coverstory *By Brian Dipert, Technical Editor*

Seeking clarity

Image sensors peer into a blurry future

ARE SENSORS' DAYS OF PREDICTABLE RESOLUTION GROWTH AND SIZE REDUCTIONS NEARING AN END? IF SO, WHAT OTHER INNOVATION OPPORTUNITIES EXIST, AND WHICH TECHNOLOGY WILL WIN THE CCD-VERSUS-CMOS TUG OF WAR?

YET ANOTHER OF THE MANY going-out-of-business local film-developing labs, one just down the street from my house, recently locked its doors for the last time. I can no longer find medium-format film locally; I need to mail-order it (once I use up the dozens of rolls currently residing in my freezer, bought last fall from a photographer who sold his film gear and went digital). I can't find 35-mm professional film locally, either, but I don't even try anymore; ever since last summer when I bought a DSLR (digital single-lens-reflex) camera, I've had no urge to shoot a single frame of silver halide.

Conventional photography's loss is commensurate with digital imaging's gain. Just yesterday, I received notification in my e-mail inbox of a Canon EOS Digital Rebel entry-level DSLR, which, courtesy of a firmware hack, you can turn into a Canon mid-level DSLR, for less than $850, including an 18- to 55-mm EF-S zoom lens. Full-featured 3 million-pixel point-and-shoot digital cameras sell for $200, entry-level digital cameras cost less than $100, and you can buy a single-use digital camera with an LCD from Ritz Camera for slightly more than $10—and hack it, too (**Reference 1**). CIF- and VGA-resolution camera phones are widely available. Samsung recently announced a variant with 3 million-pixel resolution, and 5 million-pixel versions are due out by year-end. And digital camcorders are now offering high-quality, multimillion-pixel still-image-capture capability.

After many years' worth of premature pundit predictions of a conventional-to-digital-photography conversion, the crossover point is clearly behind us. Both system and sensor manufacturers are benefiting, along with suppliers of other imaging building blocks, such as DSPs and nonvolatile memory (**Reference 2**). But although blue skies are overhead, storm clouds are beginning to form on the horizon. Numerous sensor manufacturers, including such notable names as Conexant, Freescale, Hynix, Intel, and National Semiconductor, have sold or shuttered their programs. The resolution treadmill is slowing and, without a fundamental technology breakthrough, may soon stop. Multifunction phones, PDAs, and video cameras are threatening to obsolete stand-alone still cameras. And privacy and confidentiality concerns, along with suboptimal implementation details, threaten to stall camera phones' adoption (see **sidebar** "System evolution, revolution, and potential stagnation" with the Web version of this article at edn.com).

OPTIONS AND COMPARISONS

As is so often the case in the semiconductor industry, the fundamentals of a technology rarely transform to a significant—that is, revolutionary—degree, even over nearly a decade, although the technology's nuances are quicker to evolve. *EDN*'s late-1997 coverage of image sensors provides a comprehensive and still largely relevant overview of the technology and details why CMOS sensors of the late 1990s delivered lower quality images than those of CCDs (**Reference 3**). This quality deviation is one thing that *has* changed over the ensuing years. As CMOS sensor technology has matured, benefiting from steadily increasing industry attention along with other factors, the gap between it and the CCD alternative has narrowed across a range of quantifiable quality-measurement criteria: sensitivity and dynamic range—that is, SNR, for example, along with quantum efficiency and pixel uniformity (**Figure 1**). CMOS-sensor suppliers have for years now been claiming this parity, but the proof is in the market. Canon, for example, sells a suite of DSLRs at varying resolutions, prices, and sensor sizes, all containing internally developed CMOS imagers. And an even more telling data point is the fact that Kodak now targets profession photographers with a high-end, 14 million-pixel DSLR that integrates a Cypress Semiconductor—not Kodak—CMOS sensor.

