Image SENSORS

by Reg Miles

Figure 1a. Stripe Filter

Virtually all electronic cameras, consumer to broadcast, movie and still, analogue and digital, use a charge coupled device (CCD) as the image sensor. The exceptions to this rule normally use a complementary metal oxide semiconductor (CMOS) image sensor (although alternative technologies are under development). Recently, CMOS performance has improved, and that, combined with the intrinsically lower cost of producing them, has meant that the technology is showing promise in certain market areas, and is likely to assume some significance in the near future.

CMOS sensors have the advantage of being just a variation on standard CMOS technology: the sensors thus benefit from the general improvements that are made to the widely used technology; and they can be produced in any of the normal CMOS foundries. CCD sensors, conversely, use a specialised VLSI process; and their production is largely limited to the major electronics manufacturers. CMOS can also integrate a variety of peripheral functions onto the one chip, whereas it would not be economically feasible to do so with a CCD. And a CMOS consumes less power than a CCD. But the CCD does achieve a very good image quality with low noise; and it is available in a wide variety of sizes, with pixel counts ranging from thousands to tens of millions.

Pixels, or discrete picture elements, are necessary for both types of sensor because neither is capable of recording a continuous electrical representation of an optical image. A pixel in this context is a photon conversion site; and is normally a p-n junction photodiode or occasionally a photogate which is just an MOS capacitor that is exposed to light. The photodiode is the more light sensitive, because the MOS gate partially absorbs light - particularly blue. Incidentally, pixels are generally rectangular for video use and square for just about everthing else.

If photons have an energy greater than 1eV or their wavelength is less than 1000nm sufficient energy can be transferred as they penetrate the silicon of the pixel to create electron-hole pairs. The electrons and holes are then separated: normally by applying an electrical field which, depending on the substrate, will either drain away the holes and leave the electrons, or vice versa. In most sensors it is the electrons that are left behind to become the charge carriers. However, to be effective they must be integrated into charge packets of sufficient quantity to provide a measurable output. Integration is achieved on the reverse-biased and electrically isolated junction capacitances. The number of charge carriers that a pixel can hold in its potential well is known as the well depth, and will vary between tens and hundreds of thousands.

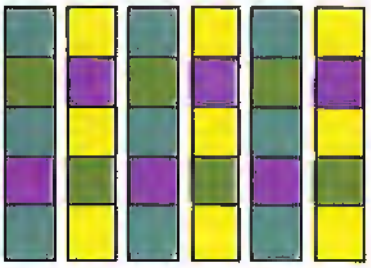

Figure 1b. Mosaic Filter

This gives only the brightness of the scene. To achieve colour reproduction the sensor must be filtered. Cameras with three sensors have dichroic filters coated onto a complex prism block so that each sensor receives only red, green or blue light. The majority of cameras have only one sensor, and use dyes coated on in stripe or mosaic arrangements (see Figure 1a/b). The colours used are normally red, green and blue (RGB); but can also be cyan, magenta and yellow (CMY); or any combination of at least three colours - the only criterion is that

