

IMAGING BEYOND PIXELS:

enhancing the elementary picture

WITH PER-SENSOR PIXEL GROWTH SLOWING, HOW ELSE CAN DIGITAL-STILL-CAMERA MANUFACTURERS DIFFERENTIATE FROM EACH OTHER? HOW DOES THIS DIFFERENTIATION BRING THEM INTO GREATER COMPETITION WITH VIDEOCAMERAS? TO WHAT DEGREE WILL CAMERA PHONES SUPPLANT THEM BOTH? AND WHERE WILL SENSOR SUPPLIERS TURN TO CONTINUE FILLING THEIR FABRS?

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

A 2004 *EDN* cover story made the then-some-what-controversial claim that image sensors' pixel-count growth rate would soon slow and that sensor implementers would subsequently need to distinguish their system designs using other measures of image quality and capability (**Reference 1**). How did the prognostication pan out? Here's one case study: In mid-2004 when I was researching the article, 8M-pixel mainstream DSLRs (digital single-lens-reflex cameras) were ramping into production; 6M-pixel predecessors had appeared approximately one year earlier. Yet, it took roughly 2.5 years for 10M-pixel

models to subsequently emerge. A number of recent analyst reports also concur with *EDN's* 2004 forecast, as did every company I contacted in researching this article.

EDN Executive Editor Ron Wilson focused one of his 2007 CES (Consumer Electronics Show) online reports on imaging, and, in it, he commented, "Today, it is arguable that the key differen-

tiator in retail camera sales is still pixel count, even though this number has become about as relevant as, say, clock frequency in PC sales" (**Reference 2**). Wilson's analogy between the camera and PC businesses is apt for several reasons. AMD and Intel encountered well-documented leakage-current issues at the 90-nm-process step, which hindered both companies' abilities to increment

their CPUs' clock speeds at historical rates. Also, consumers figured out that they no longer needed to upgrade their PCs to newer models at historic rates; their current systems, perhaps with a less costly hard-disk drive or memory midlife transplant, now delivered sufficient performance for several generations' worth of software upgrades and proliferations.

A decreased influx of consumer-tempting computer supply, coupled with decreased consumer demand for the supply that existed, has translated to the formation of a buyers' market. Nowadays, you can purchase a robust desktop PC for less than \$500 and a full-featured laptop for only slightly more money—if any. Recently published statistics suggest a similar saturation of the DSC (digital-still-camera) market. Research company IDC, for example, reports that fourth-quarter-2006 DSC shipments dropped for the first time since the company began monitoring them; at 12.1 million units, they were 3% below the year-be-

fore figure (Reference 3). And, for all of 2006, IDC estimates, the US market grew by only 5%, notably lower than 2005's 21% growth, as well as below the company's previous prediction of 8% growth for 2006 (Reference 4).

IDC also reported that only approximately 15% of cameras sold in 2006 went to first-time DSC buyers, with the bulk of sales replacing or supplementing DSC owners' existing models. IDC expects zero growth in DSC shipments for 2007, with a decline beginning in 2008. Consumers are increasingly realizing that beyond 2 million to 3 million pixels of resolution, additional captured-image detail is largely wasted when printing out 4×6-in. pictures, even if they come from cropped versions of the original photos. Interestingly, Geoff Ballew, director of product marketing for handheld graphics processors at Nvidia, confirms this fact when he states that 2 million pixels are "great for 4×6-in. prints." This "ceiling" on resolution's meaningfulness is particularly evident when the camera has aggressively JPEG-compressed the source image. And, when consumers want to create a larger print, they may have also found that high-quality pixel interpolation within their computer's image-processing software often suffices. Camera manufacturers have predictably responded to waning demand by

cutting prices. In researching this article, I collected numerous eyebrow-raising DSC prices, such as free-after-rebate, 3M-pixel cameras and 4M-, 5M-, 6M-, and 8M-pixel cameras for \$50, \$55, \$91, and \$100, respectively.

