

While this book covers many specific creative and technical issues relating to digital image processing, channels are at the core of all of these issues.

Channels are at the beginning and end of many important image manipulation tasks, and everywhere in-between. Just about everything discussed in this book relates to channels in one form or another.

Within Photoshop, people often take channels for granted. For example, when you scan an image with a desktop color scanner, you might not pay too much attention to the individual RGB channels of the image once it's been scanned. Crop, adjust brightness and contrast, run that new filter you just got in the mail—but mess with the channels? That's only for advanced geeks! Well, the fact is that you can learn quite a bit from understanding a little bit about *what* the different channels of that RGB scan actually represent. Even though you can adjust brightness and contrast as a global process to the entire image, by addressing specific channels of information you can limit your actions to only occur *when* and *where* necessary, yielding better-looking results. And who doesn't want that?

In order to get over the initial hesitation that many folks share when getting into understanding *why* things work, along with the *how*, it's important to understand some of the theory that supports the actual practice. In the case of Photoshop, or any software you're likely to use in your computing adventures, the fact is that there are always two sides to the coin: the face of the program that you, the user, see and interact with as you do your work, and the stuff under the hood—the guts of the software that do the actual nitty-gritty processing work. While you might be looking at a full-color image, unaware of the utility and power of channels, the graphics program that you're using is acutely aware of channels—in fact, it uses them when processing just about every command or feature you invoke.

Like so much terminology that overlaps the world of digital technology and the world of everyday life, channels have a specific set of meanings when we talk about digital imaging. Alpha channels are one expression of the potential functionality of channels in image processing. But it's critical to remember that, in discussing image processing software, the term *alpha channel* is likely to be virtually interchangeable with *masks*. If you speak to an experienced airbrush artist about masks, you're likely to hear a diatribe about *friskets*—a mask by any other name. And none of this addresses the fact that channels are the basic building blocks of any image, regardless of color space or graphics program or whether or not you use or care about alpha channels or masks. (After reading this book, you *will* care about these things. We promise.)

This book is about channels. Learn them. Love them. They will serve you well.

## The 8-Bit Imaging Universe

In order to start to wrap your arms around the intricacies of channels and inter-channel processing, it's important to clear up some basic core concepts relating to how microcomputers (in general) and Photoshop (specifically) work with images and visual data. The concepts that we'll cover in this section are primarily based on the Macintosh and Intel PC platforms. (The Silicon Graphics platform handles graphics differently due to SGI core graphics formats being based on more than 8 bits per channel.)

Fear not: This isn't an engineering manual. We're not going to revisit the nightmares of many college math classes. We reluctantly admit that none of us were *exactly enamored* of math during our school years—we all passed our math classes, but we definitely preferred pictures over numbers. A little theory, though, never hurt anyone, and is necessary in order to get the most from this book, as well as from Photoshop.

Perhaps the most essential atom of information in Photoshop (or any professional imaging application, for that matter)—besides pixels—are the individual channels that make up *any* image, color, or otherwise.

Everything in Photoshop is measured based on the reference of 8 bits of information. No, not the 8-bit graphics produced by converting images to indexed color mode. We're talking about the 8 bits of visual information that make up each channel of an image.

When working with images destined to be used for CD-ROM or Web applications, the file formats that tend to be used (such as PICT on the Macintosh, BMP on the PC, and GIF for Web) are usually 8-bit indexed color. Instead of the 8 bits of grayscale information we find in Photoshop channels, these indexed 8-bit (or lower) images have restrained color sets, expressed as groups of colors called palettes. While an indexed color image can have up to 256 colors, it also can have lesser bit depths (a 6-bit GIF image can have up to 64 distinct colors); discrete channels on microcomputer platforms (specifically, Photoshop image channels on the Macintosh and PC versions of the software) tend to always be expressed in 8 bits of grayscale data, or 256 levels of gray.



In order to clarify the meaning of an 8-bit channel, let's take a look at the relationship of bit-depth and color (or the mechanics of how a bucket of bits actually translates into useful information).


Every pixel (short for Picture Element—hey, don't ask us where the X comes from, we didn't invent the term!) on a computer screen is associated with some physical memory (that's the job of the VRAM in your computer, if you've ever wondered). On a black-and-white screen, each pixel has a single bit of memory. A bit can possess one of two states—on or off, black or white. This concept is the most basic core of binary arithmetic and, while relevant to many aspects of how computers work, is made very clear by placing it in the context of how a pixel is displayed on a screen.

An RGB image consists of three channels, or components, of information.


As we devote more memory to each pixel, the number of colors that a specific pixel can choose from to display on the screen increases accordingly. If we increase the number of bits for a pixel to eight, the number of possible pixel values increases to 256. Why? Because each bit can have one of two states (on or off)—and there are 8 bits—this results in 256 combinations (or *permutations*) of possible arrangements of those eight bits. This mathematically translates to 2 raised to the 8th power, which totals 256.

