

Sigma Merrill vs. Sigma Quattro. 1. Signal, Noise, and Sensor Design

Phil Service
Flagstaff, Arizona, USA
16 February 2016

Summary

A transmission step wedge was photographed with a Sigma DP1 Merrill and dp1 Quattro. Signal mean value and signal-to-noise ratio (SNR) were estimated for each color channel from the raw (X3F) files using RawDigger. Quattro raw image signal strength is substantially more uniform across color channels. Quattro images also have higher SNR (i.e., less noise) than do Merrill raw images in the red and green channels. Quattro blue channel signal strength was substantially less, and SNR was slightly inferior to that of the Merrill, particularly in the well-exposed portions of the images. Averaged across color channels, the difference in SNR is equivalent to about 1/2 EV advantage in dynamic range for the Quattro. In practice, the Quattro may have a ≥ 1 EV superiority because, in contrast to the strongly blue-channel biased overexposure of the Merrill sensor, the Quattro blue channel is relatively resistant to overexposure; and all three color channels tend to become overexposed in concert. These differences in signal strength and SNR can be understood at least partly by reference to differences in sensor design. They suggest that the Quattro design was chosen to improve signal strength and SNR in the red and green channels, while sacrificing some signal quality in the blue channel. The net effect is substantial improvement in overall signal characteristics: in particular better balance among color channels. ISO series were made with a DP2 Merrill and a dp2 Quattro. Raw files were processed through Sigma Photo Pro, exported as TIFFs, and taken into Photoshop. As expected from the signal strength and SNR analyses, Quattro images had an approximately 1-stop advantage in high ISO image quality. That is, Quattro images exposed at ISO 1600 were similar to, or slightly better than, Merrill images at ISO 800.

Key words: Sigma Quattro, Sigma Merrill, DP1, dp1, DP2, dp2, Foveon sensor, RawDigger, signal strength, signal-to-noise ratio, image noise, Sigma Photo Pro, sensor design

1. Introduction

The Sigma Merrill sensor (DP1, DP2, DP3, and SD1) is a “conventional” Foveon sensor in which trichromatic color information is obtained by vertical stacking of photosites. The design depends on the fact that light is differentially absorbed as it passes into the sensor. In standard descriptions, blue wavelengths are absorbed preferentially in the top layer, green in the middle layer, and red wavelengths in the bottom layer of photosites (Fig. 1). The principal advantage of the Foveon sensor is that under favorable shooting conditions (ample light, relatively static subjects), it can produce images with considerably greater acuity than Bayer sensor images of similar, or even higher, resolution.[1] That is, in fact, the only compelling reason to buy and use Merrill cameras. The main disadvantages of the Merrill sensor are that color fidelity may be poor, and image acuity degrades quickly at ISO setting above base level. Base ISO is 100, and serious degradation occurs by ISO 800 (as will be shown below). Poor color fidelity and ISO performance almost certainly result from the fact that light penetration to the green and red sensitive photosite layers is considerably reduced. I infer this from the observation that when the blue color channel is properly exposed (as shown by RawDigger), a large number of raw image pixels are often underexposed in the red and green channels. Similarly, overexposure tends to occur first in the blue channel. I suggest that the red and green channels of Merrill images are chronically underexposed. All other things being equal, underexposure means less "signal" and a lower signal-to-noise ratio (SNR) — in other words, a "noisier" image.[2]

The Quattro sensor is Sigma's attempt to address the shortcomings of the Merrill sensor while still retaining the advantages of the Foveon design. Specifically, I believe that the primary design goal of the Quattro is improvement of red and green- channel signal strength, with concomitant improvement in signal-to-noise ratio — all without sacrifice of acuity.[3] This paper is the first of at least two that will test whether and to what extent the Sigma has achieved that design goal.

2. Sensor Design

2.1. Merrill Sensor Contrasted with Bayer Sensor

As already mentioned, the Merrill sensor is a "conventional" Foveon sensor in which trichromatic color information is obtained by vertical stacking of red, blue, and green-sensitive photosites (Fig. 1). *Importantly, in the case of the Merrill cameras, there is a one-to-one correspondence between each raw image pixel and a stack of three photosites.* Thus, the color of an *image* pixel is computed from the responses of three *spatially coincident receptors* — located at the same *x-y* coordinates on the sensor. Merrill raw images have a resolution of 4,704 x 3,136 = 14.75 MP.[4]

The vertical capture of color information in the Foveon sensor contrasts with the horizontal capture in Bayer color-matrix sensors that are used in virtually all other cameras (Fig.

---

1 Service, Phil. 2014. Why I Use Sigma DP Merrill Cameras: Sigma DP3M vs. Nikon D7100

2 That, in a nutshell, is why higher ISO images are also noisier (in all cameras): they are intentionally underexposed.

3 Sigma's description of the Quattro sensor design supports my position: *"While retaining the distinctive characteristics of its predecessors, it offers an even higher level of image quality. Resolution is 30% higher, and its* ***noise characteristics*** *are further enhanced."* And, *"This unique structure prevents the deterioration of the* ***signal-to-noise ratio*** *that is typically associated with an ultrahigh megapixel count."* [emphasis added] (Both quotes from here.)

4 Because each X3F raw image pixel is assembled from three stacked photosites, Sigma literature refers to the resolution of Merrill cameras as 3 x 14.75 MP, or even 44 "recording megapixels". The claim is misleading: because from the point of view of displaying, printing, or cropping images, image resolution, not sensor resolution, is the relevant metric.

1). Since any *x-y* location on the Bayer sensor records only one of three colors (red, green, or blue), the complete color information for an *image* pixel must be interpolated from surrounding photosites with different color filters. In other words, the complete color information is