Granted, not all of these cameras came from name-brand manufacturers. And they offered elementary features: limited-to-nonexistent zoom ranges, for example, along with sketchy low-light performance. However, in the past, the specified resolution may have depended upon postcapture in-camera interpolation. In all of the current research's cases, on the other hand, the promoted specifications referenced the native resolution of the camera's integrated image sensor. And collapsing prices aren't restricted to point-and-shoot models; for example, the days of spending several thousand dollars for a low-end DSLR body are over. Instead, according to my researched pricing collection, an entry-level, 6M-pixel DSLR with an 18- to 55-mm zoom lens costs \$400, an 8M-pixel DSLR with *two* lenses costs \$700 (\$500 with one lens), and 10M-pixel DSLR bodies sell for less than \$800.

RETHINKING OPTICS

In assessing what features besides pixel count DSC manufacturers and their

AT A GLANCE

- ❑ A slowing rate of increase in image-sensor resolution clobbers camera prices and forces manufacturers to diversify their products' features.
- ❑ Revolutionary lens technologies decrease or eliminate the power consumption necessary for zoom and focus functions.
- ❑ Telescope-derivative optics slim the transparent material stack necessary to implement a given focal length or range.
- ❑ You've got no shortage of processing options at your disposal: generic DSP or imaging-tuned logic block, baseband or application processor, and dedicated IC or image sensor-enclosed. Selecting the correct candidate for your particular situation is a key challenge.
- ❑ Low-light sensors, new file formats, antishake techniques, and additional connectivity options are some of the factors that will define future cameras' feature proliferations.

videocamera, camera-phone, and other competitors can harness to recapture consumers' wallets, begin by looking at the lens, within which an image's photons undergo their initial transformations. I'll concentrate on two key optics attributes: focus and focal length. Fixed-focus cameras, which some marketers alternately call focus-free, have limited usefulness. They don't, for example, allow for arms-length self-portraits, and you can, therefore, find them only in the low-end market segments. And manual focus is a hassle, infeasible for users with poor eyesight, and incompatible with fast-moving subjects. These drawbacks are true even with motor-assisted—and