8-bit color depth can be used in two ways:


1. To designate a set of 256 color (different hues, as well as brightness and saturation) values to represent an image.
2. To represent 256 shades of gray, with no hue (color tonality) information.



A black-and-white pixel with its one bit of associated memory. The bit can be either on or off; as a result, the pixel can be either black or white.



An 8-bit pixel with associated memory. There are 256 possible configurations of the eight bits of memory, yielding 256 distinct shades, colors, or values.



The 256 possible gray values found in an 8-bits-per-channel world. They are the foundation of all image processing in Photoshop.

The second case is the one we're interested in at the moment. In an 8-bit grayscale image, the 256 permutations are used to represent the complete brightness range, from pure black to pure white, with 254 shades of gray in between. Solid black is given a value of zero, and the value of pure white is expressed as 255 (zero to 255 still expresses a total of 256 permutations).

This range of gray values is critical to the purpose and mission of this book. These 256 gray values constitute the entire foundation of channel processing in Photoshop, as well as any other imaging program you're likely to use on the Macintosh or PC.

## Photoshop Channel Fundamentals: Understanding Multi-Channel Documents

Since its inception, one of the most distinctive traits of Photoshop was that it understood and dealt with the concept of multiple-channel documents, additionally providing the digital artisan with direct editing access to the channels. Most other graphics software at the time, which allowed you to work with 8- or 24-bit RGB images (such as PixelPaint or ColorStudio), only let you go so far as seeing and adjusting all of the channels simultaneously. (While



early versions of ColorStudio enabled a restrained level of access to individual color channels, it was greatly overshadowed by Photoshop's channel access from the moment the latter program was released.)

Almost every document in Photoshop consists of at least a single 8-bit channel of information—the specific exceptions are “bitmap” images, which contain only pure black-and-white pixels, containing no intermediate gray values whatsoever. While the color component channels of an image are directly used to represent the actual image data (such as the RGB channels in a standard 24-bit color image), channels also are used for information other than the actual image data; alpha channels are perhaps the best example. Alpha channels are used for masking and compositing (which is discussed in more detail later in this book). These masks might have the task of isolating a specific element or portion of an image from the rest of the background, or might be used to create special effects for image blending and compositing.



Alpha channels can be used to create overlays of comments for production situations without affecting the actual image data. This technique is incredibly obvious, useful, and often overlooked.

Besides serving as the foundation for representing color images, different color spaces, and masks, other uses for extra channels include

- **Annotation layers**, for adding comments to an image that can be non-destructively overlaid on top of an image document, such as notes for informing a production artist where to work on an image needing retouching.
- **Positioning guides**, which were often needed for aligning multiple elements in Photoshop. The addition of guides to Photoshop 4 has reduced the necessity of creating guides in alpha channels, with the added benefit that the integrated Photoshop 4 guides can be snapped-to, much like the guides familiar to Illustrator and PageMaker users (and, of course, the built-in guides in Photoshop 4 take up less memory than an alpha channel). Guides and grids, however, don't supplant the technique of making vanishing point and perspective guides using alpha channels. The two techniques are complementary.
- **Text**. Very often, instead of creating color text directly onto a layer in Photoshop, you'll want to create text in an alpha channel and load that channel onto the color image as a selection; this enables you to fill the selection with different colors (for solid colored type), apply advanced color fills such as the multicolor gradients, and create filter effects between the text alpha and the background image. These general techniques enable you to create special effects that cannot be created any other way, outside of buying third-party filters, many of which don't deliver the degree of control attainable through less-automatic alpha channel techniques.



Two examples of using alpha channels for type effects. Notice how they interact with their backgrounds, something that's often hard to achieve by just using layers.

## Photoshop Image Modes

The Image>Mode submenu choices determine a variety of factors, the most important of which is the overall bit depth and number of channels of an image. Generally speaking, the number of channels and the bit depth increase as you move down the available mode choices, detailed in the following sections.

### Bitmap

A bitmap has 1 bit per pixel, with pure black (0) and white (255) values. No grayscale values are allowed here. While it might seem that the Bitmap mode has very limited appeal, the fact is that it delivers unique opportunities for both special effects and output applications.


The Bitmap mode has a variety of useful (and some rather esoteric) applications in everyday work. As far as special effects are concerned, you can create custom patterns as halftoning screens for achieving attractive (or distinctly ugly, based on the specified textures) mezzotint effects. Bitmap mode also is very handy for printing images with custom halftone screens on non-PostScript printers (such as many models of the popular Hewlett-Packard LaserJet line), which typically lack advanced halftoning handling in their imaging engines.