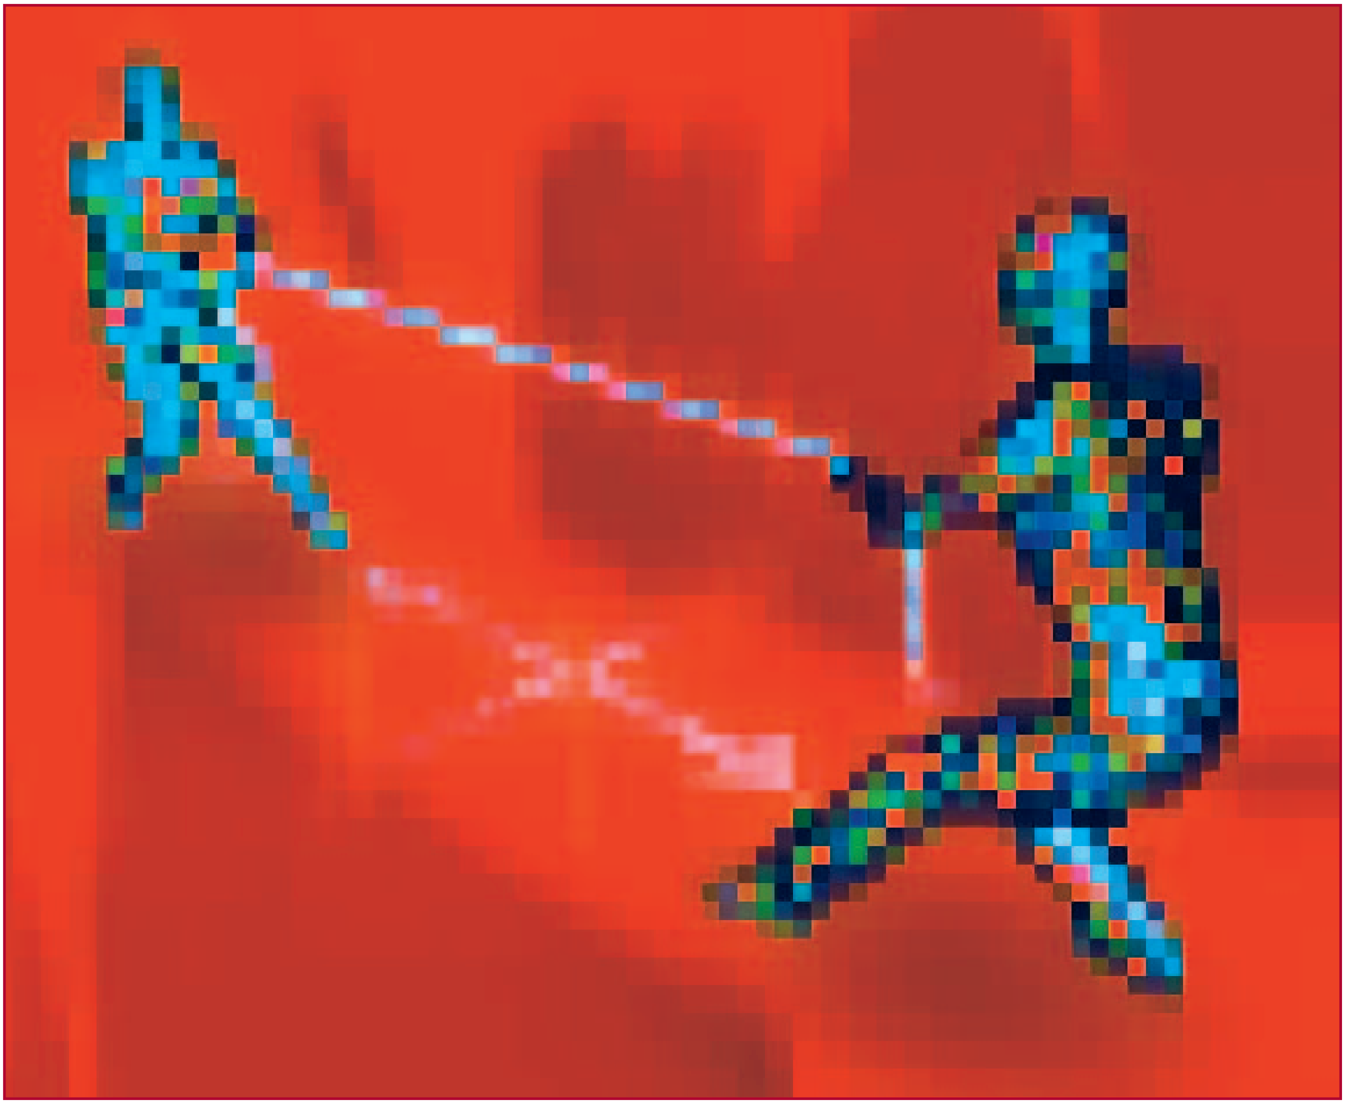

Back in the late 1990s, some analysts predicted that CMOS sensors would quickly obsolete CCDs. This scenario has not happened; coexistence has instead occurred, and the stalemate will likely continue into at least the immediate future (see **Table 1** with the Web version of this article at www.edn.com). The reasons, though, have as much to do with business relationships as they do with fundamental technical factors. The oft-touted cost benefit of CMOS sensors over CCDs has proved to be dubious at best; CMOS sensors, despite what their name implies, do not use standard CMOS processes. For example, the shallow epitaxial-layer-deposition step that conventional CMOS often uses to mitigate latch-up problems adversely blocks the transmission of red-wavelength light to the sensor's photodiodes. And shallow, heavily doped junctions, enabling dense, short-gate, conventional CMOS devices, result in suboptimal green-light response and high dark current in CMOS sensors.

The costly microlens-integration step is common to both CCDs and CMOS sensors, as is the often-required antialiasing "blur" filter ahead of the sensor. And, turning attention to the business side of the equation, it's important to keep in mind that CCDs have been around since the early 1970s, when manufacturers initially considered them for use in semiconductor memories, and today represent a mature, formidable business that companies such as Sharp and Sony would like to see continue to flourish. Aside from the cost-effectiveness that CCDs' maturity enables, Japanese CCD suppliers reportedly wield significant pressure on their sibling camera and phone divisions, along with their system partners at other companies, to employ CCDs instead of CMOS sensors.

Vendors often tout CMOS sensors as delivering longer battery life than CCDs; although this claim is potentially true, the longer life is often the result of higher level systemic factors. A CMOS sensor in and of itself may not have significantly lower power consumption than a CCD, but a CMOS sensor doesn't require the CCD's companion analog processor, handling the complex biasing, clocking, and analog-to-digital-conversion steps. You can, for example, construct a CMOS-sensor-based camera using only NuCore Technology's SiP-1280 digital-image processor; if, however, you choose a CCD, you'll also need to include NuCore's NDX-1260 analog-front-end chip. The comparatively high integration of CMOS sensors also has board-space and bill-of-materials-cost ramifications.

AT A GLANCE

▷ CCD-versus-CMOS-sensor market success increasingly has less to do with the technologies' respective technical characteristics.

▷ Shrinking pixel dimensions result in sensors that are increasingly unusable in real-life image-capture situations.

▷ Resolution one-upmanship isn't the only opportunity a sensor vendor has to differentiate itself.

▷ An assortment of competing color-detection and -interpretation techniques offer varying strengths and shortcomings.

▷ Application diversification is critical to sensors' continued success.

The relative consumption of other high-current camera subsystems complicates the CMOS-versus-CCD power-consumption comparison. If you, for example, put an optical viewfinder in the camera, the sensor will be in use for a small percentage of time. Choose an electronic viewfinder or an LCD, though, and the sensor will more frequently power up. Alternatively, if the camera user spends a lot of time reviewing his photos using an LCD, its power consumption may greatly surpass that of the sensor.

Other factors aside, developers generally consider CMOS sensors as easier to design-in and potentially offering more features than CCDs. CMOS devices also neither require complex clocking to extract pixel information nor employ multiple nonstandard bias voltages. The system can access individual pixel data in a memorylike, random fashion, leading to a simpler realization of subsampling to preview an image on a low-resolution LCD before capturing it, for example; multiposition windowing; and digital-zoom functions. Conversely, CCDs and CMOS sensors are equally adept, generally, at implementing mirror-imaging, 90 and 180° image rotation, and interlaced-versus-progressive-scan output switching. (Videocameras that also support high-quality still-image capture require progressive-scan output switching.)

