Version 1.0

for

Adobe After Effects

Key Correct Pro for After Effects Manual

www.redgiantsoftware.com

Manual Content: Lee Croft, Sean Safreed, Stonewall Ballard

Product Engineering: Stonewall Ballard, Micah Sharp

Product Design: Stonetics Inc., Red Giant Software LLC.

Product Installation

Key Correct Basics

Before We Begin

Alpha Cleaner

Color Matcher

Deartifacter

Denoiser

Edge Blur

Edge Finder

Light Wrap

Matte Feather

Matte Feather Sharp

Rack Focus

Smooth Screen

Speedy Blur

Spill Killer

Wire & Rig Remover

ZoneHLS

Product Installation

Welcome

Key Correct Pro is a plug-in designed for Adobe After Effects 5.5 and later to run on Mac OS X® (10.2 and later) and Windows® XP (Home and Pro editions). Key Correct Pro for After Effects lets you create perfect keys from an image shot against a colored background.

Installation

Key Correct Pro for After Effects comes packaged as complete software installer. You must have installed After Effects 5.5 or later before installing Key Correct Pro. The installer will automatically place the required components on your hard drive.

The installer will guide you through the process of installing the plug-in. You will be prompted for a serial number on installation of Key Correct Pro.

Activation

The final step in the installation before placing the plug-in on your system is activating the plug-in. YOU MUST ENTER A SERIAL NUMBER TO USE Key Correct Pro for After Effects. Your serial number is provided in email form automatically when you purchase the product as a download. You may also find the serial number clearly applied to the inside of the Key Correct Pro package.

The serial number for Key Correct Pro appears in the following format:

####

All the characters are digits. You must enter all characters in the serial number entry dialog before the installation can proceed.

Product Support

Product support is available to registered customers only. You can register with Red Giant Software by clicking the link below.

Register Here

You can also register with us by opening your web browser and using the following URL: www.redgiantsoftware.com/register.html.

The best way to reach support is through email. You can email your support issue to support@redgiantsoftware.com. We offer a 24-hour turnaround time on all support email received during normal business hours.

Key Correct Basics

Key Correct from Red Giant Software is a suite of tools that enables you to create world-class composites on your desktop computer system, even on footage that other systems will find problematic. When combined with professional-level keying software like Primatte Keyer™ from Red Giant Software, this powerful effects tool gives you superior edge control, element color correction, and professional compositing features.

Applying Key Correct

Key Correct is accessed like all filters from within After Effects. Make sure that your foreground element is highlighted in the timeline. Key Correct is accessed by selecting "Key Correct" from under the Effect menu.

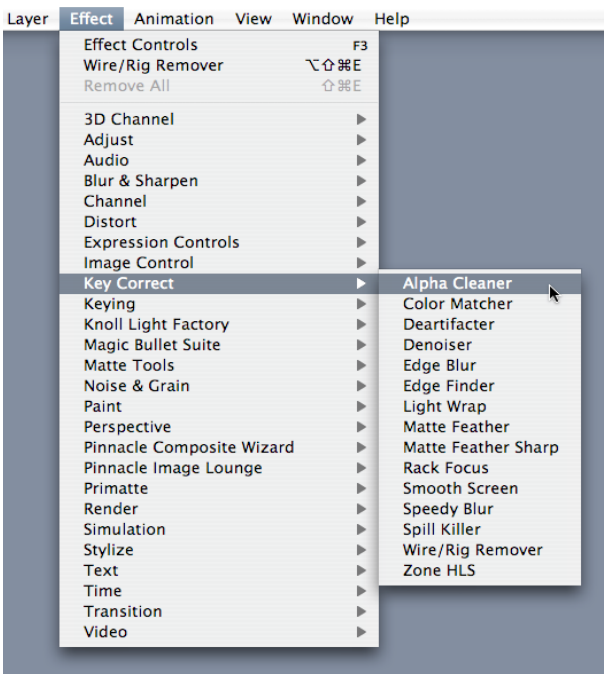

Key Correct is applied to the layer, and the interface for the individual effect appears in the Effect Controls palette.

Before We Begin

In this section we will go over a couple of interface and usage issues which concern multiple plug-ins in Key Correct.

16-Bit Mode

All Key Correct filters work in both 8-bit and 16-bit modes. Some filters, such as Deartifacter, will yield superior results in 16-bit mode, even if the original source footage is 8-bit.

Blend With Original

All Key Correct filters include a Blend With Original control at the bottom of their interface. This control sets the amount of the effect applied to the layer by enabling you to blend an effected layer with the original unchanged layer. The default value is 0%, meaning that none of the original layer is visible. At 100% the original layer is 100% visible, meaning that the effect is effectively disabled. This value is fully keyframable, enabling you to set a gradual fade-in or fade-out of the individual Key Correct Pro filter.

Blurring Versus Feathering

There are two filters in Key Correct (Edge Blur and Matte Feather) which involve softening the edges of your masked or keyed element. While both edge softening operations (blurring and feathering) might appear to be similar, it is important to understand that feathering and blurring are two completely different functions.

To illustrate this difference we have created a high-contrast image consisting of horizontal lines, to which we have applied a spline layer mask, shown below.

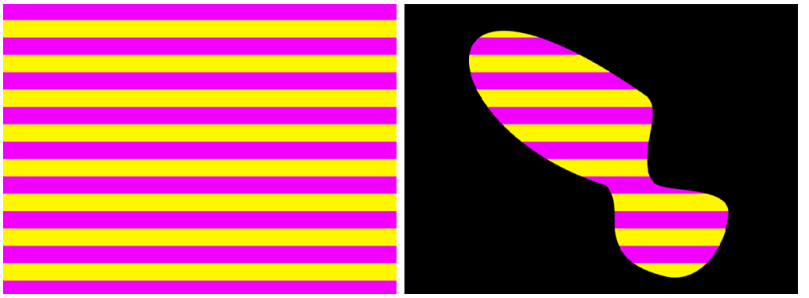

By zooming in on the edge of the shape we will be able to clearly see the difference between the two softening functions.

Blurring

In Alpha Blur the blur function is applied to both the RGB and Alpha channels. You get a soft edge in the alpha, and the detail in the RGB is blurred.

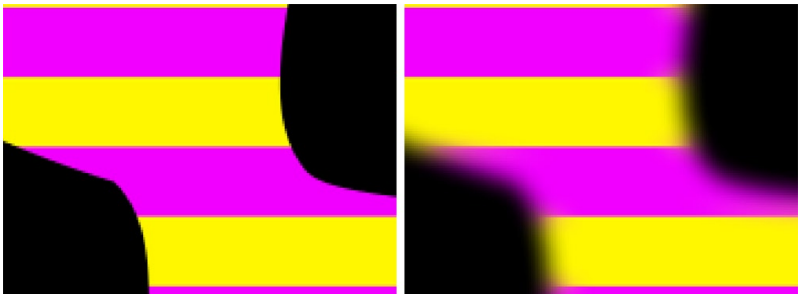

As you can see above, in the blur areas the line between the yellow and magenta stripes is very soft.

Feathering

Conversely, the image below shows the same element with Matte Feather applied. Notice how there is no blurring done to the edges of the RGB.

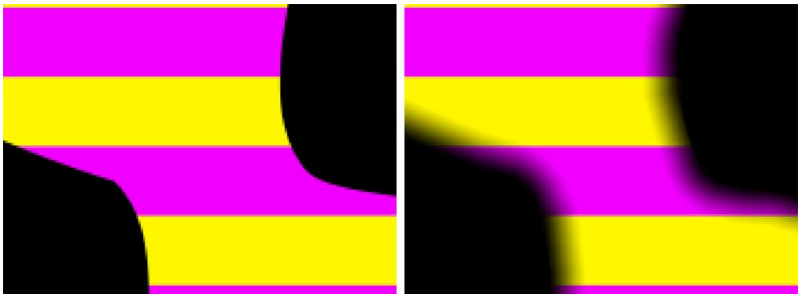

The lines between yellow and magenta are just as crisp and distinct in the edge areas as they are in the rest of the RGB.

Side By Side

To illustrate the differences even more clearly, the image below shows an extreme close-up of the edge of the masked element.

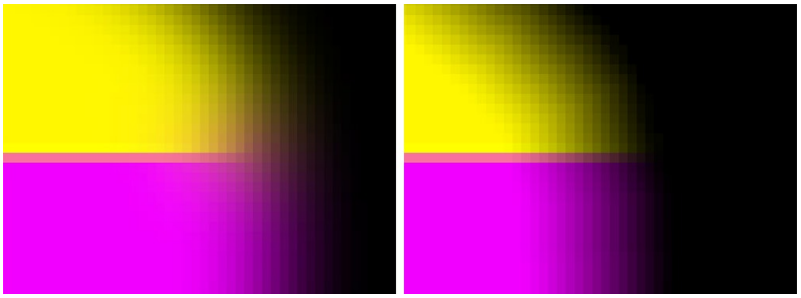

The left side has Alpha Blur applied, and the right shows Matte Feather. The blurring in the edge detail in the left image is obvious.

While both these operations are similar, knowing the differences between them is essential to getting the most out of the powerful edge treatment filters in Key Correct.

Alpha Cleaner

A large part of getting a successful key involves having good source footage to begin with. Unfortunately, that doesn't always happen. Quite often your source image has wrinkles, discoloration, uneven lighting, or other problems, all of which can make for a problem key. While your keying software may able to handle a good deal of these problems on its own, in some cases you will still run into issues, and Alpha Cleaner is designed to help in these cases. Simply put, Alpha Cleaner allows you to zap grain or noise, fill in holes in the foreground matte, and fix many other problems that would otherwise require a separate manual matte touch-up in a paint program. Alpha Cleaner works by applying a clean-up pass to your matte, identifying dirt and other artifacts based upon a combination of values, such as their size, alpha level, and proximity to the principle element in your matte.

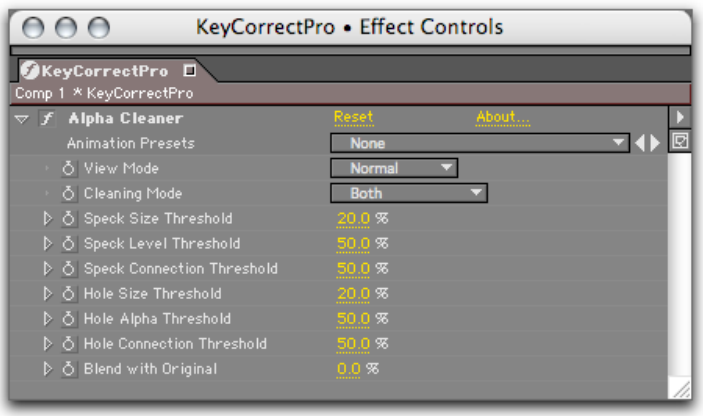

Cleaning Mode

Cleaning Mode allows you to designate which cleaning operation you wish Alpha Cleaner to perform.

Remove Specks Remove Specks tells Alpha Cleaner to remove specks and noise from the background. (White noise in black areas.)

Plug Holes Plug Holes removes holes from the foreground object. (Black holes in white areas.)

Both

Both allows you to both Remove Specks and Plug Holes in one pass. This mode has the same effect as making two passes, but is somewhat more efficient, since some internal buffers can be reused, speeding up the render process.

Threshold Values

The next six controls concern the threshold values for controlling the operations of Alpha Cleaner. The controls are **Speck Size Threshold**, **Speck Level Threshold**, **Speck Connection Threshold**, **Hole Size Threshold**, **Hole Alpha Threshold**, and **Hole Connection Threshold**. As you can see, there are three pairs of operations, with each pair concerning either specks or holes. Since the operations are identical except for the type of object they affect, we will discuss the operations below as three pairs, rather than six individual operations.

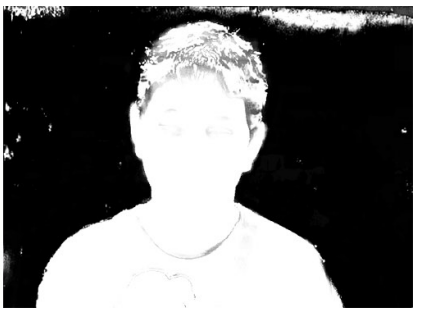

To illustrate the threshold values we will use the image above, of a boy filmed against a pond. The matte we will clean was generated using a standard color difference keyer. As you can see, there are numerous holes in the foreground due to the similarity of colors between the foreground and background, as well as a number of white, noisy areas in the background because of the varying colors found in the pond itself.

Size Threshold

The Size Threshold value determines which alpha pixels will be removed based upon their proximity to each noise or chatter area. The slider value controls the size of the clean-up area, with 100% representing the largest area of clean-up (roughly 1/5 of the image size) and 0% representing no clean-up. Increase this value to get rid of larger areas of noise (i.e. larger specks in the black areas or larger

holes in the white areas). Noise outside of the threshold range is left unaffected.

With the Cleaning Mode set to Remove Specks, we can see that simply applying the filter has a dramatic effect on the background noise.

However, there are two large noise areas in the top left corners which need to be removed. Increasing the Speck Size Threshold value will make Alpha Cleaner remove larger objects.

Increasing the value from the default 20% to 62.5% completely cleans the background while leaving the foreground unchanged.

Level/Alpha Threshold

The Level/Alpha Threshold values determine which pixels will be removed based upon the pixel's luminance value. Noise pixels with a luminance value below the Level/Alpha Threshold will be set to 0 (black) or 100 (white), depending on whether you are removing specks or filling holes. The higher the Level/Alpha Threshold value, the more pixels will be set to either solid white or solid black.

When you look at your matte, if you have areas of your background or foreground (i.e. pixels which should be 100% black or white) which appear gray, this control will allow you to remove them.

However, be advised that if you have any areas of your matte that are supposed to be semi-transparent, such as hair, smoke, or motion-blurred edges, this operation will remove the transparency in these areas. In these type of cases you would be best served by isolating the areas that need to be cleaned with a garbage matte, then using the Alpha Cleaner filter on only those areas.

Note that only pixels within the range of the Size Threshold value will be affected by Level/Alpha Threshold.

Connection Threshold

The Connection Threshold values determine which alpha pixels to remove based on their proximity to major objects in the alpha. The higher this value, the more non-essential pixels will be removed. Increase this value if you find your alpha growing "tendrils" after cleaning.

