Measuring image-
quality parameters
of integrated CMOS
active-pixel-sensor
digital camera on a
chip is more difficult
than with a CCD
system.

CMOS digital cameras need new figures of merit

Eric R. Fossum

Complementary-metal-oxide-semiconductor (CMOS), active-pixel-sensor (APS), digital camera-on-a-chip technology has progressed rapidly in the six or so years since it was invented by scientists at the NASA Jet Propulsion Laboratory (Pasadena, CA; see Fig. 1). While performing the same general image-capture function as its predecessor, the charge-coupled device (CCD), this highly integrated "system on a chip" produces new types of figures of merit for comparing performance. Some of these figures of merit can be compared to CCD imaging systems by tracing a path from the photons in a scene to the bits coming out of the camera on a chip.

Optical elements

Before entering the camera on a chip, photons are gathered by optics. Although optics typically are not considered part of a sensor, in lower-cost systems the ultimate resolution and image quality are often dominated by the optics and not the sensor itself.

Modern on-chip microlenses are formed by an inexpensive single-mask step in the backend of the silicon-wafer-fabrication process. They act as funnels to direct light incident across an entire pixel toward the sensitive portions of the pixel and not as imaging optics. Microlenses increase the responsivity of some low-fill-factor, small-pixel CCDs by a factor of two to three. Unfortunately, however, microlens performance is at its worst when the lenses are most needed.

Under lower light levels when small f-numbers are used, the rays of light striking the sensor surface come from a wide range of angles. Many of these non-normal incident rays are inefficiently funneled by the microlenses and are lost. Because the microlenses are monolithically integrated with the sensor, it makes sense to specify the sensor with its microlens in place, although responsivity at lower f-numbers will drop.

Color filter arrays (CFAs) select photons of a given wavelength range for a particular pixel. Both complementary-color (cyan, yellow, or magenta) and primary-color (red-green-
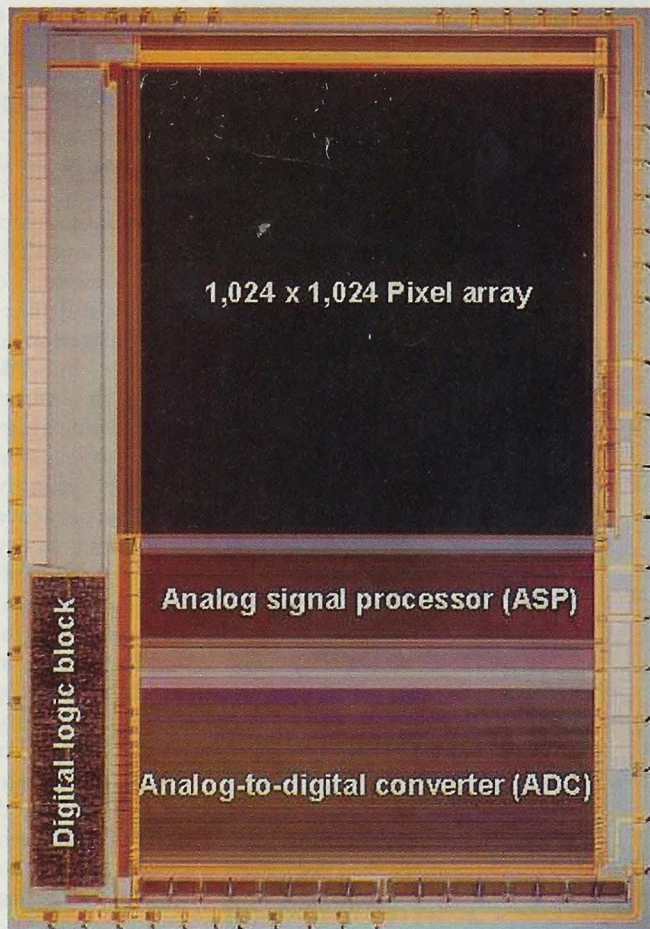
ERIC R. FOSSUM is chief scientist at Photobit Corp., 135 N. Robles Ave., Pasadena, CA 91101; fossum@photobit.com

FIGURE 1. A 1024 × 1024-element color complementary-metal-oxide-semi-conductor, active-pixel-sensor camera on a chip was designed for NASA.

blue) arrays can be used. Global uniformity of the CFA is important because changes in its absorption across the pixel array will appear as a color shift across the image when color processing is performed (see Fig. 2).

Photons to electrons

After passing through the various optical layers, the photon enters the silicon. The quantum efficiency measures the ratio of collected electrons to incident photons over a single pixel and is always less than unity for visible light. The effects of the fill factor and the optics are included in the quantum efficiency. Fill factor is the ratio of optically sensitive silicon area

to total silicon area in a particular pixel.

Unlike interline-transfer CCDs that need careful shielding in the pixel to reduce smear and have concomitant low fill factor, CMOS APS devices are immune to smear and have much larger effective fill factors—typically 30%–40%. Larger fill factors, aside from allowing more light to enter the silicon, also reduce the effects of aliasing. The quantum efficiency is important in determining the signal-to-noise ratio (S/N) of the sensor at a given lighting level. A two-fold increase in quantum efficiency can result in a 3-dB improvement in S/N under most lighting conditions.

