2012 Data Compression Conference

Progressive-to-Lossless Compression of Color-Filter-Array Images using Macropixel Spectral-Spatial Transformation

Henrique S. Malvar and Gary J. Sullivan

Microsoft Research
One Microsoft Way
Redmond, WA 98052, USA

Abstract: We present a low-complexity integer-reversible spectral-spatial transform that allows for efficient lossless and lossy compression of color-filter-array images (also referred to as camera-raw images). The main advantage of this new transform is that it maps the pixel array values into a format that can be directly compressed in a lossless, lossy, or progressive-to-lossless manner by an existing typical image coder such as JPEG 2000 or JPEG XR. Thus, no special codec design is needed for compressing the camera-raw data. Another advantage is that the new transform allows for mild compression of camera-raw data in a near-lossless format, allowing for very high quality offline post-processing, but with camera-raw files that can be half the size of those of existing camera-raw formats.

1. Introduction

For the compression of images captured using a color filter array (CFA), such as the typical Bayer CFA [1], an appealing avenue of research is the so-called “compression-first” workflow [2]–[5]. While the typical workflow for image processing consists of performing CFA demosaicing first and then compression, in the alternative workflow the compression is performed prior to demosaicing.

A clear advantage of the compression-first approach is that it avoids or reduces the data expansion produced by the demosaicing process. The quantity of this expansion is substantial, and is especially questionable when the compression is intended to operate with high fidelity, such as in a lossless or near-lossless manner. For the typical Bayer-pattern CFA in Fig. 1, demosaicing produces an RGB color triplet per pixel from the original one-color-per-pixel CFA input, thus causing a three-fold expansion of the data.

In lossless or near-lossless operation, the “demosaicing-first” approach could actually cause a net expansion rather than compression of the raw data. Although demosaicing does not increase the inherent entropy, the image compression that follows a demosaicing process typically cannot fully account for the type of redundancy that demosaicing introduces. Ordinary compression schemes would not recognize the difference between true source data and interpolated samples produced from demosaicing, and will thus waste bits on representing the interpolated samples.

Moreover, even if the compression-first approach does not provide a net reduction of the number of bits needed for the compressed picture, it can provide a benefit in computational requirements. The compression-first approach can reduce the amount of

1068-0314/12 $26.00 © 2012 IEEE
DOI 10.1109/DCC.2012.8

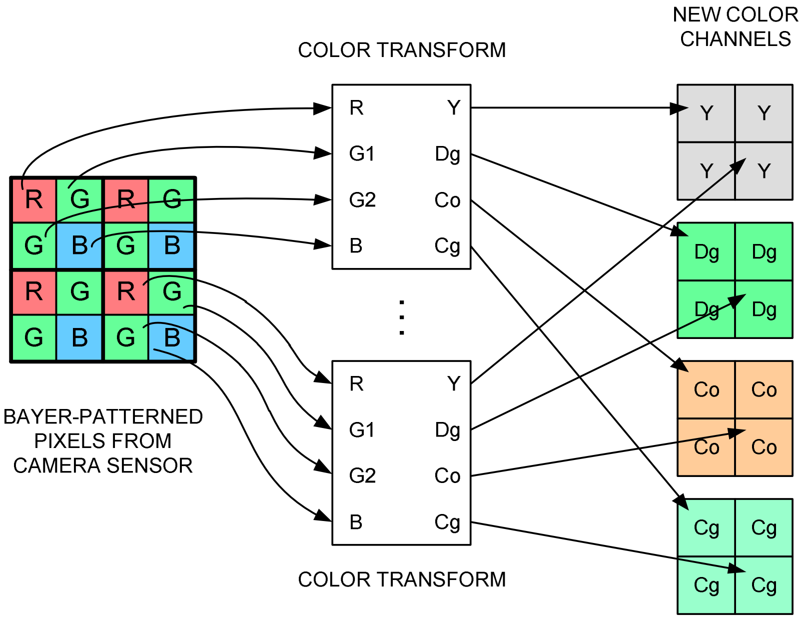

Figure 1. Four macropixels of a typical color filter array (CFA) using a Bayer mosaic (also called RGGB) pattern for single-sensor color image capture, and the proposed mapping of the original CFA image into four color channels {Y, Dg, Co, Cg}, each 1/4 of the original image size.