Color Matcher

One of the biggest problems that arises when compositing a shot comes in making the foreground and background elements appears as if they coexist in the same space. Balancing an element's tonal exposure to make it fit into a new environment can often involve multiple applications of standard color correction filters.

Color Matcher enables you to easily integrate an element into a composite environment by balancing the element's colors to match the general color balance of the environment into which it is being placed. This is accomplished quickly, and with just a few mouse clicks can yield results superior to those achieved with hours of work with conventional color correction filters.

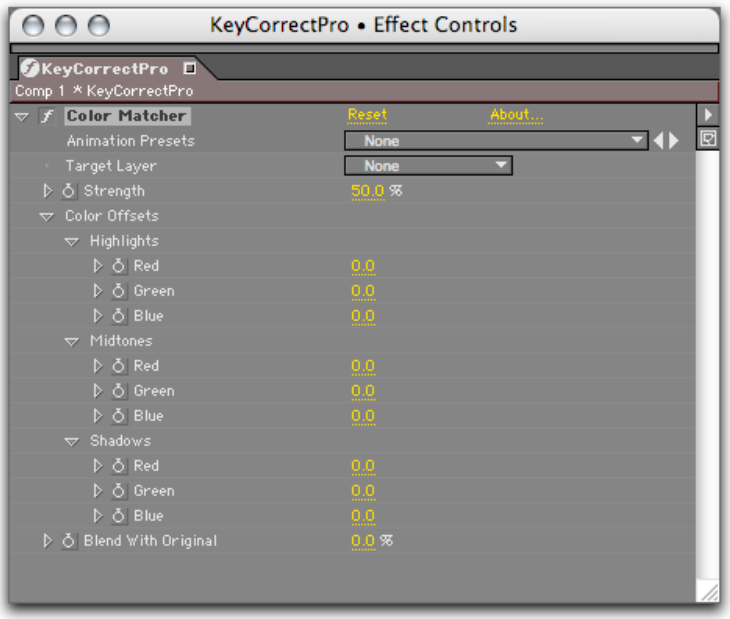

The Color Matcher algorithm automatically samples the specified background layer and remaps the foreground to match. In addition, you now have far greater control over the color correction process with the introduction of individual RGB controls for shadow, midtone, and highlight ranges.

Target Layer

The Target Layer pop-up designates which layer in your composite you are trying to match. Usually this will be the background plate. The tonal range of this element will be sampled and then applied to the foreground layer. The target layer does not, however, have to be visible, which enables you to bring in an element solely as a color correction source.

Strength

The Strength slider determines how much the original colors of the foreground element are adjusted towards the background colors. A value of 100% means that the foreground colors are completely replaced.

The image above shows Strength values of 25% and 75% respectively. That the color remap can be as subtle or as bold as you choose.

Color Offsets

These three sliders, as their names suggest, allow you to color correct the highlight, midtone, and shadow ranges of the image. Each tonal range offers a set of RGB sliders, enabling you to manually add or remove red, green, or blue in each range to achieve an accurate match.

Deartifacter

One of the unique filters found in Key Correct Pro is the Deartifacter. When performing a key on source footage from DV or HD cameras, it creates a superior result by removing blocky artifacts caused by low-resolution color sampling in the original image.

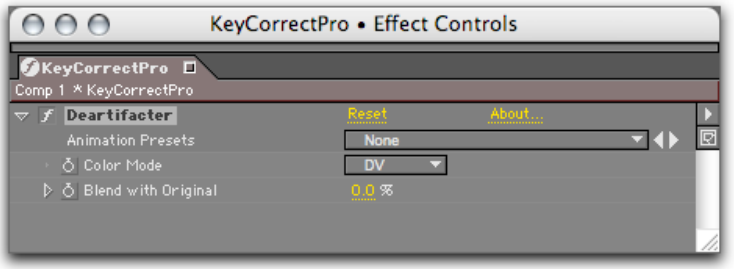

Overview

In the film world footage is shot on traditional celluloid and scanned into uncompressed RGB image sequences at 8-bit or greater color depth. With footage shot on video, however, the output is usually compressed, and saved in reduced 4:2:2 or 4:1:1 YUV color space. This compression and reduced color depth can often cause problems when keying objects with soft edges, such as hair or smoke, or on objects with a great deal of motion blur. Key Correct Pro's Deartifacter allows you to get a superior key by pre-processing the image with algorithms specifically designed to address this issue.

Applying Deartifacter

From a workflow standpoint, Deartifacting is a pre-process. These algorithms are designed to be applied to the image before it is keyed. If you plan on using Deartifacter, apply it first before your keying filter, as is shown below.

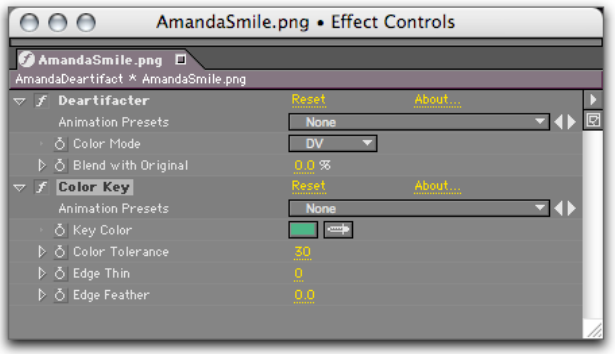

If you pull a key on your footage without first applying Deartifacter, then subsequently notice that there are jagged edges on your objects, you can go back afterwards and apply Deartifacter after you have done your key, or indeed at any other point in the process. Simply apply Deartifacter, then move it up to the top of the Effect Controls palette, before your keying filter.

You will find that virtually all keyed footage shot on video will benefit from having Deartifacter applied, regardless of the specific video format.

Using Deartifacter

Consider the following example. The left image below shows an actress in front of a greenscreen. The footage was shot on MiniDV using a mid-range consumer level camera. On the right is the same footage with a simple color key applied.

If you look along the edges of the actress' face you can see blocking and artifacting caused by the compression in the DV source. The image below shows this same footage zoomed in for closer edge

inspection.

The left side image shows the original footage with the basic key. The right image is the exact same key pre-processed by Deartifacter. The edge of the hair is significantly smoother and more realistic, which makes for both a superior key and final composite.

Once the keyer generates the matte the blocking and artifacting become quite noticeable. With Deartifacter applied you will see a remarkable difference in the smoothness of the edges in the keyed output. This smoothness is even more apparent when you view the matte.

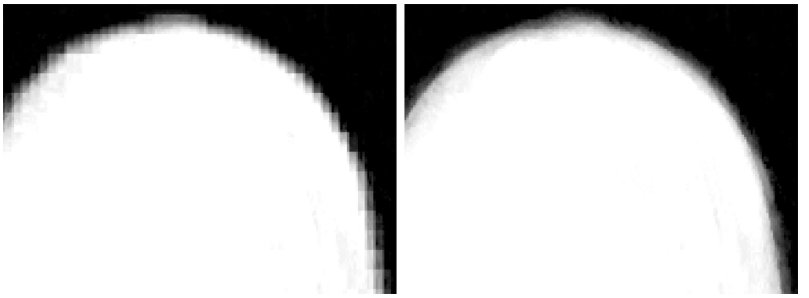

The matte images above show the same close-up of the model's head as in the previous image. Even with a simple key, the superiority of the edge detail in the Deartifacted element is clear.

Mode Menu

The Deartifacter interface is quite simple, consisting solely of a pop-up menu to designate the Mode of the input source.

DV/HDV	The DV/HDV setting is for any footage shot on a DV camera.
HDCAM	The HDCAM setting is for any footage shot on a HiDef (HD) camera.
Other	The Other setting is for when you have footage that comes from a video format not previously mentioned, or on any footage that has these types of blocking issues, no matter what its source.

Denoiser

Denoiser applies an intelligent averaging function across a sequence of frames, effectively reducing or altogether eliminating film and video noise. This can be very useful to clean up dark areas in images that pick up noise from film emulsion or video camera sensors. Once an element has been processed by Denoiser, images will exhibit sharper details and will be less troublesome when applying additional image processing effects, such as keyers or color correction.

Denoiser is more than just a simple frame average. It is able to distinguish between areas in the image that are moving or changing, and areas that are locked off and static, thereby effectively denoising only the static areas while preserving the full detail of areas of motion.

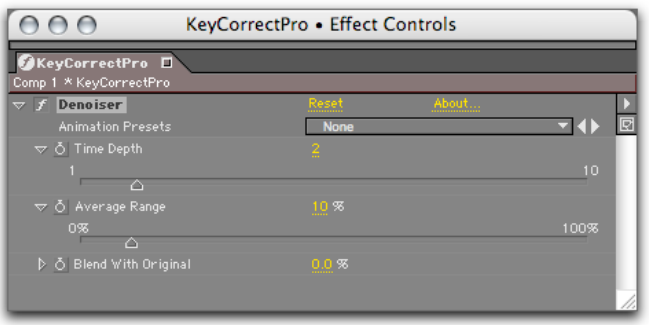

Time Depth

Time Depth specifies how many frames Denoiser will consider when determining what is noise. Higher numbers mean more frames and more accurate results, but slower rendering.

If there is a lot of motion in your image, a higher Time Depth will be wasted effort. If the scene is relatively static a high Time Depth value will remove practically all the noise.

Average Range

Average Range determines the range of values around a pixel's average that are treated as noise. Higher numbers result in more noise being removed.

Tips

Denoiser works best in subtle situations. With higher settings, detail may be smoothed out in an undesirable way.

Because of the nature of the After Effects pipeline, some layers may need to be precomposed before applying Denoiser. Time referencing works like viewing an external target layer, and bypasses masking, etc. If you find that your layers aren't behaving as you would expect, try precomping them first.

Using Denoiser

One common situation where video noise is prevalent is when filming takes place in areas of low light. Consider the following element, filmed inside the Oval Office.

Okay, so it's not the *real* Oval Office. It's a full scale reproduction at the Reagan Presidential Library in California. The room is actually in the lower level of the facility, underground, so the light coming in through the back windows is simulated. The library also has restrictions against using any kind of external light source with cameras, so the footage was shot using available light. As you can see it's quite dark and therefore incredibly noisy.

Simply applying Denoiser at the default values yields a remarkably cleaner image.

By zooming in on the image even further, then increasing the Time Depth to 9 and the Average Range to 72%, you can see that we now have a virtually noise-free image.

Edge Blur

Edge Blur is a simple Gaussian blur filter which, as its name suggests, blurs the edge of the layer it is applied to.

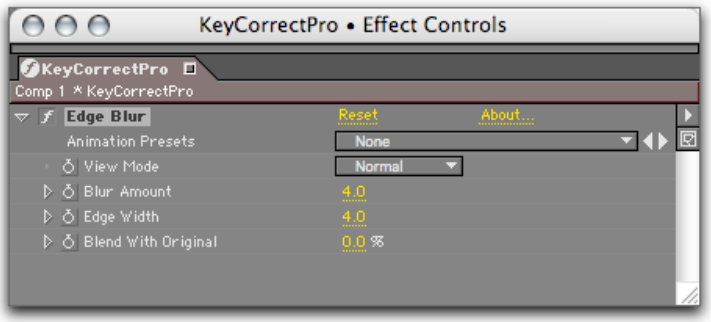

(For a detailed explanation of the difference between feathering and blurring, please see *Blurring Versus Feathering* on page 11.)

Consider the following simple composite, comprised of an image of a sunset with the Key Correct Pro logo placed over it.

View Mode

The View Mode pop-up consists of two options.

Normal Normal shows the standard composite.

Alpha Alpha shows the alpha channel of the element to which the element is applied.

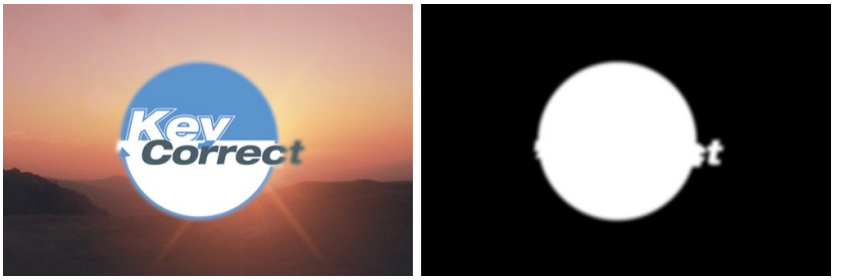

The left image above shows the composite. Alpha view mode shows the alpha channel of the layer being blurred. If you were to use the alpha view button in the composite window you would get a solid white alpha, since it would be showing you the combined alpha of all the layers in the composite.

Blur Amount

The Blur Amount slider determines the amount of the blur applied to the alpha. Higher numbers mean more blur.

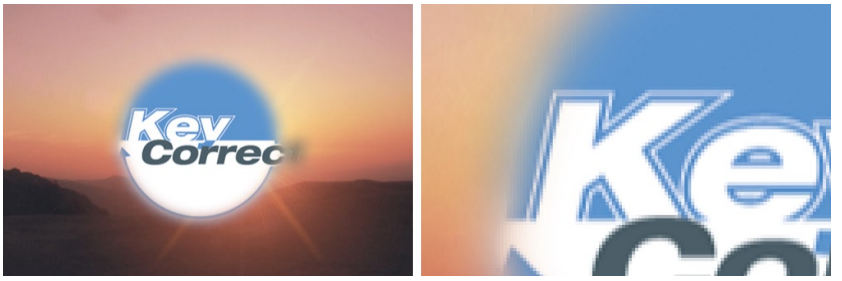

Edge Width

The Edge Width slider determines the width of the blurred edge.

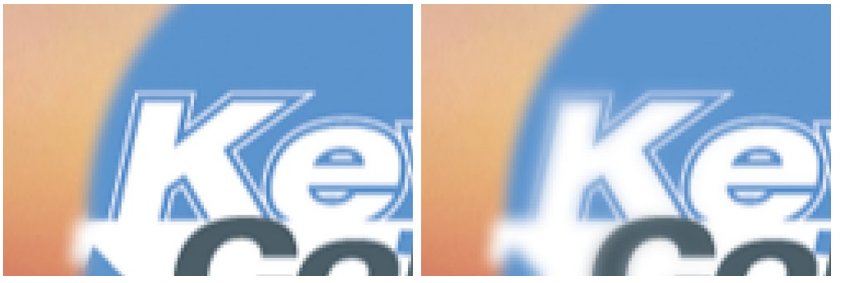

The images above show Edge Widths of 4 and 40 respectively. Note

that there is no additional blur applied to the second image, the only difference is the wider edge.