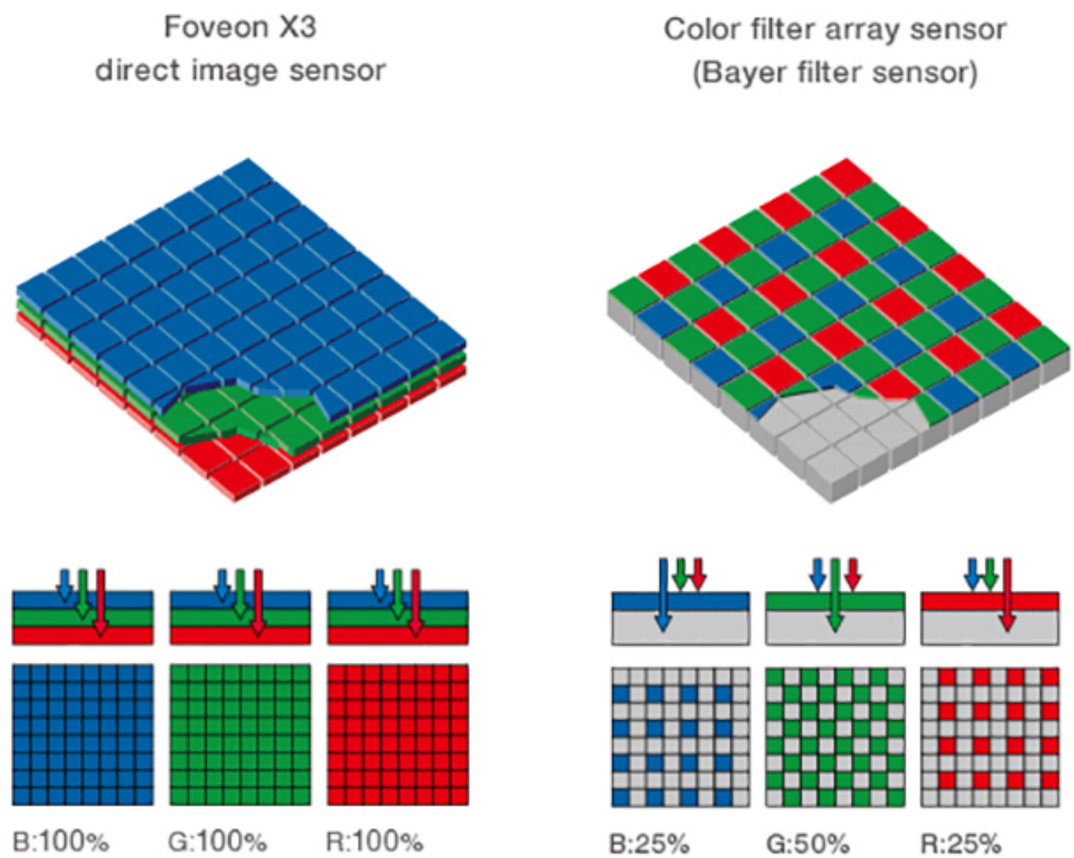

Fig. 1. Contrast between Sigma Merrill sensor(left) and a standard Bayer array sensor (right). Image source: http://www.sigma-global.com/en/cameras/dp-series/technology/

computed from several photosites that are *not* spatially coincident. Interpolation degrades image acuity (resolution and sharpness) because each image pixel is no longer an *independent sample* of the image field. Also, at certain spatial frequencies, detail in the image may interact with the color filter matrix to produce artifacts such as moiré. In order to avoid or reduce moiré, Bayer sensors often have antialiasing filters that intentionally blur high frequency detail, and further reduce acuity. AA filters are not necessary with Foveon sensors.[5]

[5] Many current Bayer sensors dispense with AA filters because moiré is less of a problem as photosite sizes become smaller — *i.e.,* the spatial frequency of the color filter matrix increases.

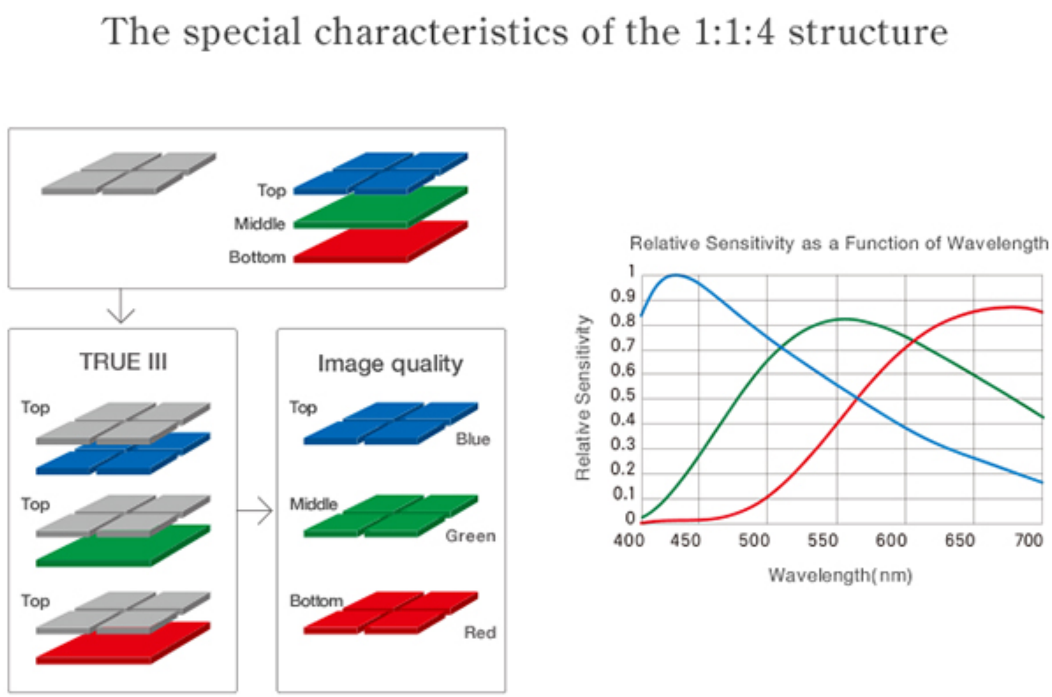

Fig. 2. Layout and image processing (TRUE III) of Sigma Quattro sensor, together with spectral sensitivities. Image source: http://www.sigma-global.com/en/cameras/dp-series/technology/

2.2. The Quattro Sensor.

The Quattro sensor is a variation of the Foveon concept in which the linear resolution of red and green sensitive photosites is one-half that of the blue-sensitive photosites (Fig. 2, top left panel). As a result, there are four times as many blue-sensitive as red or green sensitive photosites, and each red or green site presumably has four times the area of a blue site. Sigma refers to this as a 1:1:4 structure. The resolution of the top (blue-sensitive) layer of photosites is 5,424 x 3,616 = 19.6 MP. The remaining two layers are 2,712 x 1,808 = 4.9 MP, each.[6] All other things being equal, the SNR of red and green-sensitive sites should be twice that (√4) of blue-sensitive sites. Of course, all other things are not equal: we know that less light penetrates to the green and red-sensitive layers. Therefore, we do not expect the red and green color channels be half as noisy as the blue channel. But, we should expect the Quattro red and green channels to be less noisy than the corresponding Merrill channels.

Given that the sensor layers have different resolutions, and that only the blue-sensitive layer has the same resolution as the raw image files, the question arises as to how trichromatic color information is computed for each raw image pixel. Sigma literature makes possible an educated guess. First, Sigma makes a point of saying, rather cryptically, that the top layer

[6] Sigma claims 29 "effective" MP (19.6 + 4.9 + 4.9). The Quattro cameras can output 39 MP JPEG files, and in some instances Sigma claims 39 effective MP resolution. Nevertheless, the raw image resolution is 19.6 MP.