The photodetector is a nonequilibrium device, so there is net thermal generation of electrons, in addition to the optical generation. Thermal generation depends strongly on temperature (doubling every 10°C) and will result in a signal after some integration time, even in the dark. Average values for this dark current typically range from 100 to 1000 pA/cm^2. So in an image sensor with 1/30-s integration time and 5-μm-pixel pitch, the dark signal is on the order of 5–50 electrons and contributes a noise between 2 and 7 electrons rms, which is negligible. Dark-current nonuniformity is quite important, however, and broad distribution of values can lead to color aberrations as well as white spots—pixels with relatively high levels of dark current.

Electrons to voltage

The conversion gain measures the ratio of output voltage to the number of collected electrons in a pixel and is usually measured in microvolts per electron (μV/e-). A typical value is 10–30 μV/e- in state-of-the-art CMOS APS devices. Large conversion gain is good for amplifying signals above readout noise levels, but comes at the expense of dynamic range. This is because maximum signal swing is typically 1–2 V at the pixel, and the full well for a CMOS APS is given by the maximum swing divided by the conversion gain, which, for a 2-V swing and 20 μV/e-, for example, results in a full well of 100,000 electrons.

Linear full well defines the maximum number of signal levels that still preserve a certain degree of linearity in the output (roughly 2%) and is typically 80% of full well. The responsivity of an integrating, voltage-output pixel is defined in volts per lux-second, where a
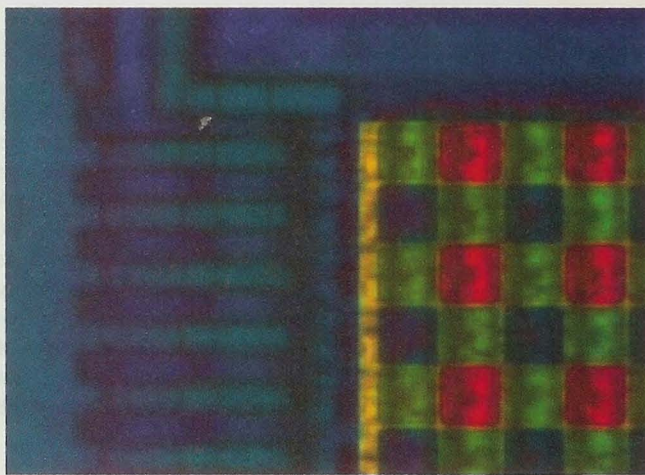

FIGURE 2. Color filter array resides on a color complementary-metal-oxide-semiconductor, active-pixel-sensor chip

lux-second exposure represents a light level of one lux illuminating the sensor directly for 1 s.

In some CMOS image sensors, a programmable gain amplifier (PGA) scales the pixel signal into a range useful for analog-to-digital conversion and reduces the impact of noise introduced prior to the analog-to-digital converter (ADC). This amplifier typically provides a gain between 1 and 16, although the actual PGA gain may differ from the setting value. If the PGA operates with low noise, high linearity, and at relatively high data rates, it can dissipate a lot of power in the image sensor. Nonlinearity can introduce difficulties into subsequent color processing.
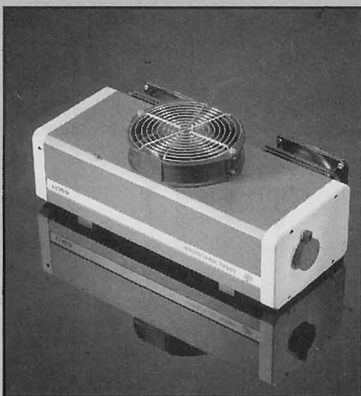
Volts to bits

The ADC converts the analog-sensor signal into a digital representation. The resolution of the ADC determines the number of significant bits in the output

word of the ADC. Most on-chip ADCs provide between 8- and 12-bit output, where 8 bits is typical for low-end applications such as teleconferencing and 12 bits for high-end applications such as digital still cameras. Higher resolution can be achieved, but typically at the cost of slower throughput (conversions/s) and higher power. The resolution of the ADC differs from the accuracy of the ADC, however. As with discrete ADCs in CCD systems, the on-chip ADC should have low integral nonlinearity and low differential nonlinearity.

All ADCs use a reference voltage to perform the conversion and map a given input voltage into a digital representation. For example, a 1-V reference in an 8-bit ADC yields a digital value of 128 bits when an input of 0.5 V is converted, where the term "bits" means the number of least significant bits. When examining figures of merit, the ADC reference voltage is an important parameter (see Fig. 3).

Signal processing

The degree of on-chip digital-signal processing can vary significantly from one application to another and may involve changing the pixel values, as in color preprocessing or autoexposure control. The output multiplexer takes words from the digital-signal processor and delivers them to output pads. The output multiplexer may deliver serial data, nibble-mode data (4 bits at a time), full

words, or parallel output words, depending on the chip-design goals. The total pixel throughput is an important figure of merit for high-speed imaging and equals the pixel-per-second output of the chip. The range of possible values includes 100 Kpixel/s for slow-interface applications, 60 Mpixel/s for HDTV-type applications, and 500 Mpixel/s or higher for high-speed motion-analysis applications.