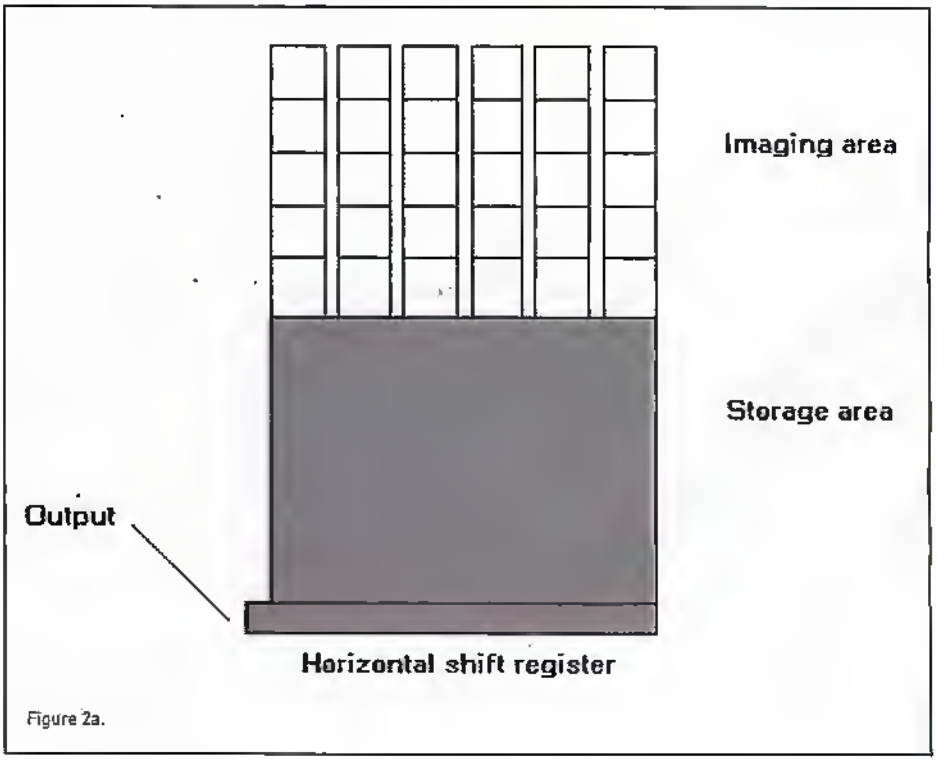

Figure 2a.

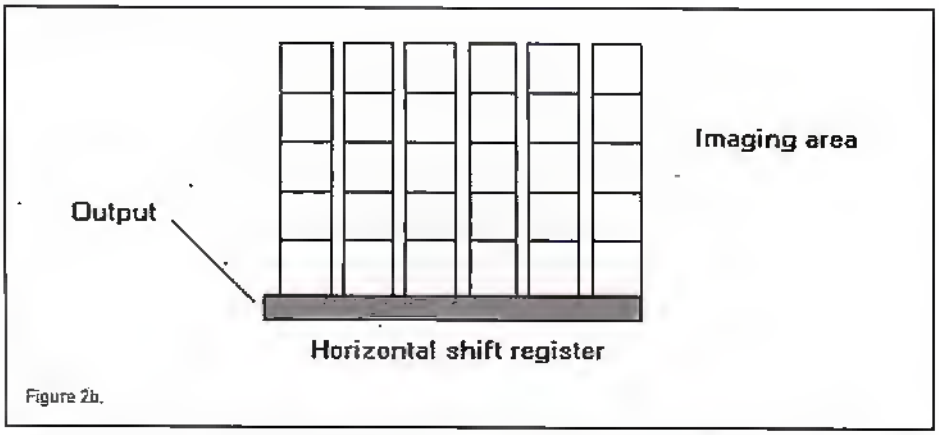

Figure 2b.

after processing they give the correct range of colours for the application.

Up to that point the CCD and CMOS are identical. Once the charges have been integrated though the resemblance ends.

To get the charges out of a CCD the horizontal rows of charge packets are moved down in parallel by clock voltages applied to gates consisting of a fine mesh of polysilicon electrodes that cover the CCDs surface. As each row reaches the bottom the charge packets are then clocked serially along a horizontal shift register. Each charge packet is then clocked via the output gate onto a floating n+ region ('floating diffusion'), and the voltage changes are sensed by means of a source-follower. After each charge packet has been sensed it is cleared in readiness for the next by applying a positive pulse to the reset gate which flushes it out the reset drain.

This is the basic principle. The details vary according to the type of CCD.

If a frame transfer (FT) type, Figure 2a, the charge packets are moved rapidly down through the imaging area, pixel by pixel, to a similar, but opaque, storage area during the field blanking period. From there they are clocked out a line at a time, through the horizontal shift register, simultaneously with a new charge being integrated in the imaging area. The disadvantage of using the imaging pixels as a shift register is that new electrons produced by areas of bright light can infiltrate the charge packets as they pass causing smearing (although modern designs minimise this).

