your script. To stop the analysis at any time, press Esc.

   After the analysis has been performed, it's time to choose the mode with which you want the *SmoothCam* node to process the image.

To smooth a shot using the SmoothCam node:

1 Once the analysis has concluded, set steadyMode to smooth.

2 Open the smooth subtree.

3 Adjust the translationSmooth, rotationSmooth, and zoomSmooth sliders to increase or decrease the amount of smoothing that is attempted. At 0, no smoothing is attempted in that dimension. At higher values, *SmoothCam* attempts to smooth a wider range of variations in the movement of the image from one frame to the next.

- *translationSmooth:* Smooths the X and Y motion of the shot.
- *rotationSmooth:* Smooths rotation in the shot.
- *zoomSmooth:* Smooths a zoom within the shot.

Important: It's faster, and will yield more accurate results, if you set the smoothing parameters of dimensions in which you know the camera is *not* moving to 0. For example, if you know for a fact that the camera is neither rotating nor zooming, set rotationSmooth and zoomSmooth to 0. In particular, if the camera is not zooming, increasing the zoomSmooth value may actually cause unwanted transformations in the image.

To lock and match shots using the SmoothCam node:

1 Once the analysis has concluded, set steadyMode to lock.

2 Open the lockDown subtree.

3 The setting of the inverseTransform parameter depends on what you're trying to do:

- If you want to lock the analyzed shot itself, leave inverseTransform set to lock.
- If you want to apply the motion from this shot to another image, attach this node to the other image, then set inverseTransform to match.

4 Turn on the locking parameters for the dimensions of motion you want to affect.

- *translateLock:* Locks the subject in the X and Y dimensions.
- *rotateLock:* Locks the rotation of the subject.
- *zoomLock:* Removes zoom from the shot, maintaining the relative size of the subject.
- *perspectiveLock:* Locks the perspective of the subject, performing the equivalent of a four-corner warp to maintain the size and shape of the subject.

Troubleshooting SmoothCam Effects

If the output is unsatisfactory, there are several things you can try to improve the result.

Change the smoothing parameters

If you're trying to smooth the motion in a shot, you should first try changing the smoothing parameters. This can be accomplished without having to reanalyze the clip.

Reanalyze the media at a higher quality

Next, try changing the analysisQuality to high, and reanalyze the media.

Try editing the analysis data
If neither of the prior solutions helps, try loading the confidence parameter into the Curve Editor, then look for frames where the confidence parameter falls to 0. If the image transformation at these frames stands out, you can try loading the translationX, translationY, rotation, and/or zoom parameters within the motion subtree into the Curve Editor, then delete any keyframes that create unusual spikes at those frames.

Removing Black Borders Introduced by Smoothing and Locking
When you use the *SmoothCam* node, the resulting transformations that are made to the input image either to smooth or lock the shot cause moving black borders to appear around the edges of the output image. While this is necessary to achieve the desired effect, you probably don't want these black borders to appear in the final shot.

There are several ways you can choose to handle this border.

Using the clipMode Parameter
The clipMode parameter provides different ways you can treat the size of the output frame in order to include or exclude this black region. After you've analyzed the input image and picked the settings necessary to smooth or lock your shot, you can choose one of three clipMode parameters.

union
This option expands the rendered frame to include the full area within which the input image is transformed. If you scrub through the *SmoothCam* node's output with this option turned on, the image appears to float about within a black frame larger than itself. The frame size of the output image is larger than that of the input image.

intersection
This option contracts the rendered frame to exclude any black area. The result is a stable output image that fills the frame, but with a frame size that's smaller than the input image.
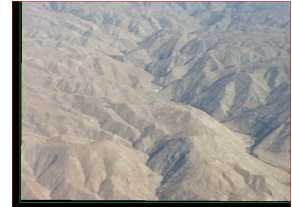
in

This option maintains the frame size of the input image as that of the output image. The result is a moving black area that encroaches around the edges of the output image.

Union   Intersection   In

You can use one of the above three clipMode options to produce the output image most useful for your purposes.

Note: Whichever clipMode you use, areas of the image that end up being clipped are preserved by Shake's Infinite Workspace, available for future operations.

Scaling the Output Image to Fit the Original Frame Size

If you need to output the resulting image at the same size as the original, the quickest fix is to leave the clipMode Parameter set to "in," and use a *Scale* node after the *SmoothCam* node to enlarge the image to the point where all instances of black borders fall outside the edge of the frame. The disadvantage of this method is the resulting softening of the image, depending on how much it must be enlarged.

Painting the Edges Back In

A more time intensive, but possibly higher-quality solution would be to paint the missing edge data with a *QuickPaint* node. A variety of techniques can be employed, including using the clone brush to sample other parts of the image, or using the Reveal Brush to paint in parts of a secondary clean-plate image that you create separately.

Warping the Edges

One last suggestion would be to experiment with the different warping nodes, to stretch the edges of the image to fill any gaps. For example, you can experiment with the *LensWarp* node to stretch the edges of the image out. This solution is highly dependent on the type of image, and may introduce other image artifacts that may or may not be acceptable.

Parameters in the SmoothCam Node

This node displays the following controls in the Parameters tab:
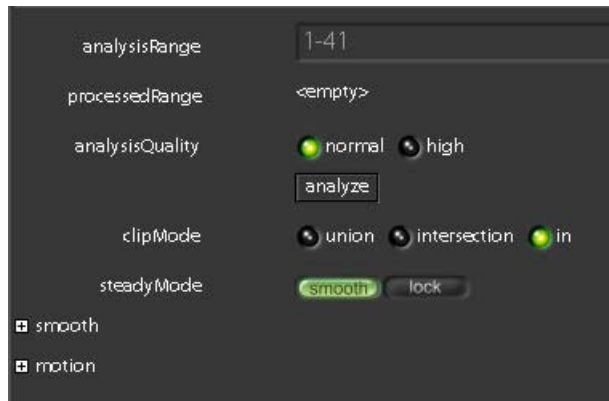
analysisRange
The range of frames of the input image to analyze. By default, this parameter inherits the number of frames in the source media file represented by the *FileIn* node to which it's connected, not the timeRange in the Globals tab.

processedRange
This is not a user-editable parameter. It indicates the overall frame range that has been analyzed.

analysisQuality
Defines the level of detail for the motion analysis. There are two levels of quality, "normal" and "high." In most situations, "normal" should produce an excellent result. Should you notice any artifacts, set this parameter to "high."

analyze
Click this button to analyze the frame range defined by the analysisRange parameter. Pressing Esc while Shake is analyzing stops the analysis.

clipMode
The clipMode parameter provides different ways you can treat the size of the output frame in order to include or exclude black regions introduced by the *SmoothCam* node's transformations. The clipMode parameter has three options:
- *union:* Expands the rendered frame to include the full area within which the input image is transformed. If you scrub through the *SmoothCam* node's output with this option turned on, the image appears to float about within a black frame larger than itself. The frame size of the output image is larger than that of the input image.

- *intersection:* Contracts the rendered frame to exclude any black area. The result is a stable output image that fills the frame, but with a frame size that's smaller than the input image.
- *in:* Maintains the frame size of the input image as that of the output image. The result is a moving black area that encroaches around the edges of the output image.

steadyMode

The *SmoothCam* node has two modes: Smooth and Lock.

- *Smooth:* This mode smooths the apparent motion of the camera, while allowing the general movement in the frame to proceed. It's useful for removing jitter from a camera movement. When enabled, this mode has three sliders for each of the dimensions that can be smoothed.

The amount of smoothing can be increased or decreased along a sliding scale.

- *translationSmooth:* Smooths motion in both the X and Y dimensions.
- *rotationSmooth:* Smooths image rotation.
- *zoomSmooth:* Smooths an uneven zoom.

  Note: Don't turn on zoomSmooth unless you're positive that the image is being zoomed. Otherwise this parameter will not have the desired effect.

- *Lock:* This mode attempts to lock the motion of the principal subject in the shot to eliminate motion. As a result, the background will appear to move around the subject being tracked.

- *InverseTransform:* This control lets you invert the effect of this node, so that other images can be matchmoved using the same motion.
- *translateLock:* Locks the image in both X and Y dimensions.

- *rotateLock:* Locks the rotation of the image.
- *zoomLock:* Locks an image that is being zoomed.

Note: Don't turn on zoomLock unless you're absolutely positive that the image is being dynamically zoomed.

- *perspectiveLock:* Locks an image experiencing a change in perspective, similar to a reverse corner-pin.

Motion

The values within the parameters in the motion subtree are not meant to be editable, or even directly intelligible. They're available to be loaded into the curve editor so that you can find possible spikes in the analysis data that, along with a low confidence value, might indicate an error. In this case, you have the option of deleting the offending keyframe, in order to see if the resulting smoothing or locking operation becomes more acceptable.

- *confidence:* Unlike the tracking nodes, the confidence value is restricted to one of three possible values. You can load the confidence parameter into the Curve Editor to quickly find problem areas in the analysis. The values are:
  - *1:* Indicates an analysis in which Shake has high confidence. A keyframe is generated at these frames.
  - *0.5:* Indicates an uncertain analysis. Shake also generates a keyframe at these frames.
  - *0:* Indicates that Shake has no confidence in the analysis. No keyframe is generated.
- *translationX, translationY:* Contains the X and Y translation data.
- *rotation:* Contains the rotation data.
- *zoom:* Contains the zoom data.

The AutoAlign Node

The *AutoAlign* node is unique among the transform nodes in that it can take multiple image inputs and combine them into a single output, similar to a layering node. This node is designed to take two or three images or image sequences that overlap, and align them in different ways to create a single, seamless output image.

Unlike similar photographic tools, the *AutoAlign* node works with both stills and image sequences. As a result, you could film three side-by-side shots of an expanse of action, and later turn these into a single, extremely wide-angle background plate. Similarly, you can set a single still image to align with the same image features in a second shot with camera motion.

Stitching Images Together

When you connect two or three overlapping images or image sequences to the *AutoAlign* node, they're assembled into a single composite panoramic image. The input images can be either vertically or horizontally oriented—the *AutoAlign* node automatically determines the areas of overlap and aligns them appropriately.

Note: The order in which you connect images to the *AutoAlign* node doesn't matter; their position is automatically determined.

For example, if you connect the following three side-by-side shots of a skyline to the inputs of an *AutoAlign* node:
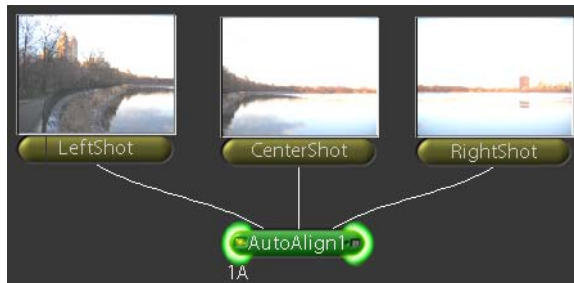
The overlapping areas are analyzed, matched, and the images are repositioned and manipulated to fit together.

As you can see in the above example, the resulting image may have an irregular border, depending on the amount and position of overlap, and the warping required to achieve alignment. If necessary, the border can be straightened with a Crop node.

Aligning Overlapping Images

You can also use this node to align two images that almost completely overlap. An obvious use of this is to align an image with something in the frame that you want to remove with a second "clean plate" image. After you've aligned the two images, you can use the aligned output image as a background plate for use in a paint or roto operation to remove the unwanted subject.

The primary advantage of the *AutoAlign* node in this scenario is that it allows you to precisely align a clean plate image to an image sequence in two difficult types of situations:

- If the clean plate image is from a different source, or has different framing or other characteristics that make it difficult to align using regular transformations.
- If the target image is in motion, and alignment difficult using regular trackers.

For example, if you have the following shot of a man jumping across a bottomless chasm, with the camera moving dramatically to follow the action, it may be difficult to line up a still image clean plate.

Moving camera with unwanted safety rigging

Still image clean plate

Using the *AutoAlign* node to align the second image with the first, you can quickly match the clean plate still to the moving image sequence, and then output the newly aligned and animated clean plate for use by a paint or rotoscoping operation to paint out the safety line.

AutoAlign Limitations

Successful use of the *AutoAlign* node is highly dependent on the content of the input images, and results will vary widely from shot to shot. In general, moving images with a great deal of perspective and parallax can be difficult for the *AutoAlign* node to analyze properly, and may cause unexpected results.

Blending at the Seams

Two parameters help to automatically manage the seams between images that have been aligned with the *AutoAlign* node:

- *blendMode* automatically adjusts the overlapping border between two aligned images, intelligently selecting which portions of each image to retain for a seamless result.

blendMode set to none   **blendMode set to smartBlend**

- *matchIllum* automatically adjusts the brightness of each aligned image so that the exposure matches.

AutoAlign Image Requirements

Although the *AutoAlign* node is a very flexible tool, it produces the best results with material that is produced with this node in mind.