Edge Finder

Edge Finder creates a precise border along the edge of a layer's alpha channel, complete with subtle softening and blurring of the edge. The resulting edge matte can be used to limit various effects to the edge areas of the element, allowing precise control over the exposure, color balance, sharpness, blurriness, or transparency of the composited element. Edge Finder can also be used in a creative manner, to place stylized borders around your moving foreground elements.

When you first apply Edge Finder it will, by default, be using the alpha channel as the control layer for the edge effect. As such, Edge Finder will do nothing if applied to a layer with a completely solid alpha channel. If you apply the filter and see no results, check the alpha channel of the source layer and make sure it is not completely white, or change the Channel pop-up menu to a value other than alpha.

Using Edge Mattes

Before we continue, we need to go over the concept of using edge mattes, and to do so we will use an example of multilayer blur. All compositing packages, including After Effects, contain a number of functions to blur elements. Key Correct Pro includes a number of options itself. Because of the layered nature of compositing, however, when you apply a blur function to the edges of a foreground layer, you are simply blurring the foreground pixels over the background. In other words, the pixels in the background are not changed.

The Basics

Consider the image below, of a boy with a pond in the background. If we zoom in on the edge of the image you can see that there is not a harsh edge between the foreground and background; pixels from the boy's flesh are intermingled with green pixels from the pond water.

To illustrate the principle of multilayer blur, we have created a rotoscope matte to isolate the boy from his surroundings. The images below show the masked boy and the corresponding roto matte.

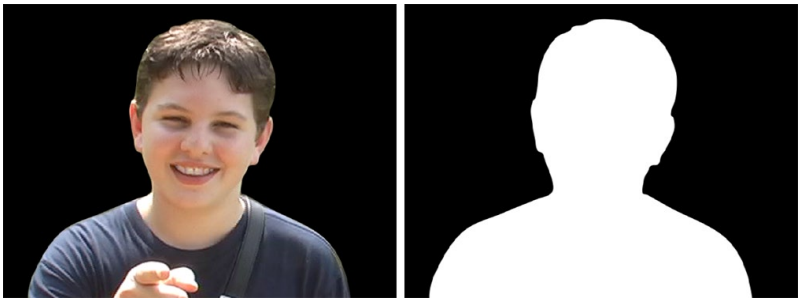

Below we see the basic composite for the boy. The new background is starkly different in color and tone from the original background.

If we zoom in on the image we can see that the edge is unnaturally harsh and smooth, which is often the case with rotoscoped elements. To fix this problem you might apply a simple edge blur to the element.

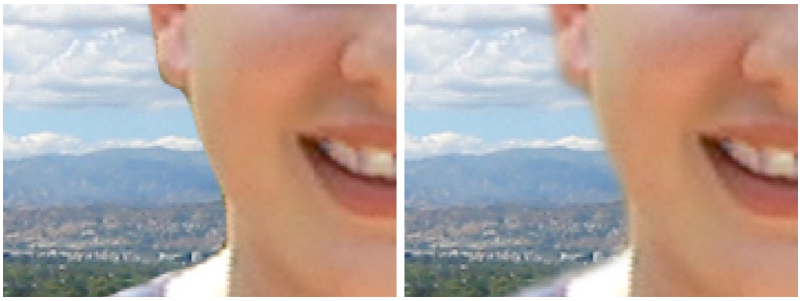

The image on the right has a simple four pixel edge blur applied. While this softens the edge nicely, note that the background pixels which come into contact with the edge are still sharp. They have not been affected by the edge blur, because the blur has only been applied to the foreground layer. The solution to this problem is to isolate the edges of the foreground layer, combine the foreground and background elements into their own composite, and then blur them together.

Generate the Edge Matte

The first step is to generate a matte for the edge using Edge Finder. The easiest way to do this is to precompose your foreground element with the corresponding matte, whether it is keyed or rotoscoped, then apply Edge Finder. (We will go over the specific procedures for the use of the tool later in this chapter.)

Adjust Edge Finder's properties as desired, and eventually you will end up with a matte for the edge of the element, which will be identical in both the foreground and alpha channels.


Isolate the Comp Edge

The next step is to create a new composite of your foreground over your background, then isolate the edge area. In this case the composite we need is the boy foreground element over the panorama background.


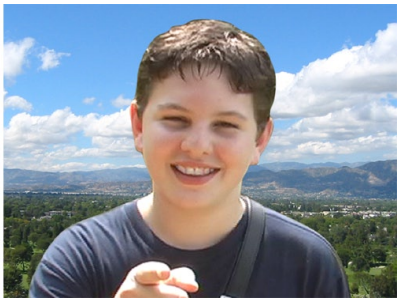

You then want to create a third composite, into which will go your precomposited foreground/background, and your edge matte element. These are shown below.

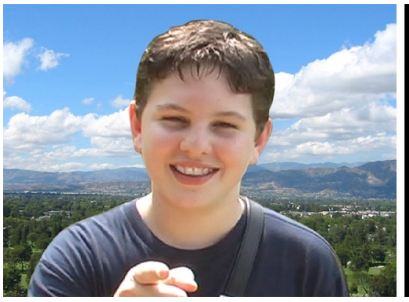


Here the edge matte will be applied as a mask (Alpha Matte or Luma Matte) to the precomped background layer.

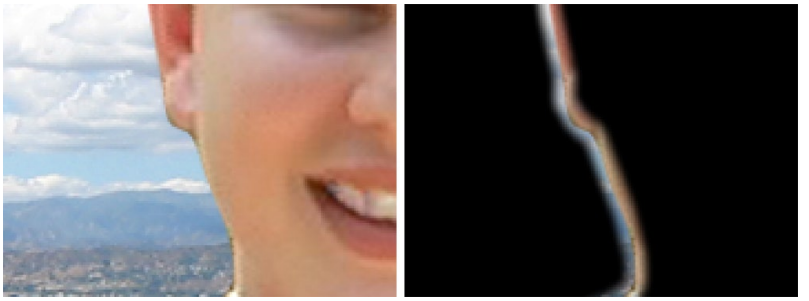

We now have both the foreground and background elements combined together, with the edge matte masking out everything except the area around the foreground edge.

Combine the Elements

The next step in our process is to combine the edge element with the foreground and background elements. Create a new composite, and place into it your foreground, background, and edge elements. Make sure that the edge element is above the other layers.

If we disable the foreground and background elements for a second, we will be able to see the next step more clearly. Select your edge matte element and apply a blur filter, such as the Speedy Blur filter included in Key Correct Pro. The effect is immediate and obvious.

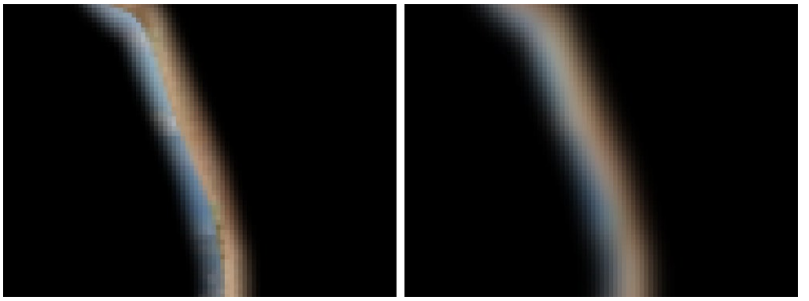

Rather than the foreground being blurred over an unchanged background, the foreground and background pixels are blurred together, more closely mimicking the way the edges in filmed elements actually appear.

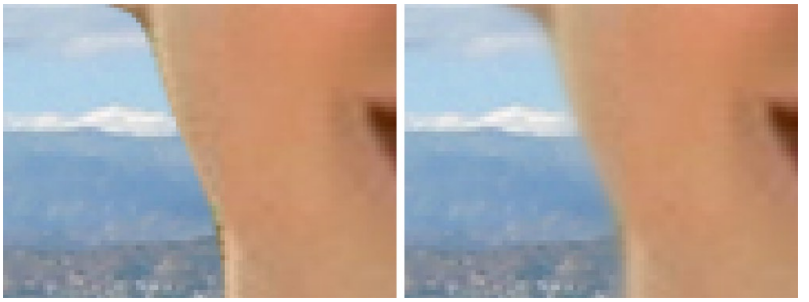

While this may seem like a minor detail, this type of technique is often the difference between a good composite and a perfect composite. When combined with the other tools in Key Correct Pro, there is literally no limit to the degree of accuracy you can achieve in your projects.

Summation

Blurring the edges of composited elements is by no means the sole use of Edge Finder, it is but one of many. As we proceed to explain the Edge Finder interface below, we will cover some more of these additional functions. No matter what use you have, it will be performed in a manner similar to the multilayer blur example above.

Now that we have gone over the fundamental logic behind Edge Finder, we will now go over the specific functions and tools available in the interface.

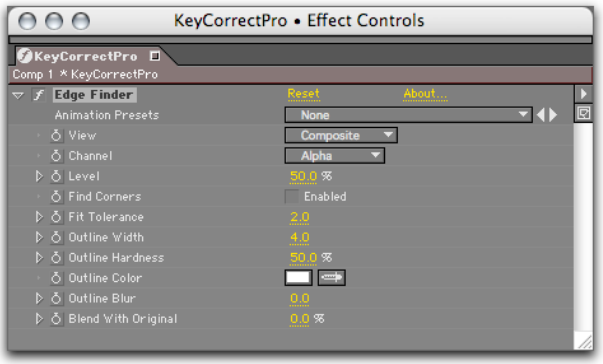

View

The View menu contains the options for what is visible in the Composite window. The examples below assume that Edge Finder is being applied to an element that is the only layer in the Composite window.

Composite

This option displays the outline over the foreground layer. The alpha channel generated represents the original foreground plus the edge matte.

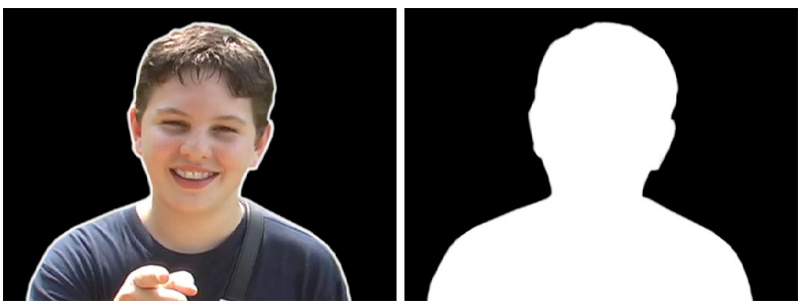

Outline

This option displays only the edge in the RGB, with a corresponding alpha channel.

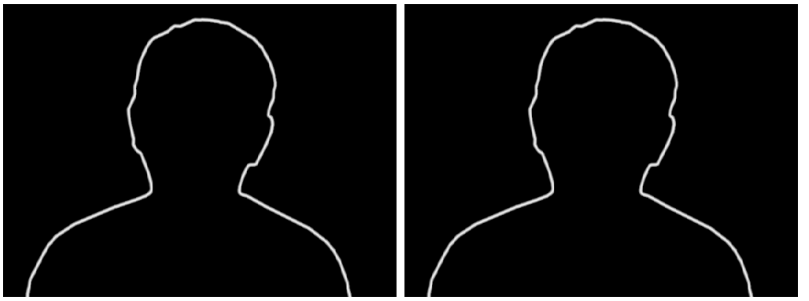

Note that the color of the generated edge is white by default, which will make the RGB image identical to the alpha unless you manually change the color of the edge. (See below for more on the Outline Color.)

Outline Mask

This mode generates the outline filled with black in the RGB, as well as a corresponding solid alpha for both the outline and the black fill.

Channel

The Channel value determines which channel Edge Finder analyzes to generate the edge.

Red/Green/Blue

These three controls allow you to designate either the Red, Green, or Blur channels. If you are trying to isolate part of the image to reduce a specific color value, choose the appropriate channel to use as a mask for a color correction filter.

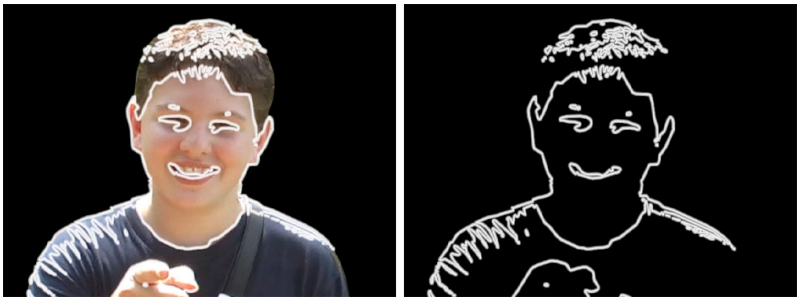

Use a wide blur value to match the natural gradients in the foreground object.

Luma Luma generates an edge matte based on the luminance values of the image.

Lightness Lightness generates an edge matte based on the lightness values of the image.

Brightness Brightness generates an edge matte based on the brightness values of the image.

Level

The Level value determines the level of the source channel used for the contour which is outlined. To easily illustrate this effect we will use a circular white-to-black gradient.

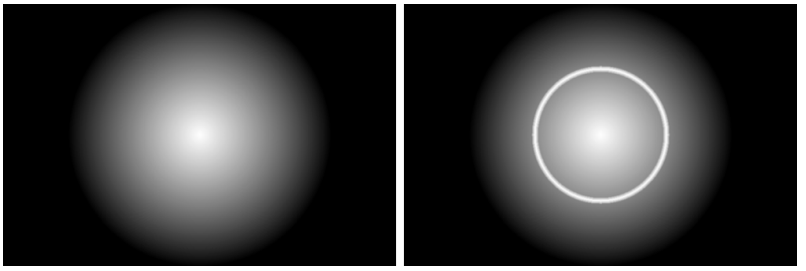

At the default value of 50% the line is drawn at pixels half way between the brightest and darkest found in the image.