MOORE'S LAW: NOT IN THE PICTURE

Until recently, when static-power-consumption issues began in earnest to rear their ugly heads, most semiconductor products greatly benefited from the cost-shrinking, speed-boosting, and voltage-lowering side effects of Moore's Law (**Reference 4**). Both CCD and CMOS image sensors, however, play by different rules. Fundamentally, the photodiodes or photocapacitors in image sensors capture and measure electrons that light photon collisions and consequent electron-hole pair creations generate. The smaller the per-pixel light-collection area, the less sensitive to light it becomes. Vendors create a smaller light-collection area by shrinking the sensor to reduce its cost, squeezing more pixels onto a sensor, in-

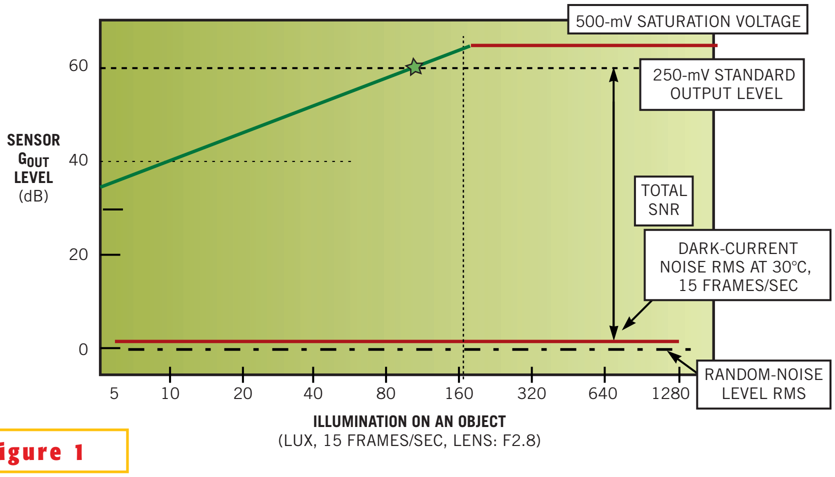

Figure 1

Intense, diverse, and ever-expanding industry attention now enables CMOS sensors to deliver specifications once achievable only by CCDs (courtesy Toshiba).

creasing the amount of overhead on-chip circuitry for each pixel, or using a combination of these techniques. Increasing the overhead circuitry reduces the pixel's so-called fill factor, the percentage of each pixel available for photon collection (see **sidebar** "Deciphering size" with the Web version of this article at edn.com).

Increasingly exotic image processing can to some degree counterbalance this decreased inherent sensitivity. But, to get to the bottom line, just "do the math," as Foveon's vice president of marketing communications, Eric Zarakov, puts it. For example, compare the company's 20.7×13.8-mm F7X3-C9110 sensor in Sigma's SD10 DSLR with the 7.1×5.3-mm FO18-50-F19 in Polaroid's upcoming x530 digital point-and-shoot camera. The F7X3-C9110's pixel dimensions are 9.12×9.12 microns, its fill factor is 67%, and it contains 3.4 million pixel locations with three photodetectors per location. The FO18-50-F19, conversely, contains only 1.5 million pixel locations with three photodetectors per location, but its per-pixel dimensions are 5 microns on a side, and its fill factor is approximately 50%.

This disparity in sizes means that the FO18-50-F19's per-pixel light-gathering area is only about one-quarter that of the F7X3-C9110. Leading-edge modern CCDs have approximately 2.5-micron pixels, and advanced CMOS sensors have approximately 3.2-micron pixel pitches, evidencing an industrywide trend, not a Foveon-specific shrinkage phenomenon. The outcome of the trend is evident to anyone who compares the ISO—ASA (American Standards Association) to you photography old-timers—specifications of cameras at different sensor sizes and, hence, prices and resolution specifications (see **sidebar** "Hands-on analysis" with the Web version of this article at edn.com). And keep in mind that the more complex the image processing, the longer the sustained shot-to-shot delay and the higher the battery drain. (To use the video analogy, the more complex the image processing, the slower the frame rate.) In the SD10, image processing occurs in the computer, because the camera outputs only raw-formatted files, but, in the x530, it occurs in the camera, as is usually the case (see **sidebar** "Interconnect controversy," also with this article's Web version).

Image processing's task would be simpler if it had to handle only the measured signal. Unfortunately, noise also factors into the real-world equation. The photodetector itself is one noteworthy noise source. It measures the amount of per-pixel accumulated electron charge and cannot distinguish between electrons photon collisions create and those that thermal effects, or "dark current," generate. To suppress thermal noise, custom cameras for long-exposure astrophotography and other specialized applications employ refrigerant cooling subsystems to operate at low temperatures. The photodector outputs a signal that the amplifier boosts, causing additional noise. The A/D converter, which digitizes the amplifier's output, is another potential noise source. As CCD- and CMOS-sensor technology matures, vendors are achieving fewer improvements in reducing noise. Boost the signal either in the analog domain using the amplifier or after the ADC stage using digital-domain multiplication, and you also boost the noise. If you couple this fact with ever-decreasing pixel dimensions and consequently lower signal strength, you may conclude that the sensor industry is about to hit a brick wall in pixel and sensor size.

To delay that scenario, sensor manufacturers are selectively employing increasingly exotic microlens structures on each pixel to collect and focus as much light as possible onto the photodetector. Sony, for example, includes a DIL (dual-internal-lens) structure in its latest-generation 8 million-pixel CCD (**Figure 2a**). Deciding whether to implement microlens inclusion and whether to customize them is a delicate balancing act, however. For one thing, the microlenses add significant cost to the sensor. Also, if the microlens is too aggressive in its operation, it will intercept and redirect light rays—particularly those entering the sensor at acute angles—that should instead be going to other photodetectors. This redirection is inherently an obstacle to accurate luminance measurement; including a Bayer filter (named for its inventor, Kodak scientist Bryce Bayer) or another

