

Digital Decisions: Considerations in High End Digital Camera Design

Ken Boydston
Goleta, USA

Three questions ought to be asked fairly early on the process anything:

1. "What it should do and how will it do it?"
2. "Who is it for?"
3. "How much will they pay?"

This paper is primarily about question 1, for that is where most of the interesting design tradeoffs occur, a number of which are influenced by answers to 2 and 3. We will answer 2 and 3 first:

"Who is it for?"

MegaVision's products are nearly all for professional photographers. The designs discussed are intended to generally be used by the same group of photographers who use large and medium format cameras. The photographic tasks fall into 2 groups:

1. taking pictures things that sit perfectly still (like products).
2. taking pictures of things that don't (like people).

MegaVision makes digital cameras and digital capture camera components that take pictures of both kinds of things. The cameras that only take pictures of still things take pictures that may be printed larger without loss of quality, take longer to capture (several ones of seconds), take less time to process, and have fewer artifacts. The cameras that take pictures of moving things can also, of course, take pictures of still things. These cameras capture very quickly. The pictures they take cannot be printed as large without loss of quality, require more processing time, and have more artifacts.

"How much will they pay?"

MegaVision customers pay between $15K and $30K for cameras and camera backs made by MegaVision, except in the very early days (10 years ago), when the price exceeded $100K. We expect in the near future to continue creating products in this price range as price/performance ratios continue to improve. We expect that the bottom end of this price range will go down. Given the price of very large silicon devices, this price range limits their use. This price range also allows considerably better performance than can be achieved with much smaller silicon devices.

"What does it do and how does it do it?"

In short, each MegaVision digital capture device competes with professional film as an image capture medium. Each must be either:

a. **Better** without being much slower or much more expensive

or

b. **Faster** without being much worse or much more expensive

or

c. **Cheaper** without being much worse or much slower.

Even the earliest expensive cameras MegaVision produced were purchased because they were cheaper and faster than film. Faster than film has never been disputed, and has often been a compelling reason for adoption. The first digital images captured on very expensive MegaVision cameras were cheaper than film in the high production studios where they were used. At today's costs, images captured on MegaVision digital cameras are cheaper than images captured on film in a much greater number of studio applications. As digital capture costs trend downward, digital capture will become cheaper than film capture in an increasing number of applications. The quality of digitally captured images exceeds film in many cases, and with each generation of digital camera produced, the quality improves.

S3 and T32 digital backs, MegaVision's most recently developed products, capture images that can be used to create professional quality 8X10 prints at least as good as those that can be made with film from medium format cameras, and Professional quality 16X20 prints comparable with film capture prints that can be sold at the same price as print made from film. The images can also print with unquestionable quality as 8 1/2 X 11 inch full page bleeds on coated stock on commercial presses, and as 11X 17 spreads with acceptable commercial quality.

Overall Design Considerations

Should the design be for a complete camera or a component of a camera system?

Answering this question was easy for the T32, a camera designed to take pictures of only still things. The T32, which is never hand held, is always mounted to a rigid fixed stand, is almost always used in a studio, and is almost always used to take pictures of still things was designed as a back for large format view cameras and as a stand-alone camera which included a shutter. The somewhat narrow usage of this camera made these design considerations fairly easy.

The S3 is intended for a much broader user base: commercial, portait, studio, location, tripod mount, hand held. Designing a complete camera system would have been well beyond the capacity of MegaVision's design capabilities, so we chose to make a component: just the capture part of a camera; ie, a camera back.

Given the currently available sensor sizes (closer to 35 mm than to 60mm cross sections) that could be used to make a camera back to sell in our target price range, the optimum camera optical fit would be 35 mm cameras.

However, making a back for 35 mm camera back posed a big problem:

Since sensors are packaged with a glass cover and material around the edges of the active sensor area, a full size 35mm size sensor won't fit between film guides, so a smaller chip would have to be used, or extensive modification to the camera would be required. Fiberoptic filmplane translation was too expensive and moving the plane back would make the 35mm camera too bulky and ugly.