- There should be at least a 15-to-20-percent overlap between any two images for the *AutoAlign* node to work properly (the amount that's necessary may vary depending on the image).
- You may still be able to successfully line up images even when limited regions within each image don't line up, although this may affect the overall quality of the resulting alignment. Examples include clouds, waves, pedestrians and birds. If you can arrange for it in advance, exact matches from simultaneously captured images is usually preferable.
- The warping performed by the *AutoAlign* node is designed to align panoramic images that have been shot by dollying the camera to the next position best. Multiple images created by panning the camera from a single position may result in undesirable curvature in the final result.
- Interlaced images *must* be deinterlaced prior to analysis.
- When aligning image sequences, the *AutoAlign* node aligns the input frames in time according to each image sequence's timeShift parameter, as represented in the Time View. Offsetting an image sequence in the Time View offsets which frames will be aligned with one another.
- When aligning overlapping images, an object or person in one image that is not in the other may prevent successful analysis if it occupies too large an area.
- Images with an animated width or height can not be aligned.

Image Stabilization Within the AutoAlign Node

You must select one of the input images to be used as the positional reference for *AutoAlign's* analysis, using the lockedPlate pop-up menu. If the shot you want to designate as the lockedPlate is not already steady, the other input images will be transformed to match the motion in this shot.

Example 1: A Procedure for Simple Image Alignment

This example covers how to set up a panoramic stitching operation using the *AutoAlign* node.

To automatically align two or three images:

1   In the Time View, adjust the offsets of the image sequences that you intend to use.

2   Create an *AutoAlign* node, and attach up to three images to it.

    The order in which you attach them doesn't matter.

3   Load the *AutoAlign* node's parameters into the Parameters tab.

4   If necessary, set the analysisRange to the frame range you want to analyze.

5 If you're analyzing images for the first time, set the mode parameter to Precise.

   If you're not satisfied with the result later in the operation, change mode to Robust and re-analyze the images.

6 Click the "analyze" button.

   Shake steps through each frame in the image sequences and analyzes each set of aligned images at every frame in the designated analysisRange.

   Note: The analysis can be interrupted at any time by pressing Esc.

7 Once the analysis has concluded, set clipLayer to All in order to expand the Region of Interest (ROI) to the overall width of the newly generated panorama.

8 Choose an input layer from the lockedPlate pop-up menu.

   For purposes of creating a panorama, you should use an image that's either already stable, or one that has been locked down using the *SmoothCam* or *Stabilize* nodes.

   Note: If the combined image is very wide, it may also be a good idea to set the centermost image as the lockedPlate.

9 To output the entire panorama as a single image, set ouputFrame to All.

10 If the seams between each aligned image are visible, set the blendMode parameter to smartBlend to make them invisible.

11 If the aligned images have different exposures, turn on matchIllum.

   If you can still see differences in color or exposure, you can attach additional color correction nodes between the upstream image and the *AutoAlign* node to make the necessary adjustments.

Example 2: Aligning a Clean Plate Image With a Moving Shot

This example shows how to align a clean plate still image with an image sequence in which the camera is moving.

To align a clean plate image with an overlapping image sequence:

1 Connect the clean plate image and the moving shot (in this example, the mountaineer image) to an *AutoAlign* node.

The order in which they are connected is not important.
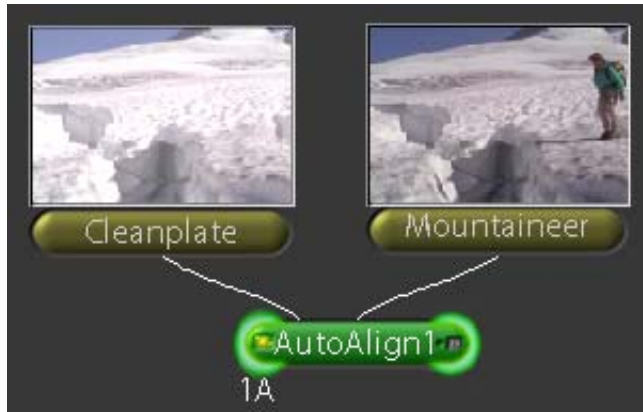
2 Use the clipLayer and lockedPlate pop-up menus to choose the input to which the clean plate image is connected.

In this example you'll be choosing Input1. Leave mode at the default setting of Precise.

3 Set the analysisRange to the number of frames you want to align.

This parameter defaults to the maximum number of frames within the longest image sequence that's connected to the *AutoAlign* node.

4 Click "analyze."
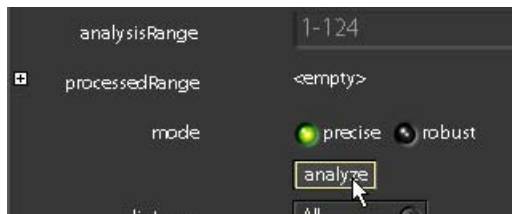
Shake begins the image analysis. The first frame may take longer than the other frames in the analysis. The analysis should speed up after the first frame. As the analysis is performed, the playhead steps through the Time Bar frame by frame, creating a keyframe at each analyzed frame.

Note: The analysis can be interrupted at any time by pressing Esc.

5 Once the analysis has concluded, changing blendMode to mix and scrubbing through the Time Bar shows you how well the resulting alignment works.

The Mix setting, in the case of two images that almost completely overlap, results in an a fifty percent blend of both images. In the above image, the ripples in the snow appear to align perfectly.

6 In order to use this result in a paint operation, set the following parameters:

a Set clipLayer to Input2, so that the moving shot defines the resolution of the output image.

b Set lockedPlate to Input2 as well, so that the clean plate image moves along with the background.

c Set outputFrame to Input1, and set blendMode to none, so that the only image that is output is the newly animated and aligned clean plate image.

The resulting output image is a transformed version of the clean plate image that matches the position of matching overlapping features in the mountaineer image.

Original clean plate image

Auto-aligned clean plate image

7 The auto-aligned clean plate can now be used in paint or rotoscoping operations to remove unwanted rigging from the actor in the mountaineer shot. For example, you can connect the original *Mountaineer* image to the first input of a *QuickPaint* node, and the output from the *AutoAlign* node to the second input.
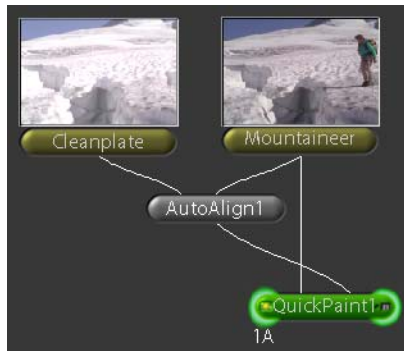
8 With this setup, you can use the Reveal Brush to paint out the rigging against the backdrop of the clean plate.

Original Image

Reveal Brush exposing auto-aligned image

AutoAlign Parameters

The *AutoAlign* node displays the following controls in the Parameters tab:
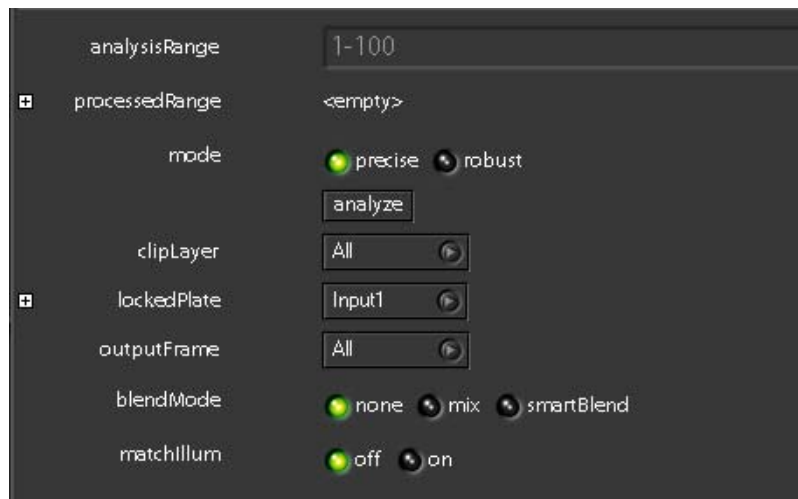
analyze
Click this button to perform the analysis that is the first step in aligning two or three images. This analysis creates a preprocessed data set that is used to perform the actual alignment. Images only have to be analyzed once, and the resulting transform data is stored inside the Shake script. The other parameters in the node act upon this transformation data—changing the parameters has no effect on the analysis data.

If, at any time, you're not satisfied with the results of the analysis, you can change the mode and re-analyze the analysisRange.

analysisRange
The range of frames you want to process. By default this parameter inherits the maximum number of frames in the *FileIn* node at the top of the tree.

processedRange
A non-editable display field that shows the number of frames that have already been analyzed. When you first create an *AutoAlign* node, this parameter reads "<empty>."

Opening the processedRange subtree reveals the following subparameters:

• *confidence:* The confidence value indicates the level of matching accuracy that resulted from the analysis. Each frame's confidence is restricted to one of three possible values. You can load the confidence parameter into the Curve Editor to quickly find problem areas in the analysis. The values are:
  • *1:* Indicates an analysis in which Shake has high confidence. A keyframe has been generated at these frames.

- *0.5:* Indicates an uncertain analysis. Despite the uncertainty, Shake has generated a keyframe at these frames.
- *0:* Indicates that Shake has no confidence in the analysis. No keyframe has been generated at these frames.
- *mMatrix, nMatrix:* These parameters contain the data accumulated by clicking the "analyze" button.

mode
Two options—"precise" and "robust"—let you change the method used to perform the analysis. In general, set this mode to "precise" the first time you analyze the input images. If the results are not satisfactory, change the mode to "robust" and re-analyze.

clipLayer
Lets you to choose which image defines the resolution of the output image—Input1, Input2, Input3, or All, which sets the resolution to the final size of the aligned and merged images.

lockedPlate
This parameter lets you choose which input image to use as a positional reference. The other two input images will be stabilized and aligned to match the position of the lockedPlate. The lockedPlate also affects how the other two images are warped to match the overall perspective of the final image. If the final image is going to be a wide-angle shot, it may be a good idea to set the center image as the lockedPlate.
- *lockedPlateOffsetX, lockedPlateOffsetY:* These subparameters of lockedPlate contain the transform keyframes used to align the plates. Changing the lockedPlate also changes the keyframe set displayed by these parameters.

outputFrame
Lets you choose which image is output by the *AutoAlign* node. The options are identical to those of the clipLayer parameter.

blendMode
This parameter provides options for automatically blending the seams where the input images meet. There are three choices:
- *none:* The seams are not blended.
- *mix:* Both sides of each seam are simply feathered together.
- *smartBlend:* The highest quality mode. The best pixels from either side of each seam are used to patch together a seamless image.

matchIllum
Turning this parameter on automatically adjusts the brightness of the input images so that the lighting is consistent across the entire output image.

Note: If necessary, you can preprocess images connected to the *AutoAlign* node with other color correction nodes to even out differences in gamma and contrast that aren't addressed by the matchIllum parameter.

The LensWarp Node

This node lets you make subtle or large adjustments to an image to either correct for, or simulate, different types of film and video lenses. As a corrective measure, you can use this node to remove barrel distortion in an image. You can also use this node to simulate the lens used in one image, in order to warp another image you're compositing over it so that they appear to have been shot using the same lens. This node affords more appropriate control over the result than the *PinCushion* node.

Before   **After LensWarp undistort**

There are two ways you can use the *LensWarp* node.

Manual Adjustments

You can manually adjust the factor and Kappa values to create the desired effect. If you're making a simple correction, this may be the fastest way to proceed.

Automatic Calibration

You can also draw shapes along the edge of any distorted features within the image that are supposed to be straight, and use the calibrate button to automatically correct the image. This provides the most accurate result, especially for moving images.

To automatically calibrate the LensWarp node:

1 Attach the *LensWarp* node to an image with wide-angle distortion.

   By default, the *LensWarp* node is in Add Shapes mode.

2 In the Viewer, identify one or more curved features that should be straight, then trace each curve with an open shape.

   Important: Make sure the *LensWarp* node is turned off in the Parameters tab before adding calibration shapes.

Each feature should be traced with at least four points. The more curved a feature is, the more points you should use to trace it with. On the other hand, you should rarely need to use more than ten points per feature.

To finish drawing an open shape, double-click to draw the last point and end the shape.

Note: The *LensWarp* node does not use Bezier curves.

3 If you're correcting an image sequence, scrub through the frame range to find more curved features, then trace these as well.

   The more features you identify in different areas of the frame, the more accurate the final result will be.

4 When you're finished, click the "analyze" button in the *LensWarp* parameters to calculate the result.

   The analyze state changes to "analysis up to date."

5 Do one of the following:
   • If you're correcting distortion in an image, set the *LensWarp* node to Undistort.
   • If you're applying the distortion from this image to another one, detach the *LensWarp* node, reattach it to the other image, and set the *LensWarp* node to Distort.

LensWarp Viewer Shelf Controls
This node has a subset of Shake's standard shape-drawing tools that you can use to identify distorted features for automatic calibration.

Button		Description
	Add Shapes	Click this button to add shapes with which to trace distorted features that you want to correct. You can only draw polygonal shapes with the *LensWarp* node.
	Edit Shapes	Click this button to edit shapes that you've made.
	Import/Export Shape Data	Lets you import or export shapes.

Button		Description
	Delete Control Point	Select a point and click this button to remove it from the shape.
	Enable/Disable Shape Transform Control	Lets you show or hide the transform control at the center of each shape. Hiding these controls will prevent you from accidentally transforming shapes while making adjustments to control points.