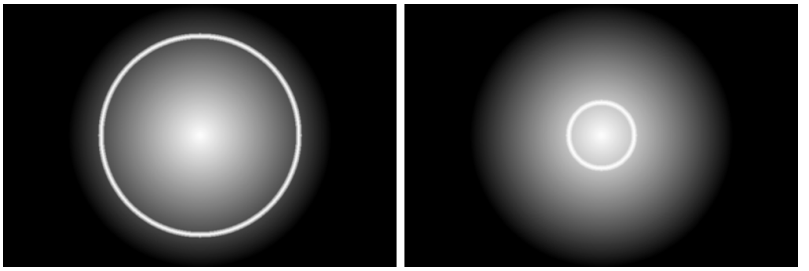

Here we see values of 25% and 75% respectively. To move the line towards the darker pixels in your layer, decrease the Level value.

Find Corners

Find Corners enables extra processing to find corners and sharp edges in the generated edge. Enable Find Corners to generate an edge on text or geometric shape layers.

Consider the text element above. This is what the edge looks like with the default Edge Finder effect applied.

These images show the same text element with Find Corners enabled. Notice the sharper edges along the corners of the T and the inside edge of the E.

Fit Tolerance

Fit Tolerance designates how far the outline can vary from the contour that created it.

The images above show Fit Tolerance values of 2 and 20, respectively. As you can see, the lower the number the more closely it adheres to the true edge of your image.

Outline Width

The Outline Width value designates the width of the generated edge.

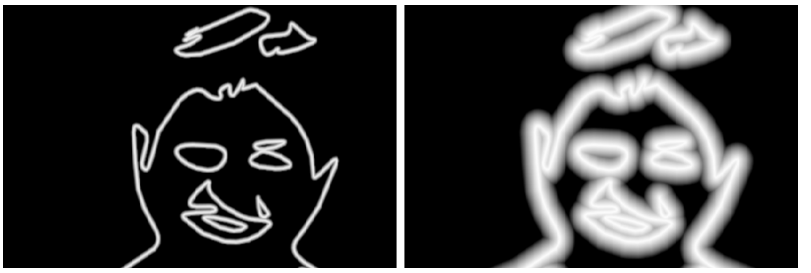

Higher numbers yield a wider edge.

Outline Hardness

If you think of the outline as being a brush stroke, the Outline Hardness value determines the flatness of the brush used to paint the outline

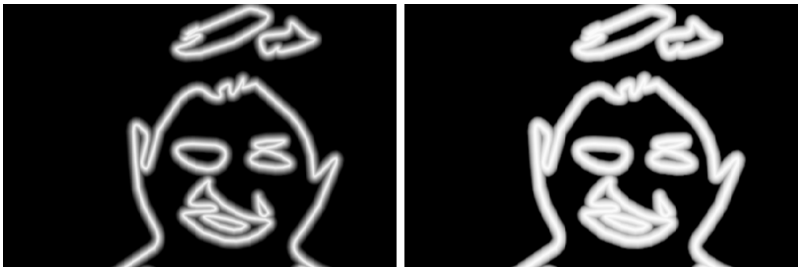

Higher numbers yield a flatter brush.

Outline Color

Outline Color determines the color of the outline. The default value is white.

Outline Blur

Outline Blur designates the amount of blur applied to the outline. Increase this value to soften your edge. Try a value of 2.0 to 3.0 to smooth out the edge of the outline.

Light Wrap

Light Wrap is the secret to seamless layer integration. It works by wrapping the edges of a foreground layer with soft light from the background layer, producing a powerful level of integration by providing the illusion that light from the background layer is reflecting along the edges of the foreground.

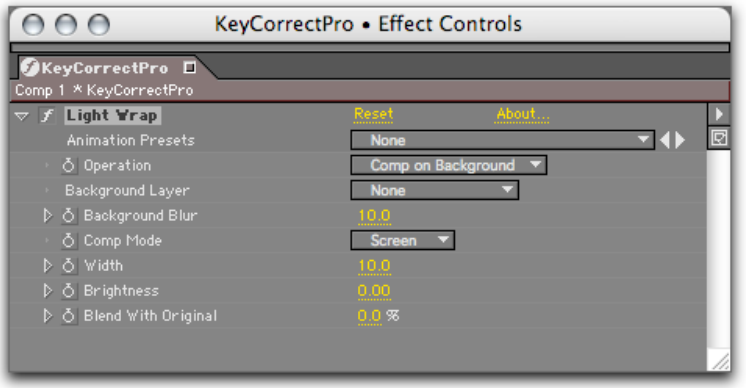

Operation

The Operation pop-up determines the output from Light Wrap. There are three modes.

Comp On Background

Comp on Background is the default output, and the one you will use most of the time. It displays the light-wrapped foreground composited over the background element.

The left image above shows the keyed and color corrected foreground, and the right shows the same image with Light Wrap applied. In this mode the output alpha is solid white.

Because the effect is often very subtle, here is a close-up of the shot. Note the diffuse wrapped edges, especially in the area around sun in the background.

Wrap Only

Wrap Only displays only the background; the foreground element disappears. However, the edge area is placed in the matte. This mode allows you to set up the Light Wrap effect as a separate element, which can be rendered out on its own to be used in another composite.

The image above shows the RGB and Alpha channel generated when this mode is selected.

Wrap On Black

Wrap on Black works virtually identically to the Wrap Only mode, except it precomposites the masked RGB layer over black and renders it out with a solid white alpha. The output from this mode allows you to bring in the element and composite it over the background by setting its transfer mode to Screen.

Background Blur

Background Blur

The Background Blur control designates the amount of internal blur applied to the background layer before the light wrap effect is

applied. This is a pre-process only; it merely softens the wrap effect, and will not alter the background element in any way.

Comp Mode

The Comp Mode menu determines the compositing mode used in Light Wrap. These modes are identical in function to the transfer modes found in After Effects. For detailed explanations of these transfer modes, please refer to your After Effects documentation.

The two most common Composite Modes you will use with Light Wrap are Screen and Add. To illustrate these two modes we will use our keyed foreground over black, as shown in the images below.

Screen multiplies the inverse brightness values of the colors in all layers.

Add combines the color values of the layer and underlying colors. The resulting color is lighter than the original.

Width

The Width value controls the width of the light wrap effect. Higher settings create a bigger backlight wrap on the layer. The image below shows Width values of 10% and 25%.

The Width value operates independently of Background Blur, allowing you to set a high blur value on the background element while still retaining independent control of the width of the effect.

Brightness As its name suggests, Brightness controls how bright the wrap effect will be. Higher settings are brighter, negative numbers are darker.

Matte Feather

Matte Feather creates a gradual change, or feathering, from opaque to transparent at the edge of a layer's alpha. It can be used to correct dirty or ragged layer edges where the keying function did not do a precise enough job. It can also be used to soften the hard edge from a computer generated object to help it make it fit into an organic scene.

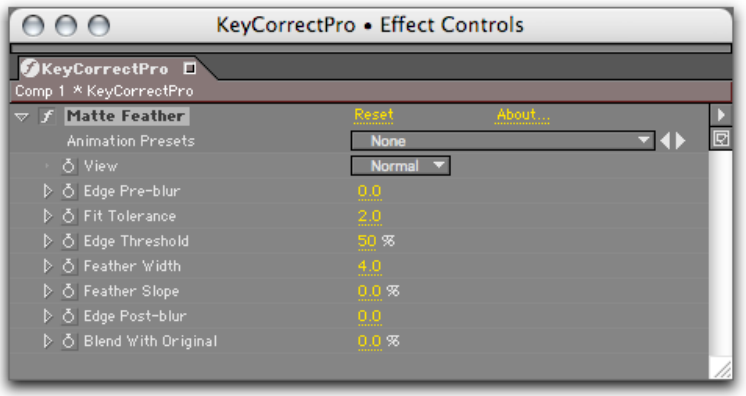

(For a detailed explanation of the difference between feathering and blurring, please see *Blurring Versus Feathering* on page 12.)

View

The View Menu provides you with the option of viewing either the default Composite Window content, or the alpha channel of the layer to which the filter has been applied.

Edge Pre-Blur

The Edge Pre-Blur value slightly blurs the edge of the alpha before applying the feathering effect, which is a good way of removing jaggies from the edge of your object. Low settings are recommended, as it changes the shape of the edge and might cut into the layer if it is set too high, which can often significantly alter the edges of your foreground element.

Consider the edges of the boy above. As you can see in the left image the edges are quite jagged and unnatural. The right image shows the edges with the default Matte Feather values applied. As you can see, we have lost quite a bit of edge detail as is. This can be fixed with subsequent tweaking. To illustrate the effect of Edge Pre-Blur, however, consider the following two images.

On the left we see an Edge Pre-Blur value of 2. Note that a significant amount of edge detail has been lost around the ears. On the right we see the result of an Edge Pre-Blur value of 15. The head has lost almost all its detail and is virtually round.

Edge Pre-Blur works by finding the edge and then tracing it internally with a sculpting operation. Blurring the alpha before this sculpting operation takes place makes the sculpting function in essence "wander away" from the edge.

To put it simply, use this tool sparingly. A little Edge Pre-Blur goes a long way.

Fit Tolerance

The Fit Tolerance function smooths the alpha edge into a rounder shape. A larger tolerance results in a smoother curve, but the trade

off is that the calculated edge will deviate from the true edge by a higher degree.

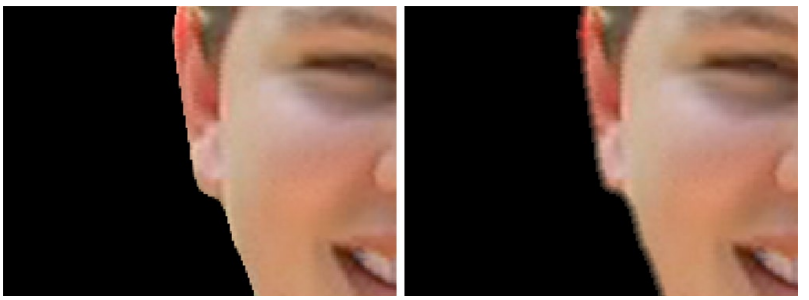

For example, consider the images above. On the left we see the original jagged edge. On the right we see the Matte Feather filter applied, with the Fit Tolerance value set to 0. (The default Fit Tolerance value is 2.) As you can see the edge is softened, but the jagged edge is still there

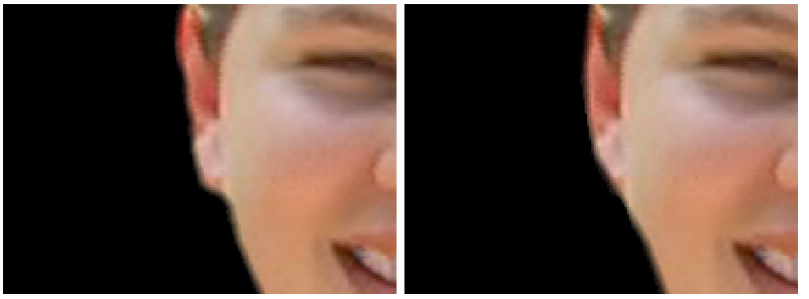

In the images above we see the default value of 2 on the left, and a value of 4 on the right. Note that the right image, while significantly more smooth, has deleted a lot of important shape and edge detail.

Edge Threshold

The Edge Threshold value designates where the feathered edge begins to become transparent, relative to the absolute edge of the alpha.

Higher numbers move the feathered edge inward. A value of 0 places the feathered edge right at the actual absolute alpha edge.

Feather Width

The Feather Width value controls the width of the feathered edge applied to your element.

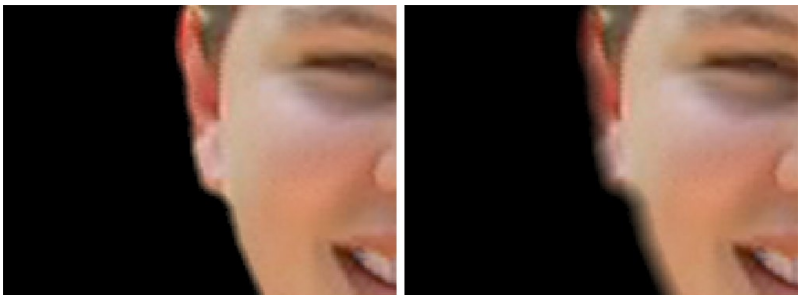

The higher the number, the thicker the feathered edge.

Feather Slope

The Feather Slope value specifies the abruptness of the transition from opaque to transparent. Lower numbers make the transition more gradual.

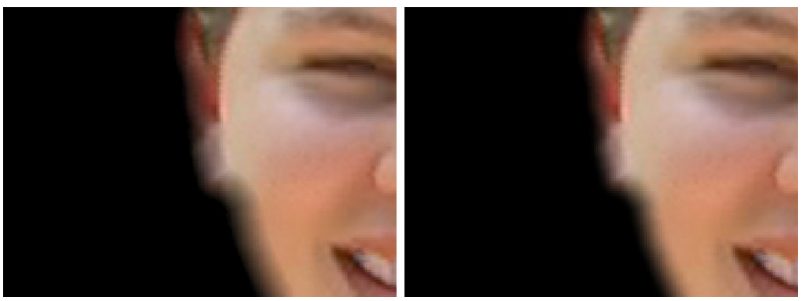

Often times the Feather Slope value can be very difficult to see on thin edges. To illustrate its effect we will give our edge a really high Feather Width value. The right image shows the edge with a Feather Slope value of 90%. Note that the inside edge of the feather does not move, the outside edge of the feather comes inward as the value is increased.

Edge Post-Blur

The Edge Post-Blur simply adds a little blur to the edges of the alpha. The RGB channels are not affected by this blur.

Matte Feather Sharp

Matte Feather Sharp is very similar in function to Matte Feather, except it uses an algorithm optimized for complex, hard-edged shapes. Using other blur or feather filters on images with sharp edges usually results in a softening of any hard angles . Matte Feather Sharp can retain these hard angles while feathering them.

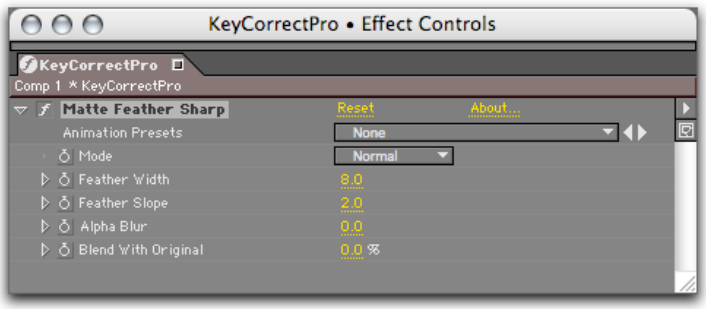

For the purposes of illustrating how this filter works, we have created a simple brushed metal texture containing a very bizarre alpha channel, one with a number of very sharp edges.