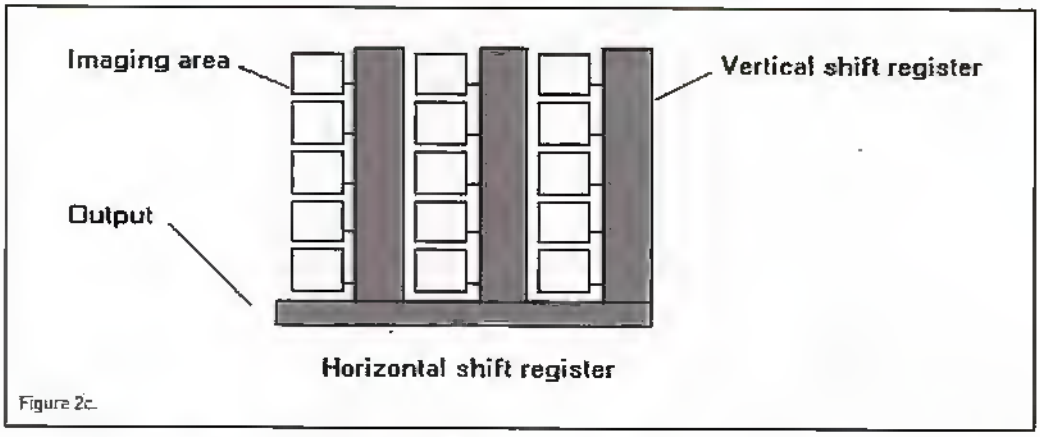

Figure 2c.

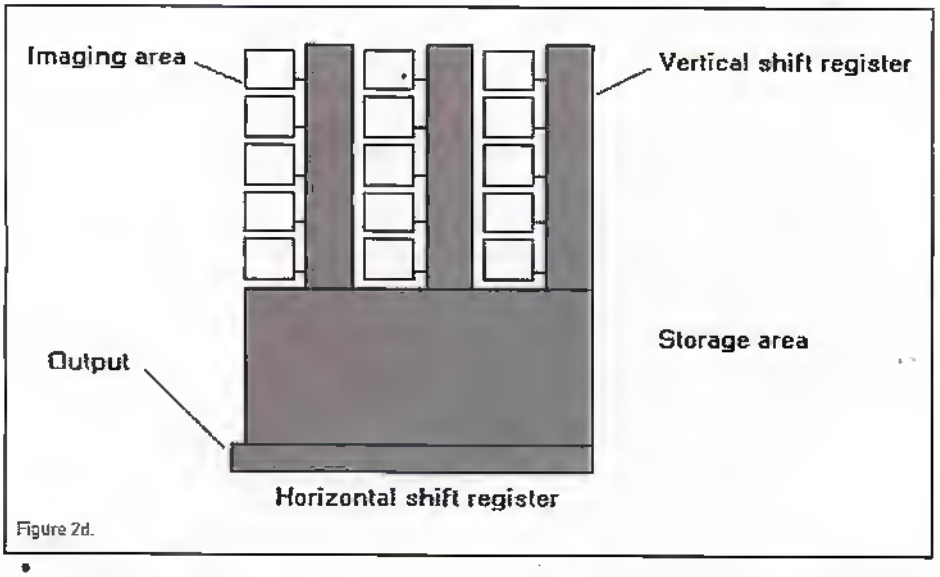

Figure 2d.

This type of charge transfer is also used for the full-frame (FF) imaging type, Figure 2b; but here the whole area is light sensitive with no opaque storage area (the signal will either be displayed directly or go to an external store). FF is thus smaller and cheaper. However, because the whole charge must be read out before a new charge can be integrated its frame rate is lower. And this makes it even more susceptible to smearing, requiring optical shuttering or strobe lighting to prevent it.

The interline transfer (IT) type, Figure 2c, reduces the time from pixel to storage area to virtually zero by having adjacent opaque shift registers. The charges are read out sideways from the pixels, via a signal path, to the vertical shift registers. They are then clocked down the vertical shift registers, and out through the horizontal shift register, during one full field period - again, simultaneously with a fresh image being integrated. Smearing is thus much reduced. The disadvantage is that the imaging area is no longer filled entirely with light sensitive pixels.

The frame interline transfer (FIT) type, Figure 2d, is a combination of the FT and IT. The charges are moved rapidly down the adjacent shift registers to an opaque storage area during field blanking, and output while a new image is being integrated. It is thus as big as an FT type, but with the reduced imaging area and added complexity of the IT - the worst of both worlds, but for non-existent smearing.

The timing will also vary according to whether scanning is interlaced or progressive, and whether it is providing still or moving images (an increasing number of CCDs can be switched between both states).

With a CMOS the pixels are individually addressed, and the charge packets are switched to charge sensing amplifiers. There are two basic types of CMOS sensor - passive pixel and active pixel (active pixel sensor, or APS). The former has a buffered, charge sensing amplifier for each column of pixels, and a single transistor in each pixel that acts as a charge gate. In operation an entire row of pixels is switched to electrically isolated column bus lines for the amplifiers to sense the individual charges and convert them to an output voltage. With the active pixel type there is a source follower amplifier in each

pixel and the charge packet is switched directly to it before being output to the column bus. In both types the vertical shift register clocks a line at a time to the sample and hold circuits and the horizontal shift register reads them out.