Parameters

This node displays the following controls in the Parameters tab:

analyze

After you've created one or more shapes to trace distorted features within the image in the Viewer, clicking calibrate automatically sets the factor and Kappa parameters to the values necessary to either correct or match the lens distortion.

analyze state

This parameter indicates whether or not the analysis has been made.

Off, Distort, Undistort

Use this control to toggle the *LensWarp* node between three modes:
- *Off:* Disables the effect of the *LensWarp* node. Make sure the *LensWarp* node is turned off before you add calibration shapes.
- *Distort:* Warps the image to introduce lens curvature to the image. This mode is useful for matching the lens curvature in another image.
- *Undistort:* Warps the image to remove lens curvature.

factor

A two-dimensional parameter that lets you adjust the amount of vertical and horizontal curvature. Opening the subtree lets you edit the X and Y dimensions independently.

kappa

The amount of warping applied to the image. Typical adjustments fall between 0 and 1.5, but higher adjustments can be made if necessary.

Negative values produce opposite effects—if the *LensWarp* node is set to Distort, negative kappa values undistort the image. If its set to Undistort, negative values distort the image. This enables you to animate changes from one state to another.

center

The center of the warp performed to the image. By default, xCenter is set to `width/2`, and yCenter is set to `height/2`.

overSampling
An integer value that represents the numbers of samples per pixel that are taken into account when performing a warp. This parameter is set to 1 by default, which results in faster rendering times. However, extreme warping effects may introduce aliasing artifacts that can be reduced or eliminated by increasing this value, up to a maximum value of 4. Increasing this parameter may cause render times to increase dramatically.

Note: Although the slider is limited to a range of 1 to 4, you can enter larger values into this parameter's value field.

Compositing With the MultiPlane Node

3

The MultiPlane node provides a simple 3D compositing environment within Shake. This environment can be used as a way to arrange and animate images within a 3D space, or as a way to integrate generated or tracked 3D camera paths into your scripts.

An Overview of the MultiPlane Node

The *MultiPlane* node provides a compositing environment for positioning 2D layers within a 3D space. A virtual camera controls the scope of the output image, similar to that found within 3D animation packages. This camera can be animated via keyframed parameters, or by importing 3D camera and tracking data from other 3D animation or 3D tracking applications. As with the *MultiLayer* node, the *MultiPlane* node accepts unlimited input images.

The *MultiPlane* node has two primary uses:

- You can composite background or foreground elements against a moving background plate using 3D camera tracking data, imported from a variety of third-party applications.
- You can arrange multiple layers within a 3D coordinate space for easy simulation of perspective, parallax, and other depth effects.

Note: There is one important limitation to positioning layers within the 3D space of the *MultiPlane* node—there is currently no support for intersecting planes. Planes that intersect appear either in front or behind, depending on the position of each layer's center point.

Additionally, the *MultiPlane* node provides transform controls for each layer connected to it. Layers can be moved, rotated, and scaled within the *MultiPlane* environment.

See Tutorial 3, "Depth Compositing," in the *Shake 4 Tutorials* for a lesson on using the *MultiPlane* node.

Viewing MultiPlane Composites

When you double-click a *MultiPlane* node to open it into the Viewer, the Viewer switches to a multi-pane interface, unique to the *MultiPlane* node.
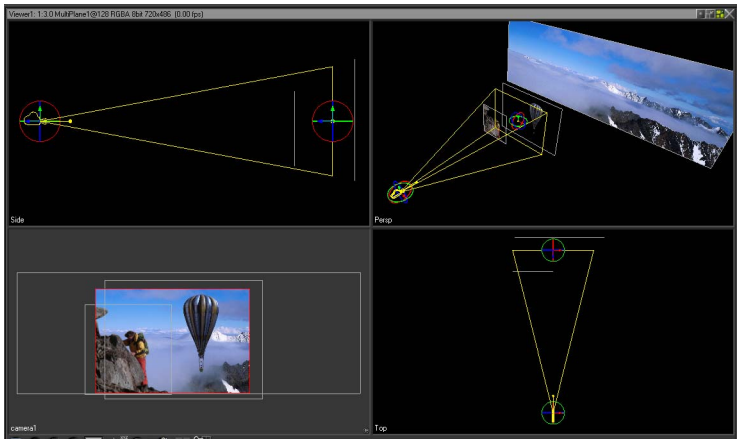
Each pane of this interface can be toggled among single, double, triple, and quadruple-pane layouts. Each individual pane can be set to display any available camera or angle in the 3D workspace, to help you position and transform objects from any angle.

Hardware Acceleration in the MultiPlane Node

The *MultiPlane* node supports OpenGL hardware acceleration of images displayed in the Viewer. Hardware rendering provides a fast way of positioning layers and the camera in space, at the expense of rendering quality.

Whichever pane of the Viewer is set to display the currently selected renderCamera (the default is camera1, although any camera or angle can be in the renderCamera pop-up list) can be toggled between hardware and software rendering using the Render Mode button in the Viewer shelf.

Render Mode Button	Description
Hardware Rendering Mode	This mode is the fastest for arranging your layers in 3D space, but doesn't provide an accurate representation of the final output. In hardware rendering mode, every layer is composited with an *Over* operation, regardless of that layer's selected composite type.

Render Mode Button	Description
Hardware Mode While Adjusting	This mode sets the Viewer to use Hardware rendering while you're making adjustments using onscreen controls. Once you've finished, the Viewer goes into Software rendering mode to show the image at the final quality. To turn this setting on, click and hold the Render Mode button, then choose this option from the pop-up menu that appears.
Software Rendering Mode	This mode displays the selected renderCamera as it appears at its actual quality. All composite types are displayed properly.

The Render Mode button only affects the image that's displayed in the Viewer when the *MultiPlane* node is selected. The image output from the *MultiPlane* node to other nodes in the tree is always at the highest quality, as are *MultiPlane* images that are rendered to disk.

MultiPlane Node Parameters

The Parameters tab of the *MultiPlane* node is divided into two subtabs, the Images and Camera tabs:

- The *Images* tab contains all the parameters for the individual layers that are being arranged in the 3D space of the *MultiPlane* node. For more information on Image tab parameters, see "Parameters in the Images Tab" on page 100.
- The *Camera* tab contains the positioning and optical simulation parameters for each of the cameras and angles used by the *MultiPlane* node. For more information on parameters in the Camera tab, see "Parameters in the Camera Tab" on page 110.

Using the Multi-Pane Viewer Display

Similarly to a 3D application, the *MultiPlane* node's multi-pane Viewer interface displays several angles of the 3D *MultiPlane* workspace. When you first open a *MultiPlane* node into the Viewer, the Viewer switches by default to a four-pane layout, which displays the Side, Perspective, Camera1, and Top angles.

Whether or not the multi-pane interface appears depends on how you open the *MultiPlane* node.

To view the final output from a *MultiPlane* node only:

▪ Click the left side of the node to display it in the currently selected Viewer without also loading its parameters.

The Viewer displays the final output from that node, just like any other node.

To work within a *MultiPlane* node using the multi-pane interface:

- Double-click a *MultiPlane* node to load its image into the Viewer and its parameters into the Parameters tab. Now, the Viewer switches to the *MultiPlane*'s multi-pane Viewer interface.

  Once the multi-pane interface is displayed, you can toggle among four different layouts.

To change the *MultiPlane* Viewer layout:

- Click the Viewer Layout button in the Viewer shelf. Keep clicking to cycle among all the available layouts.

The relative size and orientation of each pane in all four layouts is fixed, although you can zoom and pan within any pane using the standard methods.
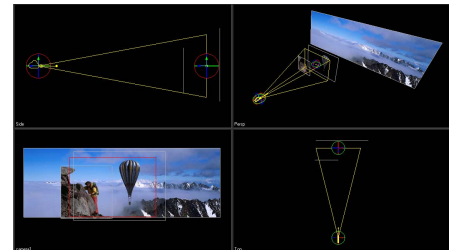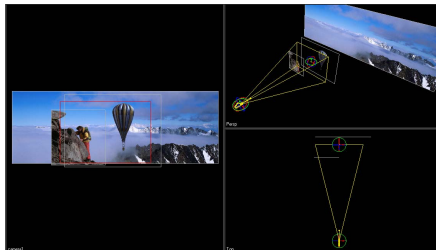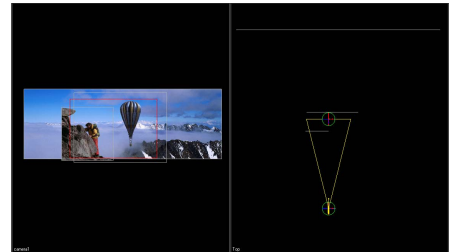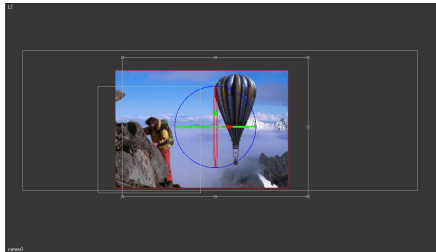
Using Favorite Views in the Multi-Pane Viewer Display

You can use the standard Favorite Views commands within the Viewer to save and restore the framing of each individual pane. In addition, when you right-click in the Viewer shelf to access the Favorite Views commands that are available from the shortcut menu, you can save and restore the state of all visible panes at once.

Changing Angles Within a Pane

Although the multi-pane layouts are fixed, you can change the angle displayed by each pane at any time. The assigned angles appear in white at the lower-left corner of each pane.

To change the angle displayed by a single pane, do one of the following:

- Right-click within a pane, then choose a new angle from the shortcut menu.

- Position the pointer within the pane you want to switch, and press one of the numeric keypad keyboard shortcuts (0-5) to switch layouts.

The following table lists the available keyboard shortcuts (using the numeric keypad only) that are available for changing angles in a pane.

Numeric Keypad	Description
0	Cycles through every angle.
1	Displays the currently selected renderCamera.
2	Front
3	Top
4	Side
5	Perspective

Using and Navigating Isometric Display Angles

The various angles that are available are intended to help you position layers and the camera within 3D space. Internally to Shake, each angle is actually an invisible camera. The first three angles are isometric views—if you set one of these angles as the renderCamera, you'll notice that the focalLength parameter is set to 0.

You can pan and zoom within any pane using the middle mouse button, and the same keyboard modifiers that apply to other areas in Shake.

The isometric angles are:

- *Front:* Useful for transforming a layer's X and Y pan, angle, and scale parameters. You can pan and zoom to navigate within the Front view.
- *Top:* Useful for transforming a layer's X and Z pan, angle, and scale parameters. You can pan and zoom to navigate within the Top view.
- *Side:* Useful for transforming a layer's Y and Z pan, angle, and scale parameters. You can pan and zoom to navigate within the Front view.

Using and Navigating Within the Perspective Angle

In addition, there is a single non-isometric angle, Perspective (Persp). This is the only pane where you can transform a layer's X, Y, and Z parameters all at once.

In addition to panning and zooming to navigate within this angle, you can also orbit the Perspective angle.

To orbit the Perspective view:

▪ Move the pointer within a pane displaying the Perspective angle, press X, and drag with the middle mouse button held down.

The Perspective angle rotates about the perspective view's orbit point.
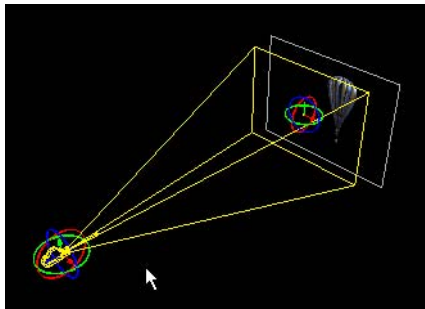
Before rotating the Perspective angle   After rotating the Perspective angle
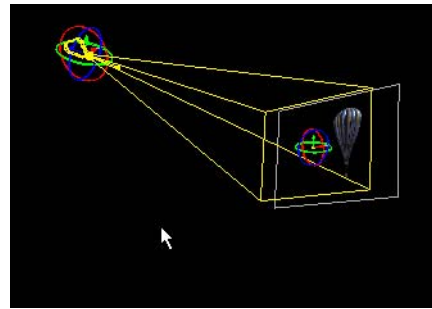
To center the Perspective view on a selected object:

▪ Select an object that you want to be the new position of the orbit point (the camera or a layer), and press Shift-B.

The Perspective view's orbit point is moved so that it is centered on the selected object, and the view within that pane is repositioned to a default position.

Note: You can also use the Shift-B keyboard shortcut in the camera view. However, Shift-B only changes the interestDistance parameter of the camera to best match the position of the selected object.

The renderCamera Angle

The camera or angle that's assigned to the renderCamera parameter in the Camera tab is special, since it represents the final output of the *MultiPlane* node. Each *MultiPlane* node is created with one default camera, named camera1. However, you can use the copy button to create duplicate cameras, or import additional cameras into a *MultiPlane* node by importing one or more .ma (Maya ASCII) files.

Whichever angle is assigned as the renderCamera has the following additional properties:

- It's the only angle that can be switched between viewing the software-rendered final output, and the hardware-rendered preview.
- It's the only angle with a compare control (in software-rendering mode only).
- The image border ROI (Region of Interest) appears only for the renderCamera angle.
- There are additional keyboard shortcuts available for transforming a camera. For more information, see "Manipulating the Camera" on page 105.