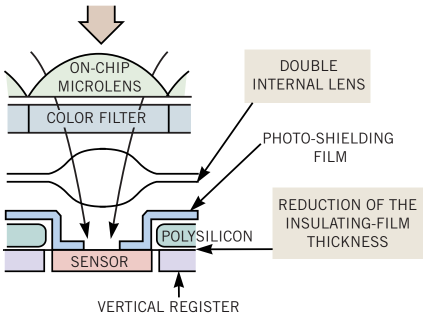

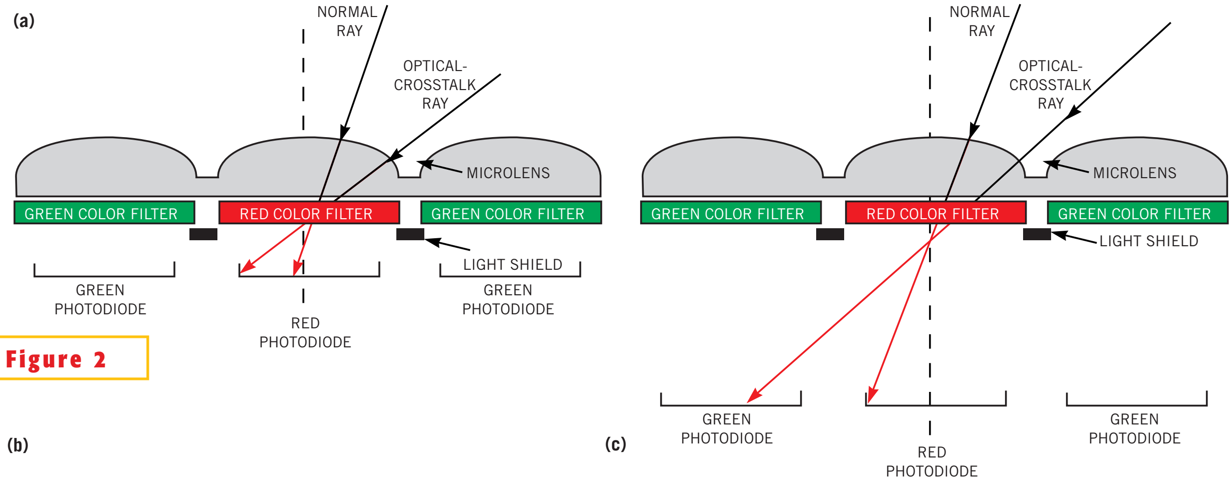

Figure 2

Dual-lens structures (a) and shallow-depth photodetectors in low-height pixels (b) and in competing "tall" pixels (c) are some of the methods vendors promote as they strive to boost sensor-light sensitivity without unduly degrading image quality (courtesy Sony and Micron Technology, respectively).

matrix-filter pattern or a cheap lens that's intrinsically susceptible to color aberration can also distort color reproduction (**Reference 5**).

Micron Technology, which in 2002 publicly announced its acquisition of Photobit, has developed low-height, or shallow-depth, photodetectors that reduce the effects of microlens-induced light-ray distortion (**Figure 2b**). Alternatively, sensors such as Foveon's products with their variable-pixel-size capability can automatically sum together the measurements of multiple closely spaced pixel sites, trading off resolution for light sensitivity as necessary. In some cases, an approach such as VPS (variable pixel size) can preclude the need for microlenses; many sensor manufacturers supply lenseless sensors upon request. The Foveon sensor in Sigma's first-generation SD9, for example, employed no microlenses. As a further cost-reduction move, especially as sensor resolutions increase and individual pixels become increasingly harder for the eye to distinguish, you might chose to disregard the antialiasing blur filter, which slightly softens the image but traditionally was necessary to suppress undersampling-induced moiré patterns and other image artifacts, along with jaggy stair-step patterns on diagonal edges (**Reference 6**).

INNOVATION CONTINUES

Despite these dismal prospects, image-sensor improvements are ongoing. Even if a refocusing of effort in directions other than resolution needs to occur, plenty of other application-tailored enhancements will for some time to come keep both vendors and implementers busy. (Analogies to CPU vendors' recently refocusing their metric of microprocessor merit on factors other than clock speed are apt.)

Electronic-versus-mechanical-shutter trade-offs exemplify these application-tuned improvements. In a conventional full-frame CCD, incremental electron collection occurs whenever you expose the sensor to light and the sensor's not in its reset state. A mechanical shutter to prevent light exposure when you don't intend it is necessary in this case. When you want to omit using a separate shutter because of cost or module-height issues—such as with camera phones—you can chose from two electronic-shutter approaches. One is the less common frame-transfer CCD, containing a separate light-shielded storage array to which the collected photo-site charge transfers at the conclusion of the exposure. This approach is silicon-costly; produces asymmetrical, rectangular die that are difficult to manufacture at high yields; and suffers from lengthy array-to-array transfer delays that restrict its practical usefulness.

The alternative and more common approach, the interline CCD, employs redundant light-shielded storage elements alongside each row of photodetectors with transfer between them taking only a few microseconds. This approach allows a subsequent exposure to begin in parallel with the extraction of the previous exposure's pixel values from the sensor. With CMOS sensors, you can use the active-pixel approach, placing a memory-based storage element and an ADC—keeping in mind the fill-factor trade-off—at each pixel location. Micron's TrueSnap technology and Pixim's Digital Pixel System are two examples. On the other end of the complexity spectrum, the passive-pixel CMOS sensor externally implements all processing functions. The interim alternative, which Micron's ERS (electronic rolling shutter) exemplifies, captures and transfers image data from the sensor one row at a time. It doesn't explicitly require an external shutter, although, without one, it creates blurred and distorted images of objects in rapid motion, conceptually similar to but even more egregious than the motion artifacts that interlaced-sensor videocameras generate (**Reference 7**).

SMaL Camera Techologies' Autobrite technology adaptively alters the traditional straight-line illumination-versus-signal-strength response of a conventional sensor in a manner similar to the logarithmic response curve of the human visual system. It performs this alteration on an as-needed basis to retain as much detail as possible in the bright and dark areas of the image (**Figure 3a**). The company is unique among imaging-building-block suppliers in that it offers not only