Images meant to be reproduced in print are always prepared by creating halftone screens. Without spending thousands of words on the subject (which is covered in extreme depth in many other Photoshop books, especially those oriented towards printing images from Photoshop), we can describe the halftoning process as breaking an image down into clusters of different shaped and spaced dots. In the case of black-and-white printing—with newspaper and laser printing as prime examples—there is no such thing as gray ink values; instead, black ink is laid onto white paper, and the illusion of intermediate shades of gray is achieved by creating clusters of black dots that exploit various deficiencies in the way that the human eye resolves detail in a visual image.


In the traditional methodology of preparing continuous tone images for black-and-white printing, a fine mesh screen was literally placed over the image and a special type of photographic exposure of the image-screen sandwich was produced, resulting in a version of the image that was broken down into the dots made by the holes in the screen. Even in the digital version of this method, the term *screening* has been retained.

Perhaps the most basic example of custom halftone screens is easily achieved by using the standard Photoshop Noise filter to create noise to be used as a halftone pattern.







Create a new document and fill it with 50% gray (128 on the color picker).



Using the Add Noise filter, create some monochromatic noise, with a setting of 50. Either Photoshop Uniform or Gaussian are fine; the differences are minimal for this exercise.




Now we'll apply the Motion Blur filter (Filter>Blur>Motion Blur) to the noise pattern to create a streaked-looking metal texture. We'll use an angle of 45 degrees and a 20-pixel distance for the Motion Blur filter.



Apply the Equalize command (Image>Adjust>Equalize), which will expand the dynamic range of the motion-blurred noise.



Choose Select>All and define the document as a pattern by selecting Define Pattern from the Edit menu.



It's time to apply the custom-defined pattern as a halftone. Choose an image to process, and start the process of converting it down to Bitmap mode. You'll probably have to convert to grayscale as an intermediate step (using the Image>Mode>Grayscale command), and once you choose Bitmap mode, you'll be presented with an options dialog box. Make sure to choose Custom Pattern as the method. (If this option is dimmed out in the Bitmap dialog box, you probably forgot the previous step of creating the custom pattern.)



Using this technique, any image or texture can be used to create a custom screen effect. Why is this cool? Simple: Printers that lack PostScript are typically challenged in the custom screening department. With this custom screening technique, you can squeeze custom halftones from any printer, or take advantage of the obvious special effects potential of creating endless varieties of custom screening effects. All built into Photoshop with no third-party filters required!

## Grayscale

With 8 bits per channel and 256 fixed shades of gray, Grayscale is most useful for printing images to black-and-white output devices (such as black-and-white laser printers) and processing grayscale scans. Grayscale is also the standard mode for single-channel mask documents. Of course, you'll need to switch to a color mode (such as RGB) in order to tint grayscale images with color/hue variations.



**Tip:** In previous versions of Photoshop, adding an alpha channel to a grayscale image would turn the image into multichannel mode. This is no longer true in Photoshop 4.0; adding alpha channels to images in grayscale mode simply creates grayscale-mode documents with multiple channels.

## Duotone

Duotone has 8 bits per channel, with up to 4 channels. Each channel represents a solid ink. This mode is only partially properly named, due to the fact that it represents one-channel (monotone), two-channel (duotone), three-channel (tritone) and four-channel (quadtone) documents. Each channel holds a representation of a standard process or spot ink type. Once converted to any of the duotone formats, the separate channels of the duo/tri/quadtone can be individually viewed or edited within Photoshop by converting the document to Multichannel mode, which places each ink plate into a separate channel.

## Indexed Color

Indexed Color has 8 bits per pixel, one channel. An 8-bit image in Indexed mode, unlike Grayscale mode, can be in color; a specific 8-bit color palette (sometimes referred to as a color lookup table, or *CLUT*) attached to the file determines what colors are represented. Colors that aren't part of that palette are either dithered or discarded. You can directly access, view, and edit the derived color palette (Image>Mode>Color Table).

## RGB

RGB is the native working mode in Photoshop, and has up to 24 channels of 8-bit information. Three of the channels are always dedicated to the RGB components, while the rest are alpha channels. Many file formats, however, store only one alpha channel with an RGB image; Photoshop's native format stores all 24 channels, if necessary.



**Tip:** If you delete an individual RGB component layer from an RGB document, the image is automatically converted to Multichannel mode.

## CMYK

CMYK is the native color space of four-color process color reproduction. CMYK files contain four color channels, with 8 bits each; only the Photoshop and TIFF file formats support more than four channels per document. The primary channels are cyan, magenta, yellow, and black, representing the four colors of ink used on a printing press. This is based on the subtractive color model. We don't spend too much time with this color space in this book, for two reasons.