How the renderCamera Defines Image Output

Even though the renderCamera angle shows the region defined by the camera target's ROI, the image data appearing outside of this area is not cropped. Shake's Infinite Workspace once again preserves this information for use by downstream nodes elsewhere in your script.

Customizing the MultiPlane Default Camera

You can customize the default angles that appear for new *MultiPlane* nodes.

Use the following declaration within a .h preference file:

```
DefDefaultMPCamera( Camera(0, "v4.0", "Front", imageWidth/(2*settings-
    >yPixelUnit), imageHeight/(2*settings->yPixelUnit), 3000,   0, 0, 0,
    "XZY", 0, 0.980, 0.735, 2, "Fill", 0, 0, 0, "default", xFilter, 0, .5,
    0 ) );
```

MultiPlane Viewer Shelf Controls

When you open a *MultiPlane* node into the Viewer, a group of controls specific to the *MultiPlane* node appear in the Viewer shelf. These controls let you toggle the visibility of various onscreen controls in the Viewer.

Button		Description
	Point Cloud Display	Displays/hides the individual locator points that display the imported point cloud data. Displaying this data can make it easier to align layers with the approximate locations of features that were tracked. Hiding this data can make it easier to manipulate the individual layers.
		For more information on locator points, see "Viewing and Using Locator Points" on page 86.
	XYZ Angle	Displays/Hides the X, Y, and Z angle controls, which can make it easier to reposition objects without accidentally rotating them. A third option is available by clicking and holding down the mouse button to reveal a pop-up menu.
	Path Display	Shows/hides animation paths for image plates.

Button		Description
	Rendering Mode	Toggles the Camera View pane of the Viewer between hardware (HW) and software (SW) rendering. Hardware rendering is faster, but the color and quality of the image is less accurate; this mode makes it easier to position objects. Software rendering is slower, but allows you to accurately see the final image as it will be rendered.
	Viewer Layout	Toggles the Viewer among four different multi-pane layouts, single, double, triple, and quadruple. Each pane can be set to display a different view of the 3D layout.
	Keyframe All or Selected Layers	When this button is turned on, adjusting a single layer using the *MultiPlane* node produces a keyframe that only affects that layer. Animation applied to any other layer is not affected by this new keyframe.

Global Parameters That Affect MultiPlane Display

Two subparameters in the Globals tab let you adjust the quality of hardware-accelerated images displayed within the multi-pane Viewer, and the relative scale of distances represented by imported locator points.

textureProxy

Located within the useProxy subtree, textureProxy sets the proxy level at which texture-rendered images that are used by the *MultiPlane*'s hardware-rendering mode are displayed in the Viewer. This is similar to the interactiveScale setting, in that the proxy level set here is used to generate on-the-fly Viewer images.
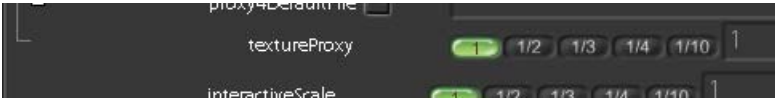
multiPlaneLocatorScale

Adjusts the size of locator points that appear in the Viewer when you load data from a .ma file into a *MultiPlane* node. This lets you scale them up to make them easier to select, or down to get them out of the way.

Connecting Inputs to a MultiPlane Node

Like the *MultiLayer* node, the *MultiPlane* node accepts a variable number of inputs. Drag input noodles from other nodes onto the + sign on the top of the *MultiPlane* node that appears when the pointer passes over it. The + sign always appears to the right of all other previously connected knots.
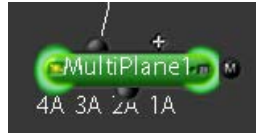
Before   **After**

You can also attach several nodes to a single *MultiPlane* node simultaneously.

To connect several nodes to a *MultiPlane* node at once:

1 Select all of the nodes you want to connect to the *MultiPlane* node.

2 Shift-click the plus sign input of the *MultiPlane* node.
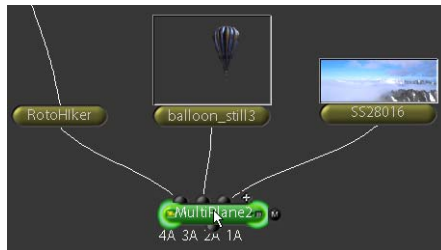
All selected nodes are connected to the *MultiPlane* node. By default, all new layers you connect appear centered on the Camera view; rotate to face the camera's current position.

Using Camera and Tracking Data From .ma Files

The *MultiPlane* node supports .ma (Maya ASCII) files, allowing you to import 3D camera paths from a variety of 3D animation packages, or use 3D tracking data clouds generated by third-party tracking applications. Every new *MultiPlane* node is created with one camera, named camera1. Importing a .ma file adds a camera to the renderCamera pop-up menu in the Camera tab of the *MultiPlane* node's parameters. You can add as many cameras as you like to a single *MultiPlane* node.

Important: When exporting camera path data from Maya, you must bake the camera data for it to be usable by Shake.

A single .ma file can also contain data for multiple cameras. Shake imports every camera that's found within a .ma file, adding each one to the renderCamera pop-up menu. Choosing a camera from the renderCamera pop-up menu displays that camera's parameters within the Camera tab.

Note: There is currently no support for the culling of 3D tracking data from within Shake. Any manipulation of point cloud data should be performed before that data is imported into Shake.

Importing .ma File Data

Data from .ma files is managed using a group of buttons at the bottom of the Camera tab within the *MultiPlane* Parameters tab.

To import a .ma file into a *MultiPlane* node:

1 Load a *MultiPlane* node into the Parameters tab, then click the Camera tab that appears.
2 Click the Load button, choose a .ma file, and click OK.

The data from the .ma file appears in the Viewer as a cloud of points.
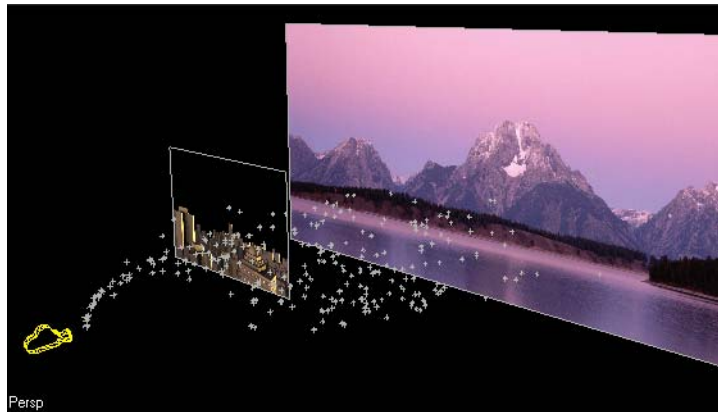
The camera or tracking data populates the parameters of the Camera tab, and a new camera appears at the bottom of the renderCamera pop-up menu.
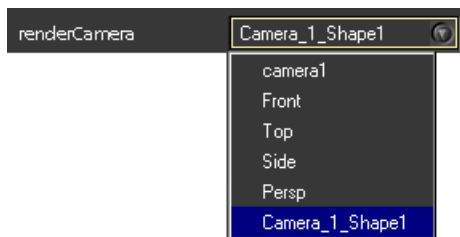
Important: Many 3D applications export camera paths with timing that begins at frame 0. Shake scripts begin at frame 1, which can result in a one-frame offset. To correct this, edit the timing of each camera path point as described in "Editing Locator Point Data" on page 87. (This is not an issue for 3D tracking point clouds—they have no timing data associated with them.)

Once a 3D camera path or tracking data cloud has been imported into the *MultiPlane* node, you can attach the layer that produced the tracking data to the camera. Doing so forces the attached layer to follow along with the camera as it moves to conform to the tracking data. As a result, the attached layer itself doesn't appear to be moving in the camera output. You can tell that the correspondence is correct if the imported tracking points conform to the matching features within the attached layer.

Unattached layers that are positioned within the 3D workspace should now appear as if they're moving along with the features within the attached layer. For more information about attaching a layer to the camera, see "Attaching Layers to the Camera and to Locator Points" on page 94.

Importing Data Clouds From Maya

In Shake, the aspectRatio parameter of a layer within the *MultiPlane* node is determined by the width and height values of the filmBack parameter in the Camera tab. These values are obtained from the imported .ma file.

When tracking a scene in Maya, do one of the following to make sure that the resulting point cloud matches the features of the tracked media:
• Set the film back aspect ratio in Maya to match the render resolution aspect ratio.
• Turn on the "lock device aspect ratio" option in the render globals—this also sets the device aspect ratio to match the film aspect ratio in the camera settings.

Deleting and Duplicating Cameras

To delete a camera, use the Delete button at the bottom of the Camera tab.

To delete a camera and its data:
1 Choose the camera you want to delete from the renderCamera pop-up menu of the Camera tab.
2 Press the Delete button (at the bottom of the Camera tab).

You can duplicate any camera or angle in the *MultiPlane* node with the Copy button. For example, you might want to create a custom viewing angle, or make some changes to a camera without losing the original camera path.

To copy a camera, creating a duplicate:
1 Choose the camera you want to copy from the renderCamera pop-up menu of the Camera tab.
2 Click the Copy button.

A duplicate camera appears in the renderCamera pop-up menu.

Linking to a Camera in Another MultiPlane Node

You can link a camera in one *MultiPlane* node to a camera in a different *MultiPlane* node, so that both cameras move together. When you do so, each camera parameter in the current *MultiPlane* node is linked via expressions to the same parameter in the second *MultiPlane* node. By default, each linked parameter is locked to prevent accidental deletion of the link.

Thus, you can set up animated *MultiPlane* composites using multiple *MultiPlane* nodes, instead of connecting all your images to a single *MultiPlane* node. For example, you might set up a composite using three different *MultiPlane* nodes—one for background layers, one for midrange layers, and one for foreground layers. This way you can apply separate color correction and *Defocus* nodes to the output of each sub-composite.
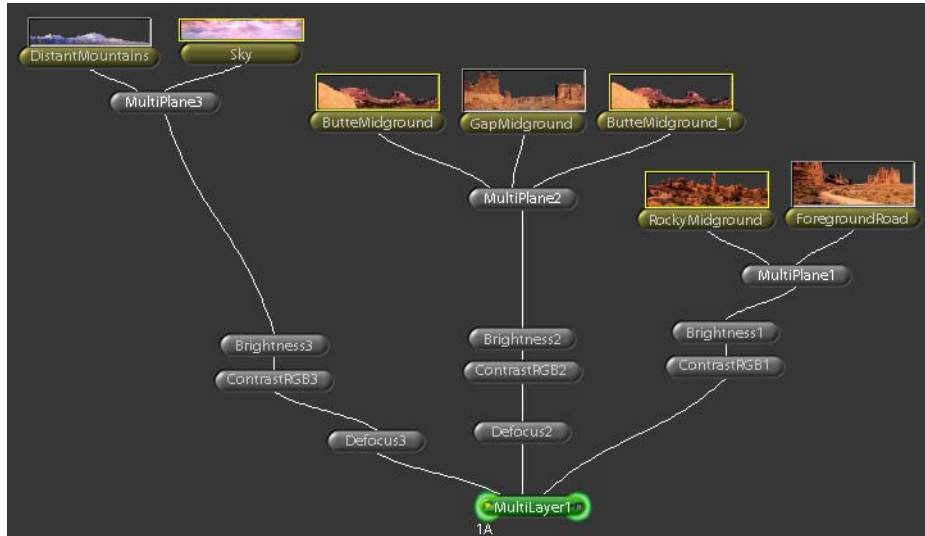
In the above example, one of the foreground image sequences has a matching 3D tracking point cloud that has been imported into the *MultiPlane1* node. For this composition to work, you need to link the cameras in all three *MultiPlane* nodes together.

To link a camera in one *MultiPlane* node to a camera in another:

1 Load a *MultiPlane* node into the Parameters tab.

2 Open the Camera tab.

3 Click the Link button, at the bottom of the Camera tab.

The "Link camera" window appears. The "Link to" pop-up menu presents a list of every camera and angle within every *MultiPlane* node in your script.
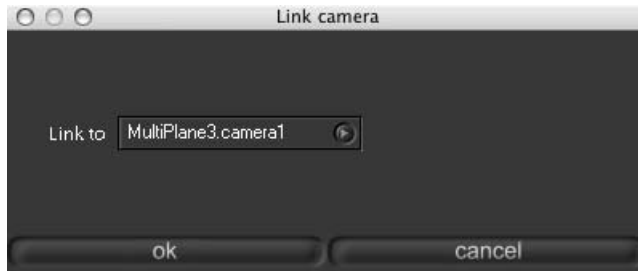
4 Choose the camera you want to link to from the "Link to" pop-up menu, then click OK.

The camera in the current *MultiPlane* node is now linked to the camera you chose. Every parameter in the Camera tab is linked, with expressions, to the matching camera parameter of the *MultiPlane* node you chose.

To unlink a camera:
▪ Unlock each parameter in the Camera tab, then clear each parameter's expression.

