Chapter 6
The Digital Zone System and Black-and-White

In the digital world, we don't shoot black-and-white photos the way we did with film. It might surprise you to know that digital sensors are actually monochrome. The color comes from the filter array that sits on top of the sensor. That aside, we still consider that we're shooting in color. There have been a few dedicated monochrome digital cameras produced, but they never sold well and didn't stay around for long. Personally, I think they're a great idea and would love to have a dedicated black-and-white digital camera that I could use with color contrast filters, the way I shot with film. Until those cameras are made available to the masses again—and I'm not holding my breath—we'll continue to work with color images that are converted to black-and-white. Which is not inherently bad. Using color information to do black-and-white conversions can give us great flexibility that we didn't have with black-and-white film.

We can shoot in grayscale mode with most digital cameras. It's a JPEG-only option, however, and it's not the best option. First, we lose the inherent quality of shooting RAW. Second, the conversions to monochrome in-camera are typically not very good. Your best bet for black-and-white in the digital world is to shoot RAW in color and then later convert to grayscale.

In the relatively short history of digital photography there have been many ways to convert from color to black-and-white. Some methods are better than others. The tools available for doing black-and-white conversions have greatly improved over the years. Compared to the tools available in ACR and Photoshop today, the Channel Mixer method that was popular just five or six years ago is downright archaic. Many people have their favorite ways to go about it. Most of the methods involve adjusting the lightness of different colors to achieve a pleasing mix. A good deal of dodging and burning can be involved as well, depending on the desired result. What these methods essentially do is use color to create tonal separation in order to generate visual interest.

Creating a Luminance Layer

Most of the conventional image-editing methods try to mimic the use of color contrast filters on black-and-white film. The goal with digital black-and-white conversion is to create color contrast. By creating contrast between colors, we get contrast in shades of gray. With the conventional methods of converting to black-and-white, this makes absolute sense. But our photos already have tonal variations in the existing luminance values. What if we could take advantage of those existing luminance differences to create our black-and-white images? Figures 6.1 through 6.6 show six examples of taking an image

▲ *Figure 6.1: The color original*

▲ *Figure 6.2: Grayscale*

▲ *Figure 6.3: Desaturate*

that originally has quite a bit of color contrast and converting it into grayscale. Five of the methods are what could be referred to as "conventional." The last is a conversion using luminance that exists within the image.

I specifically chose this image because it has a good range of colors and some subtle luminance variations that could be difficult to pick up. There are some brightness differences that the conventional methods of converting did not reproduce well. None of the conventional methods properly converted the bit of brighter blue in the middle, above the hinges. The brighter values on the far left of the image are also converted more accurately in the luminance version. The brightness variations in the large hinge are picked up better, and the brighter green of the small hinge is more faithfully converted. Remember, these are all simple default conversions with no incremental editing. The better the conversion you start with, the better the final result.

▲ *Figure 6.4: Red channel*

▲ *Figure 6.5: Black-and-white adjustment*

▲ *Figure 6.6: LAB, lightness channel*

▲ *Figure 6.7: Luminance values*

What we see as color in the world we live in gives us separation between elements in a photograph. Unfortunately, many colors convert to similar shades of gray. The classic example is a red rose and green leaves. We see the difference between the two because of the colors. Remove the color, and both produce a similar shade of gray, so we lose the separation. A similar example is red peppers with green stems.

▶ *Figure 6.8: Peppers in color*

We see contrast difference between the peppers and the stems because of the color. If we remove the color with a simple Desaturate command in Photoshop what happens?

▶ *Figure 6.9: Peppers in gray with Photoshop Desaturate command*

The contrast is gone. The difference between the red and green is gone. The two basically appear as the same in shade of gray. Figure 6.10 is a block made up of red, green, and blue. Each is set to a value of 255. We can distinguish between the blocks because of the color. What happens when we convert to grayscale?

How Do We Get Shades of Gray?

This conversion shown in figure 6.11 was done with the Channel Mixer in Photoshop at the default settings of 40% Red, 40% Green, and 20% Blue. The goal is to mix the colors so that the combined total is 100%. The green and red from the original image have become the same shade of gray. This is what happens in the classic red rose/green leaves example with film as well. The two colors reflect similar amounts of light, which leads to similar shades of gray. Using black-and-white film without a color contrast filter on the lens is a lot like doing a default grayscale conversion in Photoshop. The same thing happened with the peppers and the color blocks. The blue block is darker because of the lower percentage of blue in the mix, but if we increased the blue to 40%, it too would become the same shade of gray. Even with the blue being darker there's really no visual interest, because the tones are all so similar.