The active pixel usually contains at least three transistors - the additional ones being a reset transistor to control integration time and a row-select transistor. Because of this its fill factor (the ratio of light sensitive area to total pixel area) is less than that of the passive pixel. However, the direct amplification does reduce noise.

Whether passive or active, photodiode or photogate, CMOS is generally less light sensitive than a CCD of comparable size.

Not that CCDs are perfect: in addition to the IT and FIT types having a reduced imaging area, the channel stops dividing the pixels vertically to prevent horizontal charge

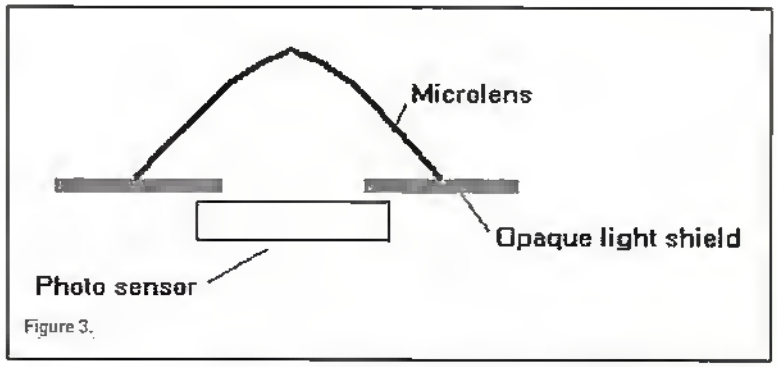

Figure 3.

insensitive areas by using finer fabrication - Philips', world's smallest, 2.4um square CCD pixels necessarily employ fine fabrication, which can equally well be applied to larger pixels. Or larger pixels can be used at the expense of resolution (smaller pixels enhance resolution, but at the expense of quantum efficiency and well depth). Or an indirect approach can be used - improving the spectral characteristics of the dyes used for filtration, for example.

As with all electronic devices CCD and CMOS sensors (particularly the latter) suffer from noise, and therefore incorporate some form of compensation.

In the case of a CCD this will probably be correlated double sampling (CDS). The output is reset to a reference value after each charge packet is flushed away, then the difference between that value and the next charge packet becomes the charge value for that particular pixel. This will largely eliminate fixed pattern noise (FPN) - the difference in dark current between individual pixels, in addition to several types of temporal noise.

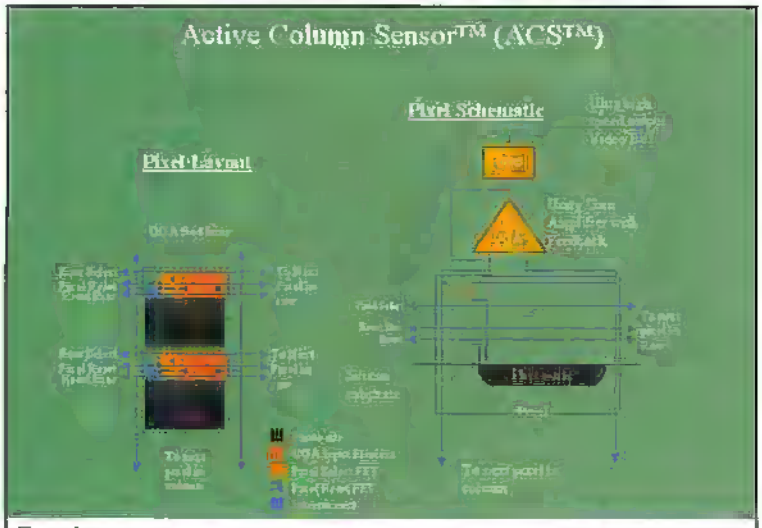

Figure 4a.

A CMOS is more likely to employ optimised fixed pattern correction, because FPN is of more concern with CMOS than CCD where it has been much reduced. The FPN shows particularly in the APS type with its multiplicity of individual amplifiers, each with its own offset and gain value. Passive pixel is less susceptible, although the pixel's position on the column bus does affect the charge level at the amplifier and the turn-on thresholds of the transistors do vary; but passive CMOS are used in products where noise is of less concern anyway - such as toys. In practice an on- or off-chip memory stores the offset values of the pixels, obtained by reading their output during reset, and subtracts the offset noise from the image. It is also possible to use this method to correct for optical fluctuations and defects.