**Reason A:** If you primarily work with CMYK scans, for output to CMYK separations, the only real importance of RGB color is to create masks—for use in the CMYK base images. Remember, it's entirely feasible to

- Duplicate your CMYK scan document (using the Duplicate command in the Image menu).
- Convert the duplicate from CMYK to RGB color.
- Use the RGB-based techniques described in this book to generate masks.
- Copy the RGB-derived masks back over to the CMYK, as additional channels. Use these new masks as desired.

This approach gives you the best of both worlds of CMYK and RGB. We have no problem confessing that the majority of our work is done for multimedia, video, and Internet applications, so we don't spend as much time in CMYK mode as we do in other color modes. Of course, like many other Photoshop artists, we often have to print our images on processed dead tree pulp, but we usually start with RGB data and leave the CMYK conversion process as one of the final steps in preparing images for print reproduction.

**Reason B:** If you want the deepest knowledge available on the theory and practice of doing deep, channels-related work with CMYK files, run right out and buy the amazing *Professional Photoshop*, by Dan Margulis (John Wiley & Sons). This book is the last word on CMYK color in Photoshop, and we felt that it would be silly to try and duplicate Dan's excellent efforts and skills. While we don't necessarily agree with absolutely everything Dan says in his book (he seems to have little interest in RGB color, and proclaims that CMYK is a better color space for many image processing tasks), we recognize genius when we see it. This guy is a bona fide CMYK genius. 'Nuff said.

## Lab Color

This is Photoshop's very own color mode: one channel of luminance (lightness) and two channels of chrominance (color), labeled A (the green-red axis) and B (the blue-yellow axis), hence the acronym *Lab*. In truth, Lab was originally developed back in the earlier part of this century. It was proposed as a color space standard by the Commission Internationale d'Eclairage (CIE) in 1931, and updated in 1976, at which time it was officially dubbed CIE L\*a\*b\*. The idea was to create a truly device-independent color modeling scheme that would preserve color values (hue, saturation, and lightness) across a variety of mechanical and digital hardware platforms (printers, displays, software, and so on).

For an excellent (and deep) technical discussion of the Lab color model, power up your Web browser and go to

<http://www.ls.com:80/cielab.html>

Ed Granger, a senior scientist at LightSource (makers of the excellent Colortron color measuring and calibration system), has posted an exceptionally relevant paper on the history—and inherent problems—of the Lab color space. Highly recommended reading from one of the great minds of the field.

While Lab looks like an arcane color space on the surface, it turns out that it's actually something of interest to any Photoshop user. Why?

- It's the native color space of PostScript Level 2. If you have a color printer with this version of PostScript, be aware that image files sent to it are converted to Lab space as part of the raster image processing (RIP) stage. Many people have reported great results converting images from RGB to Lab before printing directly from Photoshop to a PostScript Level 2 color output device.



- Lab is the closest thing left to the HSB (hue, saturation, and brightness) color mode in earlier versions of Photoshop. This mode was incredibly useful for processing the detail of an image (the brightness channel) while leaving the actual color information of the image intact.
- When you convert images between different color spaces (such as RGB to CMYK), Photoshop uses the Lab mode as an invisible, intermediary part of the color space conversion process.



Tip: Here's a time-saving Lab tip: If you spend significant amounts of time applying Unsharp Masking to large (20+ megabyte) scans, you can save some time required to apply the Unsharp Masking filter (not one of the fastest filters in the Photoshop toolbox) by converting your image to Lab and running the Unsharp Masking on the L (lightness or luminosity) channel exclusively. This also creates the desired sharpening effect on the detail areas of the image while leaving the color components intact, reducing the possibility of *halation* around strong color boundaries. The word halation is derived from halo: a particular problem with high frequency edges of an image, when being processed through contrast enhancement processes such as Unsharp Masking. Halation also can be seen with images taken with digital camera technology, especially if you do further Unsharp Masking to the digital camera image before converting it to CMYK for color output.



An image with the Unsharp Mask filter applied equally to all color channels. Note the visible halo effect around the edges of the foreground object.



An image converted to Lab color space, with Unsharp Mask filter applied only to the Luminosity channel. Note that the halation effect seen around the edges of the foreground image isn't as obtrusive as the results of the previous figure.

## Multichannel

Multichannel is the Swiss Army Knife of Photoshop's Mode menu. It's primarily used for the following reasons:

- Directly view and edit the results of a Duotone conversion. (Normal Duotone mode doesn't allow the separate plates to be viewed or edited discretely.)
- Prepare spot-color composites. For example, a grayscale image in one channel might accompany a type knock-out area designated in a second grayscale channel; the contents of the knock-out are assigned to a specific Pantone color, which prints on top of the grayscale image in the first channel. While Multichannel mode can be used to emulate spot color plates, it isn't a technique we recommend for a serious production environment, and certainly isn't a replacement for a real spot color capability in Photoshop (perhaps the most overdue feature still not found in the current release of the program).