Because the output of a digital camera on a chip is digital and many of the internal voltages are inferred, it is sensible to characterize sensors in terms of their digital output. For example, digital responsivity can be defined as bits/lux-sec for a given pixel color (for example, green), for a given PGA setting (for example, 10), and for a given ADC reference voltage. Another example is average digital dark signal, which can be described in terms of bits per second at a given temperature, PGA setting, and ADC reference voltage.

Noise parameters

Noise is another important figure of merit and also needs to be characterized digitally. In this case, digital noise is measured as bits rms at a given exposure (lux-sec), PGA gain setting, ADC reference, and total pixel throughput. The definition of noise can result in digital noise levels that are a fraction of a bit, especially at low light levels.
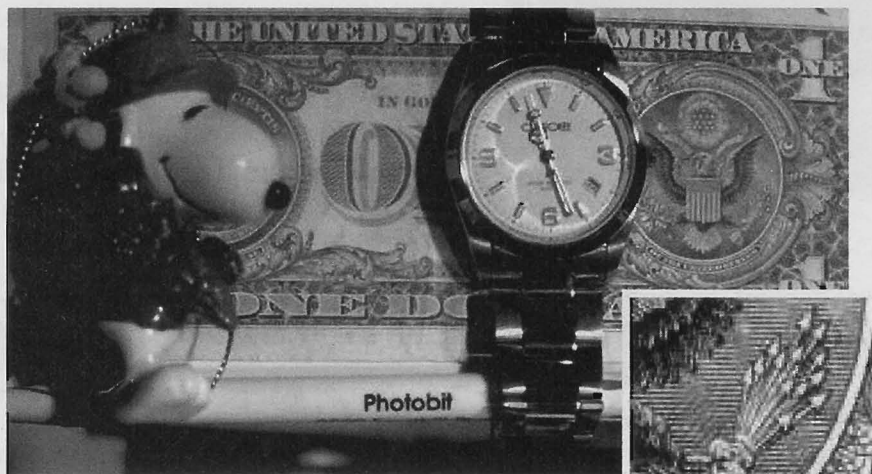
Dynamic range has traditionally been

FIGURE 3. In image captured from the Photobit PB720 sensor (inset shows enlarged portion), the 1280 × 720-element, color complementary-metal-oxide-semiconductor, active-pixel-sensor has 640 10-bit analog-to-digital converters operating in parallel and produces progressive scan imagery at 60 frame/s (55 Mpixel/s).

defined as the ratio of the maximum signal to the read noise, assuming one can see objects with a S/N of 1:1, and the intrinsic (analog) dynamic range of CMOS APS devices is typically between 70 and 80 dB. For digital output, one expects the dynamic range for an 8-bit digital output, for example, to be 256:1 or 48 dB. With a digital noise level lower than the least-significant bit, though, the digital dynamic range is extended to larger values

Perhaps the most important figure of merit for CMOS image sensors relative to CCDs is digital fixed-pattern noise (FPN), measured in bits rms, which is a pixel-to-pixel variation in offset level, independent of illumination. While CCDs and CMOS image sensors both have fixed-pattern noise, the column-parallel processing of CMOS image-sensor signals makes the CMOS image sensor vulnerable to column-to-column offset variations that can appear as faint "stripes" in the image and that are obvi-

ous to the eye. In the past few years, many CMOS image-sensor designers have overcome the FPN issue. Both global FPN rms values and vertical-horizontal FPN rms values are important figures of merit and are typically a function of PGA gain and ADC reference voltage, depending on their physical origin.

Photoresponse nonuniformity consists of both a fixed-pattern component and a gain component and can come from variations in microlenses, PGA, ADC nonuniformities, CFAs, quantum efficiency, and conversion gain—the latter two are both related to random photolithographic fluctuations in defining the pixel area. To date, photoresponse nonuniformity in CMOS APS devices is almost identical to that found in CCDs and is thus attributed primarily to random photolithographic variations.

The highly integrated nature of the CMOS APS digital camera on a chip means that measuring many of these parameters is more difficult than in a

CCD system. However, the values currently obtained in high-quality CMOS image sensors are now comparable to the best CCDs, and the market penetration of CMOS image sensors into current CCD applications, such as digital still cameras and video camcorders, can be expected to accelerate rapidly over the next few years. ❑

Bulk reprints of all Laser Focus World *articles can be ordered from Cathleen Johnson, reprints manager, at (603) 891-9415 or FAX (603) 891-0574.*
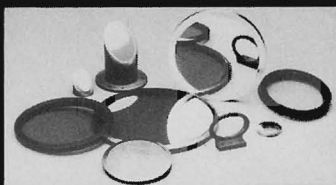
YOUR RATING, PLEASE...
Is this article of value to you? Please circle appropriate number on the Reader Service Card.
YES—502   NO—503