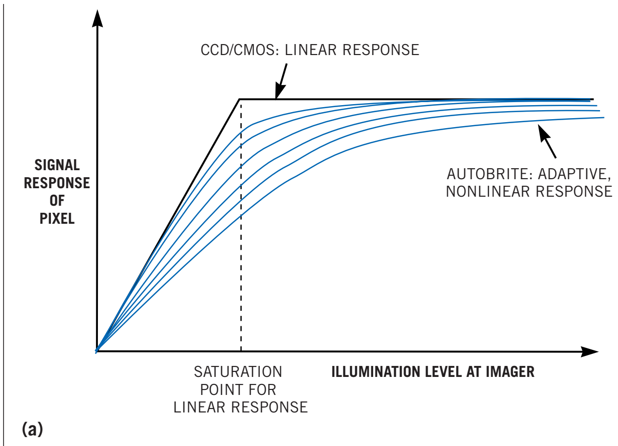

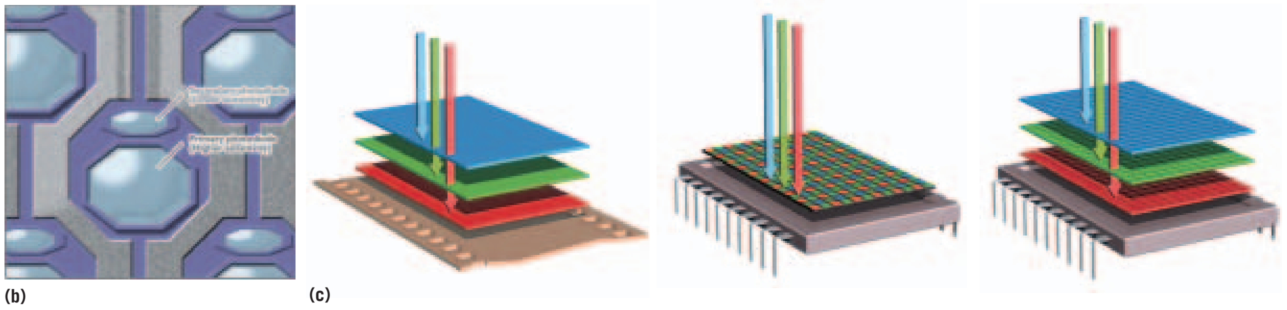

Figure 3 Autobrite (a) adjusts the illumination-versus-signal-strength curve to maintain detail in the dark and bright areas of a photograph, whereas Super CCDs (b) split the per-pixel illumination measurement task among multiple photodetectors. X3 technology dispenses with the traditional color filter array to deliver a film-reminiscent, per-pixel color-accurate representation of the captured image (c) (courtesy SMaL Camera Technologies, Fujifilm, and Foveon, respectively).

sensors, but also integrated imaging modules and even ready-for-production full-camera designs. Fujifilm's Super CCD has a honeycomb-pixel pattern, which the company claims makes more efficient use of the sensor's surface area. The fourth generation of the technology also offers dual per-pixel photodetectors, representing an alternative approach to preserving detail in dark areas of an image instead of averaging everything out to pure black or a noise-induced muddy gray. The primary photodetector, which Fujifilm explains as analogous to the woofer of a two-way speaker, captures dark- and midtone detail, and the secondary photodetector, equivalent to the photo site's "tweeter," records light at a lower sensitivity level, enabling it to capture detail in bright areas (**Figure 3b**). Fujifilm provides no detailed insights, though, on the characteristics of the Super CCD's "crossover network."

COLOR CONTRASTS

Foveon's X3 technology further expands the Super CCD multiphoto-site concept, albeit in a vertical-versus-planar fashion and focused on multicolor rather than single-color measurement. To understand it, first step back and review the fundamentals (**Figure 3c**). By themselves, image sensors know nothing about colors; if a photon with sufficient energy—that is, within a portion of the light spectrum to which the sensor is sensitive—strikes the photodetector and creates an electron-hole pair, the sensor accumulates that electron. To "force" the sensor to measure only a frequency or frequency range of light, you must externally filter the light shining on the sensor to only that or those frequencies.

One initial approach to solving the problem, which harks back to the first-generation color-television system and also is similar to the technique the single-wheel DLP (digital light processor) employs, involves rapidly rotating a tricolor RGB (red, green, blue)-primary or CMY (cyan, magenta, yellow)-complementary wheel in front of the sensor, which sequentially captures images that correlate to the three colors. And, as with today's advanced DLPs and ink-jet printers, you can use more than three colors in the wheel for incremental color accuracy. Aside from the obvious size, weight, power-consumption, and other issues related to the mechanics of the setup, you can use such a system only with still-life subjects; any movement during the multicolor multiexposure noticeably degrades the results.

Alternatively, you can use a prism to separate the light into red, green, and blue portions of the spectrum, directing the three outputs onto three sensors. High-end DLPs and professional digital videocameras employ this approach. Again, though, this approach involves size, weight, and power trade-offs, and a three-sensor array is also inherently significantly more expensive than a single-sensor alternative. In addition to the cost of the sensors themselves, a three-sensor configuration also requires incremental memory and DSP horsepower to process three times the amount of data in a reasonable time frame.

The third and most common approach nowadays employs a color-filter array on top of the sensor. The predominant Bayer pattern employs the RGB primary-color set and contains twice as many green filters as either blue or red ones, reflecting the fact that the human visual system is more sensitive to green-frequency light and, therefore, that capturing accurate detail is most critical in this portion of the visible spectrum. Post-capture interpolation generates approximations of the red and blue data for each green-filtered pixel, along with the remainder of the visible spectrum for blue- and red-filtered pixels. Complementary CMYG (CMY and green)-filter patterns also find occasional use, although the sensor outputs eventually must convert to RGB for display and printing. JVC's camcorders incorporate a hybrid complementary/primary matrix of clear, green, cyan, and yellow filters. And Sony's 8 million-pixel CCD uses an RGBE (RGB and emerald) color-filter array.

Foveon's X3 CMOS sensors employ a film-reminiscent alternative method of interpreting and capturing color information, one particularly notable for its all-important green-spectrum accuracy. Each pixel location contains three photodetectors at varying depths within the silicon and reflecting the varying penetration of portions of the visible light spectrum into the silicon substrate; blue is the shallowest, green is in the middle, and red is the deepest. One criticism of X3 that you commonly encounter is that light penetration into silicon is a continuum versus a "hard" cutoff—that is, that the blue-spectrum photodetector also reacts to green and red light and, conversely, that a few blue and green photons statistically penetrate the silicon lattice and interact with the red-spectrum photodetector. Foveon's Zarakov doesn't dispute this observation. However, he points out that it conceptually does not differ from the situation that occurs with a matrix-filter array. In those arrays, each filter doesn't *completely* block light outside its intended spectrum, although the shape of the response curve differs in the X3 case. He also notes that, from an implementation standpoint, the subsequent image-processing steps straightforwardly compensate for the overlap aftereffects.

Nowadays, Foveon, along with Fujifilm and its Super CCD, bases its sensors' specifications on the number of photodetectors, rather than the number of pixel sites they have. This approach inflates the sensors' claimed resolution and complicates comparisons with conventional filter-matrix-based sensors. I can think of a list of reasons that Foveon's approach seems misleading; Foveon has an equally long list of reasons that its approach makes sense. We've agreed to disagree. I concur with its observations that a per-pixel, tricolor photodetector cluster enables the sensor to better capture true resolution than the interpolated resolution that a matrix-filter counterpart of *comparable pixel-site count* offers. I also agree that color aliasing inaccuracies resulting from light that enters the sensor at acute angles are less problematic with X3 than with a matrix-filter array. Foveon claims that its X3 sensors are no more difficult or expensive to manufacture than are conventional CMOS sensors with color filter arrays. The company's biggest hurdle at this point isn't technical; it's the business challenge of convincing customers to commit to a US-based, single-sourced, and single-foundry sensor technology.