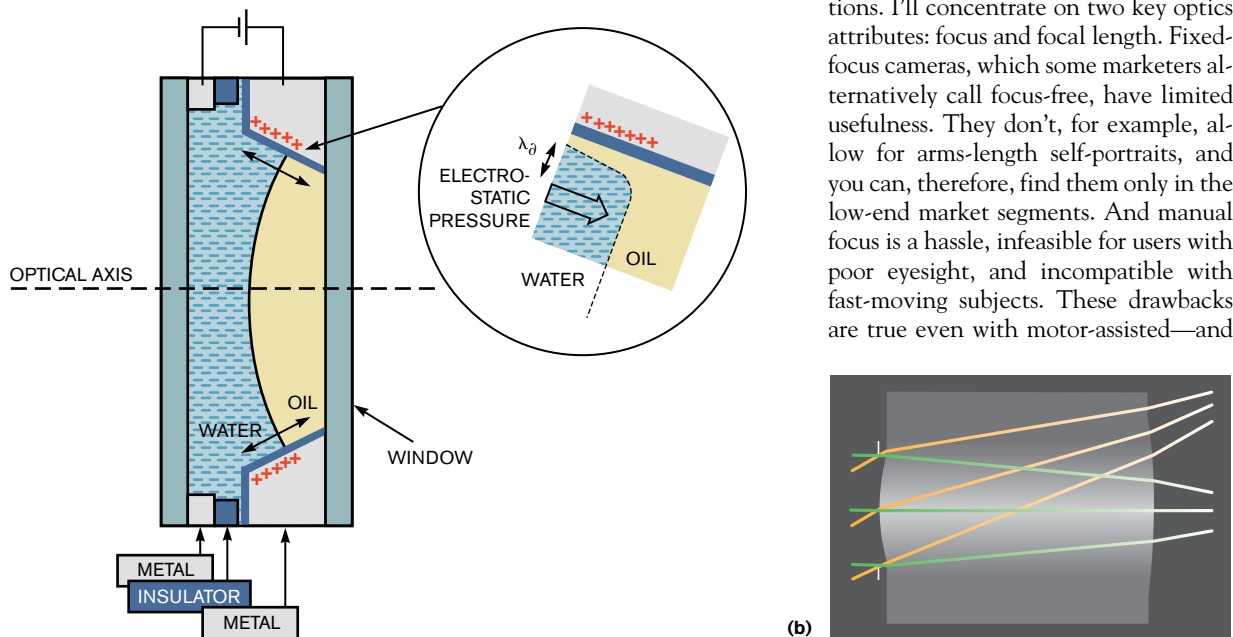


Figure 1 So-called liquid lenses implement focus and zoom functions with less power consumption than traditional motor-driven optics (a, courtesy Varioptic). Wavefront-coded lenses provide a compelling, albeit computationally intensive, no-focus path to extended depth of field (b, courtesy CDM Optics).

thus incremental power-consuming—manual focus.

Yet, consumers are equally intolerant of autofocus cameras that select the incorrect focus plane in an image or that take too long to reach that point. “Stopping down” the lens to a smaller aperture—that is, a larger aperture value—and thereby increasing the depth of focus improve the likelihood of a sharp result. However, this technique also slows the shutter speed to the point that a user would be unable to hold a camera and still capture a sharp image in low-light settings. Using this technique would thereby require the use of additionally battery-draining, artificial-looking, and often insufficiently covering flash illumination.

Deep focus may also be incompatible with a user’s desire to intentionally blur the image foreground and background to draw viewers’ attention to the primary subject. So, instead or in addition, many manufacturers are now throwing more processing horsepower at the problem. Modern cameras, particularly when users switch them to “portrait” modes, autodetect faces in the to-be-captured field of view and set both autofocus and auto-exposure to optimize for them.

One emerging solution to the battery-sapping problem of motor-assisted focus, along with, for that matter, zoom, is liquid-lens technology. Companies, such as

Varioptic, and research organizations, such as the IMRE (Institute of Materials Research and Engineering), are promoting this technology (see **sidebar** “Online addenda enhance the picture”). These lenses operate in a manner akin to that of the human eye, which is surrounded by muscles that subtly change its shape to alter its optics characteristics. Quoting from Varioptics’ technology overview: “The liquid lenses that we develop are based on the electrowetting phenomenon ... a water drop is deposited on a substrate made of metal, covered by a thin insulating layer. The voltage applied to the substrate modifies the contact angle of the liquid drop. The liquid lens uses two isodensity liquids; one is an insulator, while the other is a conductor. The variation of voltage leads to a change of curvature of the liquid-liquid interface, which in turn leads to a change of the focal length of the lens” (**Reference 5** and **Figure 1**).

Alternatively, if deep focus is acceptable in your application, perhaps you’d be interested in a five- to 10-times increase in it without any requisite decrease in light transmission, such as that encountered when you stop down a lens aperture. In such a case, check out the wavefront-coding technology that OmniVision obtained when it acquired CDM Optics two years ago and now markets under the TrueFocus moniker.

“Wavefront coding ... offers systemwide optimization, whereby specialized optics, sensors, and signal processing work closely together to provide high-quality, low-cost imaging in spite of small package requirements,” claims CDM Optics’ co-founder and president, Edward Dowski, PhD. By shifting a higher percentage of the overall image-capture burden away from the optics and toward the image processor, you can also use wavefront-coding techniques to compensate for the dubious quality of low-cost plastic lenses as well as to counteract temperature and unit-to-unit manufacturing variability.

TrueFocus is, perhaps not surprisingly, a bundled package, integrating OmniVision-defined optics and OmniVision-designed image sensors, which have on-chip image processors. However, Product Manager Michael Hepp admits that other sensor and processor suppliers are developing conceptually similar approaches. Reflecting this fact, Ping Wah Wong, PhD, Nethra Imaging’s vice president and chief imaging scientist, notes, “Digital-focus technology that uses a form of predistortion with a specially designed lens and subsequently uses image processing to undo the distortion has shown to be promising because it can eliminate the autofocus mechanism in the lens.” Echoing Wong’s comments, Clay Dunsmore, chief technology officer for Texas Instruments’ Digital Camera Solutions group, notes, “Lenses can be made more compact by transferring quality requirements from the optical domain to the digital-image-processing domain. Examples include corner illumination and geometric distortion.”