data that needs to be processed as input to the compression encoder and reproduced from the subsequent compression decoder. In addition, the compression-first approach can enable a system design in which critical resources are saved in other ways. A camera is often a portable device that operates on battery power with rather limited processing resources and a limited capacity for data transfer and storage. After capture by the camera, the images are often moved from the camera to a more powerful image manipulation environment for post-processing. Compressing the image prior to storing or transferring it can improve the speed and capacity characteristics. The compression-first approach can also allow some intensive signal processing algorithms to be moved off of the camera to the post-processing system where greater computing resources are available. Moving the demosaicing to the post-processing system not only saves on-camera resources, but also makes it feasible to use more sophisticated algorithms for the demosaicing and other processing stages. Processes such as lens aberration correction, chromatic correction, exposure and white-balance adjustment, dead pixel removal, gamma pre-compensation and tone mapping, etc., may also be best done off-camera where greater computational resources are available (and software updates are more feasible) – provided adequate fidelity has been retained.

Because the quality of the image is critical to the ability to perform such post-processing effectively, and because the ability of a compression system to produce high image quality is generally desirable in many applications, we focus on designing a flexible compression system in which reconstruction fidelity can be adjusted from providing minimal rendering-quality lossy encoding up to a fully lossless representation (that is, to make the exact values of the CFA image from the camera sensor available for off-camera post-processing). This is in contrast to various designs that may be suitable for lossy or lossless coding, but not both.

It is furthermore desirable to provide this flexibility using what is known as *scalable*, *progressive* or *embedded* encoding [6], such that a more lossy representation can simply be extracted from a higher-fidelity representation, without performing a full process of decoding and re-encoding. This scalability principle can be applied to image resolution as well as to image fidelity at a particular resolution. In a similar spirit, it should be possible to easily extract a *region-of-interest* area from a larger coded image, and to perform such operations as image cropping without full decoding and re-encoding. Two compression schemes that support this type of functionality are the international standards JPEG 2000 [7]–[8] and JPEG XR [9]–[12].

Moreover, the degree to which a compression scheme is useful for a broad range of applications is also highly important. Ideally, the method for encoding images for some particular application should require no alteration of the core technology that is appropriate for general use. Here again, we identify JPEG 2000 and JPEG XR as appropriate choices for such technologies.

The effectiveness of a color-image compression design is typically improved by performing a color conversion transformation. This step improves the compression capability by decorrelating the data of the sampled color channels. It is intuitively apparent that the data samples produced by a CFA are strongly correlated – there is even substantial correlation between samples for different sampling color elements.

In addition to providing data compaction as a decorrelation transformation, a color-conversion transformation also enables the coding system to take advantage of the differences in perceptual importance between different parts of the color spectrum.

As we require a compression design that supports lossless compression, we need for the color conversion processing stage to be fully invertible, that is, with no rounding errors. Naturally, it is also valuable to minimize the computational complexity of such color conversion process. In addition, the compression system should also minimize the introduction of visually-noticeable systematic artifacts when the coding is not lossless.

Prior work on compression of CFA images has generally not fulfilled the set of desired characteristics described above. Some schemes introduce irreversible distortion in the pre-processing steps that form the input data to the compression process. Some introduce distortion in the color conversion stage. Some schemes that do provide lossless coding capability do not have the ability to also operate in a lossy mode, as in most proprietary “camera raw” formats [13]. In schemes that can be operated either in a lossless or lossy fashion, there is typically no smooth progression between these modes of operation, no scalable/progressive encoding flexibility, and no region-of-interest or cropping support. Some schemes interleave the source data in ways likely to produce irritating artifacts if the coding is lossy. And some schemes use a design that is substantially different from ordinary general-purpose compression technology. In [5], Zhang and Wu developed an interesting approach using direct encoding of a CFA image with a wavelet image coder. However, the work in [5] did not optimize its transformation process for improved color decorrelation, and it was also specialized to the use of a wavelet image encoder.