We can look at each channel individually as well to see what happens. Figure 6.12 shows the red, green, and blue blocks, from left to right, each set to 50%.

They're all the same shade of gray. But we did have a slightly darker shade of gray in the earlier example. We can see that by adjusting the percentages of different colors we can change the appearance of tonality and luminance. But we're still "creating" that luminance difference. Why create it? The image already has luminance variation. Let's use it!

Going back to our autumn waterfall image from the previous chapter, if we do a conversion to black-and-white with a Black & White adjustment layer at the default settings, we get the result shown in figure 6.13. It's pretty blah. We've lost what separation existed between colors. The mass of trees in the background is mostly the same shade of gray. Yet we know there was color there and we know there were different areas of luminance.

▲ *Figure 6.10: RGB color blocks*

▲ *Figure 6.11: Default grayscale conversion using Channel Mixer*

▲ *Figure 6.12: RGB channels, each at 50% gray*

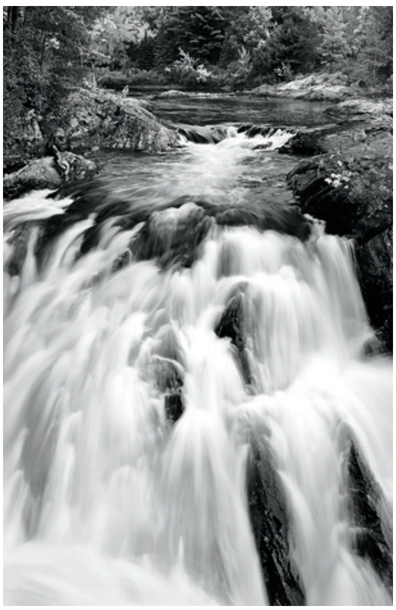

◀◀ *Figure 6.13: Default Black & White adjustment layer*

◀ *Figure 6.14: Conventional black-and-white conversion*

There's little visual interest in this black-and-white image. Now, to be fair, we could make some adjustments to the color sliders and regain some visual interest and contrast. But to really get separation, there would have to be a lot of painting in the image with your preferred tools for dodging and burning. Figure 6.14 shows just such an attempt.

This is definitely preferable to the default version. There's now some tonal separation in the trees at the back. The rocks along the left side have been lifted a bit, and overall the image has more contrast and visual snap. This took a lot of time to achieve, and there was a lot of painting to dodge and burn different areas. If you were to look at a larger version, you would see that a lot of painting was done on the trees at the back. Getting the brushed areas to have a high degree of accuracy in such small details is difficult and time consuming.

Black-and-White Conversions with the DZS

▲ *Figure 6.15: Default DZS conversion*

By making just one small change to our color DZS image, we can get a black-and-white version that retains the tonal variations in the color original, as shown in figure 6.15. The trees in the background have good tonal separation. The hints of highlight at the edges of some of the branches are retained. It's a much more appealing image at the default conversion level. All we did to create this was change the Blend Mode of the luminance layer from Luminosity to Normal. That's it. Simple. This version could stand on its own.

Because we've changed the Blend Mode to Normal, the Hue/Saturation adjustments in the color version play no part in the black-and-white version. This is simply the luminance information from the original that we extracted, plus the Curves adjustment layers and Unsharp Mask.

We probably want to do a little more to this, however. With the layer Blend Mode changed, when additional adjustment layers are created, they should be placed above the top layer that was used in editing the color version. The reason is that we're going to group the layers later so we can turn them off and on as a group. With the luminance layer Blend Mode switched to Normal, the Hue/Saturation adjustment layers below the luminance layer are no longer active. To keep the black-and-white adjustments separate from the color adjustments, any new layers should be labeled with "BW" or "Gray" or some similar notation to remind you that these are for the black-and-white version only.

The process is the same as with the color version. Load a zone mask selection and apply your desired adjustment to it. This can be repeated for as many zones as you want to work on.

You can apply different sharpening to the black-and-white version as well. Disable the Unsharp Mask Smart Filter on the luminance layer by clicking the eye icon next to it. Open Unsharp Mask or Smart Sharpen, and apply sharpening to the black-and-white version independent of the color version. Photoshop doesn't permit you to rename Smart Filters, but there is a way to know which is which. The first Smart Filter will be at the bottom and the second above it. Your black-and-white adjustments are *at the top* of your layer stack and above your color adjustments, so the sharpening Smart Filter that is *at the top* of the Smart Filter list applies to the black-and-white version.