Viewing and Using Locator Points
If you've imported 3D tracking data, it's represented in the Viewer by a cloud of locator points arranged within the 3D space of the *MultiPlane* node. Each locator point corresponds to a feature that was tracked by the camera tracking software that created it.

Note: Depending on how detailed the track is, the point cloud may visibly conform to the significant features within its corresponding image sequence. This can help you position elements you're placing within the scene.

To view an imported point cloud:
▪ Click the Point Cloud Display button to hide and show the point cloud.

To change the size of the locator points displayed in the Viewer:
- Adjust the multiPlaneLocatorScale parameter, located at the bottom of the guiControls subtree of the Globals tab.
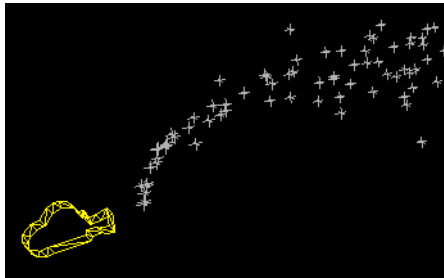
Default MultiPlaneLocatorScale of 1   MultiPlaneLocatorScale set to 3
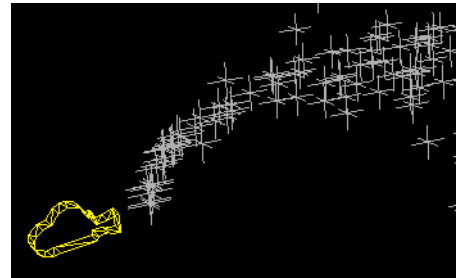
If the data cloud has individually labeled locator points, these labels are viewable within Shake. Some 3D tracking applications also allow you to add locator points by hand—labeling specific areas of the image that may be useful for layer positioning from within Shake.

To view information about a specific locator point:
- Position the pointer directly over an individual locator point in the Viewer to automatically reveal a tooltip with information about that point's name and location.
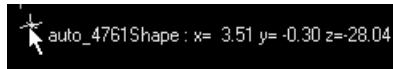
In addition to simply viewing locator points, you can connect layers directly to individual locator points. For more information about attaching layers to locator points, see "Attaching Layers to Locator Points" on page 98.

Editing Locator Point Data
You can view and edit the parameters of any locator point. You can also load points from imported camera paths into the Curve Editor to change their timing.

To reveal an individual locator point's data in the Parameters tab:
- Position the pointer over the locator point you want to edit, right-click it, then choose Expose Point from the shortcut menu.

A new subtree named pointCloud appears at the top of the Images tab of the *MultiPlane* node's parameters. Opening the subtree reveals a list of every point that's been added using the Expose command.
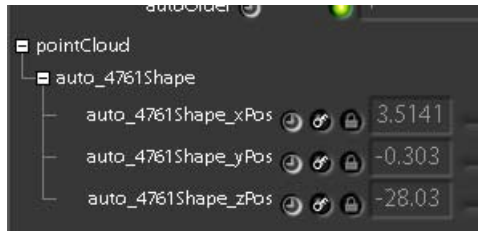
Opening an individual locator point's subtree in this list reveals its xPos, yPos, and zPos parameters. These parameters can be loaded into the Curve Editor to edit, animate, or slip their values in time.

Transforming Individual Layers

The *MultiPlane* node allows 3D transformations of all images that have been connected to it using either onscreen controls, or parameters within the Layers tab of the *MultiPlane* parameters. The onscreen controls are similar to those found within the *Move3D* node, except that these transformations occur within an actual 3D workspace.

Layer Transformation Values

Layers can be panned, rotated, and scaled. Transformations made using a layer's onscreen controls are relative to each layer's center point. When moving layers through space, the numeric values for all transformations are relative to the 0,0,0 center point of the *MultiPlane* node's world view.

Layer Onscreen Viewer Controls

Unlike other nodes that present onscreen controls in the Viewer, the *MultiPlane* node lets you select the layer you want to manipulate directly in the Viewer. A layer must first be selected before you can transform it.

To select a layer, do one of the following:

- Position the pointer within the bounding box of any layer's image in the Viewer to highlight that layer. When the pointer is over the layer you want to select, click it to expose the 3D transform controls for that layer.

- Right-click in a pane of the Viewer, then choose a layer from the Current Plan submenu of the shortcut menu.

As you move the pointer over different layers in the Viewer, the affected layer's name appears in yellow text in the upper-left corner of the screen. Once you've selected a layer, its name appears in white text in the same corner.

If you have many layers stacked up within a single *MultiPlane* node, you may want to turn one or more layers invisible to make it easier to manipulate the remaining layers. Invisible layers aren't output when the script is rendered.

Layer Controls
When you select a layer, that layer's onscreen controls appear superimposed over it.

Global axis controls

Pan controls
Center controls

Scale controls

Rotate X, Y, and Z controls
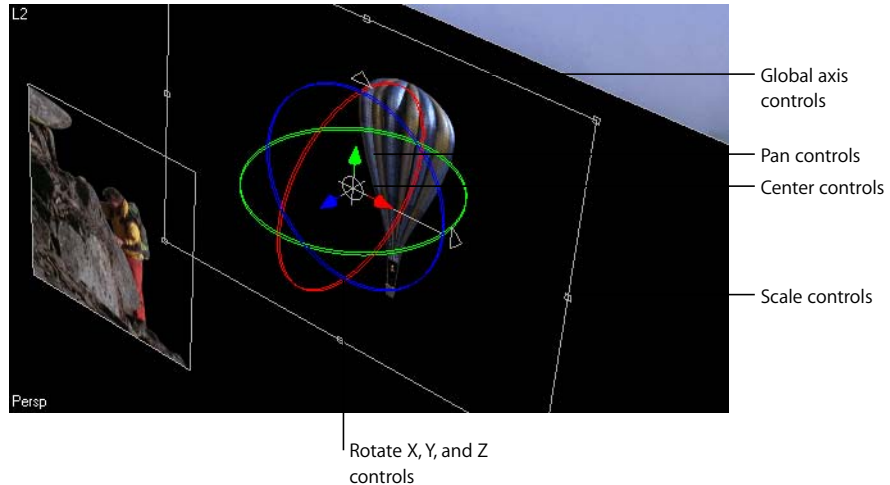
Center Controls
All transformations you make to a layer occur relative to the center point, indicated by crosshairs in the middle of the onscreen controls.
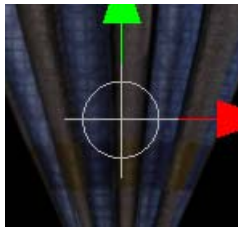
To move the center point:

▪ Hold down the Control key, and drag the center point to a new location.

Pan and Center Controls

Selected layers in the Viewer have two sets of onscreen controls for panning in 3D space—global axis and local axis pan controls.

Local axis pan control

Global axis pan control
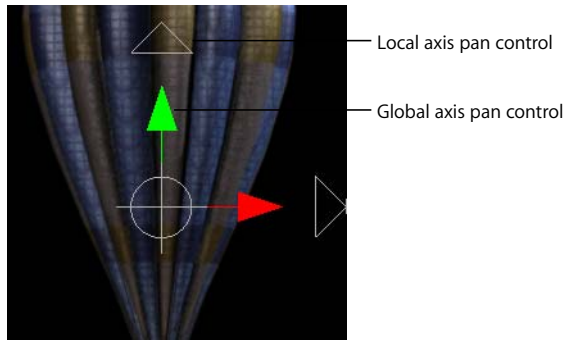
Global axis controls pan a layer relative to the overall 3D space, even if the layer has been rotated. Panning a layer up with a global axis control pans it straight up in space.
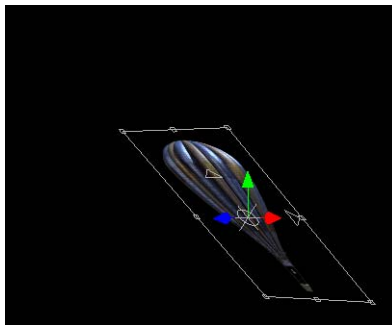
Before global axis pan
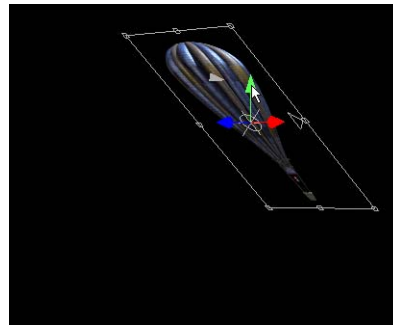
After global axis pan

To pan around the global axis, do one of the following:

- Click a layer anywhere within its bounding box, and drag in any direction.

- Drag one of the three global axis pan controls to constrain the layer's direction.

- Select a layer, and press Q or P while dragging anywhere within the Viewer to pan a layer without positioning the pointer directly over it.

The local axis pan controls pan the layer relative to its own orientation. If a layer has been rotated using the angle controls, using the local pan controls moves it along the axis of its own rotation. Panning a layer up with a local axis control pans it diagonally in space.
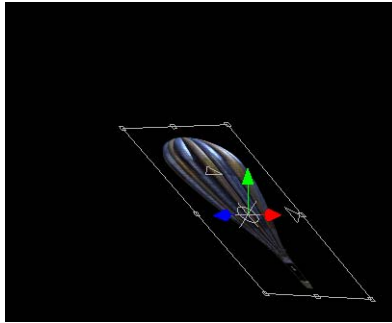
Before local axis pan   After local axis pan
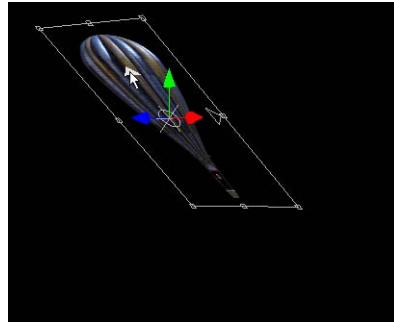
To pan along the Local Axis:
- Drag one of the two local axis pan controls to move the layer in that direction.

Angle Controls
Selected layers have three angle controls that rotate the layer around the X, Y, and Z axes of that layer's center point. These angle controls work identically to those found in the *Move3D* node. The color of each angle control corresponds to the color representing each dimension of the global axis pan controls.
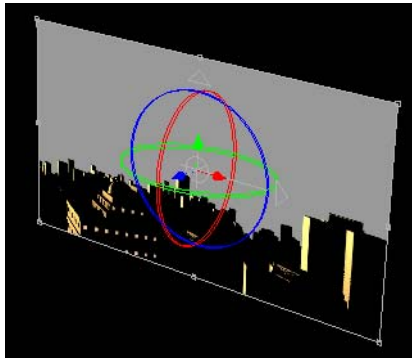
The visibility of the angle controls can be toggled by clicking the XYZ Angle button in the Viewer shelf. Hiding these controls may make it easier to use the pan controls in some situations.

To rotate a layer in the Viewer without using the angle controls:

1 Select a layer.

2 Press W or O and click in a pane of the Viewer.

3 When the dimension pointer appears, move the pointer in the direction in which you want to rotate the layer. The colors in the pointer correspond to the angle controls.

4 When you move the pointer, the axis in which you first move is indicated by a single axis arrow, and the layer rotates in that dimension.
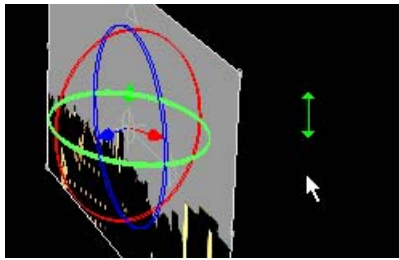
Scale Controls

Selected layers have eight scale controls located around the outer edge of the image.
• The four corner controls rescale the layer, keeping the relative width and height of the layer constrained.
• The left and right controls let you scale just the width of the layer.
• The top and bottom controls let you scale just the height.

By default, all scale transformations occur about the layer's center point. Alternatively, you can hold down the Command or Control key while you drag any of the scale controls to scale a layer relative to the opposite scale control.

To scale a layer in the Viewer without using the scale handles:

1 Select a layer.

2 Position the pointer over a pane of the Viewer and hold down the E or I key.

The dimension pointer appears.

3 Drag the dimension pointer in the direction in which you want to scale the layer. The colored arrows correspond to the pan controls.
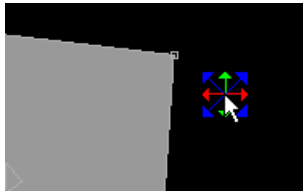
4 When you move the pointer, the direction in which you first move is indicated by a single axis arrow, and the layer scales up and down in that dimension.
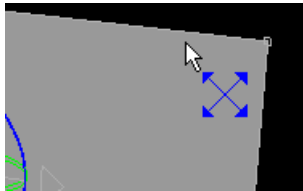
Creating Layer Hierarchies

You can create hierarchical relationships between layers within a *MultiPlane* node using the parentTo parameter. You can lock one layer's transformations to those of another by assigning it a parent layer from this pop-up menu, which contains a list of every layer that's currently connected to the *MultiPlane* node. By default, this parameter is set to off.

Note: The *MultiPlane* node only supports one level of parenting.

To assign a layer a parent layer:

▪ Choose a parent from that layer's parentTo pop-up menu.