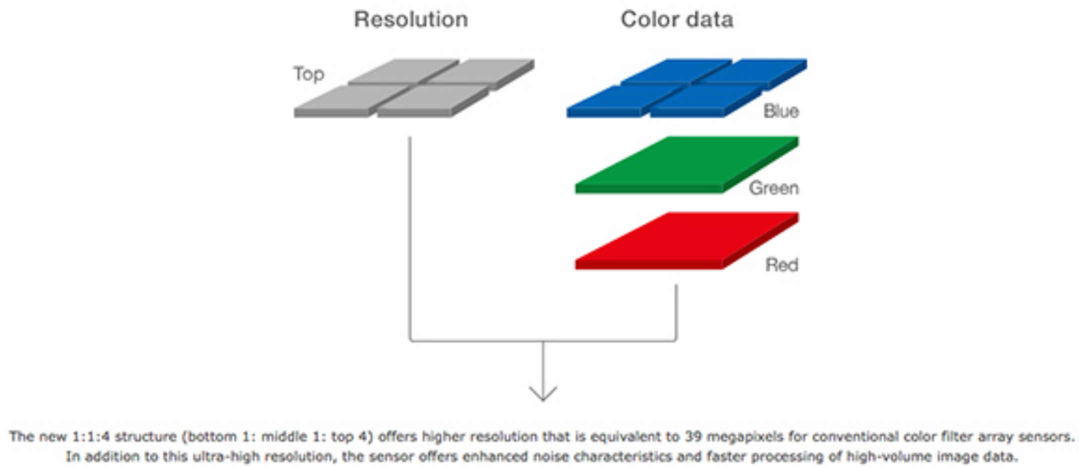

Fig. 3. Resolution information combined with color information in Quattro sensors. Image source: http://www.sigma-global.com/en/cameras/dp-series/technology/

captures *luminance* as well as color information. I say "cryptically" because Sigma's own literature reveals that all three layers have broad spectral sensitivity (Fig. 2, right panel). In other Sigma figures (Fig. 3), the top layer is labeled as capturing resolution information (rather than luminance).[7] Further, Sigma's figures and accompanying text make it clear that information from the top layer (whether it be "luminance" or "resolution"), is used in conjunction with information from the green and red layers to compute the final RGB values for each raw image pixel (Fig. 2, lower left). Data consistent with this interpretation can be extracted from RawDigger and Sigma Photo Pro. Specifically, RawDigger reveals that the X3F raw files have a structure that consists of 2 x 2 pixel blocks. Within each block, green channel values are the same for all four pixels, as are the red channel values. However, the four pixels within each block will have different blue-channel values (Table 1). After processing through Sigma Photo Pro, the block structure becomes obscured because each pixel in a block will have unique green and red channel values. The exact procedure by which the green and red channel values are adjusted is not specified by Sigma. It is certainly possible that it does *not* involve interpolation in the sense used for de-mosaicing Bayer sensor images. Rather, it is possible that the process requires information only from the blue-channel photosite that corresponds to the same raw image pixel (as implied in Fig. 2). If that is the case, pixels in the processed images have only limited non-independence that results from shared red and green channel information within each block of four photosites. The effects on image acuity may be relatively minor, or even imperceptible.

[7] To a certain extent "luminance" and "resolution" are interchangeable in this context. This is easy to see by exploring Lab color mode in Photoshop: most of the image detail is conveyed by the Lightness (i.e., luminance) channel. A nice demonstration can be found here.

Table 1. Color channel values for four contiguous blocks of four image pixels each.*					
Block	RawDigger Pixel Coordinates		X3F Raw File (14-bit values)		
	x	y	Red	Green	Blue
1	2820	1610	1921	2494	2524
1	2820	1611	1921	2494	2417
1	2821	1610	1921	2494	2484
1	2821	1611	1921	2494	2372
2	2822	1610	1919	2487	2407
2	2822	1611	1919	2487	2527
2	2823	1610	1919	2487	2494
2	2823	1611	1919	2487	2399
3	2820	1612	1941	2536	2485
3	2820	1612	1941	2536	2535
3	2821	1613	1941	2536	2478
3	2821	1613	1941	2536	2479
4	2822	1612	1976	2472	2461
4	2822	1612	1976	2472	2399
4	2823	1613	1976	2472	2478
4	2823	1613	1976	2472	2411
Sample size, n			4	4	16
Mean			1939.3	2497.3	2459.4
Std. Dev.			26.4	27.4	51.5
SNR**			73.4	91.1	47.7

* These image pixels are a sample from Step 5 of the step chart image used for Quattro noise analysis: see Materials and Methods.

** Given that the pixels sample a region of presumably uniform density and color, variation among the color channel values is a measure of noise. Signal-to-noise ratio (SNR) is the ratio of signal mean divided by signal standard deviation.

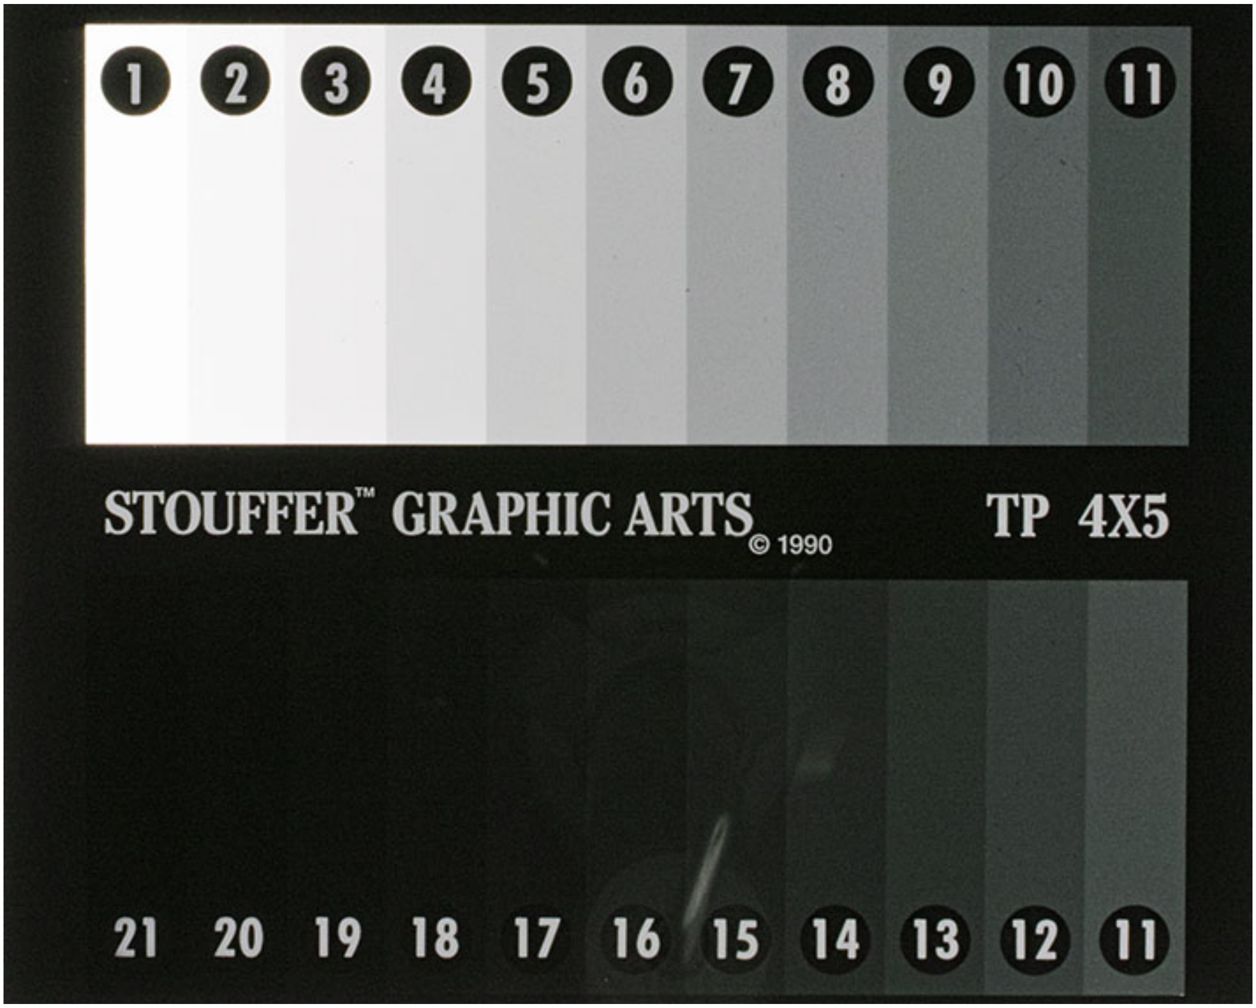

Fig. 4. Transmission projection step wedge. Stouffer Industries, Inc. TP4x5-21. Each step is 1/2 EV. The "blemishes" in zones 14-15 are reflections from a tripod leg.