▲ *Figure 6.16: Final DZS conversion*

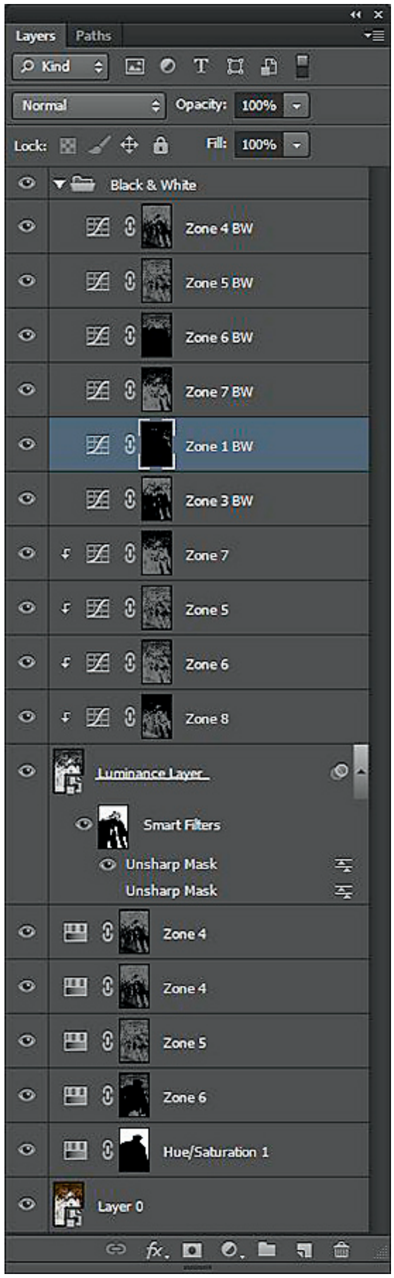

▲ *Figure 6.17: DZS layer stack*

The final black-and-white version is shown in figure 6.16. The differences from the default image are evident. We've enhanced overall image contrast by concentrating on local contrast in a few zones. Some additional tonality has been added to the deep shadow area in the rock face by the water, on the right side at the top of the waterfall. The existing lighter tones in the trees at the back of the image have been lifted slightly as well. In terms of time, this took less than half the time of the conventional conversion method. The layer stack in figure 6.17 shows which zone masks were worked on.

Zone 1 and zones 3 through 7 have had Curves adjustment layers applied. You can see the second, separate Unsharp Mask Smart Filter in place to apply different sharpening to the black-and-white version. At the top of the layer stack, you see the folder called

Black & White. That's the grouping of the six layers that were adjusted in creating the black-and-white version. By grouping them, we can turn their effect on or off all at once, rather than having to turn each one on or off individually. The other thing you'll note is that the black-and-white zone mask layers are not clipped to the luminance layer. The reason is that because the Blend Mode of the luminance layer has been changed to Normal, nothing below that layer is having an effect on the image, and the black-and-white adjustment layers need not be clipped to the luminance layer. Whether or not you clip them to the luminance layer would make no difference to the look of the end result.

You do want to leave the zone mask adjustment layers for the color version clipped so that you can switch between versions and not have to reclip each of them. These layers cannot be grouped because doing so would unclip them from the luminance layer and a layer group can't be clipped to a layer. These can be turned on or off as you wish, to render the black-and-white result you want.

You can see the difference between the default DZS conversion (figure 6.15) and the final image (figure 6.16) in the histograms of the two images, shown in figure 6.18. The histogram for the default DZS version is on the left and the one for the final version is in the middle. For comparison, the histogram for the conventional conversion is on the right.

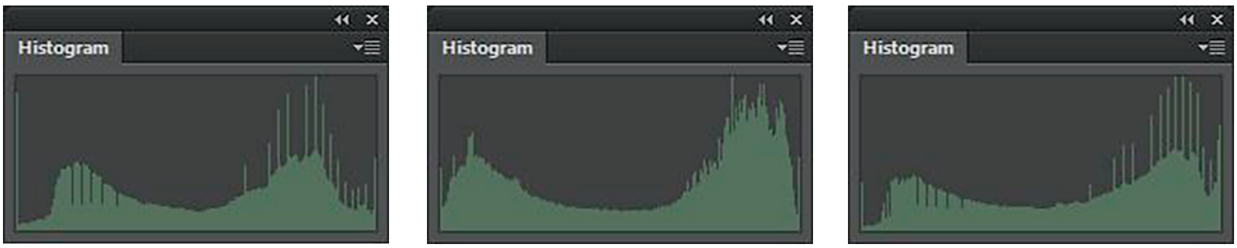

▲ *Figure 6.18: Default DZS (left), final DZS (middle), conventional (right)*

The change is not great, but it does make a visible difference. This illustrates that you don't necessarily need to make large changes to an image to have a significant impact on the final result. Darker values on the left of the graph have been affected. In particular, note the shorter spike at the left edge, indicating less shadow detail has been clipped off. On the right of the graph there is separation between the lighter tones in the image, while the midtones are largely unaffected. Note that in the histogram for the conventional method there is more clipping in the highlights and the graph is more compressed. The conventional version is flatter in midtone contrast but has more highlight clipping. Not a great result.