Traditionally, two primary focal-length factors have contributed to optics bulk: a wide zoom range, versus a fixed focal length, and a long telephoto- or ultrawide-angle endpoint to that range. Yet, consumers wishing to pocket only a single camera desire a unit with as versatile as possible an optics subsystem, and, speaking of “pocket,” they’d also like the resultant camera to be neither too heavy nor too large. OmniVision’s Hepp notes that Motorola’s ultrathin Razr phone, a “once-in-a-lifetime” success story, has heavily influenced users’ expectations for all portable-electronics devices, especially in the United States. Unfortunately, supplemental wide-an-

ONLINE ADDENDA ENHANCE THE PICTURE

Invariably, some of the interesting information that I uncover during my research ends up on the cutting-room floor, at least from a print perspective, due to page-count limitations and other factors. Fortunately, though, the *EDN* Web site provides another publishing outlet, and, this time around, it’s particularly packed with material. “Form-factor transformations” predicts to what degree camera phones and still/videocamera hybrids will erode traditional DSC (digital-still-camera) dominance in the future, and “Application expansion” covers other uses for sensor suppliers’ wares. This article discusses some of the imaging breakthroughs that have emerged from academic- and industry-research projects; “Future forecasts” expands on the list.

More imaging news will come from mid-February’s 3GSM (Third Generation Groupe Spéciale Mobile) World Congress and this month’s PMA (Photo Marketing Association) shows; “Conference updates” will keep you abreast of all the breaking developments. “Education next steps” provides suggestions for further cultivating your imaging knowledge. And “Where in the world is Foveon?” documents my thoughts on this once-promising image-sensor start-up and shares a review of a book about the company. Visit the Brian’s Brain blog at www.edn.com/briansbrain for these and other relevant imaging postings.

gle and telephoto lenses that, like a filter, screw onto the front of the primary optic haven't been largely successful for a variety of reasons, including inconvenience, added cost, and degraded quality. And the so-called digital-zoom—that is, pixel-interpolation feature—is thankfully fading from prominence; although it's arguably better than no zoom at all, its quality results are subpar, especially in the fast-capture-expectation and, therefore, limited-processing environment of a still or videocamera.

A tall order? No doubt. But several techniques can help address the seemingly divergent customer requests. Kodak's EasyShare V705 tackles the optics problem with two lenses and two 7.1 million-pixel sensors—a 23-mm, 35-mm-equivalent ultrawide-angle fixed lens and a 39- to 117-mm zoom lens (Figure 2). For another example of the range-versus-reach trade-off, look at Kodak's P712 and P880 Performance Series cameras, which have similar form fac-

tors. The P712 includes a 12× optical zoom, in a 36- to 432-mm focal-length range. Conversely, whereas the P880 offers “only” a 5× optical zoom, its reach extends from 140 mm all the way down to an ultrawide, 24-mm focal length.

Taking optics developments to a more revolutionary level, researchers at the University of California—San Diego recently unveiled an “origami” lens, which they based on reflective-mirror techniques that astronomical telescopes use but built on a single 5-mm disc of calcium fluoride. It promises to reduce the camera thickness required to implement a given focal length and focus range. “Traditional camera lenses are typically made up of many different lens elements that work together to provide a sharp, high-quality image. Here, we did much the same thing, but the elements are folded on top of one another to reduce the thickness of the optic,” says Eric Tremblay, an electrical- and computer-engineering doctoral candidate

at UCSD's Jacobs School of Engineering, in the university's newsletter (Reference 6). The initial prototype eight-fold imager, which implements a 38-mm focal length and focuses on objects 2.5m away, was at least one-seventh the thickness of a conventional multielement lens cluster.