3. Materials and Methods

3.1. Step Wedge Analysis

A transmission projection step wedge (Fig. 4) was photographed on an LED light box[8] with a Sigma DP1 Merrill and a Sigma dp1 Quattro. Both cameras were set at ISO 100 and f/8. A series of exposures was made for each camera. Images were inspected with RawDigger[9] to find the best ETTR exposure. In the case of the Merrill, the longest exposure that did not

8 Gagne Porta-Trace Stainless Steel LED Light Box. Model 1012-2. The color temperature is not specified.

9 RawDigger 1.2.6. http://www.rawdigger.com

produce overexposure in any color channel was 1/40 second. The corresponding 1/40 sec Quattro image was chosen for comparison.[10] A tripod was used.

The selected images were further analyzed with RawDigger. For each zone of the step wedge, a rectangular sample of pixels was selected. In most cases, samples included about 12,000 pixels, although smaller samples were used in a few zones in order to avoid image artifacts, such as reflections and false overexposure warnings (see Appendix). For each color channel, RawDigger computes the mean value and standard deviation for the sampled pixels. I take the mean value to be the *signal strength,* and I calculated the *signal-to-noise ratio* (SNR) as the mean divided by standard deviation.

3.2. ISO Performance

An ISO exposure series of a single subject (Fig. 5) was made with a DP2 Merrill and dp2 Quattro. Pictures were taken from the same tripod position within a few minutes of each other. In both cases, the in-camera white balance setting was "Daylight/Sunlight". The in-camera color mode was set to "Standard" for the dp2Q, but inadvertently set to "Neutral" for the DP2M. The raw (X3F) images were processed through Sigma Photo Pro 6.3.1, where in all cases white balance was changed to "Shade" (the actual shooting condition), and further adjusted by

Fig. 5. Subject used for ISO exposure series.

[10] The best ETTR exposure for the Quattro was 1/25 sec. For purposes of comparing signal characteristics of the two cameras, I wanted to use the same exposures.

reference to the gray portion of a Spyder Cube. For the Merrill images, the color mode was changed to "Standard." I accepted the default Photo Pro sharpening and noise reduction. No other adjustments were made in Sigma Photo Pro. Raw images were exported as 16-bit ProPhotoRGB TIFFs, and minimally adjusted for exposure and color balance (if necessary) in Camera Raw. Sharpening in Camera Raw was disabled.

4. Results

4.1. Signal

In order to make comparisons between cameras, it was necessary to standardize the mean color channel values. This was a two-step process. First, Quattro raw data is 14-bit, and Merrill data is 12-bit. The Merrill data was arithmetically scaled to 14-bits.[11] Second, the mean value for each channel and each step-wedge zone was converted to a percentage of the zone 1 Quattro blue-channel mean (6,320).

Fig. 6 (top) gives quantitative confirmation that the Merrill green and, especially, red channels are underexposed relative to the blue channel. In fact the red-channel signal is less than one-fourth that of the blue channel across the entire step wedge. (Note the $\log_2$ scale of the y-axis.) This amounts to about a 2 EV difference in exposure between blue and red channels, for a white illuminant. For the Quattro, on the other hand (Fig. 6, bottom), color channel values are much more similar — blue and green diverge only in the darkest zones, and the red-channel value is about 75 – 80% of the blue-channel value (less than 1/2 EV exposure difference).

Direct comparisons between sensors for each color channel are shown in Fig. 7. Note that the Merrill blue-channel signal is about 150 – 160% that of the Quattro (Fig. 7C). In return, the Quattro has consistently superior signal in the red and green-channels, although the green-channel advantage is not large. When signal strength is *averaged* across color channels, the two sensors have essentially the same performance (Fig. 7D). I noted in the Introduction that Merrill overexposure tends to be biased toward the blue channel. And, conversely, when the blue channel is properly exposed, the red and green channels tend to be underexposed. Fig. 6(top) makes the same point.

4.2. Signal-to-Noise Ratio

Photon noise scales with the square root of the signal mean.[12] Thus, SNR will be higher when the mean signal is higher, all other things being equal. Therefore, it is no surprise that the SNR results (Figs. 8 and 9) closely mirror the signal-strength results (Figs. 6 and 7).[13] For the Merrill, the blue channel has consistently higher SNR than the other two channels, the only exception being approximate equality with the green channel in the very brightest zones (Fig. 8,

---

[11] This required attention to the maximum values that each channel can take. According to RawDigger, these are 14,336 and 4,079 for the Quattro and Merrill, respectively. The ratio is 3.5146. Thus, a Merrill channel value of 1,349 = (1,349 x 3.5146) = 4,741 when scaled to 14-bit.

[12] This means that noise increases at a slower rate than signal: doubling the signal increases noise by $\sqrt{2} = 1.41$.

[13] Note that the y-axis in Figs. 8 - 9 is linear, not $\log_2$ scaled as in Figs. 6 – 7.