In this paper we present an image coding scheme that uses a low-complexity integer-invertible macropixel spectral-spatial transformation (MSST) for each macropixel of a CFA as a pre-processing stage, to convert the original Bayer-sampled image data into a four-channel representation, as shown in Fig. 1, which can be handled readily by standard compression designs such as JPEG 2000 and JPEG XR. We refer to such as pre-processing step as a spectral-spatial transformation (a term previously used in [5]) because it performs a transformation across both the color component domain and the spatial location domain as well – transforming the data elements of each elemental block of a CFA pattern into an equal number of integer-valued data elements that are each processed as a color channel by the subsequent general-purpose image encoder. The resulting scheme then supports all of the flexibility and scalability features of the underlying general-purpose image coding scheme. A specific MSST suitable for Bayer CFA images (usually referred to as "camera raw" images) is described, and experiment results illustrate the use of this transformation within the JPEG 2000 and JPEG XR image coding designs, enabling their direct support of camera raw images.

The MSST approach described herein was actually implemented as part of the Microsoft "HD Photo Device Porting Kit" [14], first released in Nov. 2006, on which the design of the JPEG XR standard was originally based. However, although this aspect of the design was discussed within the JPEG committee during the standardization process, it has not previously been described in scientific publications.

2. Four-channel spectral-spatial transform

In the JPEG XR image coding system [9]–[12], in order to allow for lossless and lossy compression in the same architecture, we use a new integer-reversible color transform [15] that maps a set of original {R,G,B} values into one approximate luminance value and two new kinds of chrominance channels: {Y,Co,Cg} (Y = approximate luminance, Co = "excess orange", and Cg = "excess green"). Using ladder/lifting operators, we define the YCoCg-R transform, which is fully-reversible in integer arithmetic. We have shown that the YCoCg color space leads to higher coding gains than either the commonly-used YUV (YCrCb) transform or the RCT transform used in JPEG 2000 [15].

For CFA raw images, e.g. as in the Bayer pattern shown in Fig. 1, each individual pixel position has only one color value, corresponding to a specific color depending on the pixel position. Thus, each sample of the pixel data array contains only one channel, consisting of one of three types of signals that are subsampled in the spatial domain and filtered in the chromatic domain. Adjacent pixels come from different chromatic filters and thus do not have cross-correlation characteristics similar to traditional luma/chroma image channels. As one would expect, feeding this interleaved data directly to a single-channel image compressor does not, in general, lead to the best results [5].

An efficient approach towards compression of raw images is to map the original CFA pixel array into a set of images, each with correlation characteristics similar to those of RGB or YUV images. We propose the spectral-spatial transform shown in Fig. 1 [16]. For each 2x2 Bayer cell (or "macropixel"), we compute the 4-color transform by the direct-inverse transform pair as follows:

$$\begin{bmatrix} Y \\ Dg \\ Co \\ Cg \end{bmatrix} = \begin{bmatrix} 1/4 & 1/4 & 1/4 & 1/4 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 1/2 & 1/2 & -1/2 & -1/2 \end{bmatrix} \begin{bmatrix} G_1 \\ G_2 \\ R \\ B \end{bmatrix} \Leftrightarrow \begin{bmatrix} G_1 \\ G_2 \\ R \\ B \end{bmatrix} = \begin{bmatrix} 1 & -1/2 & 0 & 1/2 \\ 1 & 1/2 & 0 & 1/2 \\ 1 & 0 & 1/2 & -1/2 \\ 1 & 0 & -1/2 & -1/2 \end{bmatrix} \begin{bmatrix} Y \\ Dg \\ Co \\ Cg \end{bmatrix} \quad (1)$$

This is an extension of the YCoCg transform as defined in [15] and used in JPEG XR. The YDgCoCg color space has several interesting properties:

- The Y channel is just an average of all four original values in a macropixel, with a 50% green contribution and 25% contributions of R and B, just like in the original YCoCg sapce. Thus, Y is approximately a luminance channel, that is, it contains an approximation of the grayscale brightness information of the image. The dynamic range of Y is the same for each of the original R, G_1, G_2, and B pixels.
- Dg, Co and Cg are all chroma channels, that is, they do not carry luminance information. If a macropixel has values R = G_1 = G_2 = B, then the macropixel is just a gray level, for which Dg = Co = Cg = 0.
- Dg is a “difference green” channel. The smoother the original image, the smaller the values of the Dg pixels, as desired.
- As in our previous YCoCg color space design [15], Cg is an “excess green” channel; in Eqn. (2), the original green values can be reconstructed from just the luminance Y, difference green Dg, and excess green Cg. Co is like an orange channel (although not quite), because the value of Co is largest when the input pixels have R = maximum value and B = minimum value, corresponding to a hue between red and yellow (depending on the G value) whose mid-point is orange.
- The direct and inverse transform matrices only have entries with magnitudes of 0, 1/4, 1/2, or 1. This reduces computational complexity, since the multiplications can be implemented by right-shift operators for integer pixel values.

Fig. 2 shows an example of the application of the YDgCoCg transform, with the four resulting quarter-size images produced from the original CFA image. The Dg image is of a highpass nature, as it contains green differences within each original macropixel.

2.1. Integer-invertibility of MSST

For lossless compression applications, we need to convert the direct and inverse transforms in (1) to ladder/lifting-based operators that perform rounding via shifting, in the appropriate orders for exact reversibility in integer arithmetic. Using the same techniques as described in [15], we arrive at the following:

- YDgCoCg-R Direct Transform:

```
Co = R - B;
Dg = G2 - G1;
 u = B + (Co >> 1);
 v = G1 + (Dg >> 1);
Cg = v - u;
 Y = u + (Cg >> 1);
```

(2)

Figure 2. Example of images produced with the 4-color spectral-spatial transform. Top left: original image, with three colors (RGB) per pixel location. Top right: simulated CFA (one color per pixel) using the Bayer sampling pattern of Fig. 1. Bottom left: a zoom of the CFA image; we can clearly see the discontinuous intensities from each of the R, G, and B locations. Bottom right: the sub-images formed by the 4-channel spectral-spatial color transform (clockwise from top left: Y, Dg, Cg, Co).

- YDgCoCg-R Inverse Transform:

```
 u = Y - (Cg >> 1);
 v = u + Cg;
G1 = v - (Dg >> 1);
 B = u - (Co >> 1);
G2 = G1 + Dg;
 R = B + Co;
```

(3)

For exact reversibility under integer arithmetic, each operation in Eqns. (2) and (3) should be computed in the order shown. Computational complexity is very low, as only a small number of additions and shifts are needed.

3. Application to non-Bayer CFA patterns

The MSST approach can also be readily generalized for application to non-Bayer CFA pattern designs, as follows:

- Any full-spectrum brightness (sometimes called "whiteness") samples should not be modified by transformation with samples of other colors, in order to preserve their nature as a luminance representation – since the luminance signal is perceptually very important and tends to contain much of the signal energy.
- Samples for different color filters within macropixels should be processed by a (simplified) decorrelating transformation to enhance compression performance.
- Samples that form a rectangular grid both within and across macropixel boundaries, such as the full-spectrum (sometimes called "whiteness") samples of the newer Kodak color patterns B and C [17] can be processed subsequently as a larger rectangular array (e.g. as a double-height or double-width spatial array) by the subsequent spatial compression transformation step.
- Other cases of multiple same-color samples within macropixels, such as the Y_1 and Y_2 samples of the Bayer-like Cyan, Yellow_1, Yellow_2, and Magenta (CYYM) sampling grid, and various samples of the newer Kodak color patterns [17], should be transformed by spatial transformations within macropixels (as was done with G_1 and G_2 for the Bayer case).
- Transformation steps should be performed using ladder/lifting operations to enable integer invertibility.