FORECASTS

Antishake systems in cameras have historically taken either a lens-housed mechanical approach, which requires the purchase of esoteric, expensive optics, or a mostly electronic image-stabilization system, which requires motion-sensing transducers and which produces passable results only when you use it with an oversized image sensor (**Reference 8**). Konica Minolta is rolling out yet another stab at solving the problem of the shakes in its latest cameras. The company employs conventional optics, but, instead of including an excessively costly, oversized sensor, it mounts the sensor on a movable bracket that shifts the sensor's location, using a piezoelectric element, in synchronicity with the transducers' feedback.

Olympus' E-1 camera also employs sensor-centric movement but aims to solve a different problem. The company's catchy Supersonic Wave Filter moniker refers to the fact that a movable, transparent membrane protects the image from dust. The membrane connects to an ultrasonic transducer and vibrates during the camera's power-up cycle to shake off dust, which sticky tape subsequently captures. Sigma's SD10 protects its Foveon sensor with a simpler approach: a transparent, passive dust shield that, being outside the lens' focal plane, doesn't noticeably degrade image quality.

One of the long-touted benefits of CMOS sensors is that they let you include ever-increasing amounts of conventional logic and memory circuits on the same die as the sensor. Most image-processing tasks, including multipass noise detection and subtraction, defective-pixel compensation, color interpolation, auto-white balance, auto-exposure calculation, and image scaling and compression, now occur in a separate image-processing chip (see **sidebar** "Integration trends" with the Web version of this article at edn.com). However, numerous recently published academic papers posit the feasibility of on-sensor per-pixel image processors to handle some or all of these functions, and manufacturers have fabricated several test chips. Start-up founding and funding will invariably follow.□

REFERENCES

1. www.maushammer.com/systems/dakotadigital/DakotaDigital.html.

2. Dipert, Brian, "Pick a card: card-format options," *EDN*, July 8, 2004, pg 53.

3. Kempainen, Stephen, "CMOS image sensors: eclipsing CCDs in visual information?" *EDN*, Oct 9, 1997, pg 101.

4. Dipert, Brian, "Heat wave: FPGAs confront increasing, evolving power consumption," *EDN*, Aug 2, 2004, pg 61.

5.Dipert, Brian, "A crash course in color conversion," *EDN*, June 7, 2001, pg 46.

6. Dipert, Brian, "Display technology's results are compelling, but legacy is un 'clear,'" *EDN*, Oct 26, 2000, pg 63.

7. Dipert, Brian, "Video quality: a hands-on view," *EDN*, June 7, 2001, pg 83.

8. Schweber, Bill, "Image stabilization shows diversity of engineering approaches," *EDN*, Oct 26, 2000, pg 79.

AUTHOR'S BIOGRAPHY

Technical editor Brian Dipert doesn't miss all those shoeboxes filled with negatives, slides, and prints one bit, and he's even bought a 12-in. Apple PowerBook to speedily run Photoshop while he's on the road. Now, all he needs is a bigger hard drive. Reach him at 1-916-454-5242, fax 1-617-558-4470, bdipert@edn.com, and www.bdipert.com.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in *EDN*.

Agilent Technologies
www.agilent.com

Atsana Semiconductor
www.atsana.com

Canon
www.canon.com

ChipWrights
www.chipwrights.com

Conexant
www.conexant.com

Cypress Semiconductor
www.cypress.com

Eastman Kodak
www.kodak.com

Epson
www.epson.com

Foveon
www.foveon.com

Freescale Semiconductor
www.freescale.com

Fujifilm
www.fujifilm.com

Hynix Semiconductor
www.hynix.com

Intel
www.intel.com

JVC
www.jvc.com

Konica Minolta
www.konicaminolta.com

Micron Technology
www.micron.com

MIPI Alliance
www.mipi.org

National Semiconductor
www.national.com

Nokia
www.nokia.com

NuCore Technology
www.nucoretech.com

Pentax
www.pentax.com

Pixim
www.pixim.com

Polaroid
www.polaroid.com

Qualcomm
www.qualcomm.com

Renesas Technology
www.renesas.com

Ritz Camera
www.ritzcamera.com

Samsung
www.samsung.com

Seiko Epson
www.epson.co.jp

Sharp
www.sharpusa.com

Sigma
www.sigmaphoto.com

SMaL Camera Technologies
www.smalcamera.com

Sony
www.sony.com

Starzen Technologies
www.starzen.com

STMicroelectronics
www.st.com

Texas Instruments
www.ti.com

Toshiba
www.toshiba.com

VESA
www.vesa.org

SYSTEM EVOLUTION, REVOLUTION, AND POTENTIAL STAGNATION

Photography magazines and online discussion forums often debate the phenomenon of the ever-decreasing light sensitivity—in other words, the ever-increasing noise at a particular ISO setting—with increasing resolution. Momentarily stepping off the seemingly unending resolution treadmill, it's useful to honestly ask, given these and other tradeoffs, just how much resolution the average camera user needs. You might cynically suggest that the comprehension behind this question is beyond the capability of average consumers, who historically find their primary—if not only—motivation is "bigger is better." But, just as more consumers have figured out that yesterday's PCs can run any software they might require, more consumers are deciding that the cameras they own or the cheapest, lowest resolution models now available are sufficient for their purposes.

Consumers don't typically enlarge most photos larger than 4×6 in. For a full-frame print of these dimensions, a 2 million-pixel camera is more than adequate. Even if a consumer wants to crop and enlarge a portion of the frame, today's 3 million-pixel cameras are up to the task. And 4 million pixels (for full-frame images) to 6 million pixels (for cropped and enlarged cameras) are sufficient for $8\frac{1}{2}\times 11$-in. prints. Don't let the 1200/1440-dpi (dots per inch), 2400/2800-dpi, and even higher dot-per-inch claims of latest generation ink-jet printers fool you. As the term "dots per inch" implies, they refer to the application of a single ink dot to the printed page, and typically they refer to either the length or the width of the print; the other dimension is a much smaller dot-per-inch value. Multiple complementary-color dots are necessary to construct a full-color cluster. And multiple clusters combine to create an accurately shaded pixel. An ink-jet printer delivers much fewer pixels per inch than the dot-per-inch specification would suggest, and the figure is often less than 200 ppi.