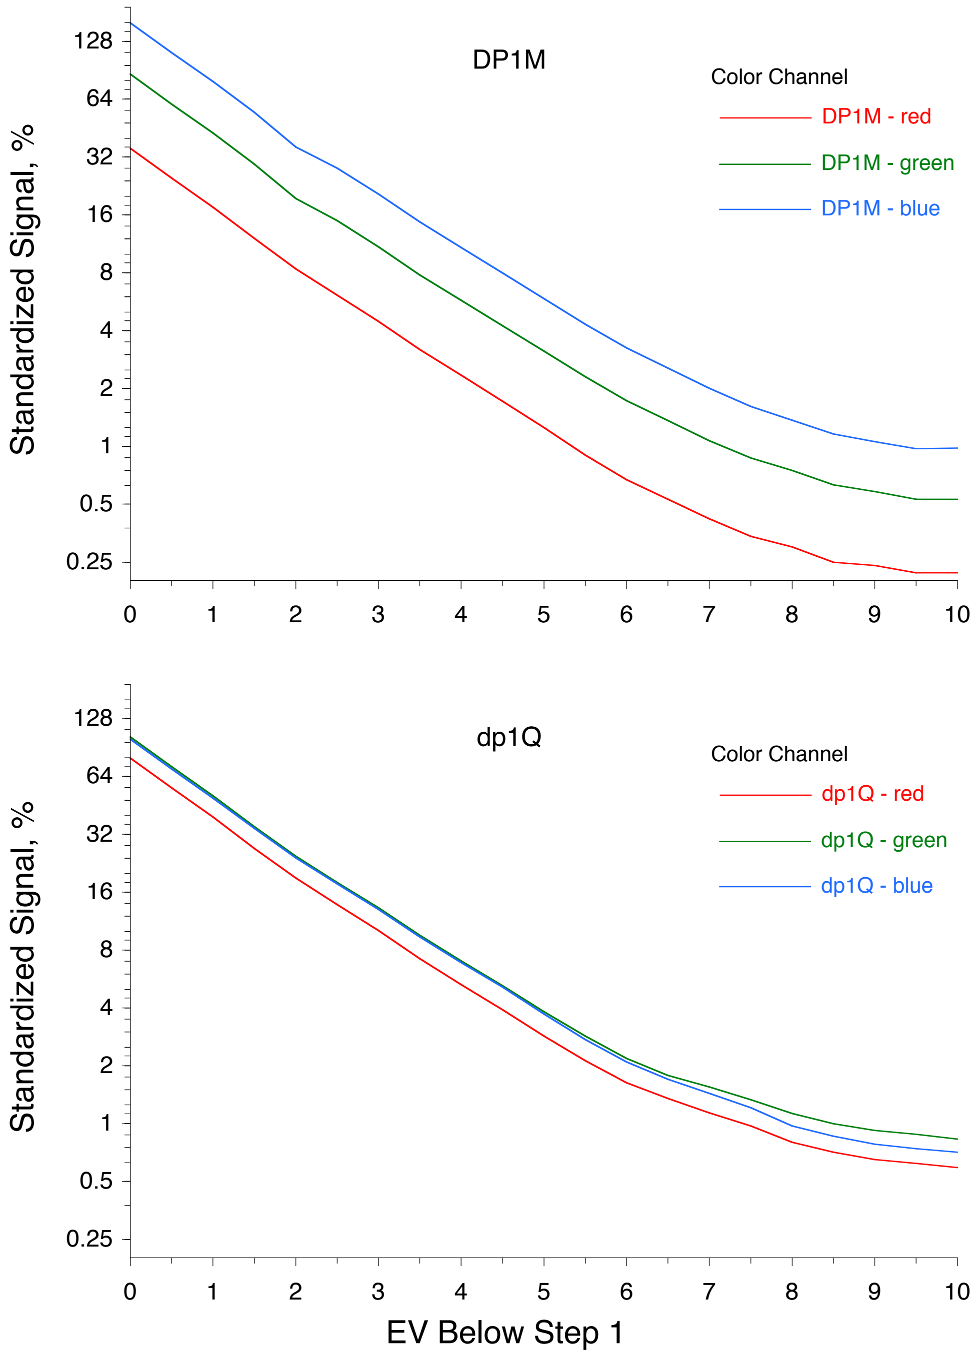

Fig. 6. Standardized signal strength for each color channel. Top: DP1 Merrill; bottom: dp1 Quattro. Note $\log_2$ scaling for y-axis. See text for explanation of "standardized signal strength".

top). The red-channel has substantially lower SNR than either of the other two channels throughout the exposure range. The Quattro exhibits a more complex pattern in which the blue channel has a *lower* SNR than both red and green channels in the brightest parts of the image (Fig. 8, bottom). From about the eighth zone of the step wedge, and into the darker zones, the red channel has the lowest SNR. Blue-channel SNR doesn't match that of the green channel

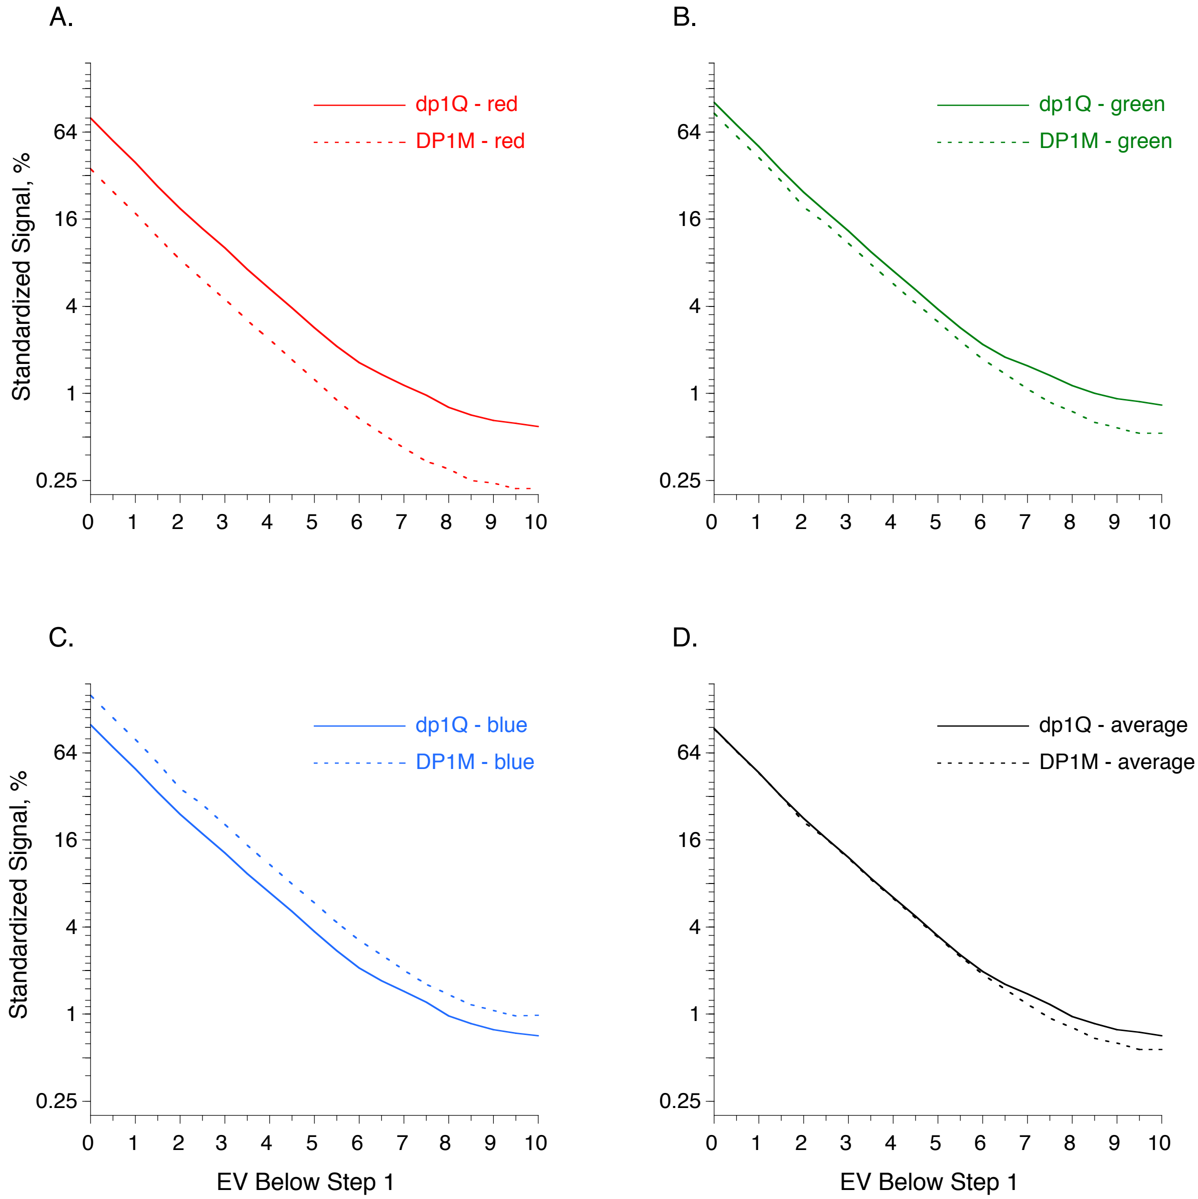

Fig. 7. Signal strength comparison between DP1 Merrill and dp1 Quattro. A: red channel; B: green channel; C: blue channel; D: channel average.

until about zone 12, after which they remain essentially equal. Perhaps more significant than color-channel ranking is the fact that the Quattro exhibits greater similarity of SNR among channels in the dark zones.