Photon Vision Systems has further reduced FPN by the invention of an alternative to APS, called Active Column Sensor. In addition to reducing FPN, ACS increases signal strength because the gain of each amplifier in an active pixel is typically only 0.84. Unfortunately, the solution to the low and varied gain, a unity gain amplifier (UGA), would require at least six transistors in each pixel. What Photon Vision Systems has done is to compromise with a shared UGA, with one dual input transistor per pixel and around four shared column transistors (Figure 4a shows the arrangement, Figure 4b shows a conventional APS layout). Typically, the ACS pixel will have over twice the fill factor and full well capacity, providing at least twice the S/N ratio. Integrated test chips, with 25 micron pixels, have only 0.08 FPN and a S/N ratio greater than 86d8 as measured on the output

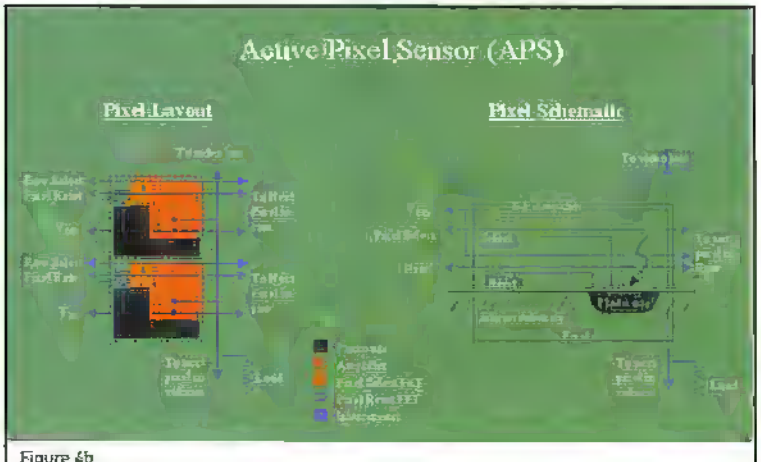

Figure 4b.

leakage are opaque, and the polysilicon electrodes absorb some of the blue light. In applications where sensitivity and full spectral response are necessary, such as astronomy, FF CCDs are backside thinned and used back to front allowing photons to impinge on the base of the pixel wells. This process of thinning is a straightforward one, the silicon is dissolved away by acid to a depth of about 15 microns, and the CCD is mounted on a rigid substrate. The quantum efficiency (response to different wavelengths of light) of these back illuminated CCDs is therefore high by comparison with when they were front illuminated.

For more general applications microlenses are being increasingly employed on both CCD and CMOS to direct the photons onto the active areas of the pixels (see Figure 3). These are produced by a photolithographic process. Another method of increasing pixel sensitivity is to reduce the size of the

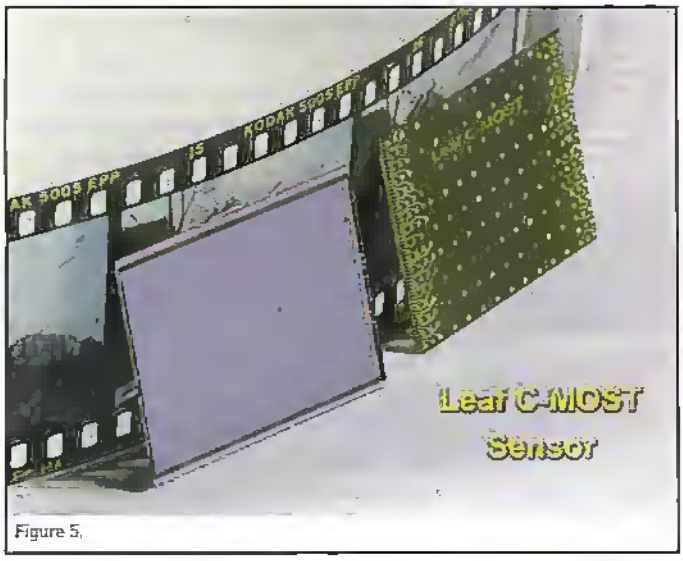

Figure 5.

of the sensor. A VGA format array with 10 micron pixels will have in excess of 70dB S/N at full video rates on the output of the sensor. Coupled with a CDS circuit to minimise offset variations, the company is claiming CCD performance at CMOS cost.

A photogate APS has noise control in each pixel. This is achieved by using the active circuitry to perform correlated double sampling. However, this is achieved at the expense of a reduced fill factor and greater complexity.