Once a layer has been assigned a parent layer, it is removed from the parentTo pop-up menu. When you select a layer that is parented to another, the parent layer appears with a red border in the Viewer.

You can make local transformations to a layer after you've assigned it a parent. This lets you create an offset between it and the parent layer. Transformations that are applied to a parent layer also affect all layers that are assigned to it.

Important: Always assign a parent layer before making any local transformations to a layer itself. Otherwise you may encounter unpredictable results.

Deleting Parent Layers

When you disconnect a layer that's being used as a parent, a warning appears:
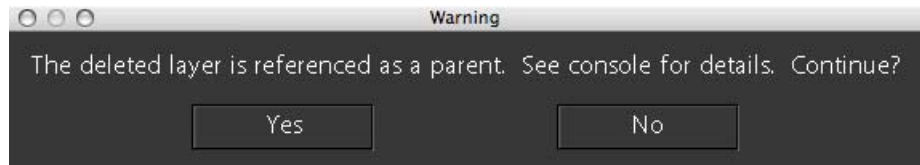
Clicking Yes removes the parent layer, eliminating the parent-child hierarchy. The remaining layers' parentTo parameters still show the original layer, in the event you decide to reconnect it later on.

Showing and Hiding Layers

You can toggle the visibility of layers. For example, if you need to transform a layer that's hard to select because it's obscured by other layers in the composition, you can "solo" it to hide all other layers, or simply hide the individual layers that are in the way. HIdden layers are not rendered.

To show or hide layers, do one of the following:

▪ Right-click a layer in the Viewer, and turn it on or off in the Plane Visibility submenu.

 The Plane Visibility submenu also has commands to Hide All Planes and Show All Planes.

▪ Open the Images tab, then click the Layer Visibility button for that layer.

To solo a layer:

▪ Open the Images tab, and turn on the Solo button for that layer.

 When you solo a layer, every other layer in that *MultiPlane* node is hidden.

Animating Layers

Layer transformations within the *MultiPlane* node can be animated similarly to the parameters within any of Shake's transform nodes. For more information on keyframing parameters, see Chapter 10, "Parameter Animation and the Curve Editor," on page 291 of the Shake 4 User Manual.

Attaching Layers to the Camera and to Locator Points

When you import camera path or 3D tracking data into a *MultiPlane* node, you gain the ability to attach layers to the camera, or to one of the locator points distributed within the 3D workspace.

Attaching Layers to the Camera

To use a *MultiPlane* node to matchmove one or more images into a scene using 3D tracking data, you need to do three things:

- Import data from a .ma file.
- Position the images you want to matchmove.
- Attach the originally tracked image sequence to the camera.

Attaching the layer that produced the tracking data to the camera forces the attached layer to follow along with the camera as it moves according to the tracking data. As a result, the attached layer itself doesn't appear to be moving in the camera output, while the unattached layers that are positioned within the 3D workspace appear as if they're moving along with the features in the attached layer.

To attach a layer to the camera:

▪ Click the Attach to Camera button of a layer in the Images tab—the lock icon closes when Attach to Camera is on.

That layer's image is automatically locked to the full area of the renderCamera angle in the Viewer.
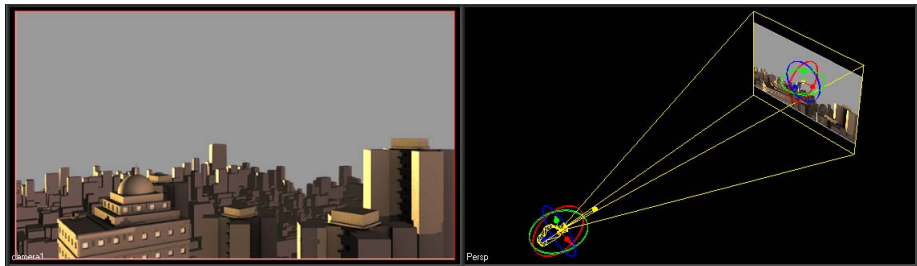
When you turn on a layer's Attach to Camera button, the faceCamera, parentTo, pan, angle, scale, center, and aspectRatio parameters all disappear from that layer's subtree in the Images tab, replaced by a single parameter—cameraDistance.

The cameraDistance parameter lets you adjust the relative spacing between layers that are attached to the camera, and the other unattached layers that are arranged within the 3D workspace. This lets you determine which layers appear in front of and behind attached layers.

Decreasing the cameraDistance value brings the layer closer to the front of the composition, while increasing the cameraDistance pushes the layer farther away. Attached layers move back and forth along the lines that are projected from the camera itself to the four corners of the frustum surrounding the camera target.
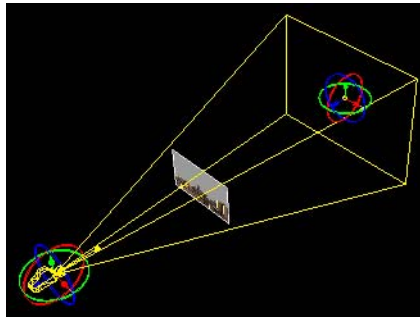
cameraDistance reduced   cameraDistance increased
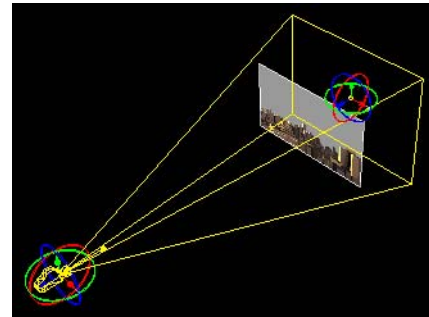
Regardless of how the cameraDistance parameter is set, a layer that's attached to the camera always lines up with the locator points imported from its matching .ma file. Similarly, attached layers always appear, within the camera angle, at their original scale, regardless of their position in the 3D workspace.

You can attach multiple layers to the camera. A typical example is when you need to isolate one or more foreground elements in a tracked shot, so that a matchmoved element can be placed behind it. In the following example, a hot-air balloon is inserted into a tracked cityscape background plate.

Example: Isolating an Element in a MultiPlane Composite

1   First, duplicate the cityscape image sequence that's used as the background plate, and create a rotoshape to isolate the front building.

2   Next, attach the original cityscape image, the isolated building, and the hot-air balloon images to a *MultiPlane* node.

They appear within the Images tab as separate layers.
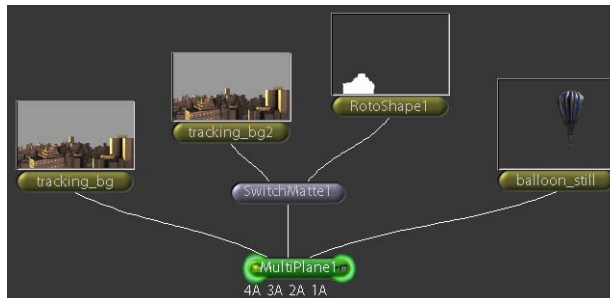
3 Turn on Attach Layer to Camera for the background and isolated building layers to lock both images to the full area of the renderCamera angle in the Viewer.


4 Open the subtree for each attached layer, and adjust the cameraDistance parameters to create the necessary spacing between each element for your composition.

In this case, the front building is moved forward so there's room between the building and the rest of the background plate for the hot-air balloon.
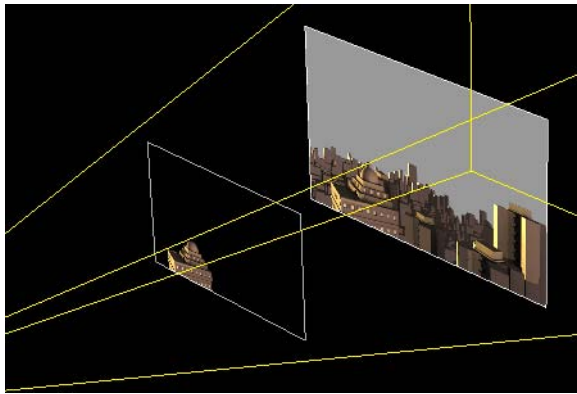

5 Connect the hot-air balloon image to the *MultiPlane* node.

A third layer appears in the layer list.

6 Using the onscreen controls, you can now position this new layer between the isolated building and the overall cityscape plate.

As a result, the balloon appears positioned between the front building and the rest of the city in the camera view.
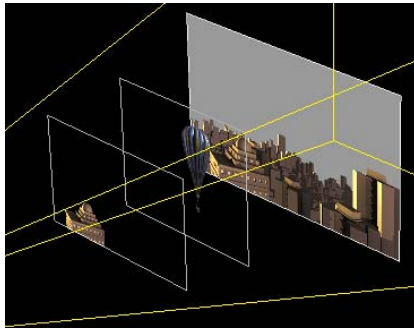
A balloon image inserted between. the front building and cityscape
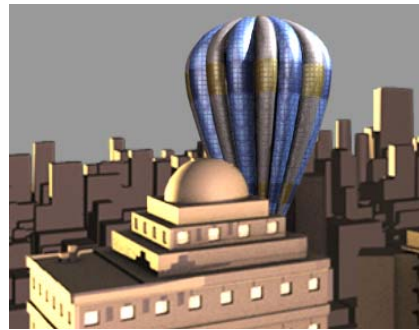
The resulting camera angle

Because both the building and city layers are attached to the camera, they look identical to the original input image from which they're derived, regardless of each layer's position within the 3D workspace.

Attaching Layers to Locator Points

In addition to attaching layers to the camera, you can also attach a layer to any locator point in the data cloud. This lets you quickly set a layer's position in space to match that of a specific feature in a tracked background plate.

When you attach a layer to a locator point, the layer is transformed to match the position of the locator point using expressions. As a result, operations that change the position of locator points, such as changing the sceneScale parameter, also change the position of any layers that are attached to locator points. In addition, animated locator points (from a tracking application capable of tracking moving subjects in addition to backgrounds) will transform any layers that are attached to them as well.

You can attach a layer to either a single locator point, or to a group of three locator points. Attaching a layer to a single locator point only pans the layer, it is not rotated.

To attach a layer to a locator point:
1 If necessary, move the layer's center point to a position at which you want the layer to be attached to the locator point.
2 Right-click a locator point in the Viewer, then choose a layer from the Attach Plane to Point shortcut menu.

The layer is panned to the position of the locator point, and attached at the layer's center point. Expressions are added to that layer's pan parameters that maintain the relationship between the attached layer and the locator point.

To attach a layer to three locator points:

1 If necessary, move the layer's center point to a position at which you want the layer to be attached to the locator point.

2 Shift-click three locator points in the Viewer.

3 Right-click one of the selected locator points, then choose a layer from the Attach Plane to Point shortcut menu.

The layer is panned so that its center point is at the center of the triangle defined by the three selected locator points. In addition, it is rotated to match the orientation of the plane defined by the three points. The order in which you select the locator points determines the orientation of the attached layer. If you select the same three points in the reverse order, then choose Attach Plane to Point, the layer will be flipped.

Note: You can attach a layer to three locator points to orient its rotation, then attach a layer to a single point afterwards to nudge its position in space, while maintaining the current rotation.

Adjusting sceneScale

The sceneScale parameter, at the top of the Camera tab, lets you scale the relative distribution of locator points making up a camera path or tracking data cloud within the 3D workspace of the *MultiPlane* node. This lets you increase or decrease the space between different planes within your composition, to make it easier to arrange layers.

If you adjust the sceneScale parameter in the Camera tab, layers that are attached to locator points move along with the expanding or contracting point cloud:

• Lowering sceneScale contracts the distribution of locator points, bringing any layers that are locked to locator points closer to the camera. The locator points themselves appear to bunch up.
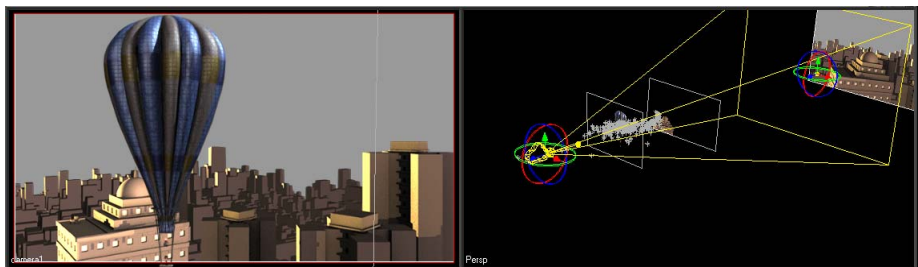
- Raising sceneScale expands the distribution of points, moving any layers that are locked to locator points away from the camera. The locator points themselves appear to stretch out.
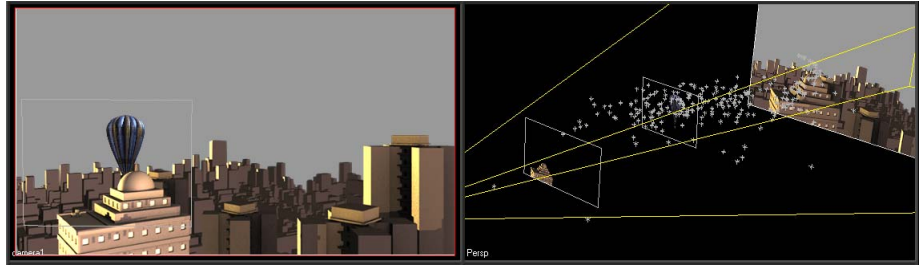
Changing the sceneScale parameter has no effect on layers that are attached to the camera, nor does it affect the position of layers that are not attached to locator points. It does, however, affect the size of the frustum, increasing or decreasing the area that is seen by the camera. It also changes the position of the camera—if you make a big enough change to the sceneScale parameter, the camera may move past layers in the 3D workspace.