Furthermore, the camera would have to compete with other 35mm cameras, which are relatively cheap compared to medium format cameras.

Making a back for medium format cameras has several advantages. The mechanical design is relatively easy, since the film plane is defined in removable backs well behind the focal plane shutter or barn door. No modification to existing platforms were required. And medium format cameras are in wide use among our intended market; they are, in fact, the preferred choice of our target market.

The downside to this choice: The size of the medium format image plane is a roughly 6 cm. cross section. Existing large sensors approaching this size could not used because of price, packaging, and performance limitations. Read out time at optimum S/N would result in several ones of seconds per capture, yet clearly closer to 1 frame per second is required by the target market.

A smaller (roughly 30mm cross section) sensor has the advantages of 1. good price 2. good package, and 3. good performance.

The clear disadvantage is that its wrong the size for the range of lenses designed for 60 mm image plane cross sections. However, since our target market is mostly portrait & commercial photographers, wide angle is not so important. Wide angle lenses become normal, so the optical range for our target market did actually exist. Thus, made what appears to be an odd choice; we chose a small format sensor for a medium format camera.

Note: a minor improvement is that since coverage is in center of existing lenses, sharpness is optimal out to the edge of the image area.

For the S3 digital back, we chose a 24mmX36mm, 6MegaPixel, 12micron pixel CCD sensor from Philips. Some very attractive features of this sensor are very good S/N at room temperature, very good anti-blooming, and somewhat attractive quantum efficiency curve.

Mechanical configuration requirement

The back must not exceed the size of a film back. This requirement, together with the complexity and cost of multisensor designs for large sensors, ruled out the development of a multisensor camera. Multichip designs necessarily occupy space between the lens and the focal plane, and force considerable design complexity.

The mechanical configuration requirement did not come about arbitrarily; the ergonomics of handheld cameras are important to the creative process of the photographer and should be given very careful consideration. In considering this requirement, we asked "should the camera be considered as one box or as two boxes?"

Since most of the time, most of the customers would connect directly to a computer, we considered the computer to be a part of the camera, and therefore 2 boxes evolved naturally. For portability, given power consumption, required battery life, and available battery technology, it was not possible to include the battery into same enclosure as the sensor, so a 2nd enclosure which held the battery was natural. A second enclosure allowed most heat generating components to be removed from the enclosure containing the chip, resulting in a power dissipation of less than 1 watt in the camera back and a negligible increase in CCD operating temperature. Since the CCD includes a charge dumping feature, and dissipation while not in read down mode is very low, power to the CCD did not need to be removed between exposures, which results in more stable operation.

Capture rate:

Since it was determined that the camera platform to be supported was medium format, the capture rate would naturally be compared with the auto-winder capture rate of the host camera. These capture rates are in the range of 1 frame per second, and this capture rate was set as a target design point. Capture rate is fundamentally limited by how fast the pixels can be read from the array. In a CCD, the rate is determined by:

1. Clock speed
2. number of pixels
3. number of paths (amplifiers, possible ADC's)

Clock speed is related to S/N; on the Philips sensor the optimum clock speed is about 6MHz. To maximize S/N, it was determined that the sensor should operate near this frequency. 6M pixels at 6MHz translates to 1 second (plus some overhead) to read the array. A second path is available which could double this rate, but at the expense of considerable complexity and cost. Since a single path was somewhat close to the design target, it was deemed sufficient. A single path also has the added benefit that the CCD can be installed normally or 180 degrees rotated, and the better of the 2 orientations can be used. For example, flaws can be minimized by orientation or the better of the 2 output amplifiers can be selected.