Red channel SNR is much improved in the Quattro relative to the Merrill: a 69% increase when averaged over zones (Fig. 9A). The Quattro advantage is less pronounced in the green channel, but is still appreciable: about 38% on average (Fig. 9B). For the blue channel, the Quattro is actually *inferior* to the Merrill except in the darkest zones: by about -7% on average (Fig. 9C). Nevertheless, considering average SNR across all three channels, the Quattro is

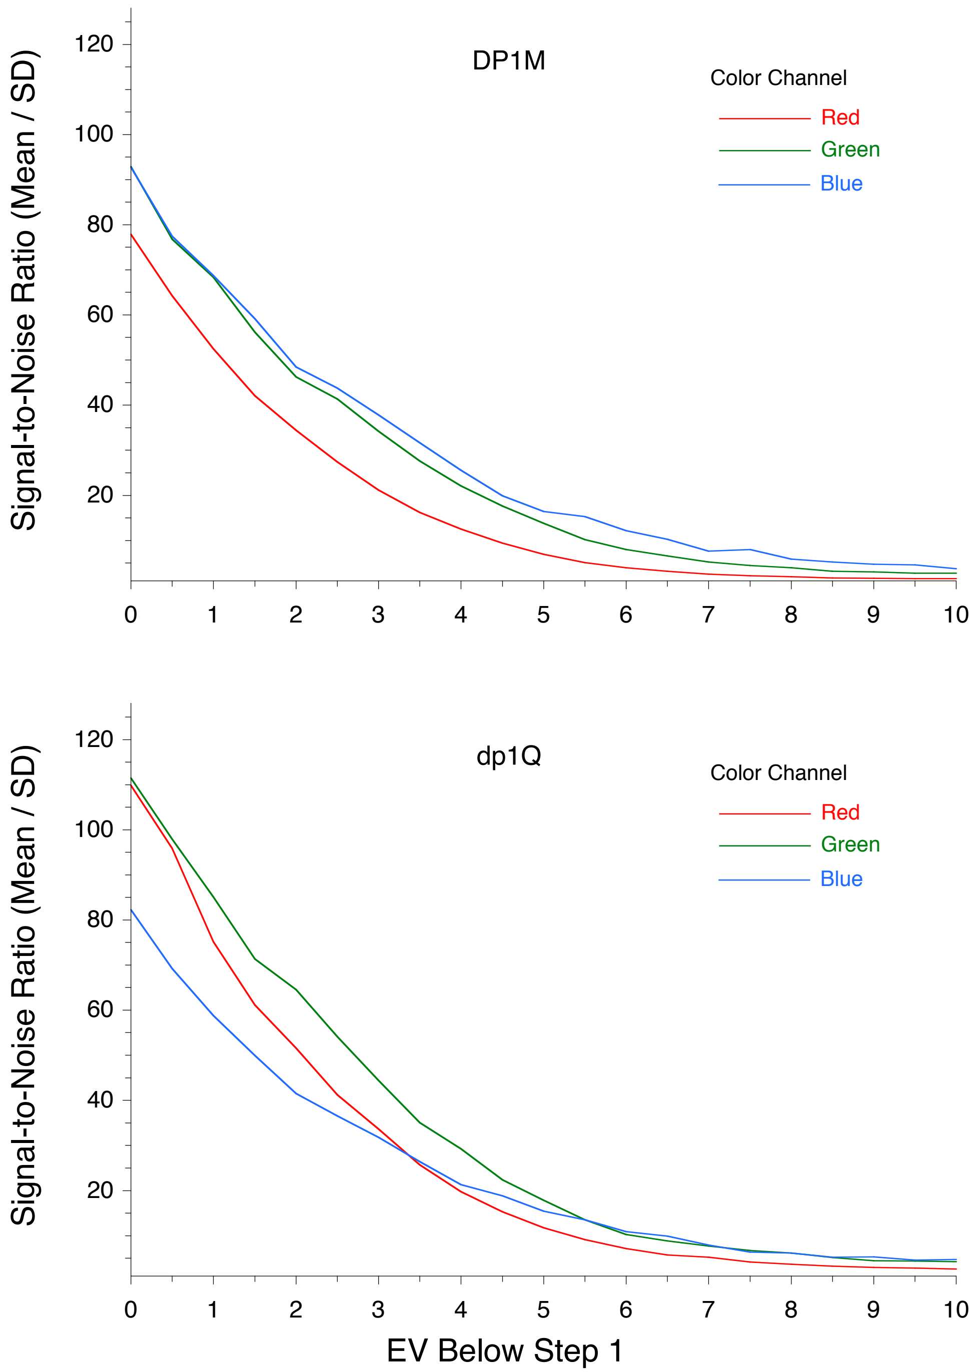

Fig. 8. Per channel signal-to-noise ratio from raw images of transmission step wedge. Top: DP1 Merrill; Bottom: dp1 Quattro.

consistently superior to the Merrill: by about 24 – 33%, depending upon how one chooses to do the calculation (Fig. 9D).