Parameters in the Images Tab

The first three parameters in the Images tab determine the overall image that is output from the *MultiPlane* node.

ClipLayer

Defines the output resolution of image data produced by the *MultiPlane* node.

postMMult

When turned on, the postMMult parameter premultiplies the output image produced by the *MultiPlane* node.
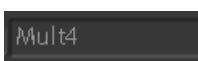
autoOrder

When turned off, layer order is determined by the position of layers in the Parameters tab, much like with the *MultiLayer* node. When autoOrder is turned on, layer order is determined by each layer's position in 3D space.

Note: The layer order of coplanar layers (occupying the same coordinate in the 3D workspace) is determined by their position in the Parameters tab.

Individual Layer Controls

Each image that you connect to a *MultiPlane* node is represented by its own set of layer controls and subtree parameters in the Images tab. These controls are similar to those found within the *MultiLayer* node, and are labeled in the order in which the layers were connected to the *MultiPlane* node. For example, L1 is the name of the first connected layer, followed by L2, and so on.
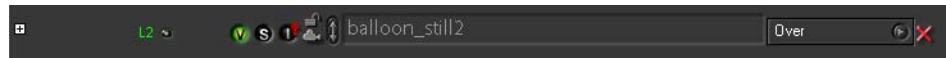
Layer Button or Control	Description
L1	The number of each layer corresponds to the input knot that image is connected to. L1 is the first knot at the left.
Input	Clicking this button shows the input image for that layer in the Viewer.
Layer Visibility	Toggles the visibility for that layer. Layers with visibility turned off are not rendered.
Solo	Turns off the visibility of all other layers. You can only solo one layer at a time. Click an enabled solo control to disable solo.
Ignore Above	Turns off all layers above the current layer, keeping only the current layer and those behind it visible.
Attach Layer to Camera	If a particular layer corresponds to a tracked sequence that's imported from a .ma file, turn this button on to attach the image to the camera, so that it's not transformed when the camera moves to follow a track. As a result, unattached layers appear locked relative to the tracking information when composited against the attached layer.
Reposition	Dragging this control up and down lets you rearrange the layer order within the parameter list. This only affects layers that are coplanar, unless autoOrder is turned off.
Mult4 — Input Layer Name Field	Gives the name of the preceeding node that's actually connected to the *MultiPlane* node.
Over — Composite Mode	A pop-up menu that lets you set each layer with its own composite mode, which affects how that image's color data interacts with other overlapping layers. Certain composite modes add parameters to that layer's parameter subtree. New layers default to the Over mode.
Disconnect Mode	Disconnects that input image from the *MultiPlane* node, and removes it from the layer list without removing the input nodes from the node tree.

Individual Layer Parameters

Opening up a layer's parameter subtree reveals a group of compositing and transform parameters affecting that particular layer.
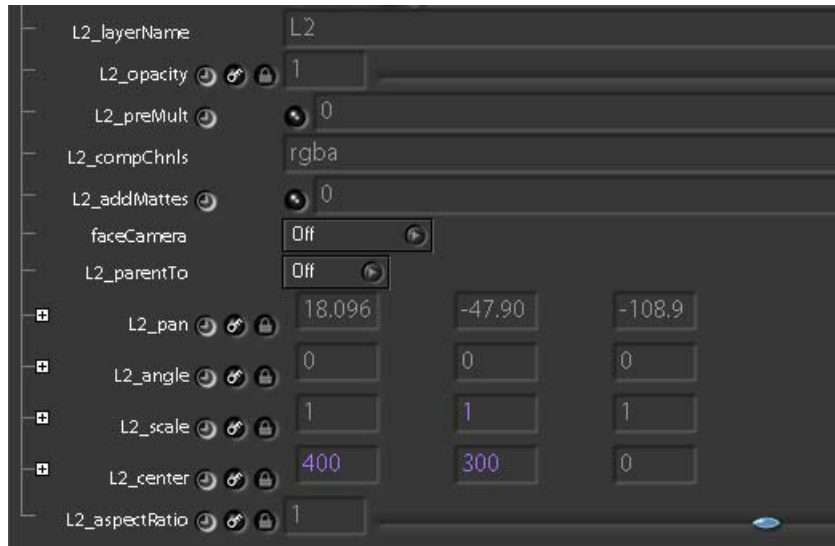
layerName

The name of the layer. All associated parameters for that layer are prefixed by the *layerName*.

opacity

Opacity of the input layer. With an imported Photoshop file, there is an additional PSOpacity parameter. See "Supported Photoshop Transfer Modes" on page 476 of the Shake 4 User Manual for more information.

preMult

When this is on (1), the foreground image is multiplied by its alpha mask. If it is off (0), the foreground image is assumed to already be premultiplied.

compChannels

Sets which image channels are passed up from below (behind) this layer.

addMattes

This parameter appears when a layer is set to Over. When enabled (1), the mattes are added together to create the composite.
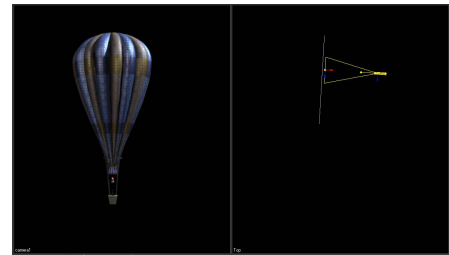
faceCamera

The faceCamera pop-up menu lets you choose a camera. Once set, the layer will always rotate to face the same direction as the selected camera. If the camera is animated, the layer animates to match the camera's movement so that the object remains facing the camera at every frame. This automatically animates that layer's angle parameters, even though no keyframes are applied. In the following example, the camera starts off facing a 2D balloon layer.
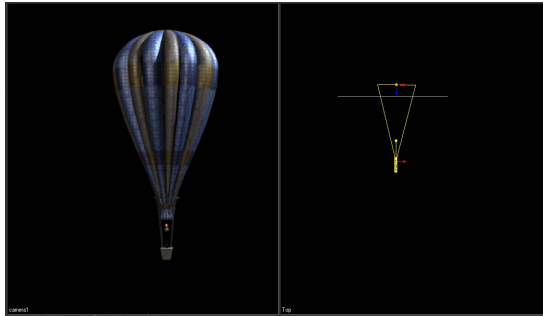
With faceCamera turned off, moving the camera around the layer results in the image flattening as the camera reaches the side of the layer. When faceCamera is set to Camera1, the image automatically rotates to face the direction the camera is facing.
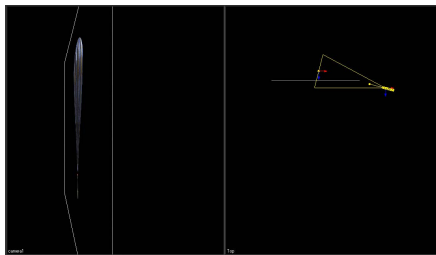

Rotated camera with faceCamera turned off

Rotated camera with faceCamera turned

Unlike ordinary rotation made with a layer's angle controls, the automatic rotation that's made as a result of faceCamera being turned on is relative to the center point of the bounding box that defines the actual image of the layer, and *not* by the layer's center point.

One useful application of this parameter is to offset a layer's center point, then use layer rotation to control layer position, even though faceCamera is turned on. In the following example, an image of a planet is arranged to the right of an image of the sun. The planet layer's center point is offset to the left to match the center of the sun layer.


The arrangement in the Camera pane looks like this:
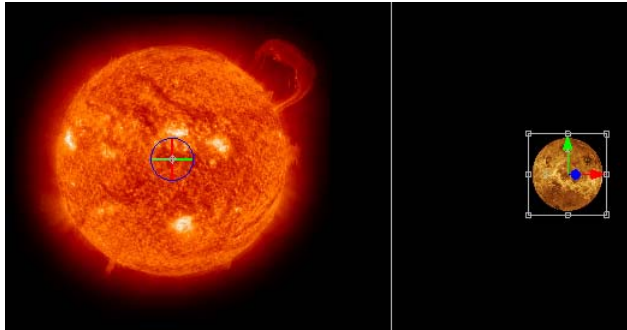
Sun and Venus images courtesy NASA/JPL-Caltech

Rotating the planet layer's yAngle parameter (corresponding to the green axis arrow) now moves the planet layer around the sun layer, so that it appears to orbit about the sun. By default, the 2D planet layer thins out as it approaches the camera. Setting faceCamera to Camera1 results in the planet layer rotating to face the camera as it moves around the sun.
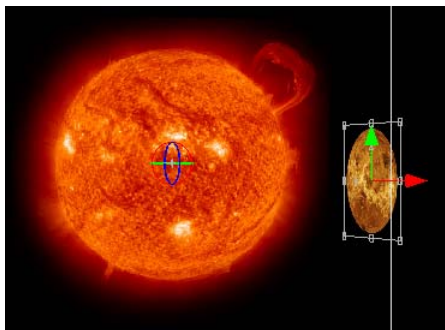
Rotated planet with faceCamera turned off
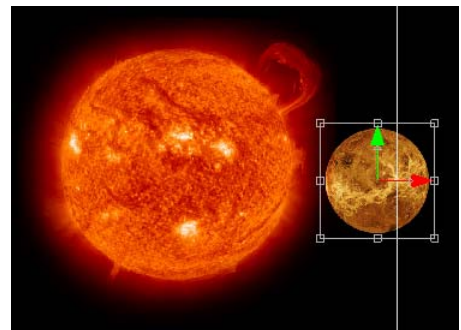
Rotated planet with faceCamera set to Camera1

parentTo

This parameter lets you create layer hierarchies within a single *MultiPlane* node. You can link one layer to another by choosing a parent layer from this pop-up menu. Layers with a parent can still be transformed individually, but transformations that are applied to the parent are also applied to all layers linked to it. For more information on using the parentTo parameter, see "Creating Layer Hierarchies" on page 93.

pan (x, y, z)

Corresponds to the onscreen pan controls, allows you to pan the layer in 3D space.

angle (x, y, z)

Corresponds to the onscreen angle controls, allows you to rotate the layer in 3D space.

scale (x, y, z)

Corresponds to the onscreen scale controls, allows you to resize the layer in 3D space.

center (x, y, z)

Defines the center point of the layer, about which all pan, rotate, and scale operations occur. By default the center of newly added layers corresponds to the horizontal and vertical center of the layer's bounding box.

aspectRatio

This parameter inherits the current value of the defaultAspect global parameter. If you're working on a nonsquare pixel or anamorphic image, it is important that this value is set correctly to avoid unwanted distortion in the image.

Manipulating the Camera

The camera can either be positioned or animated manually, like any other layer, or by importing 3D tracking or camera data via a .ma file.

3D Transform Controls
Click the camera to expose its 3D transform controls.

Camera target
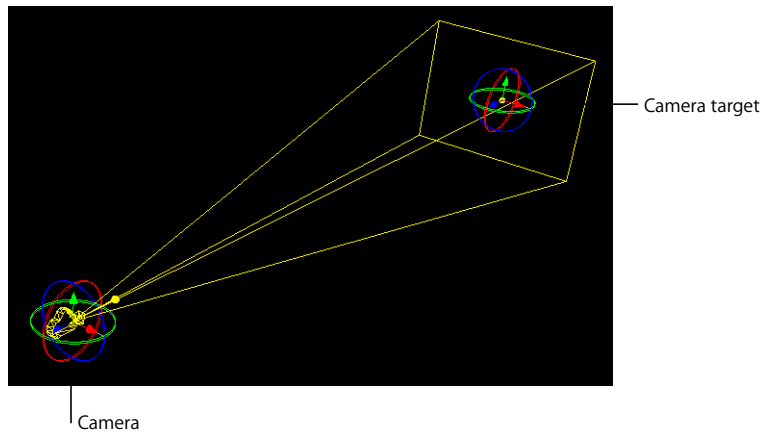
Camera

The camera consists of two groups of controls, the camera itself, and the camera target. Both controls are connected so that they always face one another—moving one rotates the other to match its position.
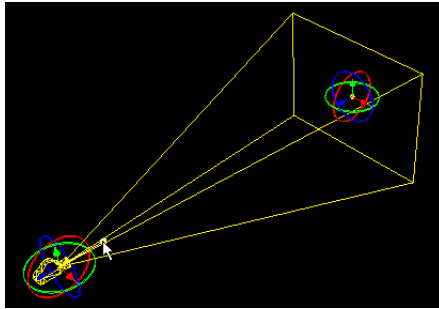
Camera Controls
Similar to layer controls, the camera has rotate x, y, and z controls, and translate (move) x, y, and z controls to constrain movement of the camera to one of these directions. Dragging the camera image itself freely moves the camera within the Viewer.
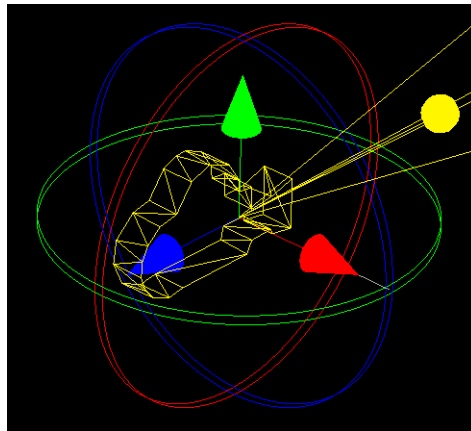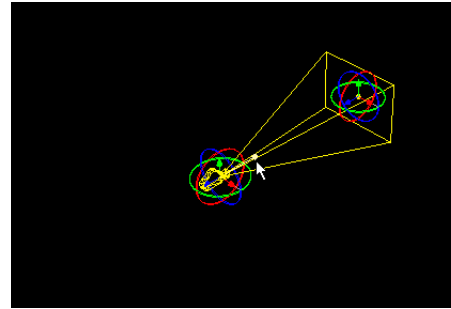
Moving the camera in relation to the camera target