Both CCD and CMOS are also susceptible to blooming: where the well depth is insufficient to hold all the electrons and they spill over into adjacent pixels, causing image highlights to spread. This is countered by draining the excess to the substrate. But with the consequence of a reduced pixel area, smaller electron well depth and lower quantum efficiency. In areas of use such as astronomy CCDs may be used without anti-blooming; countering the overspill by taking a series of short exposures, then combining them into one by using image processing software.

The dynamic range in CMOS can be increased by having pixels that operate in logarithmic mode, with direct readout and random pixel addressing. There is no integration time, the logarithmic conversion circuit continuously converts current to an output voltage that will vary with the instantaneous light intensity (like a photographic exposure meter). The result will be a range of about 140dB, by comparison with around half that figure in linear mode. However, in low light levels the response time is slow, and with low contrast images the contrast is further reduced.

A somewhat similar approach to that of anti-blooming allows an electronic shutter to be incorporated into CCD and CMOS sensors. Here the pixels are held in reset, with the charges draining away, until the pre-set exposure time begins. CMOS sensors can also vary the timing inputs to the sensor. Both can also have an increased integration period and employ a mechanical shutter.

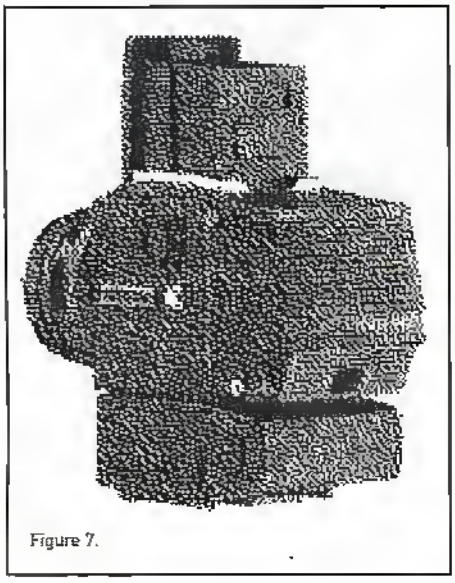
Figure 7.

There is now a rapidly growing need for sensors that can be switched between 4:3 and 16:9 aspect ratios. Some consumer camcorders offer a 16:9 facility, but it is only electronic (letterbox) masking of the 4:3 image, not a true 16:9 image from a 16:9 sensor. The CMOS with its XY addressing has no problem with switching - any aspect ratio can be easily achieved by building in the facility to select different pixels. With a CCD, conversely, it requires major retiming to read out a 4:3 panel in a 16:9 chip.

This pixel selectivity of CMOS also enables the sensor to operate in reduced resolution, sub-sampled modes for a viewfinder display. This is achieved by writing to certain serial register bits in the sensor, which will then integrate and read out every second, fourth or eighth pair of lines, and every second pair of pixels. Keeping pixels in pairs maintains the arrangement of the colour filter blocks and, thus, the colour signal. A CCD cannot do this.

However, it can bin pixels (combine them) to increase the sensitivity at the expense of resolution. Binning is achieved by clocking the charge packets of individual pixels in both horizontal and vertical directions to give one large 'super pixel': vertical binning by clocking multiple lines into the horizontal register and horizontal binning by clocking multiple pixels under the summing gate. Thus, 2x2 binning combines four adjacent pixels - quadrupling sensitivity while halving both horizontal and vertical resolutions, 3x3 binning combines nine, and so on. However, it cannot be used with CCDs that have individual colour filters on the pixels, because the colour would be mixed up.