Viewing multichannel documents in Photoshop results in only viewing one separate grayscale channel at a time. Multichannel format is often a decent choice for storing alpha channels outside of your main working document (the one that the alpha channels would actually be used on). Because Photoshop allows the saving of selections to external documents, all the masks for an image can be stored in a different document, if needed.

## Bit Depth and Color

Most people involved in imaging know that the lower the bit depth, the fewer colors you can fit into your image. We've all seen it: beautiful 24-bit images that show up on the World Wide Web looking pretty lousy in 8-bit color. Some artists wonder why their final CMYK print has weird moiré patterns and other artifacts when their scan was done in 8-bit color mode, converted to RGB, and printed in CMYK.

An indexed color image has a relatively limited number of colors. If we look at the Indexed Color dialog box in Photoshop, it's obvious that an image can have fewer than 8 bits of memory for each pixel, resulting in fewer colors (which translates into smaller overall file sizes). The relationship between the value of the pixel depth (for example, 5 bits per pixel) and the number of colors in a particular image is relatively straightforward: raise 2 (the possible state of a bit: on or off, 0 or 1) to the power of the image's pixel depth (how many bits of color data each pixel contains) to find out exactly how many colors are in it (the maximum number of possible color permutations available for a single pixel). For example, an 8-bit color image can have a maximum of 256 colors in it; raise two to the eighth power and the result is 256.

Given this exponential relationship, the difference in size and range of possible color variations between two images of different color depths is more significant than you might expect. An 8-bit color image, with 256 colors, is much more than twice as big as an image with 4-bit color, which only has 16 colors in it!

The following series of figures is an example of an image indexed down to different color depths and converted with different tools. These tools include Digital Frontiers' excellent HVS Color plug-ins for Photoshop, as well as Debabelizer.






A 24-bit original RGB color image.



An 8-bit adaptive palette, dithered in Photoshop.



A 4-bit adaptive palette, dithered in Photoshop.



An 8-bit adaptive palette, dithered using HVS Color.



A 4-bit adaptive palette, dithered using HVS Color.



An 8-bit Macintosh system palette, dithered in Photoshop.



An 8-bit adaptive palette, dithered in Debabelizer.




A 4-bit adaptive palette, dithered in Debabelizer.



An 8-bit Macintosh system palette, dithered in Debabelizer.

Let's say that you scan an image in grayscale mode with a typical desktop scanner. Even though the original image might be color, and your scanner is a color model, you can scan a single channel of information. This yields a single channel, 8-bit grayscale image in Photoshop.



A scanner plug-in module in Photoshop.



The Channels palette, displaying the channels of a 24-bit RGB image.



**Tip:** We consider the HVS color reduction technology to be one of the better investments that you can make for your Web-related imaging work. You can find the publisher, Digital Frontiers, on the Web at


<http://www.digfrontiers.com/>

Digital Frontiers HVS ColorGIF 2.0 and HVS JPEG 2.0 are excellent—the HVS (Human Visual Systems) color reduction algorithm goes far beyond the standard color-polling techniques used by programs like Photoshop and results in far more accurate color representation using far fewer colors than anything produced by any other program we've seen (not even DeBabelizer can touch it!). And, yes, they make the plug-ins for both the Macintosh and PC hardware platforms.


## Color Depth and Channels

So how do full-color, 24-bit images fit into this 8-bits-per-channel paradigm? Quite easily, in fact. Without going too deep into color and visual theory, the short answer goes something like this...

Unlike a grayscale image, which has a *single* channel (which therefore grants it 256 levels of brightness), an RGB image has *three* component color channels: one Red, one Green, and one Blue. This is called *transmissive*, or *additive*, color, wherein red, green, and blue light are displayed simultaneously on a cathode ray tube. This is how images are shown on computer monitors, television screens, and video displays of all kinds. Each of these color channels, like a single-channel grayscale image, contains 8 bits of data, or 256 different tones. Three of these 8-bit channels in a single image (eight bits per channel, multiplied by three channels) yields a single image with *24 bits of color*. When all three RGB components are displayed at full strength (a full 255 brightness readout for each), the resulting color is white. An absence of all three RGB components (0 readout for each component) results in pure black.



Photoshop's Color palette is a tiny, self-contained education in basic color theory. As you can see, when all three RGB values are set to 0, the result is black.




Pump all three RGB sliders up to their full possible value (255), and the result is pure white. And look at the actual color readouts within the slider bars themselves, which shows you how color components interact. Isn't it interesting to see that when you have all three RGB sliders set to full value, the Color palette is trying to tell us something about the relationship between RGB and CMYK.