4. Lossless and lossy coding experiments

As shown in Fig. 1, we can use our spectral-spatial YDgCoCg transform to implement an efficient compressor for CFA raw images, using an existing codec (designed for non-subsampled images), such as JPEG 2000 [7]–[8] or JPEG XR [9]–[12].

To evaluate the performance of such an approach, we used as input images the 16-bit per color JPEG XR test files, a set of 14 TIFF images known as “HD Photo test set” that was previously distributed to the JPEG committee. These are high-resolution (4.5 and 10 megapixel) images captured with high-end digital cameras. To simulate a camera-raw image, we subsampled them according to the Bayer pattern of Fig. 1. We also reduced their dynamic range to 15 bits (as the actual sensor data has at most 14-bit resolution), because Dg, Co, and Cg have one extra bit of dynamic range, compared to the input data.

Table 1 shows results for lossless compression. In Column 2 we show the result of feeding the CFA image data directly to the codec as if it were a grayscale image. As discussed in [5], that approach produces reasonable results when the codec is wavelet-based, such as JPEG 2000, because the first stage of the wavelet transform acts somewhat like a color transform and the filtering effect of the wavelet bases is akin to some level of interpolation of the CFA data. For the JPEG 2000 encoding, the reversible 5/3 wavelet transformation was applied here, as is necessary for lossless operation. From Table 1 we

HD Photo image number	JPEG 2000 on CFA image	JPEG XR on CFA image	G merge [5] with JPEG 2000	YDgCoCg with JPEG 2000	YDgCoCg with JPEG XR
01	12.09	12.76	11.87	11.86	11.68
02	10.74	12.20	10.23	10.08	9.94
03	11.33	11.86	11.20	11.16	11.04
04	11.90	12.50	11.80	11.68	11.46
05	11.38	11.90	11.36	11.13	10.96
06	11.84	13.08	11.49	11.36	11.20
07	12.00	12.26	12.04	11.89	11.76
08	10.34	11.60	9.93	9.81	9.77
09	10.29	11.38	9.99	9.83	9.74
10	12.02	13.28	11.36	11.26	11.12
11	11.87	12.44	12.01	11.98	12.02
12	10.69	11.77	10.24	10.26	10.15
13	11.46	12.23	11.22	11.15	11.01
14	10.95	11.70	10.71	10.69	10.65
Average	11.35	12.21	11.10	11.01	10.89

Table 1. Bits-per-sample for lossless coding of CFA images of "HD Photo" set, truncated to 15 bits/sample. In cols. 2 and 3, the CFA image is presented to the encoder as if it is a grayscale image. Col. 4 uses the approach in [5]. The last two columns use the proposed approach in Fig. 2.

see that with JPEG 2000 applied directly to the CFA image, the effective bit rate drops by 3.65 bits/sample (from 15 to 11.35 bits/sample)[1]. With JPEG XR applied directly to the CFA image, the rate is higher, as the length-4 core transforms don't quite provide an intrinsic CFA interpolation. In Column 4 we show the rate with the "G channel merging plus color differences" approach of [5], which improves the rate slightly. In Columns 5 and 6 we show the results with our proposed approach, with JPEG 2000 and JPEG XR compression, respectively. Thanks to the relatively smooth Y, Co, and Cg channels produced by our 4-channel spectral-spatial transform, the rates are even lower, with the best results now coming from JPEG XR, showing a 4.11 bits/sample rate reduction.

For lossy compression, Table 2 shows results at the rates of 2.0 and 1.0 bits/sample. For each rate, the first reference is the classical approach of first interpolating the CFA raw values with a demosaicer, followed by encoding. For our experiments, we used the linear demosaicer described in [18]; in practice, in-camera nonlinear demosaicers can provide an improvement of 0.5–1.0 dB in peak-signal to noise ratio (PSNR). As a desired functionality of our approach is progressive-to-lossless encoding, for JPEG 2000 we used the 5/3 reversible wavelet transform, which is the same as used for lossless encoding.