Nearly every "standard" interface design available, including SCSI and 1394 (Firewire), imposed significant capture rate limitations for capturing directly into a computer. SCSI has the further limitation of short cable length and fat cables. Since the PCI bus is a cross platform standard with bandwidth sufficient to digest data at the optimum readout rate, it was determined that a very simple PCI board with minimal buffering would provide a good solution to for optimum data transfer. Since it was determined that a 2 box design was fundamentally indicated, the PCI interface card could conveniently be considered as the 2nd box, at no extra expense. Splitting the camera into 2 parts adds a second benefit to life cycle costs. As the back housing, together with its enclosed CCD, embody the major cost of the digital back system, removing components from it reduces downtime of the major cost item in the event of component failure. Since the PCI board is relatively inexpensive, yet it contains the majority of the components, spare inventories are cheaper to maintain.

The tradeoff, and it is significant for a broad market, is that the computer must be opened to install the PCI card. As many computers still require a SCSI card or 1394 interface card, this point is moot in many cases.

We wanted enough resolution to print an exceptionally good 8X10 print and a very good 16X20. A 2K X 3K chip is not really the best use of the pixels for most formats for which the camera is used... only about 2500 pixels of the 3K available are used for an 8X10. This number of pixels, as can be seen from the examples that I have brought, is certainly adequate to satisfy the print size requirement. In fact, some of our customers routinely print 30 inch prints. However, because of the color under-sampling of the array, (bayer pattern), moire can be a problem. Most color aliasing does not moire, and adaptive median-like filtering very successfully removes such. Color aliasing can be eliminated optically at the expense of resolution. Tests suggest that with optical filtering the current CCD will still produce a very very good 8X10, and possibly and acceptable 16X20, but enough optical filtering to eliminate aliasing definitely reduces the resolution and hence the acceptable print size.

More pixels with optical filtering will certainly be a consideration in future designs, but reasonably priced CCD's with more pixels will mean smaller pixels, since silicon is priced by the acre, more or less independent of the amount of pixels on it. Smaller pixels mean less well capacity, hence less dynamic range. Lower noise floors and increased quantum efficiencies in chip designs will help offset the pixel size limitation.

Dynamic Range and Sensitivity

How much dynamic range is needed? The simple answer is, “All you that is possible.” You can never have too much. Unlike pixels, the number of which the target print size suggests practical limitation, more dynamic range can always be used, because if you've got more than you need to make a really good print, it can be used to increase the sensitivity and thus require less light, which can often be in short supply. What is reasonable: enough to make a good picture at minimum ASA 100 or possible up to ASA 400. For professional quality, we felt that we needed to have some 2,000 to one at ASA 100 and upwards to 1,000 to 1 at ASA 400. Active Cooling: at short shutter speeds, cooling the CCD does not provide much benefit. Since most of the uses required short shutter speed, and active cooling would add considerably to the power budget, package size, cost, and design complexity, we decided if we could meet our S/N requirements without active cooling, we would do so.

Negative: our practical exposure time is limited to several ones of seconds ("Usable" images up to 8 seconds are made"). Also, if the choice is between longer exposures and increased sensitivity, we chose increased sensitivity.

If the CCD is operated at close to full well capacity, its theoretical S/N limitation exceeds 4,000 to 1, or 12 bits. Operating close to full well requires increased camera complexity to maintain linear output. Thus, the decision was made to operate well down from full well (less than 1/2), which served to both reduce complexity and increase sensitivity. By so doing, we were able to use a 12 bit ADC without loss of information, reduce camera complexity, and increase camera sensitivity. I might mention that it was considered that some of the camera response curve shaping be performed in analog, which would reduce the possibility of banding due to bit limitation. However, there was no compelling evidence that this was needed, and using to a linear amplifier reduced color transform computational complexity and increased color accuracy.

Adjustable ISO

Adjusting the ISO can be done either by software or by hardware. Just as with film, opting for a higher ISO trades off some quality. Whether the sensitivity is pushed via hardware or software, some compromise is mandated. The S3 ISO adjustment is a combination of hardware and software. A hardware gain adjustment very near the CCD output increases signal without increasing data path noise, so very little noise other than that from the CCD is amplified. This switch operates over a 2 stop range. Software curves further adjust the gain in addition to shaping the response curve.

Conclusion

The S3 works quite well, and more than satisfies its original design requirements. Some of the features traded off will return in future designs as technology changes and their implementation becomes viable.