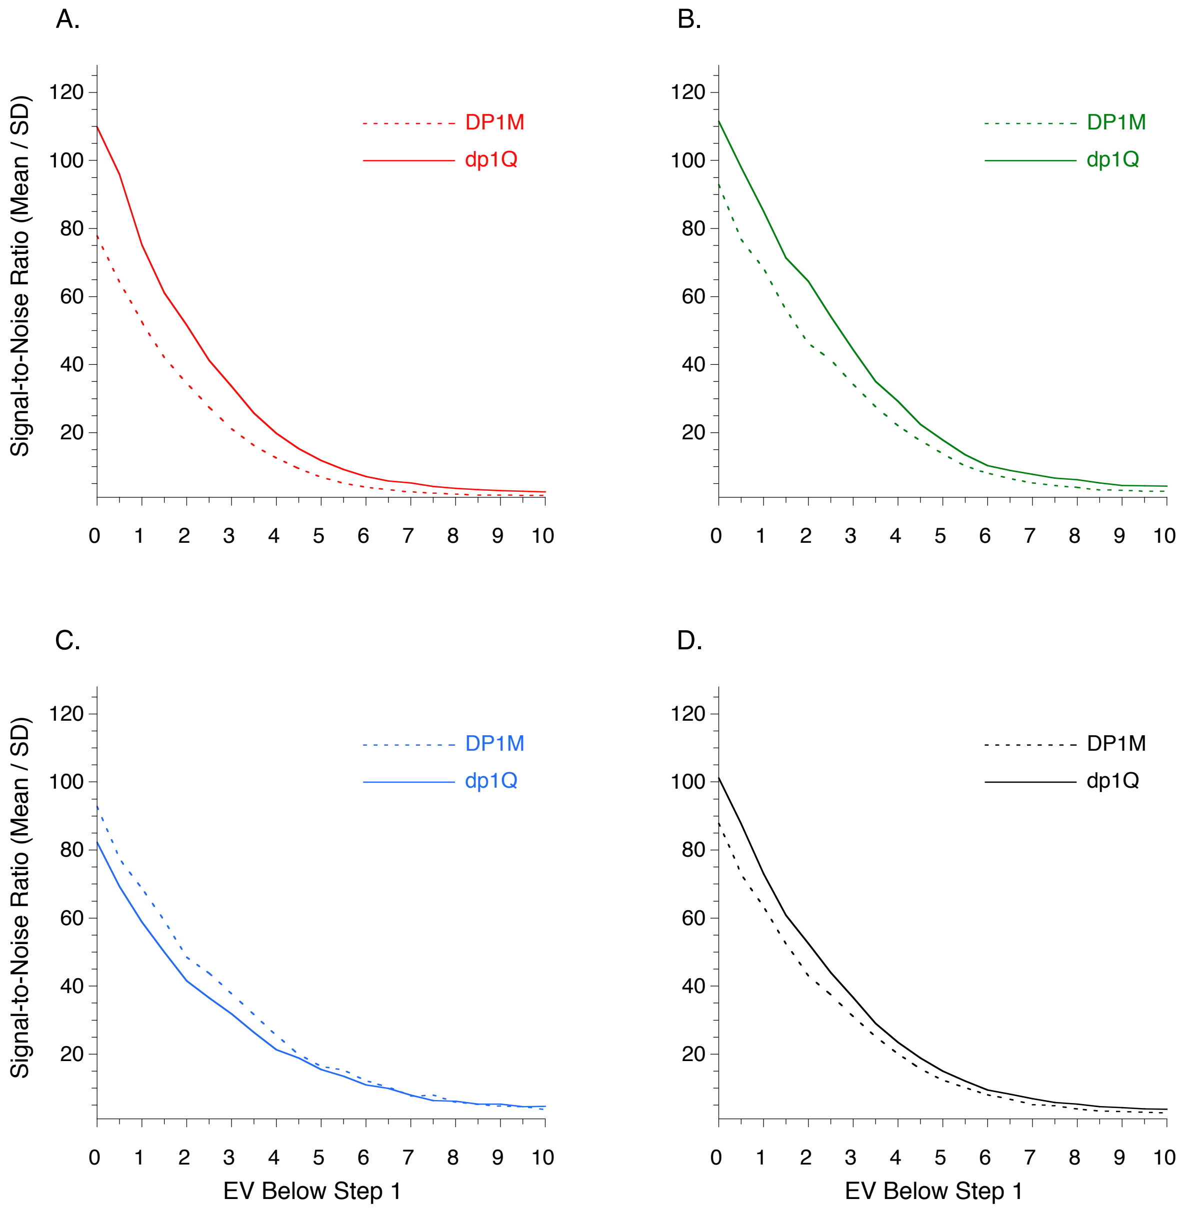

Fig. 9. Signal-to-noise ratio comparison between DP1 Merrill and dp1 Quattro. A: red channel; B: green channel; C: blue channel; D: average of all three color channels.

4.2. ISO Exposure Series

In order to present images at 100% magnification, they are provided as separate files that can be accessed by clicking on the links in the text. I show only images from processed raw files. For both cameras, they were vastly superior to in-camera JPGs at the higher ISO values. A 100% crop from the right side of the image is shown. Results for the Merrill are in Fig. 10. There is progressive degradation of fine detail starting (just perceptibly) at ISO 200, and color saturation decreases at ISO 800 and 1600. Both of these issues could be ameliorated in post-processing. Magenta-green blotching becomes apparent at ISO 800, and is pronounced at

ISO1600. This is a commonly noted defect of higher ISO Merrill images. It is much worse in in-camera JPGs.

Quattro image detail starts to deteriorate at about ISO 400, and color saturation becomes slightly reduced at ISO 1600 (Fig. 11). The Quattro series does not exhibit any magenta-green blotching. However, this should not be taken to mean that the Quattro is immune to strange color artifacts at high ISO. A series of the same fence taken with a dp0 Quattro showed small magenta blotches at ISO 1600 in some parts of the image. Overall, I judge the Quattro to have about a 1 EV ISO advantage over the Merrill. For example, the Quattro ISO 800 image is a fairly good match for the ISO 400 Merrill image. It should be noted that this is fence is a low-contrast subject. As such it does not demand much from the sensor in terms of dynamic range at elevated ISO.

5. Discussion

The results presented here confirm that the Quattro sensor represents a substantial improvement over the Merrill sensor with regard to signal quality. First, with a white illuminant, there is much greater similarity in signal strength across the three color channels. That has important consequences that will be detailed below. Second, signal-to-noise ratio is improved in the green and, especially, red color channels. The improvement in red and green-channel SNR is at least partly a consequence of the increased area of the red and green-sensitive photosites. All other things being equal, SNR should improve as the square root of area. In this case, the red and green photosites are about three times the area of the Merrill photosites. The square root of three is 1.73, which agrees reasonably well with the step wedge noise analysis. On the other hand, the Quattro's blue channel SNR is actually somewhat worse than the Merrill's. I suggest two possible reasons. First the higher resolution of the blue-sensitive top layer of the Quattro means smaller photosites — about 78% of the area of the Merrill photosites. Secondly, a deliberate decision was probably made to absorb less light in the blue-sensitive layer, thus allowing more light to penetrate to the relatively light-starved green and red-sensitive layers below. I am suggesting that a design choice was made to sacrifice blue channel signal strength and SNR in order to increase them in the red and green channels. I believe it was a good choice. The lower blue-channel SNR of the Quattro is most pronounced in well-lit portions of the image: exactly where it is less likely to have a visible effect on image quality. On the other hand, in darker portions of the image, Quattro SNR is more uniform across all three color channels. I suggest that that leads to greater color fidelity in underexposed areas, relative to the Merrill. Finally, in this regard, absorbing less light in the blue-sensitive layer means greater resistance to overexposure in that channel: the channel that is most prone to overexposure in Merrills. For most subjects, it would be an improvement to have all three color channels approach saturation at a similar rate. In fact, when I process Merrill images in Sigma Photo Pro, it is often necessary to dial down overall exposure in order to recover blown blue-channel highlights.