From the two right-most columns for each rate, we see that results with JPEG 2000 are comparable to those with JPEG XR at 2.0 bits/sample. For the higher degree of compression of 1.0 bit/sample, the results with JPEG 2000 are better than those with JPEG XR, but only by 0.1 dB, which is not statistically significant. As a second reference, we also compare the results with the naïve approach of direct lossy

---

[1] If the raw image data came from a 12-bit sensor, for example, as we would expect about the same reduction of 3.65 bits/sample for JPEG 2000, we would expect a resulting rate of 8.35 bits/sample.

HD Photo image number	Compression to 2.0 bits/sample				Compression to 1.0 bit/sample			
	Demosaic first	JPEG2K on CFA	YDgCoCg JPEG2K	YDgCoCg JPEG XR	Demosaic first	JPEG2K on CFA	YDgCoCg JPEG2K	YDgCoCg JPEG XR
01	43.99	43.45	43.70	43.68	42.54	39.99	41.02	40.87
02	52.13	51.23	51.91	51.86	51.67	48.48	50.91	50.72
03	43.63	43.59	43.63	43.54	42.94	41.59	42.42	42.35
04	42.60	42.09	42.30	42.26	42.02	40.49	41.27	41.05
05	46.59	46.37	46.54	46.46	46.16	44.56	45.44	45.30
06	43.29	42.97	43.23	43.23	42.93	41.10	42.56	42.36
07	40.38	40.25	40.26	40.17	39.69	38.36	39.01	38.91
08	47.06	46.94	47.09	46.96	46.84	45.90	46.79	46.68
09	51.27	51.08	51.28	51.16	50.72	49.47	50.52	50.33
10	44.65	43.60	44.40	44.28	44.04	40.18	43.33	43.22
11	42.41	42.39	42.05	41.88	40.53	39.14	37.97	38.18
12	50.43	50.02	50.32	50.21	49.58	47.43	49.14	49.05
13	44.96	44.95	45.02	44.93	44.19	42.71	43.77	43.56
14	46.84	46.63	46.85	46.73	46.23	44.32	45.79	45.72
Average	45.73	45.40	45.61	45.52	45.01	43.12	44.28	44.18
	Delta	−0.33	−0.12	−0.21	**Delta**	−1.88	−0.72	−0.84

Table 2. Peak SNR for the reconstructed RGB data for lossy compression of the CFA raw image in the "HD Photo" set. We compare the usual approach ("demosaic first") with applying lossy JPEG 2000 compression directly on the CFA array, and with using our proposed 4-color YDgCoCg MSST with a JPEG 2000 or JPEG XR encoder.

compression of the CFA data; we see in Table 2 that such an approach leads to a noticeable loss of PSNR, especially at low bit rates, as well as a noticeable increase in compression artifacts.

Compared to the classical approach of first demosaicing the pixel data and then applying an ordinary image encoder, our proposed approach of first applying the 4-color spectral-spatial transform and then applying an existing encoder brings the flexibility of allowing the final demosaicing (and other processing operations) to be done off-line, in a cloud service or on a user's computer. That allows for the use of a much higher quality demosaicer, which can improve PSNR by more than the ~0.2–0.8 dB loss from the proposed approach, as well as improve visual quality, with a potential for full elimination of "zipper artifacts" [18]. Camera chip design is also simplified by the removal of the need for a good quality demosaicer and the reduction of the amount of data to be processed by the in-camera encoder.

5. Conclusion

The proposed spectral-spatial pre-processing of CFA images can enable a compression-first processing workflow with progressive-to-lossless capability. With lossless compression, the proposed approach provides better compression that the other tested designs. With "medium" to "high" quality lossy compression (e.g. at around 2 or more bits per pixel), essentially no fidelity loss is observed relative to demosaicing prior to compression, thus enabling the use of high-quality compressed "semi-raw" pictures with about half the file size of today's camera-raw formats (which are typically around 6 bits per pixel). Even at higher compression rates (e.g. 1 bit/pixel), the proposed scheme can be coupled with more sophisticated off-camera post-processing to provide a workflow with improved final fidelity relative to the current practice of on-camera demosaicing. The measured compression performance when using the proposed approach is essentially equivalent when using either JPEG 2000 or JPEG XR for the core image coding technology. As JPEG XR is simpler to implement, it may be the preferable design choice.