As you'll see, it's the interaction of these three color channels that produces the ultimate full-color image. In a single color channel, black indicates the absence of that particular color, while white represents full-color presence. Let's assume a particular pixel in a 24-bit RGB image is black in the red channel, 50% gray in the green channel and white in the blue channel. This pixel's RGB color is 0% red, 50% green, and 100% blue, resulting in a light blue hue. The total possible permutations of hue in a 24-bit color image is upwards of 16 million different colors (16,777,216 colors to be exact, the result of 2 to the 24th power).

Specifically, when discussing the actual color channels of an image, the presence of brightness in a component color channel (as expressed in values getting close to pure white) indicates the presence of that component color in the overall image. Let's look at the example of a typical color photograph.

If we take a look at the blue channel of this image—by selecting it from the channels palette, or pressing Command-3 (Control-3 for Windows users) from the keyboard—it's clear that the area of the blue sky contains predominantly brighter pixels, especially when compared to the same area of the image in either the red or green channels. (In the case of a blue sky, the red channel is relatively dark, while the contents of the green channel in the sky area fall somewhere in between the brightness values of the red and blue channels.)



RGB images have three 8-bit channels, one for each color. The interactions of these channels create full-color, 24-bit images.



A typical 24-bit RGB image from our stock photo library.



The red, green, and blue channels of our sky image. In Photoshop, the name of the channel appears in the title bar of the document window. Note that the pixels in the sky area in the blue channel are significantly lighter than the same areas of the red and green channels.

The only times that RGB images have more than 24 bits of data per pixel are

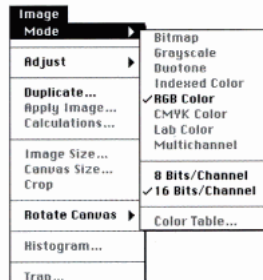
- If they have embedded *alpha channels*, 8-bit channels used for masking and compositing, effectively making them 32-bit image files (with no increase in potential colors).
- If they have been scanned on high-resolution scanners. Drum scanners, and even some flatbed models, can scan images with up to 16 bits of data per color channel. This extended bit depth primarily asserts itself in the image as extended detail in darker (shadow) regions of the image, as well as in smoother color gradations, which involve subtle variations of a base color (such as gradients between pure black and white). Images originating on SGI systems also often have more than 24 bits of data.

Photoshop has some limited support for handling 16-bits-per-channel images. The toggle for 8- or 16-bit modes can be seen in the Image>Mode submenu. Normally, the default setting is 8 bits per channel (in all color modes, with the exception of Bitmap, in which neither 8 nor 16 bits are active), with the 16-bit option appearing in both Grayscale and RGB modes.


If you have a Grayscale or RGB scan created with a typical desktop scanner, it's most likely that they will default to the 8-bit setting. Some higher-end flatbeds, as well as a number of slide and film scanners, have the ability to generate 16-bits-per-pixel scans, which significantly extends response in shadow areas (yielding greater detail in predominantly dark portions of a scan). Photoshop can actually import and display a 16-bits-per-pixel file.

What's the use of 16 bits of color data per channel? Four words: *greater dynamic color range*. If you get a 24-bit color scan (8 bits per channel) that's both dark and lacking in contrast, using Levels or Curves actually discards some image data in favor of some level of aesthetic improvement. You can visually see this by opening up the Levels dialog box, where you'll see empty bands in the image's histogram, indicating missing colors or tones.


With a 48-bit scan (16 bits per channel), doing drastic color or brightness correction still leaves a good deal of data intact. Looking at the Levels dialog box of a 48-bit scan after significant color correction still shows a histogram with a full dynamic range of tones. Of course, after initial scan tweaking you'll have to convert your images into 24-bit mode for filtering or other significant editing.




In Grayscale and RGB modes, you can choose from either 8 or 16 bits per channel.




This is the Levels histogram display for an original 24-bit scan...



...and this is the Levels histogram for an original 48-bit scan. Notice how these two histograms look very similar before any image editing is done.



This is the same 24-bit scan after Levels has been used for contrast enhancement. You can start to see the empty vertical bands where color data has been lost.



When the 48-bit scan has the same Levels command applied to it, however, the histogram looks entirely intact; when converted to 24-bit mode, it has a greater range of color data than the 24-bit scan.



## Color Spaces and Channels

Color space, in the context of this book, is an abstracted method of color representation on the computer. RGB, CMYK, CIE, Lab, Indexed—the color modes in the Image>Mode menu of Photoshop—are all individual color spaces. Despite many experts, many opinions, and untold numbers of international standards organizations, there is no single universal method of constructing a color space that can display color on all devices for all occasions. That's why there are several different standard color spaces in Photoshop and other applications. Each color space has its own intended use.