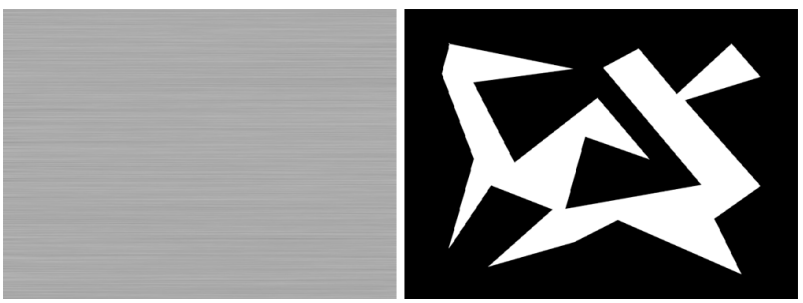

(For a detailed explanation of the difference between feathering and blurring, please see *Blurring Versus Feathering* on page 9.)

View

The View Menu provides you with the option of viewing either the default Composite Window content, or the alpha channel of the layer to which the filter has been applied.

Feather Width

The Feather Width value controls the width of the feathered edge. Higher numbers created wider edge feathers.

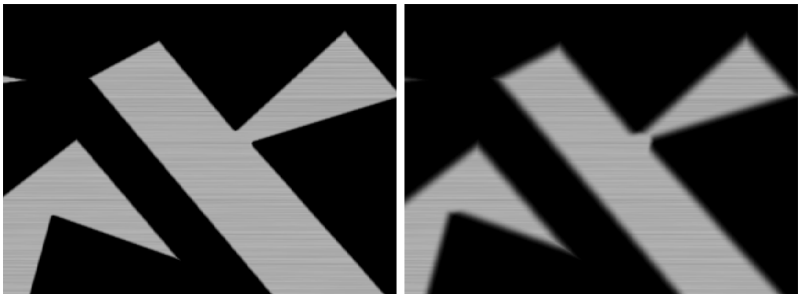

The image above shows the edge feathering at values of 4 and 16 respectively.

Feather Slope

The Feather Slope value specifies the abruptness of the transition from opaque to transparent pixels along the edge of the alpha. Lower numbers make the transition more gradual. Since the effect is easier to see on edges with a higher Feather Width, our example images will be using a high Feather Width value of 16.

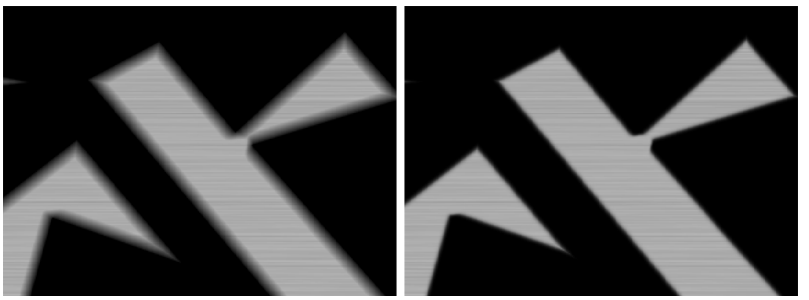

The left image above shows a Feather Slope value of 1. Note that there is a long transition from fully solid pixels to fully opaque. Increasing that value to 4, however, makes for a much sharper transition.

Alpha Blur

The Alpha Blur value applies a basic blur function to the edges of the alpha channel. The RGB pixels are not blurred. You can use this to soften up your edges, but using too much can mean that you lose some of your edge sharpness detail.

Rack Focus

Rack Focus is a z-buffer depth blur filter, enabling you to create realistic depth of field simulation where a z-buffer or "depth map" is available. Depth maps are grayscale images which can indicate the distance from the camera to specific points within the image. Rack Focus can reference this depth map to progressively blur an image as it recedes in space away from or approaches the viewer. By animating the virtual camera parameters within Rack Focus the feel of a camera operator changing focus as a scene progresses can be created with startling realism.

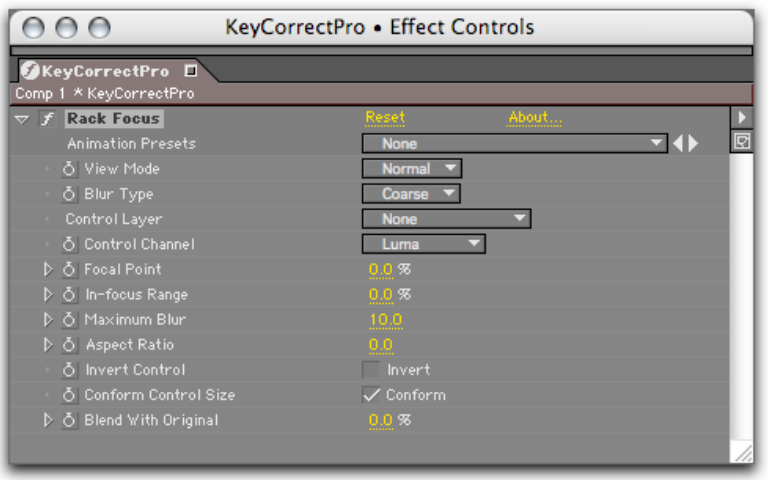

There are two ways that depth maps can be generated. The first is as part of a multi-pass render from within a 3D program. The second is a manually created image used with real-world footage shot with a regular camera. While we will demonstrate how to use Rack Focus with both types of elements and depth maps, for the purposes of explaining the function of the filter we will be utilizing an image and map created in a 3D application.

The Example 3D Environment

The images below show a side and top view of the 3D environment created to demonstrate the use of the filter. As you can see the word DEPTH is in green, and the numbers 12345 are to the right of

it in red. The letter D and the number 1 are both at the same level of depth away from the camera. Each letter subsequently moves away from the camera, and each number moves towards it.

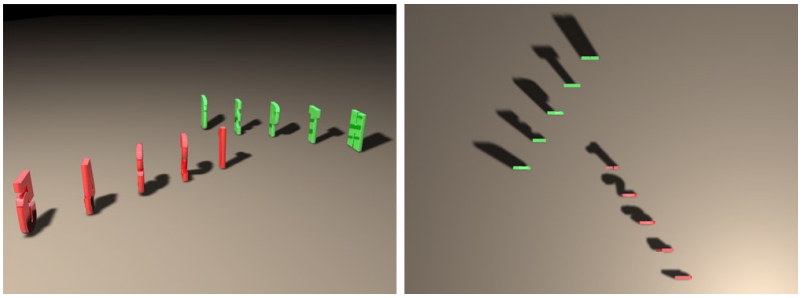

Below is the rendered perspective camera image and its corresponding grayscale depth map.

As you can see, the camera is situated in 3D space closest to the number 5. Each subsequent number is further away from the camera, and its corresponding area in the depth map becomes lighter. The number 1 and the letter D are the same distance from the camera, so they have the same grayscale value in the depth map. Each letter in turn is further away from the camera still, and become lighter and lighter. The lightest pixels are the area shortly behind the letter H, which represent everything else in the distance. Also note that the ground plane is represented by a smooth gradient, indicating each pixel's distance from the camera.

Much as an alpha channel determines transparency, a depth map determines distance from the camera. The lighter a pixel, the more blur will be applied to it.

Antialiasing

When rendering out an element from a 3D program the edge of the RGB element will undoubtedly be antialiased, and chances are the grayscale depth map will be as well. This, however, can cause unwanted artifacting around the edges of your blurred RGB layer.

When the antialias function smooths the edges of the grayscale depth image during rendering it adjusts the luminance values of the edge pixels. In an alpha channel this is a desired effect, as luminance designates transparency. However, in a depth map luminance designates distance from the camera, and the changes in pixel luminance from antialiasing indicate that the edge pixels of an object are a different distance from the camera than the object itself.

In the left image we see our depth layer with smooth antialiased edges. On the right is the same depth layer with jagged, rough natural edges. Your depth map should, if possible, look like the one on the right. Make sure that you render out your depth pass with antialiasing disabled.

The 3D images in this chapter were created using Cinema 4D, courtesy of Maxon Computing. For more information visit www.maxon.net.

View Mode

The View Mode menu contains two options for working with Rack Focus.

Normal

Normal shows the standard RGB composite view. As the blur effect is applied to the RGB it will be visible in this mode.

Focus

The Focus view mode shows the depth map. Subsequent controls in the Rack Focus interface make changes to the focus mode, and these are visible here.

When in Focus mode the visible image will be in grayscale *except* for the area that is 100% in focus.

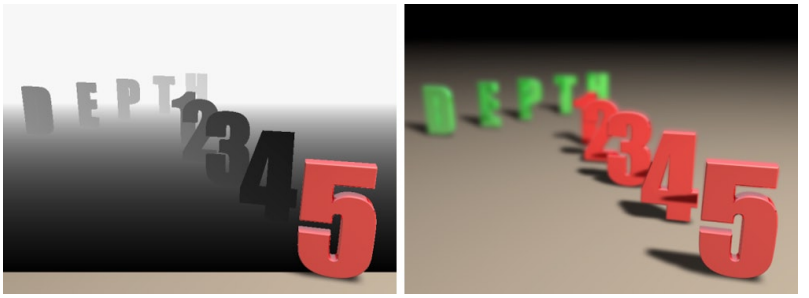

In the image above we see that the only part of the depth map in color is the number 5. In the corresponding RGB image we can see that the only part of the image completely in focus is the same part of the 5.

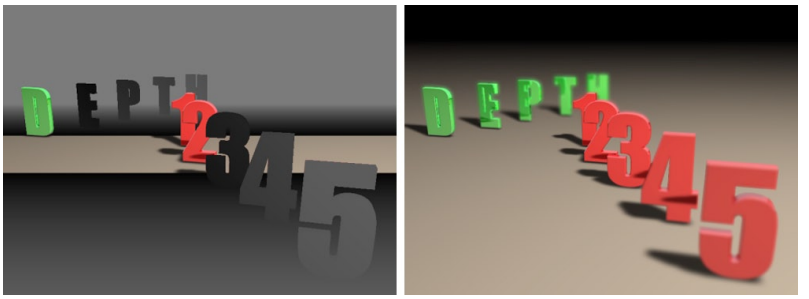

Conversely, in the above images the color section comprises the area of the letter D, and the numbers 1 and 2. In the RGB image those numbers and letters are the only part of the image in focus.

Blur Type

The Blur Type menu designates which type of blur is applied to the RGB.

Coarse

Coarse is the roughest and fastest rendering blur function. Mathematically this is known as a box blur.

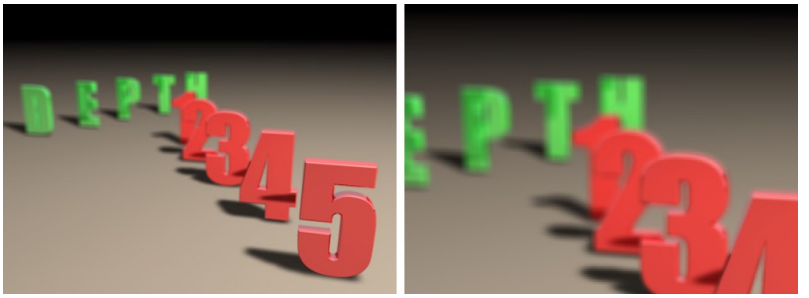

As you can see there is a large degree of square artifacting on the image. Coarse blurs are good for background elements where fine detail is not of paramount importance.

Fine

Fine is the highest quality blur. This is known as a double-pass pyramid blur, and is virtually indistinguishable from a traditional Gaussian algorithm, while still retaining a significant speed advantage.

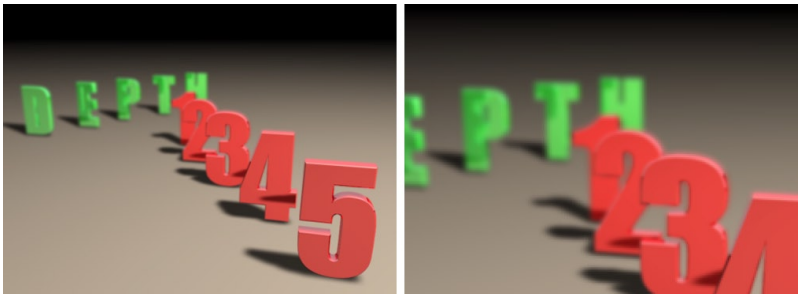

The overall detail, including the curvature on the P and the notch in the top of the H, is preserved to a great degree.

Control Layer

The Control Layer specifies which layer in the composite will control the blur effect.

Control Channel

The Control Channel specifies which attribute of channel of the Control Layer will be used to control the blur. Luminance, the default value, is the most common choice. The other attributes allow for very specific control to create artistic or stylized blur effects. Note that when using a grayscale depth map image, all of these

values will return the same result.

Focal Point/In-Focus Range

The Focal Point and in-Focus Range values work in tandem to set the blur area, and as such they will be explained together. To illustrate these values we will be viewing our element with the View Mode set to Focus.

Focal Point

The Focal Point value designates the point in the depth map where the RGB image will be in focus. In practical terms, think of this point as the distance from the camera which will be perfectly clear. Everything in front of and behind this point will be blurred.

In-Focus Range

The In-Focus Range value determines the degree of focus in the image. Where the Focal Point value sets the point in space which is in absolute focus, the In-Focus Range value determines how far in front of and behind this point is also in focus. In other words, it determines the width of the focal area, expressed in terms of the distance from the camera.

Focal Point/In-Focus Range Usage

We will now look at practical examples of how these two values work together.

The left image shows our focus zone with the default Focal Point value of 0%. Only the front of the image is in focus. The darker pixels are in the very front of the image and the lightest pixels are in the back, meaning that the blur will be applied mostly to pixels in the distance. If we increase this value to 45%, we see the right image. The focal point is roughly in the middle of the image. The least

amount of blur will be in the area of the red 1, with blur increasing in front of and behind it.

Say that we want the focal point of the image to be in the middle. With the default In-Focus Point value of 0%, we see that the actual range that is in focus (shown in color) is only a pixel or two wide. Increasing the In-Focus Point value will expand the area that is in focus.