PROCESSING PIXELS

I asked Stuart Boyd, Analog Devices' product-line director for the company's high-speed-signal-processing group to provide an overview of the imaging businesses in which Analog Devices participates. “We're seeing manufacturers place more effort into helping their customers take better pictures. This [effort] is showing up in ... face-detection-assisted autofocus and exposure, in-camera image correction at the push of a button, assistance in stitching shots together for panorama shots—stuff that used to be possible only in software that few of us bothered to deal with. What's happening inside the camera ‘engine’ is getting a lot more complex to manage and ensure [that] performance is fast enough to give the immediate results we all expect.” Speaking of consumers who inappropriately judge their photographs against professionals' work in magazines, EDN's Wilson echoes Boyd's observations: “This puts pressure on the camera architects to cram as much image quality into the acquisition process and as much postprocessing capability into the platform, as possible,” he wrote in his online CES report.

Let's be clear: This problem is a *good* one for a processor vendor to have, and




Figure 2 Kodak's EasyShare V705 harnesses dual lenses and dual image sensors to sveltely extend its optical range (a), and the University of California—San Diego's “origami” lens harnesses mirror-telescope-derived technology to the same end (b), versus a bulkier traditional multielement lens stack (c).

it shows no sign of disappearing any time soon. But which processor vendor and which architecture are optimal for your next system design? This problem has no easy answer; when researching alternatives, you'll face a spectrum of offerings—from hard-wired to fully software-programmable. NuCore Technology's processors' imaging-centric pipelines inhabit the hardware region of the spectrum, with Texas Instruments' DSC processors' combination of a generic DSP core and imaging-centric peripherals on the opposite, software-centric side. (Note that both NuCore's and TI's products embed an ARM core as the overall system processor.) Reflecting TI's image-processing approach, Kanika Ferrell, marketing manager for the digital-camera group, comments, "Because multiple functions are being run on a single core, it requires that the image processor in the system have enough horsepower to run multiple functions simultaneously. ... To keep power consumption low, manufacturers move to more advanced process-technology nodes and run the processor at higher speeds while using lower power to do so."

Nvidia's Ballew, on the other hand, suggests some image-processing areas in which dedicated-hardware support makes particularly good sense:

- **"A very fast datapath into the GPU [graphics-processing unit].** Benefit: Very fast click to capture; the sensor module can be set in full resolution for preview so there is no delay resetting the sensor from a low-resolution preview to full-resolution capture.

- **"Real-time-JPEG encode.** Benefit: Rapid multishot; the image can be compressed as fast as the sensor sends the data, so the user can capture several frames in a row with a single click of the shutter release. This [feature] helps to capture action shots or those hard-to-catch photos. Real-time JPEG also reduces memory requirements, so it helps OEMs make affordable phones. ... And full hardware JPEG encode and decode reduce the power required to compress and decompress images.

- **"ISP [image-signal processor].** Benefit: The ISP is key for autofocus, auto-white balance, and auto exposure and the variety of image features, such as sepia, black and white, antique, red-eye reduction, and edge sharpening. It also

ALL OF THE VARIOUS IMAGE-PROCESSING APPROACHES HAVE STRENGTHS AND SHORTCOMINGS.

gives the phone OEM more flexibility to choose camera modules with or without an ISP. The newest sensors typically come to market without an ISP integrated. Putting this function into the graphics processor saves cost and board space compared to an external ISP. This [feature] also reduces the power required compared to software-based ISPs."