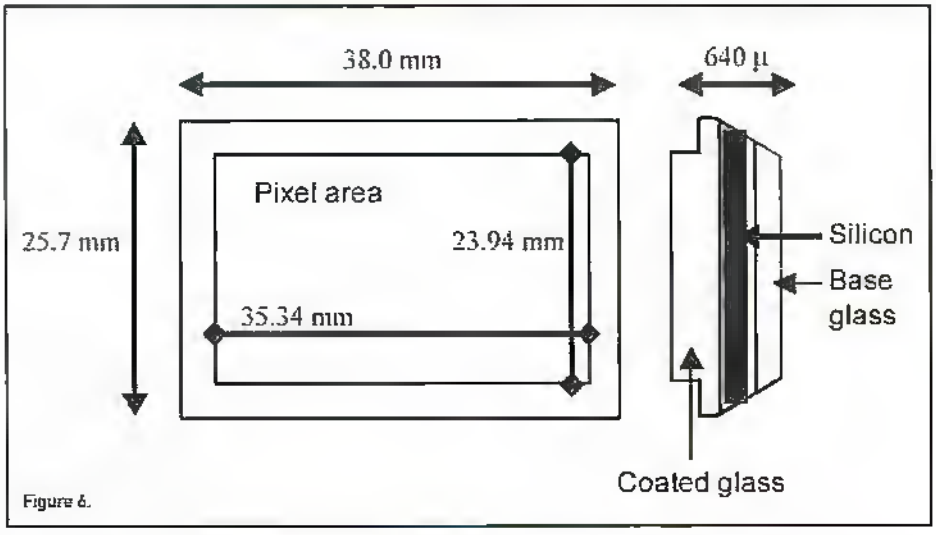

Figure 6.

This limited flexibility of CCDs is going to become an increasing disadvantage. And, with its complex clock requirement, something of an Achilles heel. The clock amplitude and shape are critical for correct operation, and this requires clock drivers using multiple supply voltages. Also, CCDs normally require secondary chips to implement the various functions such as clock drivers, timing logic, signal processing, etc.

Companies producing CMOS sensors used to be almost apologetic about their products. Now they are becoming quite bullish, with some predicting the imminent demise of the CCD and others, more cautiously, anticipating CMOS replacing the CCD at the lower end of

5	8	8	8	9	8	8	8	4
10	1	1	1	1	1	1	1	13
10	1	1	1	1	1	1	1	13
10	1	1	1	1	1	1	1	13
10	1	1	1	1	1	1	1	13
11	1	1	1	1	1	1	1	12
10	1	1	1	1	1	1	1	13
10	1	1	1	1	1	1	1	13
10	1	1	1	1	1	1	1	13
10	1	1	1	1	1	1	1	13
2	6	6	6	7	6	6	6	3

Figure 8.

the digital still camera market as a more realistic objective at this stage. In addition to a lot of small companies developing and/or producing CMOS, like VLSI Vision in Edinburgh (the first company to integrate a complete camera on a single chip), there are major names working, singly or in partnerships, on improving the performance still further, such as Canon, Kodak, Motorola, Texas Instruments and Toshiba. While NASA is experimenting with CMOS cameras to see whether they can get the performance of a CCD in a smaller, lighter camera that consumes less power - an obvious advantage in space.

NASA's Jet Propulsion Laboratory, the first to develop a practical CMOS APS, has developed the first fully digital camera on a chip, needing only five wires for operation. And has recently invented an approach that will make APS compatible with silicon-on-insulator (SOI) technology, that it is said will become the baseline for CMOS VLSI implementation. While in a separate development JPL is working on a high speed APS to be used as an integrated, 'smart sensor'.

The CMOS people are getting very excited by the prospect of these 'smart sensors'. They will combine the function of an imaging device with the integral facility for processing information acquired from those images to provide a result that can indicate position, motion, identity, etc; or be used in areas such as robotic vision; or to vary the charge integration time in individual pixels for adaptive exposure control. It will be possible to use a CCD for some of these purposes; but the facility to integrate functions onto a single chip means that CMOS generally has a significant advantage in this area.

CMOS technology is also moving gradually upmarket, as typified by the announcement of the CreoScitex Leaf C-Most sensor. This is the first CMOS with 6.6 million pixels (3150x2100); and the first that equals the frame size of 35mm film (see Figure 5). The pixels measure 11.4x11.4 microns. It is said to be particularly suitable for integration in upmarket 35mm, medium and large format photographic equipment; and was shown at Photokina 2000 in a medium format back. Its ultra-thin packaging will allow it to fit exactly into the focal plane of a standard 35mm camera (see Figure 6). CreoScitex collaborated with FillFactory for the design of the APS CMOS, with Tower Semiconductor for wafer fabrication and ShellCase for the packaging. The technology is scalable so larger and smaller devices can be expected.

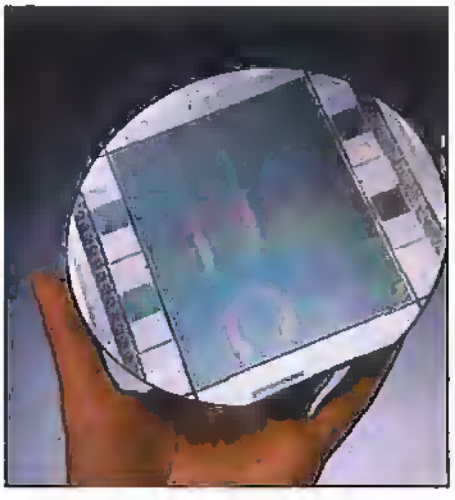

Figure 9.