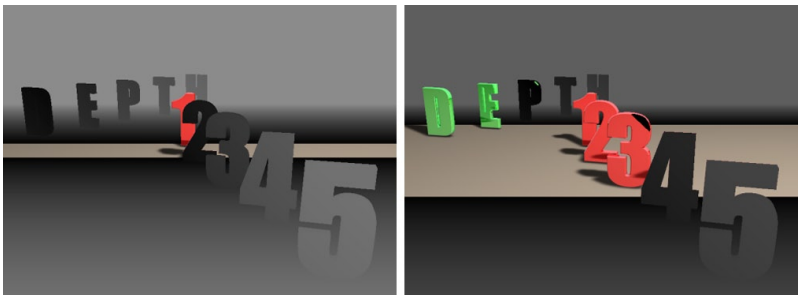

The left image above shows an In-Focus Point value of 5%. The number 1 is now completely in focus. On the right we see an In-Focus Point value of 23%. At this setting the letters D and E and the numbers 1, 2, and 3 are all in focus.

Rack Focus

You've seen it in films a million times. There's a shot of a room. In the foreground is an object, a telephone for example. In the background is a door, completely out of focus. The phone rings, and in walks a woman to answer it. In one move the focus shifts from the phone to the woman. The phone, previously in focus, is now completely defocused. This type of animated focus is known as a *rack focus*. Simulating a rack focus in this manner is as simple as animating the values for the In-Focus Range and Focal Point.

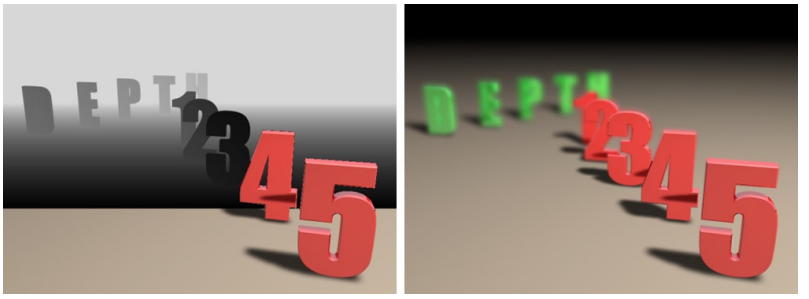

Consider the image above. We've set the Focal Point to be the very

front section of our image, and the In-Focus Range to encompass the area of about two numbers.

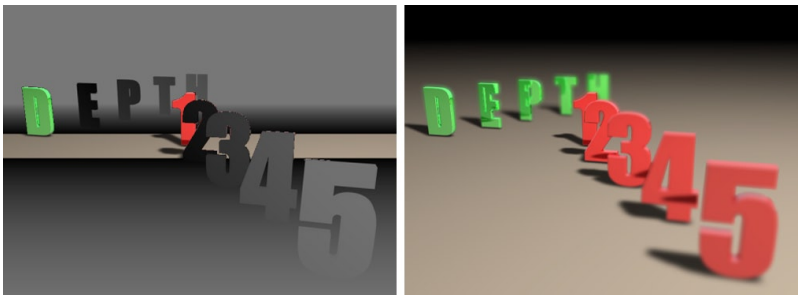

In the images above we've animated the Focal Point to be roughly the middle of the image. Note that while we haven't changed the In-Focus Range value, it is smaller than it was on frame 1, which is representative of the perspective in the shot. The numbers 4 and 5 are now defocused, while the middle of the image is in focus.

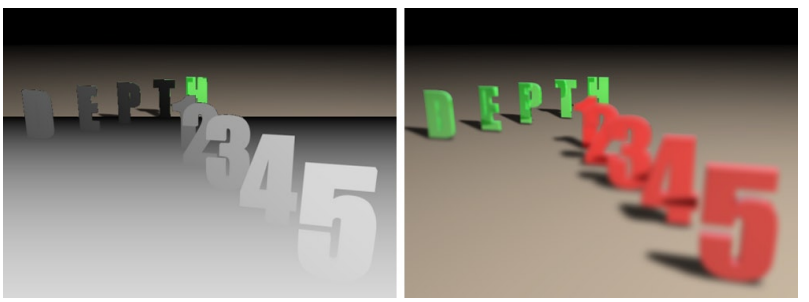

On the final frame of our animation we've animated the Focal Point to the very back of the frame. The foreground is now completely defocused, with only the letter H and the area behind it in focus.

A perfect rack focus, every time.

Maximum Blur

Maximum Blur controls the maximum amount of blur applied to the image. Because the amount of blur applied to each pixel is designated by the control layer, this value sets the high threshold amount for blur applied. Higher numbers produce more blur.

Aspect Ratio

The Aspect Ratio value sets the aspect ratio bias for the blur effect. At the default value of 0 the blur is applied in equal amounts along the X and Y axes.

The left image shows an Aspect Ratio value of -1. The blur is applied vertically along the Y axis only. The right shows a value of 1, with the blur applied horizontally along the X axis.

Invert Control

Rack Focus normally interprets a depth map with black pixels being closest to the camera and white being the farthest away. The Invert Control checkbox simply inverts the current value of the depth map. Any animated effects of changes to the map will also be inverted. This is particularly useful when working with output from a 3D program which generates an inverted depth map image.

Conform Control Size

Conform Control resizes the Z-map image to be the same size as the current layer. This has no effect unless the Z-map is a different size.

Working with 2D Elements

In the previous examples we have been working with an element rendered from a 3D program with a corresponding rendered depth-map. We'll now look at how to use Rack Focus with standard filmed 2D images.

The images above show a standard panoramic image and a corresponding grayscale depth map. The depth map was created in Adobe Photoshop™ using various selections filled with gradients. Since the panorama was a locked-off shot we only need to make one still depth map. If this shot had pan, tilt, or zoom in it then the depth map would have to be animated accordingly.

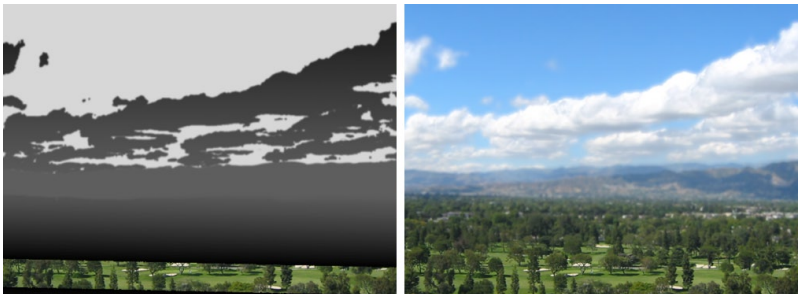

Once the depth map is created simply use Rack Focus in the usual manner. The left image above shows a focal zone right in the front of the image, and as you can see on the right this is the only area in focus.

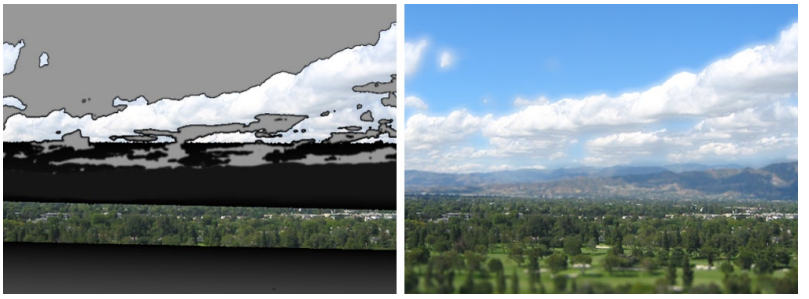

In the images above the middle section of the panorama (including the clouds) is in focus, while the rest of the image is defocused.

Smooth Screen

Video and film production is becoming more and more digital. Many major motion pictures today are filmed using HD cameras, after which the digital footage is transferred to film. Bluescreen and greenscreen elements shot on video are becoming commonplace. Additionally, many post facilities are finding that they can shoot on relatively inexpensive video equipment, then use products like Magic Bullet Suite™ to achieve astonishingly realistic filmic looks at a fraction of the cost of shooting on film. One drawback to this technique is that the reduced color of many video formats can make keying difficult.

In addition, as anyone who has ever done visual effects production at a major facility can tell you, good directors and massive budgets do not necessarily translate into well-shot bluescreen elements. Often times compositors are handed atrociously-shot elements, then expected to perform miracles to get a perfect final shot. Many times these elements are sent to the rotoscoping department to have a matte cut by hand, which is an expensive and labor-intensive process.

For all these reasons, we present Smooth Screen, which corrects poorly shot or lighted bluescreen or greenscreen elements by comparing the background pixels to a sampled color, then averaging together the background pixels to smooth out the key area. This pre-process enables even the most mediocre keying software to achieve superior results.

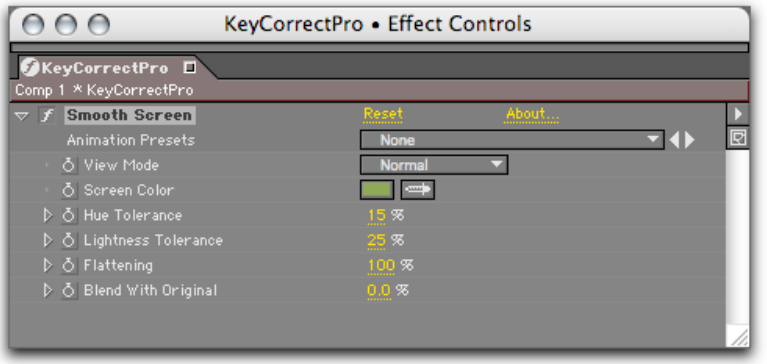

View Mode

The View Mode pop-up contains two options.

Normal Normal displays the standard Composite View.

View Strength The View Strength mode shows you a graphical representation of the area that the Smooth Screen operation will be applied to. Lighter pixels mean more smoothing.

As you tweak and adjust the settings in the rest of the filter this Strength view will update accordingly.

Screen Color

Use the Screen Color picker to sample a representative color in your background. Consider the example below. As you can see, this is an incredibly poor greenscreen, both in quality and lighting.

The left image shows the original greenscreen. The right image is the same greenscreen after a single click sample of the background. As you can see, the difference is dramatic, even with the default settings.

Smooth Screen is also extremely helpful in smoothing out areas before performing luminance or color difference keys. In the above images we see our actress sitting on some grass. One click to sample the grass area smooths it out to an incredible degree.

Hue Tolerance

The Hue Tolerance filter designates the range of hues (colors) that are included in the Smooth Screen operation. The higher the value, the more colors are smoothed.


The images above show the smoothing done with Hue Tolerance values of 15% (the default) and 40%. However, because this value concerns averaging colors in the RGB channels, pushing it too high will cause unwanted tinting in your foreground elements.


The left image shows a Hue Tolerance value of 40%. At this setting the actress' legs are unchanged. However, when this value is pushed to 50% you begin to see green tinting on her skin. As you adjust this value make sure that you are not creating unwanted tint on your foreground subject or elements.

Lightness Tolerance

Tolerance The Tolerance slider controls the upper value, defining which pixels will be corrected. The higher the setting, the more pixels will be corrected. In most cases, a setting of 100% is recommended.

Strength The Strength slider determines how much the original colors of the foreground element are corrected. A value of 100% means that the foreground colors are completely corrected.

Flattening

Flattening determines the degree of uniformity in the smoothing effect applied to the element. By default this is at 100%, and you will usually want to leave it there.

In the image above we see Flattening values of 25% and 100% respectively. Higher numbers make for a smoother key.

Speedy Blur

Speedy Blur is an incredible fast blur filter which calculates blurs at a rate many times faster than traditional blur operations. It offers you three different blur qualities, allowing you to make a trade-off between speed and quality based upon the prominence of a particular element in your composite.

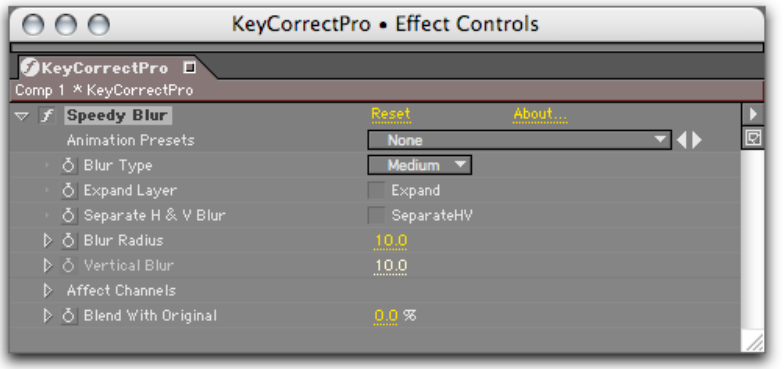

Choosing an appropriate level of blur quality can significantly speed up rendering times. For example, a background element that is mostly obscured by other layers will not require the same quality of blur as a hero element prominently visible in the foreground.

Blur Type

There are three types of blurs for you to choose from. To illustrate the blur types we will be using some footage of Notre Dame in Paris.

The building contains a good deal of intricate edge detail, which

will be good for illustrating the effects of the various blur types on large and small elements.

Coarse

Coarse is the roughest and fastest rendering blur function. Mathematically this is known as a box blur.

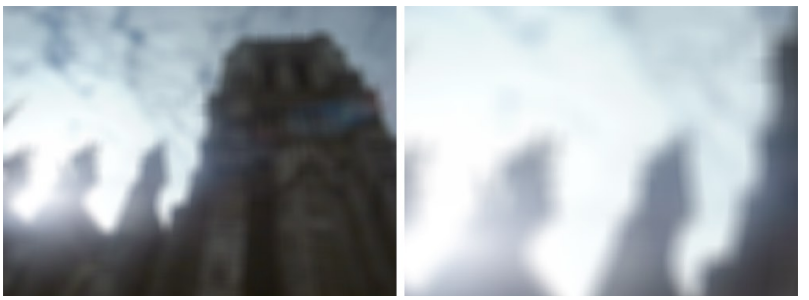

As you can see there is a large degree of square artifacting on the image. Coarse blurs are good for background elements where fine detail is not of paramount importance.

Medium

Medium is a pyramid blur function. It gives the appearance of a more traditional Gaussian-style blur, but at significantly improved render speeds.

The fine detail on the towers is preserved to a much greater degree than in the Coarse blur.

Fine

Fine is the highest quality of all the three blurs. This is known as a double-pass pyramid blur, and is virtually indistinguishable from a traditional Gaussian algorithm, while still retaining a significant speed advantage.