If resolution's upward spiral does end, how else can hardware manufacturers and their sensor partners convince consumers to keep buying new gear? One likely approach involves incorporating still-image-capture functions within other pieces of equipment, such as cellular phones, PDAs, and camcorders. Camera phones are so far the biggest success story, but they also have some of the most significant challenges to continued growth. Cost and form-factor pressures force them to include the smallest sensors, which have the lowest light sensitivity at a given resolution; this fact is contrary to camera phones' common usage in dimly light bars and restaurants. The media has well-documented the privacy and intellectual-property-protection concerns of camera phones; copyright concerns also abound. For example, many bookstores ban camera phones because some patrons photograph book and magazine pages instead of purchasing the publications.

Although some digital still cameras have added video-capture capability, the low-frame-rate, postage-stamp-sized results are to date underwhelming, and consumers are more likely to migrate in the opposite direction. A premature flattening of the digital-still-camera market at the hands of camera phones and camcorders might not significantly alter the number of sensors vendors sell, but it will have a drastic affect on, for example, a still-camera manufacturer that historically has not participated in these other markets and for which it would be too expensive to enter these markets and establish a successful brand.

Sensor suppliers can also implement other strategies as they pursue their aspiration to sell more silicon. They can advocate the incorporation of multiple sensors into each system or the addition of sensors to systems that lack them. For example, a high-end camera-phone design might include two image sensors—one for snapshots and the other for videoconferencing—much as today's high-end phones include multiple LCDs or OLEDs (organic light-emitting diodes). Vendors often tout automobiles as *the* new applications for image sensors; they provide front-, side-, and rear-vision assistance in navigation, thereby improving safety for drivers, passengers, and passersby.

DECIPHERING SIZE

As you compare manufacturers' image sensors, you'll often find several measures of the sensors' size in the data sheets. These measures include total package size, the dimensions of the active array, and the aspect ratio—typically 4-to-3 or, mimicking 35-mm film, 3-to-2. You'll likely also encounter the "optical format," an at-first-glance baffling number with values such as 1/4, 1/3, 1/2, 1/1.8 and 4/3 in. The term harks back to the standard sizes manufacturers applied to 1950s-era TV-camera tubes; the specification refers to the outer diameter of the tube's long glass envelope. This designation, clearly long obsolete in practice, is still in wide use. To translate optical format to the lenses' projected imaging plane and, therefore, to the required diagonal dimension, multiply the optical format by a two-thirds scaling factor and, if necessary, convert from US units to metric units.

Camera users do not typically concern themselves with optical formats, except indirectly when the sensor size affects light sensitivity and other operating parameters. (For example, smaller, cheaper lenses are also more prone to chromatic aberration and other optical distortions.) One important exception to this rule involves digital DSLR (digital single-lens-reflex) cameras onto which a consumer wants to mount lenses. Full-frame sensors that match 35-mm-film frame dimensions are expensive; smaller sensors lead to a focal-length-multiplication factor that may be undesirable, depending on the situation. In my case, lenses I mount on my Pentax *ist D undergo an approximate 1.6-times focal-length extension. When I'm shooting sports or wildlife photography, I welcome this transformation; it lets me get closer to the action without requiring an expensive super-telephoto lens and without the optical and light-transmission trade-offs of a teleconverter. If I'm shooting panoramic landscapes, conversely, the multiplication factor is undesirable; that's why I've held onto my Pentax 67 medium-format camera and optics. Addressing this issue for new-equipment purchasers, the so-called 4/3 format represents a coordinated effort by camera and lens manufacturers to develop well-matched bodies and lenses based on small-format sensors.

INTERCONNECT CONTROVERSY

Image sensors and camera phones continue to generate increasingly larger frame rates and sizes. Meanwhile, the phones are becoming more commoditized and cost-sensitive. As a result, phone suppliers are increasingly pressuring sensor, display, and processor manufacturers to come up with an industry-standard, high-speed, and low-pin-count interconnect bus. Several contenders are currently duking it out in standards bodies, customer briefing rooms, and the market. They include:

- National Semiconductor's MPL (Mobile Pixel Link);
- Nokia and STMicroelectronics' SMIA (Standard Mobile Imaging Architecture);
- Qualcomm's MDII (Mobile Display Digital Interface), which the VESA (Video Electronics Standards Association) is currently evaluating;
- the MVI (Mobile Video Interface) Alliance of Renesas and Seiko Epson; and
- the efforts of MIPI (the Mobile Industry Processor Alliance, which includes Nokia as a highly visible member.

HANDS-ON ANALYSIS

This spring, during the Embedded Systems Conference in San Francisco, I spent an afternoon at the city's Asian Art Museum. I photographed, among other things, an image of a small, off-white Bodhidharma figurine in front of a black background, and, because the museum prohibits the use of flash photography and the ambient illumination was low, I used a 1/45-sec shutter speed on my Pentax *ist D DSLR (digital single-lens-reflex) camera. When I viewed the enlarged image on my computer monitor, I noticed several small white spots in the dark background areas of the print, and, looking at my image archive, I found similar discolorations at the same pixel locations in other pictures.

I at first thought that my camera's CCD was dirty, but carefully cleaning the image sensor didn't improve the situation. Because I'd used different lenses to shoot the pictures, I knew that optics blemishes weren't to blame. I reluctantly concluded that my camera's CCD had defective pixels, a not-uncommon phenomenon that you also find in LCDs. However, this situation baffled me because, soon after buying the camera, I'd run test images created by it through Starzen Technologies' useful freeware, dead-pixel-test utility (www.starzen.com/imaging/deadpixeltest.htm) and found no problems. Now, however, the utility was reporting as many as 20 "hot" pixels—those with a luminance value of at least 60, with pure white being 255—but no dead pixels—those with a luminance value of 250 or higher—at various shutter speeds.

The utility uses as its input TIFF files you generate with the camera's lens cap on, thereby theoretically creating completely black images. The Pentax *ist D's noise-reduction feature, which I'd always enabled, was active only at shutter speeds of 1/4 sec and slower; that's the speed at which I'd done my initial testing and thereby not uncovered the CCD's hot pixels.

I contacted Pentax and asked them to send me a replacement body for testing and possible substitution. When it arrived, I reran my testing on both it and my original *ist D through 21 shutter speeds of 1/90 of a sec to 10 sec and with noise reduction turned on and off. The results were surprising in several respects. This time, my original *ist D reported only a few hot-pixel errors. Results on my original *ist D also varied at shutter speeds greater than 1/4 sec with noise reduction on and off, even though they shouldn't have. And in all cases, my original *ist D outperformed its replacement, which I promptly shipped back to Pentax.