On that last point, given that one of the long-touted advantages of a migration from CCDs (charge-coupled devices) to CMOS sensors is the ability to include image-processing logic alongside the sensor array, you might predict that the CMOS-sensor suppliers wouldn't necessarily agree with Nvidia's partitioning stance, and you'd be right. The latest generation of ISP-inclusive sensors even incorporates support for the JPEG-encoding function. Andrew Burt, vice president of Toshiba's ASSP (application-specific-standard-product) business unit's imaging- and communications-marketing group, comments, "There is continuing debate about whether the ISP should be integrated on the CMOS-image-sensor die or on-board the baseband processor. Many handset manufacturers tell us that a CMOS-image-sensor SOC [system on chip] with an embedded highly optimized ISP is preferable. It offers system benefits that leverage the CMOS-image-sensor designer's in-depth knowledge. This [benefit] becomes key as resolutions approach 5M pixels and beyond. However, even at 2M pixels, an SOC approach can provide end users with a better visual experience." In reality, all of the various image-processing approaches have strengths and shortcomings. To guide your selection, assess factors such as cost, performance, degree of integration, power consumption, flexibility versus the flexibility requirements of your application, and development-tool maturity and robustness.

One increasingly key camera feature

that all the companies I interviewed for this article mention is the need to deliver increasingly high-quality images at increasingly low-ambient-light levels. Said a different way, they point out the need for the camera to operate at ever-higher shutter speeds for a given subject illumination strength, thereby obviating the need for battery-draining flash operation and acting as a first-pass image-stabilization scheme. This consumer demand flies in the face of the fact that, as sensors' pixel pitch decreases, each pixel's photodiode captures fewer photons in a given amount of time, thereby inherently decreasing the photodiode's light sensitivity.

It is essential to accomplish this high ISO (International Organization for Standardization)-processing trick with low image noise, because noise directly impacts compression efficiency, and the smaller the JPEG file, the more pictures a consumer can take before filling up a flash-memory card or another storage device. (ISO 5800:1987 describes photographic film's or a sensor's sensitivity to light, also known as its "speed" and often also called its ISO number.) Nvidia, for example, has branded its proprietary JPEG scheme with the Fotopak marketing moniker. And NuCore's two-chip CleanCapture approach, which works in conjunction with CCDs, does as much image processing as possible in the analog domain with its NDX chips before passing on the data to a digital-domain SiP companion processor. (CMOS-sensor-based designs don't use the company's front-end analog processor.) Evolutionary sensor improvements can help to some degree; for example, Toshiba's Burt notes that the company is "working on new conformal-microlens technology that provides a gapless microlens with increased light-gathering efficiency instead of today's circular microlenses that can't cover the entire pixel area of the image-sensor array."

Perhaps, though, the industry needs a more fundamental breakthrough in image-sensor design to fully address consumers' low-light-with-high-quality contradictory expectations. At both this year's and last year's CES, Planet82, a small company and partner of the Nano Scale Quantum Devices Research Center of the Korea Electronics Institute, showcased the ability to capture



+ Go to www.edn.com/070315df to find this article's associated references and vendor box. At this URL, you can also post a comment on this article.

+ At the Brian's Brain blog, www.edn.com/briansbrain, you can learn more about the topics discussed in this article, provide feedback, ask questions, and peruse and respond to the postings from other EDN readers.

discernible, albeit fuzzy, images—both still and low-frame-rate video and both black-and-white and color—of objects at *extremely* low ambient-light levels. I attended a demonstration at this year's show and was impressed by the company's achievement. Granted, the sensor was low-resolution—that is, VGA—although the company expects a 2M-pixel sensor to be out of fab by the time you read this. Also, the technology's cost, manufacturing yield, stability through operating lifetime and extremes of voltage and temperature are all currently unknown.

“Planet82's new VGA-color SMPD [single-carrier modulation-photo detector] is the ... first, full-color, high-sensitive imaging chip for taking pictures or video in the dark without a flash,” says Hoon Kim, PhD, chief technology officer for Planet82. “SMPD combines the clear-image quality, high sensitivity, and wide dynamic range of existing imaging technology with powerful nanotechnology, making it 2000 times more light-sensitive and 50% smaller than traditional CMOS and CCD sensors. ... Unlike photodiode-based CMOS and CCD technologies, which require millions of photoelectrons per pixel unit to create an image, the SMPD is able to react to tiny amounts of photons in light levels less than 1 lux, the equivalent of the light from one candle a meter away.”