Over time you'll train your eye to know which zones are which, and thus know which zone mask you need to load to affect specific areas of the image. There is a way to help determine this, though. If you hold your cursor over an area of the image you want to work on, you can see the RGB numbers for that area in the Info panel. The RGB numbers will be the same because you're working with a grayscale representation in the luminance layer. You can then load the zone mask that corresponds to the RGB numbers under your cursor. For example, if your reading is 128, 128, 128, you know you need to load the Zone 5 mask. If the readings are 42, 42, 42, you know you need to load the Zone 2 mask. You can refer back to the step wedge in chapter 1 until you become used to the zone values. Recall from the previous chapter that you should make your measurements before you begin editing so that you're reading from the original luminance values and not the adjusted values. For images with color and black-and-white versions in the same image, you should turn off your adjustments for the color image as well before measuring to determine which zone masks to use.

In figure 6.13, there was quite a bit of luminance difference to take advantage of. Can the DZS be used on an image in which there may not be as much separation between tones? Yes it can. Let's go back to the red peppers and green stems example.

The straight grayscale conversion of the peppers looks pretty blah, as we can see in figure 6.20. There are some highlights on the peppers from ambient light reflection, but not a lot of distinction between the stems and the peppers. There isn't even much distinction between the peppers at the top of the pile in more light and the ones at the bottom of the pile in less light. The histogram for the image is pretty much as expected—mostly flat with spikes in the middle and at the left end, and not extending all the way to the right on the highlight end.

Now let's take a look at the DZS image in figure 6.21. We see a much better result. There is tonal separation between the peppers in different amounts of light. There is separation between the peppers and the stems. Overall contrast is much improved. It's a far more visually interesting image. The histogram shows this as well, where the tonal range has been spread out. The spike in the middle of the original has been moved to the right, and the graph extends fully to the right edge. We've sacrificed a

▲ *Figure 6.19: Peppers in color*

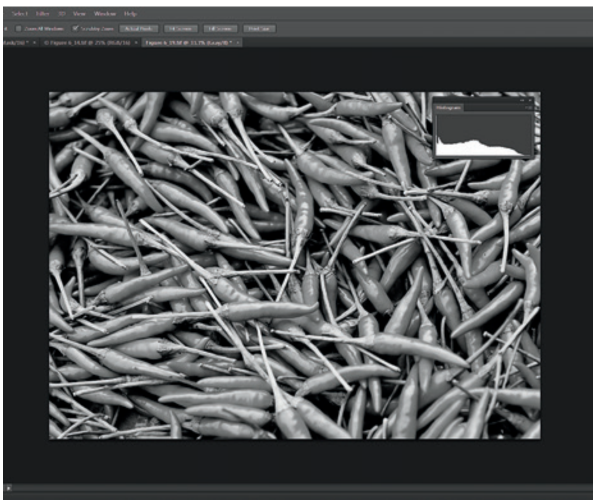

▲ *Figure 6.20: Peppers with default grayscale conversion and histogram*

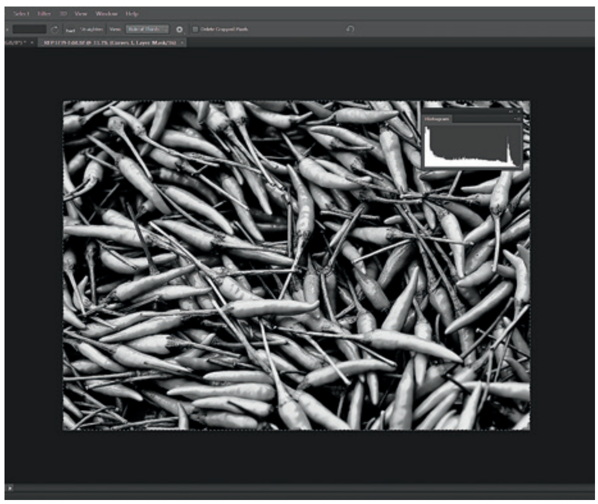

▲ *Figure 6.21: Peppers with DZS conversion and histogram*