After additional research, dialogue with Pentax, and comparative analysis with other members of the Pentax Discuss Mailing list (www.pdml.net), I think I've figured out what's happening. First, I reran my testing of my original *ist D with the eyepiece covered; light leaking onto the sensor might have swayed the results the first time. Realize, too, that hot-pixel manifestations, likely the result of dark current, vary slightly from shot to shot. Thermally generated electrons are key contributors to noise; others include the analog circuitry in the downstream amplifier and ADC circuits and the natural variability in the battery and, therefore, the system-supply voltage. By slightly altering the hot-pixel threshold in the dead-pixel test above and below its default 60 luminance value, I could dramatically diminish and increase the number of hot pixels the utility reported for each TIFF file. A subsequent retesting of the camera bodies might suggest a different conclusion.

I've concluded that hot pixels are preferable to other sources of image degradation, such as imperfections on a camera or an enlarger lens and particulate matter embedded on negatives and slides. This conclusion ties directly to hot pixels' predictable locations and behavior. I can relatively easily fix the defects in Photoshop or another image-editing tool. I'm eager to run the test on other digital cameras, but I'm resisting the urge because I know the tool's results would suggest a worse situation than exists. Our other cameras lack an option to output TIFF or raw files; they support only the lossy-compressed JPEG format. The 8×8-pixel-clustering technique JPEG employs during its DCT step would likely spread the hot pixels' altered luminance across the 64-pixel cluster, mimicking the distortion you'd find in a typical JPEG photo but overestimating the extent of the sensor's actual defect rate.

Table A details relevant specifications for various cameras. Note that, as sensor size decreases and resolution increases, the default ISO setting of the camera also drops. A recent article on the latest crop of 8 million-pixel digital cameras further validates and extends this trend of high resolution's equating to low ISO/high noise (**Reference A**). Visit the version of this sidebar on the *EDN* Web site, where you can download and analyze my test results on the two *ist D bodies, along with several example TIFF images. I also have a Sigma SD10 DSLR on the way, courtesy of Foveon, and I'll share my observations and test results in the Web-site addendum.

REFERENCE

A. Richards, Dan and Michael J McNamara, "Great '8' megapixel shoot-out!" *Popular Photography & Imaging*, July 2004, pg 75.

TABLE A—SAMPLE DIGITAL CAMERAS AND THEIR RELEVANT SPECIFICATIONS

CCD-based camera	Sensor active size	Total pixels (millions)	Effective pixels (millions)	Default ISO	ISO range
Pentax *ist D	23.7×15.5 mm (Sony ICX413AQ)	6.3	6.2	200	200 to 1600 (3200 as custom option)
Kyocera FineCam SL300R	1/2.7 in.	3.34	3.17	100	100 to 400
Kodak DX6340	1/2.7 in. (5.27×3.96 mm)	3.34	3.1	100	100 to 400
Kodak DX6490	1/2.5 in.	4.2	4	80	80 to 800

INTEGRATION TRENDS

Few TV or computer-display OEMs still buy stand-alone bare panels from LCD manufacturers. Instead, they buy prefabricated screens that incorporate the necessary backlight, column drivers, and other circuits. They might even take a full-blown reference design that includes the system interface circuitry to production. Or, as is increasingly likely, they'll simply stick their name on someone else's hardware design that uses customized firmware.

Similarly, customers will, as time progresses, ask their sensor suppliers to provide increasingly integrated and comprehensive imaging products. Using the example of a camera phone, a sensor supplier today either internally develops or partners with another company to provide the required digital and, for CCDs, analog image processors, timing and voltage bias generators, and other components. Eventually, some or all of this circuitry will likely go onto the same piece of silicon containing the sensor. Most sensor companies targeting the camera-phone market will optionally sell you a module containing the optics subsystem. And, from companies such as Agilent and Sharp, you can even buy internally developed LEDs that provide a flash with adequate illumination for those all-important dim-ambient-light photos (**Figure A**).

Figure A

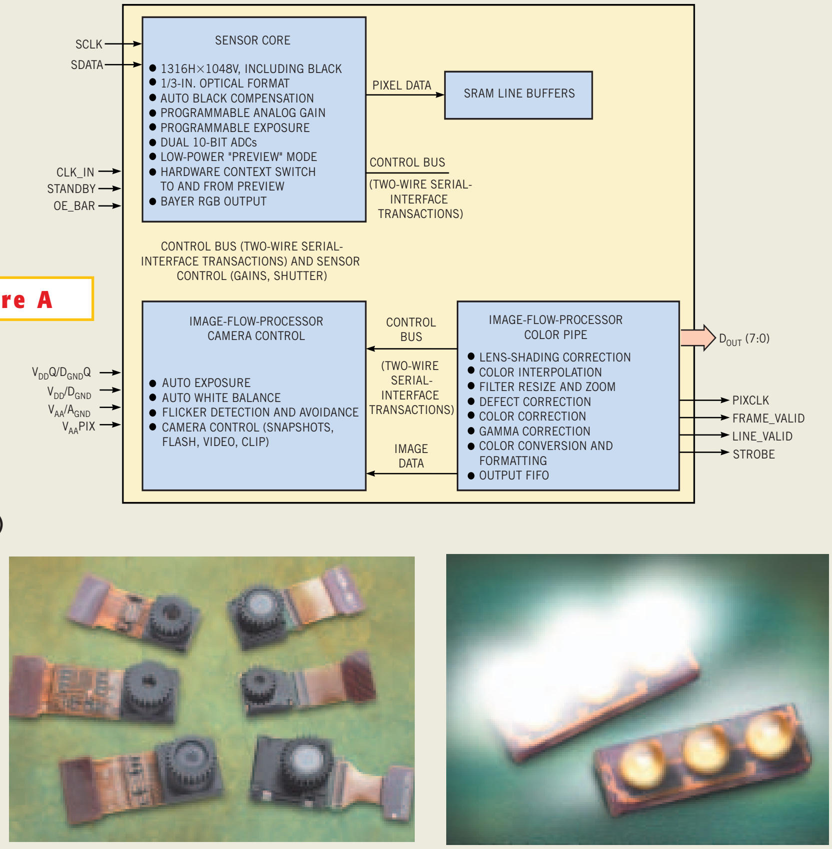

Micron and other companies bundle the sensor and image processor (a), whereas Agilent offers a one-stop-shopping approach to imaging-building-block acquisition (b).