Color images usually have a dedicated channel for each component color in the image's designated color space: one channel each for red, green, and blue for an RGB image (three channels total); one channel each for cyan, magenta, yellow, and black in a CMYK image (four channels total); and so on. One exception to this rule is the Indexed Color mode, which uses only a single channel to represent a specific set of colors; this set of colors is called a *look-up table* (LUT for short, sometimes called a *CLUT* for color look-up table).

## Channels Reveal Their Secrets

The different color space options offered by Photoshop each have their appropriate use, depending on the context of the acquisition and final application of an image. For example, images scanned on professional laser drum scanners are typically delivered as CMYK files right from the scanner, while images scanned on most desktop scanners are typically in RGB format (with the exception of a newer generation of desktop laser drum scanners, which have begun to appear on the market in the last couple of years).

## The Importance of RGB

Much of this book deals primarily with the RGB color mode in Photoshop, often at the expense of coverage of the other useful color models (such as CMYK) found in Photoshop. Why did we decide to do this?

### The Electromagnetic Spectrum

Color, brightness, and saturation—terms you're likely familiar with from working with Photoshop, as well as other imaging software—are ways to describe particular characteristics of the portion of light that human beings can see. The visible spectrum of light is part of a larger construct of physical reality (no, we didn't just go New Age on you, we're talking about actual empirical physics here), a world that's called the *electromagnetic spectrum*. Did you ever wonder how broadcast television and radio work? Or how it is that a sound of a frying hamburger reaches your ears?

### Authors' Note

At this point, you're probably wondering why we're taking this little detour into physics... "Hey, this is supposed to be a Photoshop book!" Well, the fact of the matter is image processing is largely based on the concepts of signal processing, which are more generalized than visual imagery, but can help you understand the more complex topics encountered in learning advanced image processing.

The truth is that much of David Biedny's knowledge of image processing is derived from his intense interest at an early age in audio processing, which shares many base concepts with image processing. (Ever notice the High Pass filter in Photoshop? Wonder what it does? The concept has an extremely important and useful counterpart in audio signal processing, and is an important link in the audio equalization chain.) But let's get back to channels!

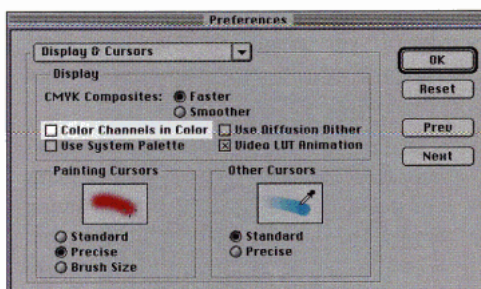
Most of the devices used by human beings to perceive and interpret color are inherently based on the RGB color model. The light gathering mechanisms of the human eye are based on RGB color, as are the computer and display monitors used by anyone reading this book.

Well, as discussed in the preceding paragraphs, it's the native color space for most input devices that you're likely to use in daily work (with the notable exception of high-end pre-press drum scanners, which typically create scans in CMYK color space). RGB also is the native color space for your computer, the monitor on your desk, and your own internal optical input devices (eyes, by any other name). Most important to the mission of this book, though, is the fact that RGB is perhaps the most revealing of the color spaces, in terms of how light works and the characteristic behavior of light when described in RGB.

Although Photoshop can display images in other color spaces, the actual colors on your screen are still physically created by the RGB guns in the Cathode Ray Tube (CRT). In the day-to-day world of image creation and reproduction, just about the only time you'll likely deal without another color model is in the four-color printing process—combinations of cyan, magenta, yellow, and black inks. Hey, face it: It's an RGB world, after all.



**Tip:** One of the options in the Display & Cursors screen of Photoshop's Preferences dialog box is the capability to display color component channels in either grayscale or the actual component color. You'll always want to make sure that you're looking at channels in Grayscale mode, given that the actual brightness values of the pixels in each component channel are of any actual value; looking at the blue channel in shades of blue can be quite distracting, making it harder to actually evaluate the real information present in the channel.



Make sure that you've unchecked the Color Channels in Color option in the Display & Cursors Preferences dialog box (screen three) as the method for displaying individual channels.

## The Red Channel

The red channel is the ultimate proof of life on Mars. Okay, maybe not. Like the other channels, the red channel displays the amount of red component in an image. Perhaps the most noticeable example of this is seen when looking at a shot of a typically lit human head. Human flesh tones have a significant amount of red component, represented by brightness in the red channel.

In a general sense, the red channel tends to contain the greatest range, or *spread*, of contrast values in an image. If you want to know a bit about the brightest and darkest values in a photographic image, you'll usually start by taking a look at the red component channel (which can be directly accessed by pressing Command-1).



The red channel is always active where human flesh is involved.



## The Green Channel

Of all of the RGB channels, green is extremely useful in general image processing work. While the red channel tends to hold the contrast spread, the green channel is the place where all high-level or high-frequency detail in a photographic image can be found.