The overall detail, including the fine edge detail in the towers is preserved to a great degree.

Expand Layer

The Expand Layer function comes into play when you are applying a blur to a foreground object that is of a smaller size than the composite environment.

The Notre Dame footage is standard NTSC DV 720x480. On top of this we have added the Key Correct logo, which is 250x208.

The image on the left shows the logo with a 150 pixel blur applied. As you can see the blur is constrained to the 250x208 extents of the source image. Enabling the Expand Layer checkbox, however,

permits Speedy Blur to draw the blur outside these extents, giving you a natural looking edge no matter what the size of the original image.

Separate H & V Blur

The Separate H & V Blur checkbox enables you to choose whether or not you wish to have separate blur values for the horizontal and vertical axes. By default it is disabled, and the Vertical Blur value (see below) is grayed out, and the Blur Radius value determines the blur amount along both axes.

Blur Radius

The Blur Radius control determines the amount of blur applied to the layer. When the Separate H&V Blur checkbox is disabled this value controls the amount of blur applied to both the horizontal and vertical axes.

Vertical Blur

The Vertical Blur value is grayed out and disabled until the Separate H&V Blur checkbox is enabled. When active, this control allows you to set a separate value for the blur along the vertical (Y) axis, with the Blur Radius value controlling the blur along the horizontal (x) axis.

Affect Channels

This control provides checkboxes for the Red, Green, Blue, and Alpha channels. All four are enabled by default, meaning that any blur functions in the Speedy Blur filter will be applied to all four channels. You may choose any combination of channels to blur by checking or unchecking the corresponding checkbox.

Spill Killer

Spill Killer mathematically analyzes and corrects for colored spill, caused by bounced light from the surface of a blue- or greenscreen. When a foreground element is contaminated by reflected spill light, Spill Killer removes the spill from the tinted pixels without altering their original color. It will also correct many spill situations that other spill suppression filters cannot, including soft edges like fur and motion blur.

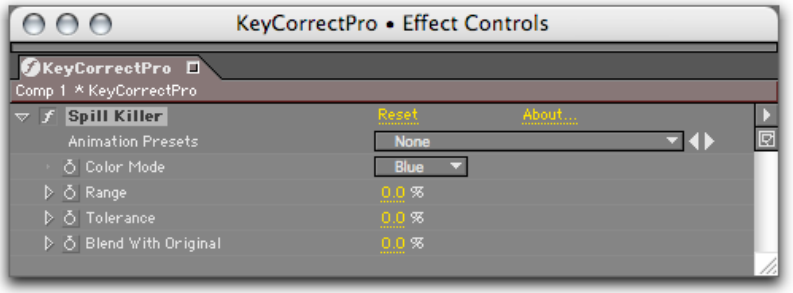

Spill Killer For the examples below, consider the following greenscreen element and the subsequent keyed footage, which has a significant amount of color spill.

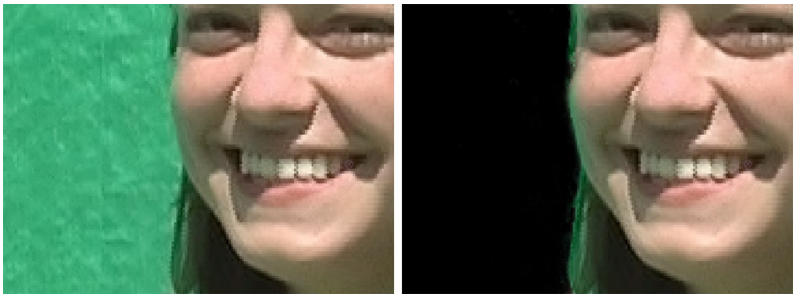

Color Mode

The Color Mode pop-up allows you to choose the default spill removal color. Choose between **Blue**, **Green**, and **Red**.

Range

The Range slider allows you to specify the lower value, designating which pixels will be corrected. The lower the setting, the more pixels will be affected. Range also defines the limit of spill suppression. The lower the setting, the more correction will occur.

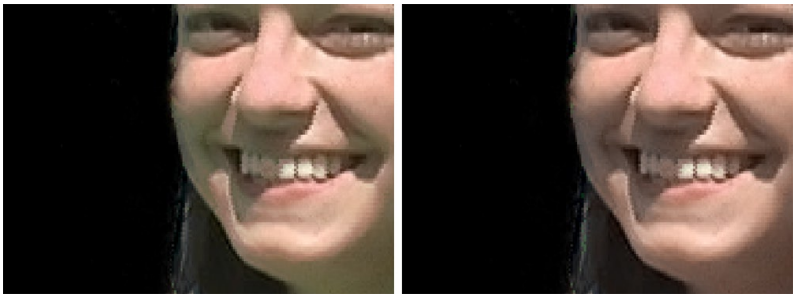

In the images above we see Range values of 100 on the left and 30 on the right. Note the magenta tint to the flesh tone areas in the image with the lower Range value. As more green is suppressed in the image a magenta tint will appear; you should keep in mind as you adjust this value in order to maintain color fidelity in the non-keyed portions of your image.

Tolerance

The Tolerance slider controls the upper value, defining which pixels will be corrected. The higher the setting, the more pixels will be corrected.

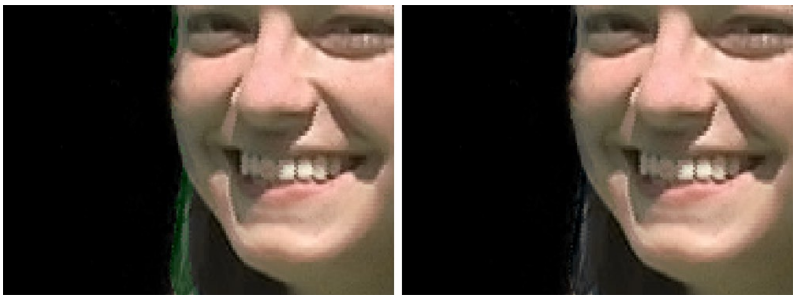

In the images above we see Tolerance values of 10% and 24% respectively.

Strength

The Strength slider determines how much the original colors of the foreground element are spill corrected. A value of 100% means that the foreground colors are completely corrected.

Wire & Rig Remover

Wire & Rig Remover replaces unwanted objects like wires and rigs with colored pixels, cloned or offset pixels, or pixels stitched from either side. Time offset cloning is also available, allowing you to pull pixels from past or future frames to cover up the unwanted object.

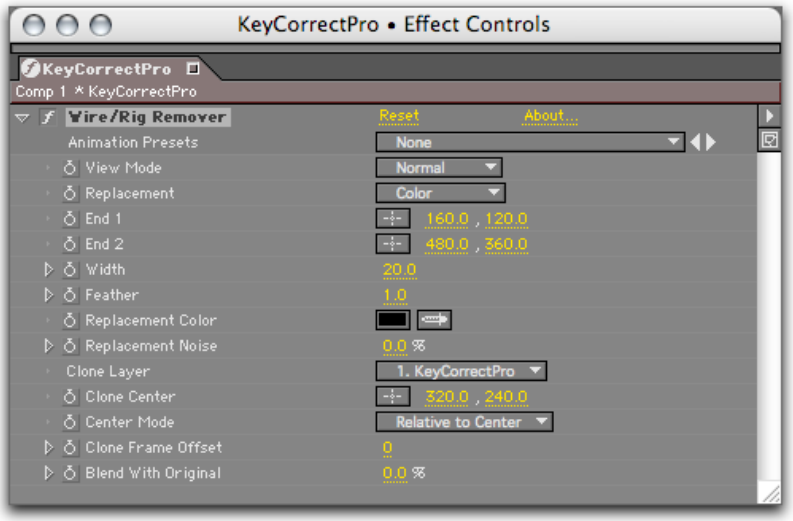

Though wire and rig removal can be accomplished by duplicating and masking layers, or by painting frame-by-frame, Wire & Rig Remover offers a fast, automated solution appropriate for most situations.

View Mode

The View Mode pop-up controls what is displayed inside the selection region.

Normal Normal displays the removal operation as applied to the layer.

Show Area Show Area fills in the selection region with red. This enables you to quickly and easily see area where the filter is being applied. It is much easier to keyframe the end point position when in this mode.

The images above show the selection range filled in with a Replacement mode (see below), and with the Show Area option enabled.

Replacement

Clone

Clone pulls pixels from any frame in the designated Clone Layer. (See below.) In addition, you can also slip the clone layer horizontally and/or vertically.

In the image above we see Clone in action. We first position the two end points over one of our tracking markers. When Clone is selected the pixels are pulled in by default from the same frame at the same location. By adjusting the Clone Offset values (see below) we tell Wire & Rig Remover to fill the Clone area with similar pixels from another part of the image.

Mask

Mask cuts a hole in the alpha, revealing the layers or background color underneath.

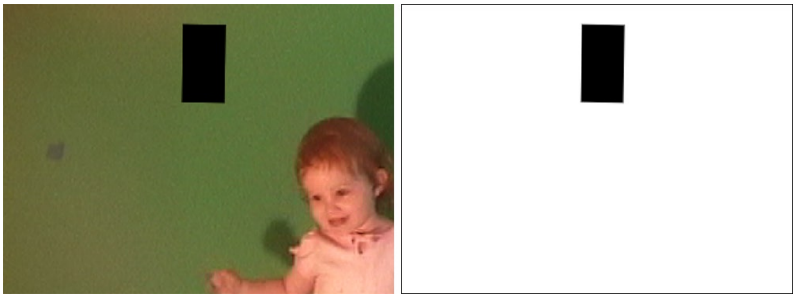

The image above shows the masked element with the corresponding alpha channel generated by Wire & Rig Remover.

Mask Clone

Mask Clone is essentially a combination of Clone and Mask. It applies the Clone effect and generates an alpha channel.

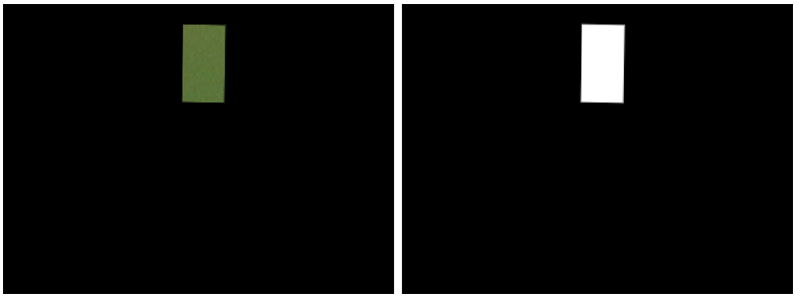

The image above shows the masked element with the corresponding alpha channel generated by Wire & Rig Remover.

Stitch

Stitch pulls in pixels from the edges of the selection area and duplicates them across it. While this will appear to be streaking, you should have no problem keying out this area since the colors used are also colors present in the greenscreen.

The examples above show the difference between Clone and Stitch.

Color

Color simply fills in the selection area with a solid color.

The example images above show the selection filled with the default solid black and with a green color sampled from the image. The color is adjusted using the Replacement Color value, which is explained below.

Ends 1 and 2

Ends 1 and 2 determine the position of the selection box. The points are located at the center of each end of the selection.

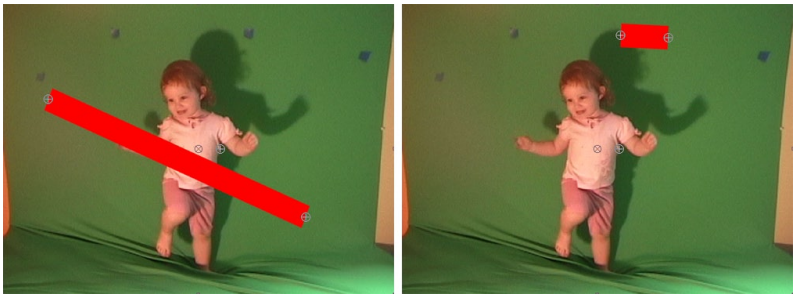

The example images above show two positions of the Ends.

Width

Width controls the width of the selection box. Higher numbers make for wider selection areas.

The examples above show width values of 10 and 45.

Feather

Feather controls the softness of the edge of the selection box.

The examples above show Feather values of 1 and 20.

Replacement Color

The Replacement Color value designates the fill color used when the Replacement menu is set to Color. (See above.)

Replacement Noise

The Replacement Noise value adds a basic noise function to the fill area, enabling you to match grain or video noise.

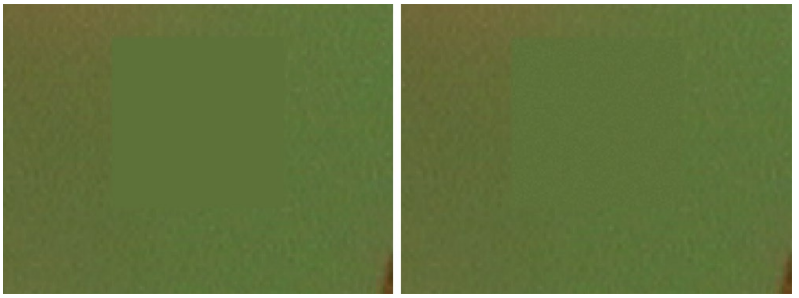

The examples above Replacement Noise values of 0 and 10%. Note that in the majority of cases the noise will not have to perfectly match the background to guarantee a successful key.

Clone Layer

The Clone Layer value is used in conjunction with Clone setting of the Replacement menu. It designates the layer from which the clone pixels are pulled from.

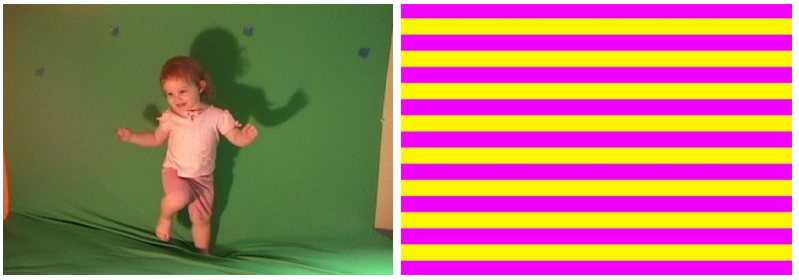

For example, consider the two layers above. If the Clone Layer pop-up is set to the be the same layer to which the effect is applied the clone source will be the same as the current visible frame. When another layer is designated the pixels are pulled from that layer.