An additional control at the front of the camera constrains its movement so that it moves either closer to or farther away from the camera target—which also adjusts the interestDistance parameter in the Camera tab of the *MultiPlane* parameters.
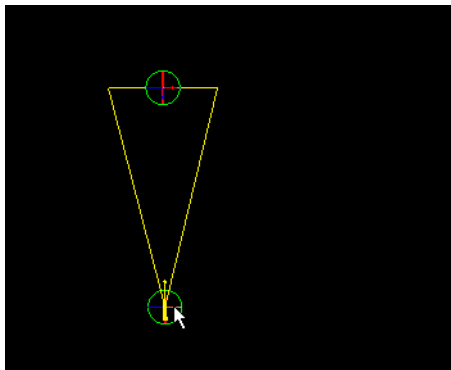

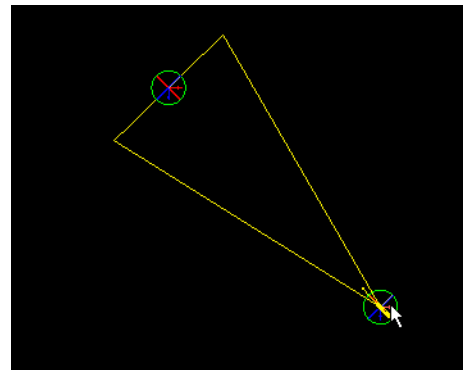**Before moving interestDistance control**

After moving interestDistance control

No matter where you move the camera in space, it rotates to face the position of the camera target. The orientation of the camera target, in turn, rotates to face the camera.
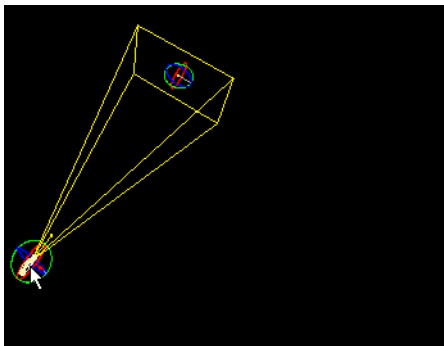

Before panning camera

After panning camera

Additionally, a special group of keyboard shortcuts lets you move the camera by dragging anywhere within the renderCamera view, without selecting the camera itself.

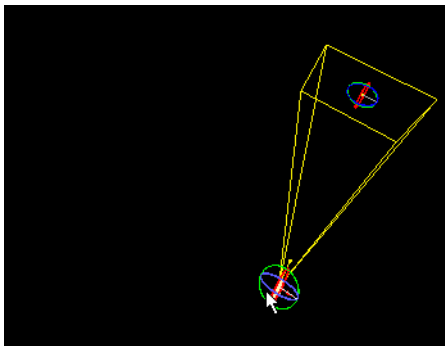*Note:* These keyboard shortcuts only work within the camera view:

Keyboard	Explanation
V-drag	Rotates the camera about its Z axis.
S-drag	Rotates the camera about the X and Y axes, about the camera's own center point, changing the position of the camera target.
Z-drag	Pans the camera in and out along the Z axis.
D-drag	Pans the camera and camera target together along the X and Y axes.
X-drag	Pivots the camera about the camera target's orbit point.

To move the camera and the camera target together in any view:
- Press T and drag the camera or camera target controls in the Viewer.

Before T-dragging camera   After T-dragging camera

Camera Target Controls and Frustum

The camera target represents the image that is displayed by the camera1 view. In addition to the standard angle and panning controls, the camera target displays the orbit point as a yellow sphere at the center of the target controls. This is the point about which the camera rotates when you move the camera independently of the camera target, or when you manipulate the camera target angle controls.

Note: The orbit point is also the point about which the Persp (perspective) view rotates when you rotate this view using the X key.


Orbit point
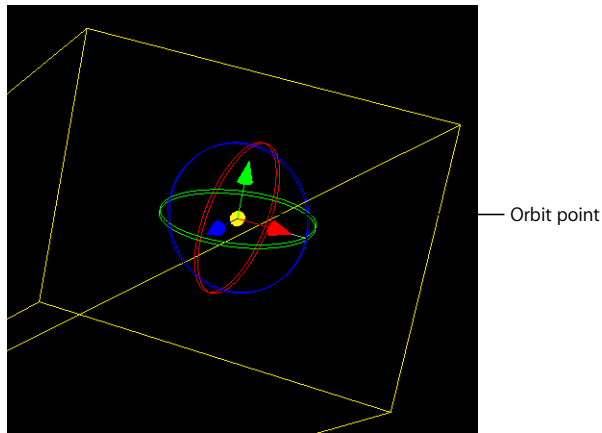
The outer bounding box represents the frustum, which defines the camera view frame that produces the output image. The size of the frustum determines, in part, the area of the 3D workspace that is output as the renderCamera image that's output by the *MultiPlane* node.

Note: Unlike frustum controls found in other 3D animation packages, Shake's *MultiPlane* frustum does not crop the image outside of the frustum boundary. Thanks to Shake's Infinite Workspace, this image data is preserved for future downstream operations in the node tree.

The controls within let you move the camera target. Using the translate (move) x, y, and z controls moves the camera target—causing the camera to rotate so that it continues to face the target.

Using the rotate x, y, and z controls rotates the camera about the camera target's orbit point.
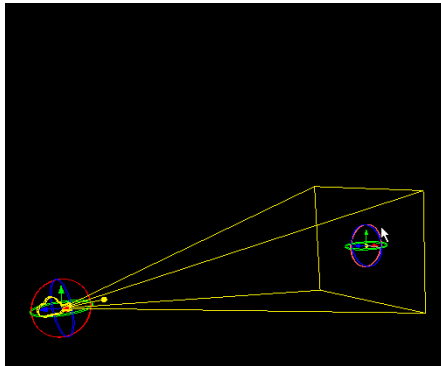
Before rotating camera target   After rotating camera target
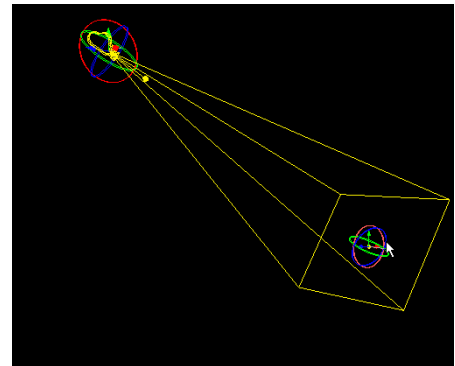
Moving the camera target closer to or farther away from the camera adjusts the interestDistance parameter in the Camera tab of the *MultiPlane* parameters.

Animating the Camera
The camera can either be animated manually, like any other object, or by importing 3D tracking or camera data via a .ma file.

Animating the Camera Using .ma Data
You can use the Copy, Load, Delete, and Link buttons at the bottom of the Camera tab of the *MultiPlane* parameters to import, use, and clear camera and tracking data from a .ma file.

Use the sceneScale parameter at the bottom of the guiControls subtree of the Globals tab to change the size of the locator points displayed in the Viewer.

Parameters in the Camera Tab
All of the parameters that affect the camera are located within the Camera tab of the parameters.

renderCamera
A pop-up menu that lets you choose which angle provides the output image from the *MultiPlane* node.

Lock, Unlock, Reset
Click the Lock button to lock all of the camera parameters at once, preventing them from being accidentally edited. Click the Unlock button to unlock these parameters again. Click the Reset button to restore the camera to its default position.

cameraName
A text field that lets you rename the camera or angle to more accurately describe your setup.

sceneScale
Scales the depth of the virtual space used to distribute the locator points that are displayed in the Viewer (which represent 3D tracking data clouds imported from .ma files). This parameter allows you to expand or compress the relative distance from the camera to the tracked background plate. Adjusting this parameter lets you more easily position layers in space when camera tracking data comes from a subject that's either very far away, or very close. This parameter is for reference only, and has no effect on the data itself. The default multiPlaneLocatorScale value is 50.

Camera Transform Data
The following parameters contain transform data for the camera, as well as parameters that determine how the image is resolved based on mathematically simulated camera attributes such as focal length and film gate resolution. These parameters are adjusted whether you simply reposition the camera within the 3D workspace, keyframe the camera manually, or import a camera path or 3D tracking data from a .ma file.

Note: Many of the parameters found within the Camera tab are identical to parameters exported by Maya. For more information, see the Maya documentation.

Lock, Unlock, Reset
Three buttons let you lock all camera parameters, unlock all camera parameters, or reset all camera parameters.

focalLength
Sets the focal length of the camera's virtual lens. Shake measures the focal length in millimeters. The default focalLength is 35mm.

angleOfView
A subparameter of focalLength. This value is provided for convenience, and represents the horizontal field of view of the frustum (as opposed to the vertical field of view that the *Move3D* node uses), based on the current focalLength. The value of the angleOfView has an inverse relationship to that of the focalLength parameter—raising one lowers the other, and vice versa.

translate (x, y, z)
Transform data for the camera's position. By default, the translate parameters are set to expressions that automatically set their values.

rotate
Transform data for the camera's rotation. If you keyframe rotation, the data is stored here.

rotateOrder
The order in which rotations are executed, by dimension. Defaults to XZY.

interestDistance
This parameter defines the distance of the camera target to the camera. By default, this parameter is set to an expression that automatically sets its value.

filmGate
The filmGate pop-up menu provides presets for setting the filmBack width and height parameters. There are options corresponding to most standard film formats.

filmBack (width, height)
Represents the width and height of the image that strikes the virtual film, in inches.

Note: If you're not importing camera data and you want to avoid squeezing the image improperly, make sure the aspect ratio of the filmBack width and height parameters matches that of the input image.

scale
Represents the size of the camera compared to the scene. Lowering this value reduces the visible area taken in by the virtual lens, increasing the size of the scene relative to the viewer. Increasing this value increases the visible area seen by the lens, reducing the size of the scene relative to the viewer.

Changing the scale parameter also changes the *effective* focal length of the camera. For example, if the focalLength parameter is set to 50, setting the scale to 2 changes the effective focal length to 25. Setting the scale to 0.5 changes the effective focal length to 100.

fitResolution
A pop-up menu with four options: Fill, Horizontal, Vertical, and Overscan. This parameter determines how differences between the filmBack aspect ratio and that of the input image are resolved.

filmFitOffset
If the filmBack resolution is different from that of the clipLayer, this parameter offsets the image within the filmBack resolution in inches.

filmOffset (x, y)
Offsets the image within the filmBack in relation to the area defined by the clipLayer. This parameter is measured in inches.

useDeviceAspectRatio
If this parameter is turned off, the camera uses the aspect ratio of the input image. If this parameter is turned on, the deviceAspectRatio parameter is used instead.

deviceAspectRatio
By default, this parameter uses an expression that computes the aspect ratio based on the filmBack parameter.

xFilter, yFilter
Lets you pick which method Shake uses to transform the image in each dimension. For more information, see "Applying Separate Filters to X and Y Dimensions" on page 863 of the Shake 4 User Manual.

motionBlur
A Motion Blur quality level of 0. 0 produces no blur, whereas 1 represents standard filtering. For more speed, use a value less than 1. This value is multiplied by the global parameter motionBlur.

- **shutterTiming**

  A subparameter of motionBlur. Zero (0) is no blur, whereas 1 represents a whole frame of blur. Note that standard camera blur is 180 degrees, or a value of .5. This value is multiplied by the global parameter shutterTiming.

- **shutterOffset**

  A subparameter of motionBlur. This is the offset from the current frame at which the blur is calculated. Default is 0; previous frames are less than 0.

.ma Load/Camera Copy, Delete, Link Buttons
Four buttons at the bottom of the Camera tab let you control how .ma camera data is used within a *MultiPlane* node.

Load
This button lets you load .ma data from a file. This creates a new camera in the currently open *MultiPlane* node.

Copy/Delete
These buttons let you duplicate or delete the currently selected renderCamera.

Link
The link button at the right lets you link the camera in the currently open *MultiPlane* node to an animated camera within another.

Delete Cloud
Deletes a point cloud, but leaves the camera angle intact. This is useful if you plan on redoing the track and you want to clear the old point cloud in preparation for importing a new one.