What is high-frequency detail? Does it have something to do with your stereo system? Is there a treble control involved? Well, not quite, but the ideas are actually related.

In discussing image processing, the concept of high frequency is placed into a specific context. Let's take the example of the most extreme case of a high-frequency edge.



The area where the two regions meet is considered a very high-frequency edge. If you sample the black-and-white areas with the eyedropper tool, it's clear that the difference between the two values is as extreme as it gets.

In the case of CMYK, for example, the distinctive, recognizable details of an image are primarily contained in the black channel (often called the black plate by those in the printing and pre-press world). This is the primary reason that Photoshop enables you to get very specific about how you want the black channel derived from the mixture of red, green, and blue when doing color space conversions from RGB to CMYK.


Any edge that displays a significant differential in brightness values is where the actual defining detail of an image can be found, and as it turns out, the green channel of an RGB document tends to contain the best representation of high-frequency edges in a color composite image. The green channel can be accessed by pressing Command-2.




This represents the most high-frequency edge you can find in image processing: a border between two regions of absolute black and white.



It's obvious that there's a 255 value spread between white (255) and black (0).



The density of the black plate of a CMYK document (converted from an RGB source image) is determined by specifying the Black Generation tonality curve and the Total Ink Limits (Total Ink Limit describes the relationship between the amount of mixture of CMYK inks and the derivation of black values from that mixture).



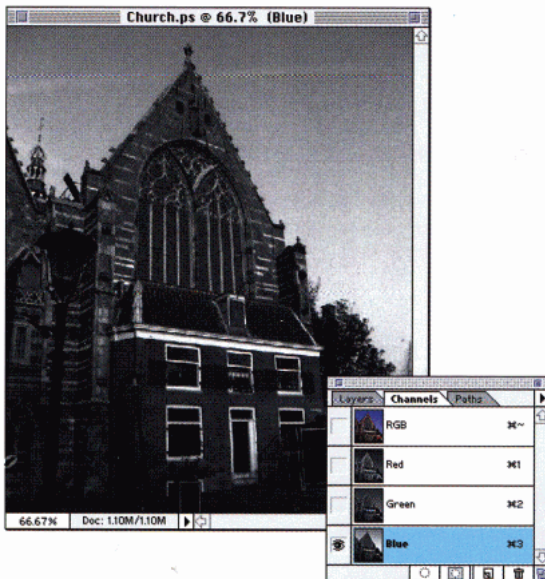
If you look closely, it becomes obvious that the best representation of the fine detail in an image can be found in the green channel.

## The Blue Channel

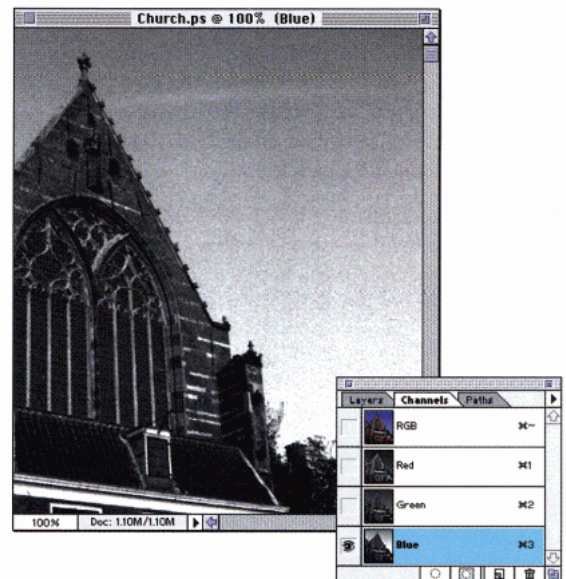
In the food chain of RGB channels, blue is at the bottom of the heap. No, that doesn't mean that the blue channel is bad—but very often the bad things about a color image try to hide themselves in the blue channel of a color image. Of course, if you're scanning lots of images of blue skies, blank blue screens, blue-colored construction paper, or a new fashion line of pure blue clothing, you would expect significant activity in the blue channel (remember, the presence of a color component in an image is represented by brightness values in that component channel). But for most natural images, the blue channel is the place where you can find

- Noise artifacts
- Film emulsion grain characteristics
- Scanner noise and undesirable artifacts
- Video jitter (in the case of images grabbed or digitized from video sources)

One way to tell how well any scanner performs (in terms of overall image quality) is to take a look at the blue channel by itself (press Command-3 to display just the blue channel in the active RGB document window).



A typical blue channel of a color image.



The emulsion characteristics of a specific type of film are always visible in the blue channel of a scan of the slide.

If you're evaluating the blue channel of an image scanned from film (such as a 35 mm slide), you'll notice a distinctive grain pattern underlying the overall image.