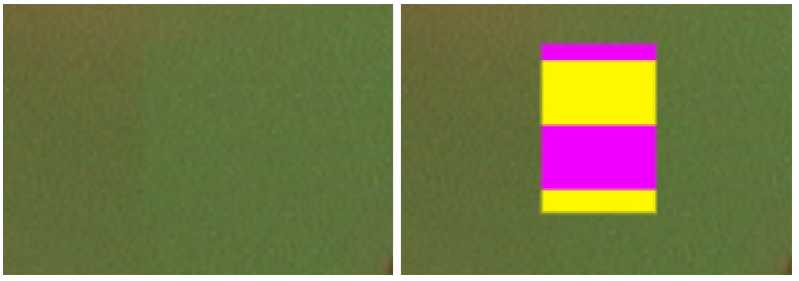

The images above show our stroke with pixels pulled from the cur-

rent frame and from the secondary layer.

Clone Center

The Clone Center values designate the X- and Y-axis pixel offsets used for cloning. Positive numbers move the offset to the right, negative numbers move it to the left.

When Wire & Rig Remover is initially applied the default values for Clone Center are set to the center point of the layer. For example, on a 640x480 image the Clone Center values will be 320,240. These values are directly related to the operation of Center Mode, explained below.

Center Mode

The Center Mode pop-up works with the Clone Center value, designating whether the values there are calculated absolutely of relatively. The two modes allow you to animate the position of the clone rectangle and have the Clone Source point move with it (absolute) or stay in place (relative). To illustrate the difference between the two modes we will use two totally different images, shown below, with the image of the girl being applied to the sky as a clone stroke.

Absolute

In Absolute mode the Clone Center is mapped to the center of the clone rectangle. Below we see the clone rectangle area, as well as the result of the clone operation with the Center Mode set to Absolute.

Note that the center of the Clone Layer is placed at the center of the clone rectangle.

If we move the clone rectangle 100 pixels to the left the center of the Clone Layer is still located at the center of the stroke.

If we increase the X value of the Clone Center value by 100 the Clone Layer shifts 100 pixels to the left.

Relative

In Relative mode, the offset between the center of the Clone Layer and the Clone Center is used for the clone source.

As you can see, the fill in the rectangle directly corresponds to the same location in the clone source.

If we move the clone rectangle 100 pixels to the left the clone source stays in the same location.

If we increase the X value of the Clone Center value by 100 the Clone Layer shifts 100 pixels to the left.

Clone Frame Offset

The Clone Frame Offset enables you to set a frame (time) offset to pull pixels from frames before or after the current frame.

ZoneHLS

ZoneHLS is a powerful color correction tool which enables you to adjust the Hue, Lightness, and Saturation (HLS) of an image by individually adjusting the parameters for its highlight, midtone, and shadow zones. This is incredibly powerful when performing precise color adjustments.

The HLS Color Model

Before we delve into ZoneHLS, it is important to have a fundamental understanding of how color correction is done in HLS color space, so we'll begin by going over the HLS color model.

The most common color space that artists working on a computer deal with is RGB. RGB color space is based on the fact that all visible colors may be formed by a combination of three different colored lights: red, green and blue. (Your computer monitor displays colors in RGB.) Colors in RGB space are defined in terms of the percentage of those three colors components found in the overall color. For example, black has an RGB value of 0,0,0, white is 100,100,100, and blue is 0,0,100. Most of you are probably well-versed in working with RGB.

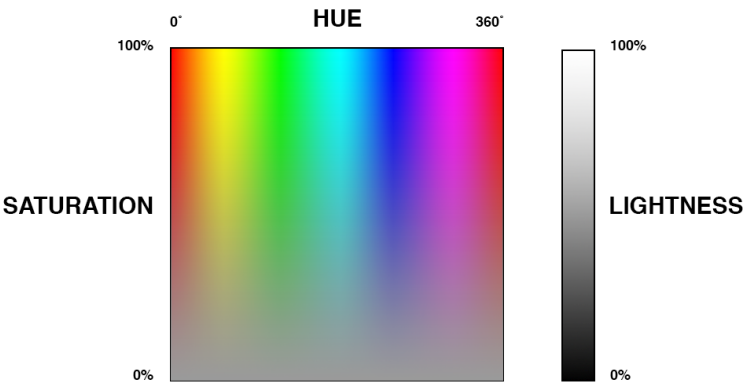

The HLS color model, however, is different in that it defines color in terms of three color characteristics: Hue, Saturation, and Lightness.

Hue Hue is what we think of when we think of an object's color. When we say an apple is red we are referring to its hue. In the diagram above the hue is expressed as the colors running horizontally in the square.

Lightness Lightness is the intensity of that color, a measure of the color's value in terms of light and dark.

Saturation Saturation is the measurement of the purity of the color. Fully pure colors are 100% saturated. Colors at 0% saturation have no hue, and appear as grayscale. If you've ever adjusted the saturation control on your TV set you will understand this. As the saturation decreases the colors on the screen become muted, and eventually you end up with a black-and-white image. In the diagram above the saturation is expressed as the colors running vertically in the square.

It's easy to understand how the components of the HLS color model work together if you think of a color in terms of a three-step process. We'll choose red as our example color.

1 The first step is to choose a hue, which is expressed as the colors running along the top of the square color spectrum. These pixels are at full intensity. The color we want is pure red, located at both 0˚ and 360˚.

2 The next step is to determine the saturation. At the top of the square the colors are 100% saturated and pure, while at the bottom the colors are completely desaturated and gray. We want pure red, so we choose a saturation of 100%.

3 As soon as we have made these two initial choices the color we choose appears at the center of the lightness gradient. Whatever combination of hue and saturation we choose is what appears at the 50% point on the lightness gradient. Because we chose pure red, the lightness gradient shows black at 0%, red at 50%, and white at 100%. To choose pure red we need to keep the lightness value at 50%, where the color is not lightened or darkened.

Top recap: Hue chooses the color, Saturation determines how intense the color is, and Lightness determines how light or dark the color is. A solid understanding of this process is essential to unlock-

ing the power of ZoneHLS, and will be elaborated upon as we go over the interface below.

The Hue Control

One other area we need to go over is the controls for the hue parameters, which are the same for all three zones.

The hue values are expressed in terms of a degree offset from 0. The circular control represents a 360˚ color wheel, as shown below.

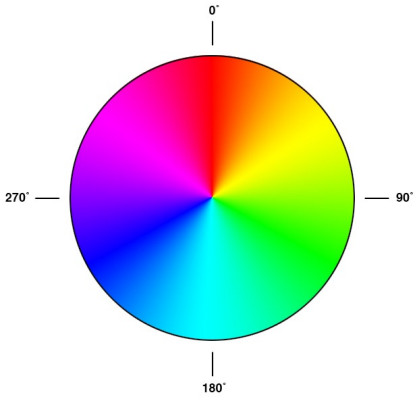

Moving the control to the right will tint the zone towards the colors found there, like yellow and green. Moving it to the left will tint the zone towards magenta and blue.

Interface

We will now go over the individual parameters found in the ZoneHLS interface. We will be using the Oval Office image that we previously used to illustrate Denoiser. The image used in this example has already been Denoised to provide superior image quality.

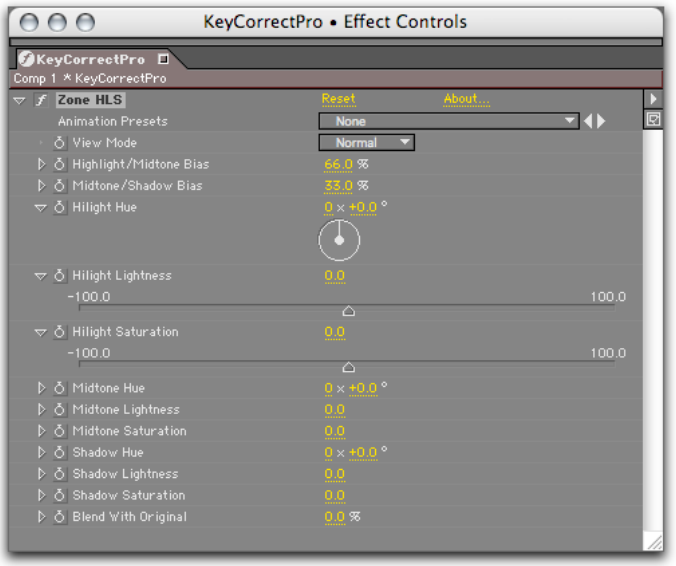

View Mode

The View Mode menu allows you to look at the image in terms of zones. For example, it's one thing to be able to look at your regular RGB image and say to yourself that you need to desaturate the highlights a little. But exactly what in the image will be adjusted? The View Mode allows you to see the exact distribution of pixels in each zone (Highlights, Midtones, and Shadows) that will be affected when you adjust its values.

Highlights

The image below shows the original RGB and the pixels in the Highlights zone.

Only the bright areas behind the desk, as well as some of the reflective areas on the desk, will be affected.

Midtones

The image below shows the original RGB and the pixels in the Midtones zone.

As you can see the majority of the image is in the midtones range. These zones work much the same as an alpha channel, in that the brighter a pixel is in the zone view the more change will be applied to that pixel when the zone's values are adjusted. Note that much of the desk is dark, whereas the walls are bright. When adjustments are made to the midtones, they will be more visible on the walls than on the desk.

Shadows

The image below shows the original RGB and the pixels in the Shadows zone.

As you can see here, the shadow areas are brightest. The walls and window are almost solid black, meaning that any adjustments to the Shadows zone will only affect the desk and other dark areas.

Level Controls

The Level controls consist of two individual sliders, Midtone/Highlight Level and Shadow/Midtone Level. To explain how these sliders work we need to go over how the zones interact with each other.

The image below shows the three zones in our image: Highlight, Midtone, and Shadow. Think of the zones as operating sort of like an alpha channel, with each pixel having a value of 0 to 100, with 0 being black and 100 being white.

A pixel can be in more than one zone. In the three images above we see each zone. Most of the pixels in the leftmost image (the highlight zone) are black. Solid white pixels are completely within their particular zone, and solid black pixels mean that that particular pixel is not represented in that zone. Pixels that are not completely black or completely white are partially within one zone or another. If you took the three images above and added them together, you would end up with a solid white image.

The Level Controls concern the pixels that lie in the area between the two zones, those pixels that are not 100% black or 100% white. They control the distribution of those pixels, whether they should be "weighted" more to one zone or the other.

Midtone/Highlight Level

The Midtone/Highlight Level control designates the distribution of pixels that occupy space in both zones. The value is expressed in terms of a percentage, with the default value being 66%.

The image above shows the Highlight and Midtone zones using the default setting of 66%. Increasing this value will bias the pixels towards the Midtone zone, and decreasing it will bias them towards the Highlight zone.

The images above show the same two zones with the value increased to 90%. As you can see, some of the highlight pixels (from the window area) have disappeared from the Highlight zone, and are now visible in the Midtone zone.

Here we see our two zones with the value decreased to 30%. The Highlight zone now contains significantly more pixels than it did previously, and the Midtone zone is far darker. By adjusting the bias of these pixels you can greatly affect how the color correction is applied to each individual zone.

Shadow/Midtone Level

The Shadow/Midtone Level control functions in exactly the same manner as the Midtone/Highlight Level control, determining the distribution of pixels between the two zones. The default value of the control is 33%.

The image above shows the Midtone and Shadow zones using the default values.

Here we see the Midtone zone before and after increasing the value from 33% to 66%. As you can see pixels are shifted out of the Midtone zone into the Shadow.

Conversely, here is the Shadow zone before and after increasing the value. The Shadow zone has a dramatic increase in value as pixels from the Midtone are biased towards the shadows.

Hue, Lightness, and Saturation Controls

We will next cover the Hue, Lightness, and Saturation controls, which are available for each individual zone: Highlights, Midtones, and Shadows. Earlier in this section we went over the HLS color model, which explained how each one of these controls work. To illustrate their practical use here, we will use the Midtone zone as an example.

Hue

The Hue Control determines the overall color of the zone. If we refer back to our color wheel you can see the various distribution of Hue colors. The control ranges from 0˚ to 360˚. At the right is 90˚.

Halfway between 0 and 90 is 45˚. At 45˚ the color on the wheel is yellow.

Here we see our image, with a Hue value of 0˚ (the default) and of 45˚. As you can see, the image has taken on a noticeably yellow coloration.

Lightness

The Lightness Control determines the intensity of the color in terms of light and dark.

Here we see our hue-corrected image. On the left is the default Lightness value of 0. On the right we see the lightness value boosted to 100. This significantly increases the overall lightness of the Midtone zone.

Saturation

The Saturation Control determines the purity of the color, expressed in terms of the relationship of the color to gray.

In the left image above we see our hue- and lightness-corrected image with a Saturation value of 0. On the right is the same image with the Saturation increased to 0. The colors are more vibrant and rich.

Conversely, here are the same two images with the Saturation value set to -100. Note that the Midtone values are devoid of color, and the image has taken on a grayish tint. Because this only affects the Midtone zone, there is still color detail in the image from the Highlight and Shadow zones.

Finishing the Shot

Since we have shown you the theory behind ZoneHLS, you might be curious to see how we could take our original footage and make a realistic shot.

Denoiser

Because the footage was filmed in a low-light environment, the first step was to use Denoiser to remove the noise and smooth out the detail in the shot. (See the Denoiser chapter for more information on this process.)

If you look at the wall, especially, you can see how much better the image quality is after the noise is removed.

ZoneHLS

The next step is to apply ZoneHLS to the image to bring out the hidden detail in the desk and background.

As you can see, the image has taken on a much more natural look, as if the lighting coming through the windows was providing actual lighting.

Magic Bullet

The final step involves using another product from Red Giant Software, Magic Bullet Suite™.

Magic Bullet gives the shot more of a filmic color and appearance. The light coming in through the windows, though simulated, now

has more of natural bloom and glow, and the detail in the desk is clearly visible yet still retains its original dark quality. The presidential flag returns to its original blue, the walls retain their eggshell tint, and the curtains in the background are a brilliant yellow and red.

If we look at the original element and the completed shot, we can see significant improvement. This particular clip was shot using a consumer-level MiniDV camcorder. Even with an input source as poor as the one in this example, the Red Giant Software family of products was able to turn out an excellent, natural-looking shot with less than ten minutes work.