Planet82 remains tightlipped about the specifics of its nanotech accomplishment, but, according to a report by MicroDesign Resources' Max Baron in April 2006, “Planet82 seems to base its development of nanotechnology on SMPD pixels that can deliver high amplification (high electron yield) by creating a semiconductor 3-D confine-

ment smaller ... than the de Broglie wavelength. [The de Broglie hypothesis states that all matter has a wavelike nature, **Reference 7.**] The mechanism of translation of light energy into charge for available CCD and CMOS devices is based on the conversion (subject to efficiency) of one photon to one electron. The SMPD uses the same mechanism, but the electron created is amplified by quantum effects, generating several thousands of electrons. CCD and CMOS devices use a PN-junction photodiode to provide the photoelectrical transformation. The detection mechanism employed by the SMPD is fundamentally different: the SMPD implements artificially made potential barriers in every electrical energy band of interest. A few injected photons will find it easy to lower the barrier, allowing the structure to generate a significant electrical charge” (**Reference 8**).

FEATURE INFLATION

Sometimes, the two-step scheme of simply cranking up the signal gain coming off an image sensor and then processing out the consequent noise is insufficient for stabilizing images. For example, Nethra's Wong points out, “Camera shake can happen easily in cell-phone cameras because of the small form factor and because the camera is usually operated with one hand.” In such cases, more elaborate image processing is necessary. Electronic stabilization incorporates an oversized image sensor. The processor identifies the pixel locations of high-contrast images, and, if they move from one frame to another, it compensates by proportionally shifting the captured image—as long as the entire scene still fits within the sensor's boundaries. Even more elaborate image-stabilization schemes involve electromechanical techniques that shift either the image sensor or the lens elements in response to an accelerometer-sensed jostling of the camera.

It might be easy to conclude that smaller pixel pitch is overwhelmingly a bad thing for imaging. That supposition would be premature, however. Applications that beg for ultrahigh resolution do exist. Ask Hasselblad, for example, which just unveiled the \$24,995, 31M-pixel H3D-31 camera, which represents a step *backward* from last year's 39M-

YOU CAN HARNESS THE MULTIPLE-SENSOR-PIXEL-PER-IMAGE TECHNIQUE TO EXPAND THE DYNAMIC RANGE OF THE RESULTING IMAGE.

pixel model. Or ask Canon, which, according to rumors, is developing a 22M-pixel DSLR. Alternatively, think of what else you could do with all those pixels, besides devoting them to resultant image pixels on a 1-to-1 basis.

Most modern cameras, with the exception of prism-based, three-sensor configurations or those derived from Foveon's three-photodetector-per-pixel-sensor approach, embed sensors that include a multicolor-filter array ahead of the photodetector array. Common filter matrices include the Bayer pattern, which contains twice as many green filters as either red or blue; the Sony-developed RGBE (red/green/blue/emerald) Bayer enhancement, which adds an emerald-green filter; and the CYGM (cyan/yellow/green/magenta) filter array. Regardless of the raw image's filtered state, postcapture interpolation derives red, green, and blue values for each pixel. Instead, though, why not harness burgeoning image-sensor-native resolutions to dedicate red-, green-, and blue-filtered sensor subpixels to each corresponding image pixel?

You can also harness the multiple-sensor-pixel-per-image technique to expand the dynamic range of the resulting image. Today, photographers who desire to generate HDR (high-dynamic-range) pictures take underexposed, overexposed, and correctly exposed shots of the same scene and then combine them in a computer using Adobe Photoshop or a similar program. This approach is not only incompatible with moving subjects, but also cumbersome and time-consuming. Instead, by placing a variable neutral-density filter array ahead of the sensor or by designing sensor subpixels with varying photon-integration characteristics, you can accomplish a similar result in the camera and with only a single shot.