References

[1] B. E. Bayer, *Color Imaging Array,* U.S. Patent No. 3,971,065, July 20, 1976.

[2] Y. T. Tsai, "Color image compression for single-chip cameras", *IEEE Trans. Electron Devices*, vol. 38, pp. 1226–1232, May 1991.

[3] S.-Y. Lee and A. Ortega, "A novel approach of image compression in digital cameras with a Bayer color filter array", *Proc. IEEE Intl. Conf. Image Proc.*, vol. 3, pp. 482–485, Oct. 2001.

[4] C. C. Koh, J. Mukherjee, and S. K. Mitra, "New efficient methods of image compression in digital cameras with color filter array", *IEEE Trans. Consumer Electronics,* vol. 49, pp. 1448–1456, Nov. 2003.

[5] N. Zhang and X. Wu, "Lossless compression of color mosaic images", *IEEE Trans. Image Proc.*, vol. 16, pp. 1379–1388, June 2006.

[6] J. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. Signal Proc.*, vol. 41, pp. 3445–3462, Dec. 1993.

[7] D. S. Taubman and M. W. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice*. Boston, MA: Kluwer Academic Publishers, 2002.

[8] P. Schelkens, A. Skodras, and T. Ebrahimi (eds.), *The JPEG 2000 Suite*, Chichester, UK: Wiley, 2009.

[9] S. Srinivasan, C. Tu, S. L. Regunathan, and G. J. Sullivan, "HD Photo: A New Image Coding Technology for Digital Photography", *Proc. SPIE Appl. Digital Image Proc. XXX*, vol. 6696, paper 6696-90, sequence number 6696 0A, pp. 1–19, Aug. 2007.

[10] F. Dufaux, G. J. Sullivan, and T. Ebrahimi, "The JPEG XR Image Coding Standard", *IEEE Signal Proc. Magazine*, Standards in a Nutshell series, vol. 26, pp. 195–199, Nov. 2009.

[11] ITU-T and ISO/IEC, JPEG XR image coding system – *Part 1: System architecture*, ITU-T Rec. T.Sup2 and ISO/IEC TR 29199-1, 2010.

[12] ITU-T and ISO/IEC, *JPEG XR image coding system – Part 2: Image coding specification*, ITU-T Rec. T.832 and ISO/IEC TR 29199-2, 2011.

[13] Wikipedia: "Raw image format," available at http://en.wikipedia.org/wiki/Raw_image_format#Standardization.

[14] Microsoft, "HD Photo Device Porting Kit," v. 1.0, available at http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=5863, Nov. 2006.

[15] H. S. Malvar and G. J. Sullivan, "YCoCg-R: A Color Space with RGB Reversibility and Low Dynamic Range," ITU-T/ISO/IEC JVT Document JVT-I014, available at http://ftp3.itu.ch/av-arch/jvt-site/2003_09_SanDiego/JVT-I014r3.doc, Sept. 2003.

[16] H. S. Malvar. System and method for encoding mosaiced image data employing a reversible color transform, U.S. Patent No. 7,480,417, appl. filed Oct. 2004, granted Jan. 2009.

[17] "Kodak High Sensitivity Image Sensor Tech," *Digital Photography Review,* available at http://www.dpreview.com/news/2007/6/14/kodakhighsens, June 2007.

[18] H. S. Malvar, L. He, and R. Cutler, "High-quality linear interpolation for demosaicing of Bayer-patterned color images," Proc. IEEE Intl. Conf. Acoustics, Speech, and Signal Processing (ICASSP), vol. III, pp. 485–488, May 2004.