If one defines dynamic range as the darkest EV step with a specified SNR, the Quattro's dynamic range is about 1 EV greater than the Merrill's, considering the red and green channel results. Taking into account the blue channel as well, the Quattro has about a 1/2 EV increase in dynamic range overall, relative to the Merrill. That may be somewhat conservative, as the ISO series suggest that the improvement may be as high as 1 EV. In regard to dynamic range, it is worth mentioning one other point: for the Quattro, the best ETTR exposure of the step wedge was actually 1/25 sec. That is 2/3 EV more than the best ETTR exposure for the Merrill (the 1/40 sec exposure for both cameras that was actually used for comparison). We can get some idea of the Merrill's blue-channel overexposure bias — and the importance of relative uniformity of signal strength across color channels — by examining the step wedge images with RawDigger. At 1/25 sec, zone 1 of the Merrll image is completely overexposed in the blue channel. On the other hand, *no* zone 1 pixels are overexposed in the red and green channels. Green channel overexposure begins at 1/15 sec (+ 2/3 EV). *No* red-channel overexposure is evident in zone 1 at 1/13, sec, the longest exposure tested. For the Quattro, zone 1 does not become overexposed until 1/15 sec, at which point all pixels are blown in *both* blue and green channels. The red channel shows complete overexposure at 1/13 sec. If ETTR exposures are compared (1/40 and 1/25 sec for the DP1M and dp1Q, respectively), the SNR advantage of the Quattro is even greater than depicted in Figs. 8 – 9. The step-wedge-on-light-box-scenario is hardly typical. In particular, it involved a white illuminant and a gray-scale subject. However, the data support the hypothesis that reduced color channel bias in overexposure makes longer ETTR exposures possible. If that is generally true for typical subjects under typical conditions, it further enhances the Quattro's advantage in dynamic range.

The issue that I have not addressed directly in this paper is image acuity. As I mentioned in the Introduction, superior acuity is really the only reason to consider using Merrills. From the data that I have presented here, it does not necessarily follow that the Quattros are better than the Merrills in this regard. In particular, one wonders if the lower resolution of the red and green-sensitive photosite layers will have adverse effects on acuity. An in-depth analysis will be presented in a subsequent paper. For the moment, however, we can examine images of the fence used for an aperture series: both images 1/15 sec, f/5.6. ISO 100. Fig. 12 shows central crops at actual-pixels magnification. You be the judge.

Appendix

The Curious Case of "Overexposed" Dark Pixels

RawDigger often reports pixels as overexposed in the red and green channels of Quattro raw images, even though the pixels are located in very dark regions. This phenomenon has been replicated over three different cameras. It does not occur in the blue channel, and is not seen in Merrill raw files. Fig. 13 shows small crops from two captures taken only seconds apart. The top two panels are the red (left) and green (right) channels from the same capture. The bottom two images are the red and green channels from another image made just seconds earlier. The magenta pixels (actually blocks of four pixels in a processed image) are "overexposed".

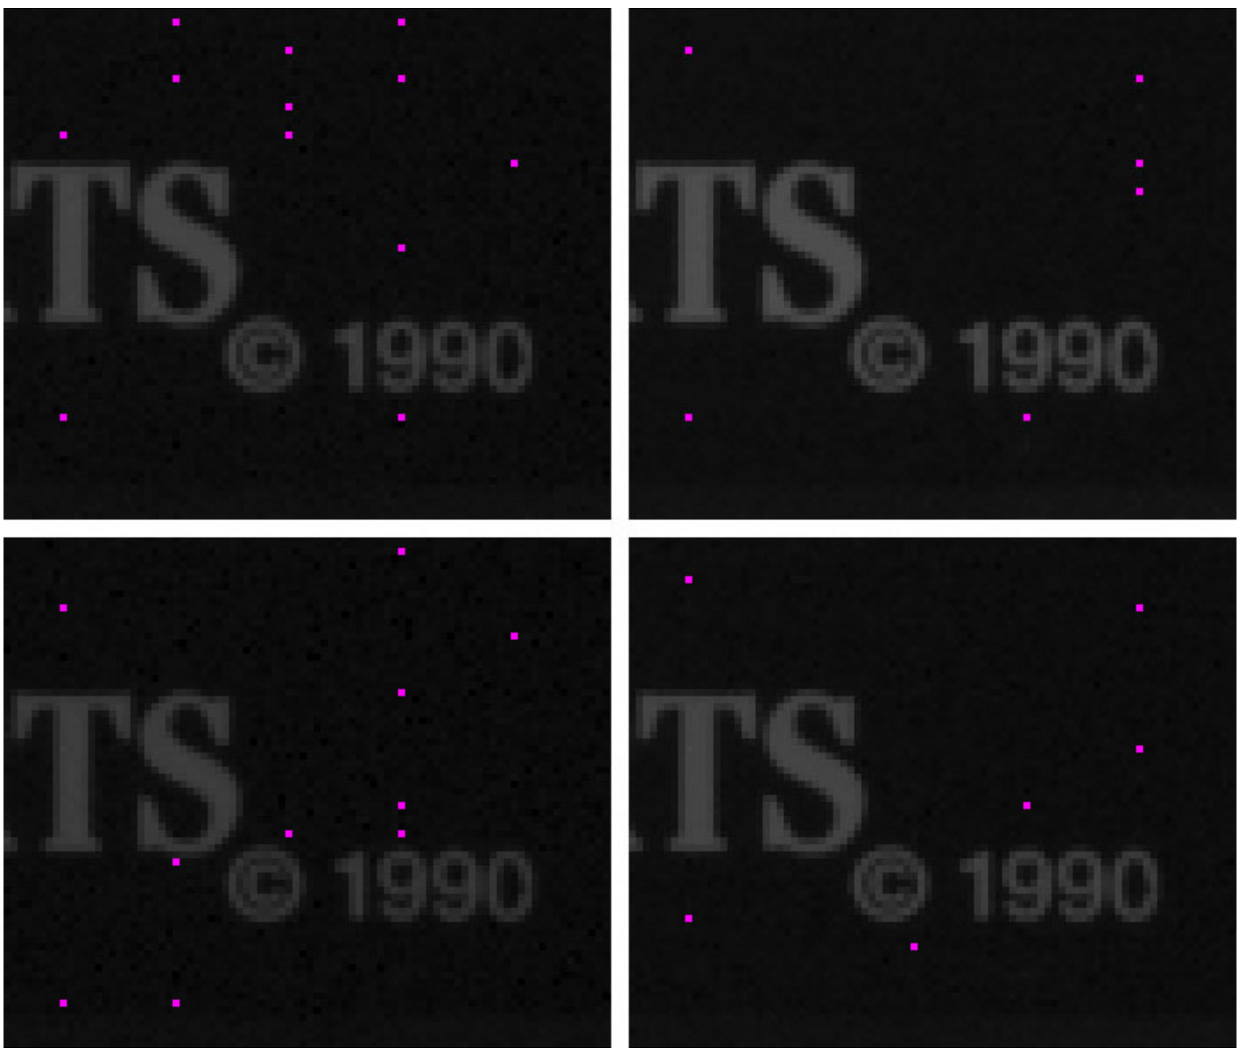

Fig. 13. Four-pixel blocks reported by RawDigger 1.2.6 as overexposed in the red channel (left-side) and green channel (right side). Top two panels from same capture at 1/40 sec. Bottom two panels from same capture at 1/50 sec. Sigma dp1 Quattro. 200% magnification.

Only a small minority of pixels are reported as overexposed simultaneously in both red and green channels: three pixels in the top pair, and one pixel in the bottom pair of Fig. 13. See, for example, the pixel under the "9" in the top pair. Similarly, "overexposed" pixels are generally different within channels from one frame to the next: only one red pixel was the same (left-hand pair of images), and only two green pixels were the same (right-hand pair). More pixels are reported as overexposed in the red than the green channel.

These "hot" pixels are not present in TIFF files produced by Sigma Photo Pro. Either they are real, and SPP corrects them; or they are simply an artifact of the way that RawDigger reads the raw files. I would put my money on the latter.