Unfortunately, the more-than-10-

year-old JPEG format, which still dominates in still-image-capture applications, has outmoded characteristics that hinder imaging advancements, such as HDR. It supports only 24-bit-per-pixel, or 8-bit-per-color-per-pixel, maximum dynamic range, and the more flexible JPEG 2000 approach hasn't achieved widespread adoption. Microsoft's Windows Vista OS and the company's latest generation imaging applications robustly support its Windows Media Video-derived HD Photo format, and, with it, the company hopes to break the JPEG bottleneck on imaging innovation. Key HD Photo attributes, according to Principal Product Manager Bill Crow, include 2× compression efficiency versus JPEG for typical photographic content; a lossless-compression option, which typically provides 2.5× compression; significantly higher color fidelity at any compression ratio, which the company primarily accomplished through the use of 4:4:4 or 4:2:2 chroma sampling; fewer additive artifacts with multiple recompressions; support for many more pixel formats, including gray scale, CMYK, and N-channel; and support for 8 to 32 bits per color, integer and fixed- and floating-point formats, HDR, and wide-gamut images.

"HD Photo uses a biorthogonal lapped transform, combined with advanced entropy coding, a high-performance, reversible color-space transform, and numerous other innovations," says Crow. "The core transform is functionally equivalent to the DCT [discrete-cosine transform] in JPEG. HD Photo's algorithm is incrementally more complex to support the innovations that significantly improve compression efficiency, but this additional processing does not require complex operations or instructions. All processing involves simple integer math, which can be easily accelerated with parallelization, pipelining, [or both]. HD Photo encoding and decoding is based on macro blocks, enabling progressive image encoding and decoding using a minimal memory footprint. For a typical camera-encoding application, the target memory-buffer requirement is 18 pixel rows.

"While HD Photo's compression algorithm is incrementally more complex than JPEG, the improved compression efficiency results in smaller compressed

bit streams, reducing the total processing required. Most important, this [approach] reduces the amount of variable-bit-length processing, which is the portion of the overall algorithm that is least able to be accelerated." A DPK (device-porting kit) is currently available in a 100%-royalty-free fashion for building support for HD Photo into an imaging application. It includes a bit-

stream specification, reference ANSI C source code for the encoder and decoder, and sample applications, and it supports both little- and big-endian processor architectures. Microsoft also plans to deliver a Quicktime codec for Apple's OS X.

According to TI's Ferrell, "Camera features are beginning to normalize. We aren't seeing many 'new' features on

cameras but rather improvements in the existing features, such as improved response time and faster processing of the advanced features already integrated into the system." Although Ferrell's perspective may be accurate in the broad sense, plenty of examples exist of cameras that make more radical feature transformations in attempting to tempt both DSC owners to upgrade and those on the fence to take the first-time DSC plunge.

CES is a prime opportunity for vendors to unveil their latest image-capture ideas. New products include Go-Pro's Digital Hero 3 for extreme-sports fanatics, which straps to a user's wrist and captures either bursts of 3M-pixel images or as much as 54 minutes of 30-frame/sec video with sound. The 4.5-oz unit is waterproof to 100 feet and has a manufacturer's suggested retail price of \$139.99. Kodak unveiled its 4M-pixel EasyShare-One at the 2005 CES. It contains a Wi-Fi transceiver that enables you to wirelessly e-mail, print, and upload photos to your PC, as well as to directly access Kodak's online-storage service. Cameras from numerous suppliers now contain Bluetooth connectivity for printing and photo transfer, supplementing traditional wired-USB links. UWB (ultrawideband)-silicon vendors are eyeing this same application as a future opportunity. Another unit, Samsung's 7M-pixel VLUU i70, includes HSDPA (high-speed-downlink-packet-access) cellular-data connectivity, along with text-messaging functions. Sanyo's Xacti HD2 hybrid high-definition still/videocamera embeds an HDMI (high-definition-multimedia-interface) transmitter for directly tethering a display. Ricoh's 8M-pixel 500SE, like several other cameras now available, integrates a GPS (global-positioning-system) module that enables users to record the location where they took a picture. Some camera phones mimic this capability, albeit with a lower degree of accuracy, through cellular-tower triangulation. **EDN**

You can reach
Senior Technical Editor
Brian Dipert
at 1-916-760-0159,
bdipert@edn.com,
and www.bdipert.com.