On the subject of smaller devices, Canon has already launched the EOS D30 digital SLR with a 22.7x15.1mm image area CMOS. This has 3.25 million pixels, measuring 10.5x10.5 microns. And incorporates a programmable gain amplifier to minimise noise.

But the CCD is far from finished. Research continues on finding new applications for it, and both furthering its performance and reducing its cost.

Phase One has shown a prototype of the LightPhase H20 digital back which incorporates a 4020x4020 (16 million pixel) CCD from Eastman Kodak (Figure 7 shows it attached to a medium format camera). It is expected to be made available in July.

But for sheer size and number of pixels nothing can beat the modular CCD from Philips. This takes the basic building block of a 1024x1024 FF chip: which, using standardised connections, can be produced in combinations of up to 7x9 - 86x110mm and 66 million pixels. The layout is produced using just four basic blocks: imaging area blocks, vertical blocks, horizontal output register blocks and output amplifier blocks (see Figure 8). The 'split' arrangement allows readout through one, two or all four amplifiers. Production takes advantage of the repetitive structure to use a step and repeat process; thus overcoming the present inability of lithography equipment to produce fine details over a large area. Joining the individual parts electrically in the silicon is achieved by a patented 'stitching' technique that is claimed not to produce any visible seams in the image. A 7x9 array is the maximum at present; the result of using 12x12 micron pixels and necessarily fitting it all onto a 6 inch wafer (see Figure 9). And there can be any number between one and the maximum to suit the customer's requirements; with the only unique mask being that used to make the interconnect between the bonding pads. The particular configuration can thus be realised in production rather than design - greatly reducing development time and costs.

Both FF and FT devices can be produced (in the latter case with half of the sensor masked for charge storage). If larger pixels are required then 2x2 or 3x3 can be obtained by binning. If the application requires colour, then a colour mask can be added to the CCD.

The full 7x9 CCD is initially being used for astronomy; and it is anticipated that it will also be used for digital photography, where the current maximum for an area array device is 6x6cm (medium format in photographic terms). A 1x2 FT CCD is already in production for progressive scan still imaging. While the modular device additionally has the potential for high speed imaging, with one block receiving the image and eight blocks around it storing sequential frames.

On a much more modest scale, Minolta has also embraced a modular approach: their RD-3000 camera uses two half CCDs to achieve almost 3 megapixels at lower cost. The image is divided by half mirror prisms and received by the two CCDs; the slightly overlapping images are then processed to remove differences between them and seamlessly joined.

Fuji has developed a significant change to the design of an IT CCD, and called it the Super CCD. They have taken the pixels, turned them through 45 degrees, changed their shape from rectangular to octagonal and arranged them in a honeycomb formation (see Figure 10). The charge packets can now be transferred directly to the vertical shift register, making the signal path redundant. The pixels are also larger, due to both their shape and arrangement, so their light-gathering capacity is improved. They are also better shaped to make the most of microlenses. With the result that the sensitivity, S/N ratio and dynamic range are all claimed to be improved by 2.3x as a result. The horizontal and vertical resolutions are also improved by the new, more closely packed, arrangement. In combination with Fuji's newly developed honeycomb signal processing LSI the effective resolution is claimed to be at least 1.6x greater than a conventional IT CCD with the same number of pixels (the new FinePix 6800 Zoom camera (see Figure 11) is said to give a resolution equivalent to a 6 million pixel IT from a 3.3 million pixel Super CCD). Conversely, the Super CCD can achieve the same resolution with fewer pixels - reducing power consumption and cost. Its colour reproduction is also claimed to be better. And, with R, G and B pixels in each horizontal line, alternate lines can be skipped for interlaced scanning - which will give better resolution than combining pairs of lines as is done conventionally for video output. It is also possible to sub-sample the pixels, with any vertical ratio possible and 1/3 horizontal ratio - speeding up the process.

Fuji envisages digital cameras with picture quality approaching that of 35mm film; combination digital cameras and camcorders giving high quality video as well as stills; and, by miniaturising the CCD, an extremely compact digital camera.

Although the CCD continues to dominate, and is likely to do so for some time to come, and CMOS is working its way up the imaging pile, there are other contenders waiting. I will deal with those, and other variations on image sensors, in the second part of this article.