

Rubber Focal Plane for Sky Surveys

John L. Tonry^a, Gerard A. Luppino^a, Nicholas Kaiser^a, Barry Burke^b, George H. Jacoby^c

^aIFA; ^bLincoln Laboratory, Massachusetts Institute of Technology; ^cWIYN Observatory

ABSTRACT

We describe progress in removing image motion over large fields of view. A camera using a new type of CCD has been commissioned and we report first results which are very promising for wide field imaging. We are embarking on a project to build a new type of astronomical CCD which should provide image motion compensation over arbitrarily large fields of view, very fast readout, autoguiding capability, good red sensitivity, and should be significantly less expensive than the present generation of CCDs.

Keywords: CCD, image motion compensation, tip-tilt correction, isokinetic angle

1. INTRODUCTION

Optical astronomical observations must contend with distortion of the images by the atmosphere and telescope which usually are much larger than the diffraction limit of the aperture. This is a very serious limitation on the detectability of objects; the signal to noise of a faint point source goes inversely as the size of the PSF, and working to a given S/N therefore requires an exposure time which goes as the square of the PSF. Atmospheric turbulence is quite well described by a Kolmogorov spectrum¹, which causes a scale-free phase distortion of the wavefront which grows in variance as length scale to the 5/3 power. This scaling means that the lowest moment of the distortion – the overall tilt of the wavefront – in angular measure is almost independent of telescope aperture, and therefore problematic for all telescopes. Roughly speaking, the total size of an integrated image has a 60% contribution from overall image motion (wavefront tilt) and 80% contribution from higher order distortions which add in quadrature.

Removal of this image motion is therefore extremely important, and many instruments have been built to do this (cf Tonry, Burke, and Schechter² and references therein). However, a complication arises for wide fields of view. Because the telescope pupil for widely separated directions samples different turbulence, the image motions are uncorrelated, and it is not possible to make a single offset in the focal plane to compensate for both motions. This “isokinetic angle” where the image motions decorrelate to some fiducial level (e.g. 70%, where correction of one location by another does no good) depends on the telescope diameter and height of the turbulent layer, and is typically taken to be $\sim 1'$. What is required to correct for image motion over larger fields of view is some sort of “rubber focal plane”.

This article has three sections. The first describes a camera we have commissioned which is capable of measuring and correcting for image motion over angles up to $12'$. The second reports on results obtained over a fair sample of conditions on Mauna Kea during Spring 2002 which suggest that estimates of the isokinetic angle may have been overly pessimistic. The final section describes a new device and camera we are in the process of designing which should make this concept fully scalable, but also improve the speed and cost to the point where it is feasible to build mosaics with 1–3 Gpixels.

Further author information: (Send correspondence to J.L.T.)

J.L.T.: E-mail: jt@ifa.hawaii.edu, Telephone: 1 808 956 8701, Address: Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA

2. THE ORTHOGONAL PARALLEL TRANSFER IMAGING CAMERA (OPTIC)

The heart of the OPTIC camera is the MIT Lincoln Laboratory (MIT-LL) CCID28. This is a CCD designed by the authors using funding from MIT-LL as well as the NSF. The CCID28 has 2K x 4K 15 μm pixels, with serial registers at both 2K ends and amplifiers at each corner. The CCD is mirror symmetric around the 4K centerline, and each of the 2K x 2K halves have separate clocks for the “lower” 512 rows closest to the serial registers and the “upper” 1536 rows closest to the centerline. The CCD, therefore, has effectively four “cells” which can be clocked independently. The CCID28 sits on a molybdenum package and is buttable to $< 0.5\text{mm}$ on the 4K sides, and buttable to within 5mm on the 2K sides. The bond wires from the CCD go to a piece of AlN which has traces running around to load resistors, follower JFETs, and a Nanonics connector all lying underneath the footprint of the CCD.

Each serial register is read out of one amplifier, for a total of four video outputs from the focal plane. We use SDSU-2 electronics for readout, with 3 clock boards providing the 32 parallel clock signals, 20 serial clocks, 4 reset gates and 4 serial dump gates. Two video boards handle the four outputs, a utility board operates the shutter, filter wheel, temperature control, pre-flash LEDs, and a timing board carries the DSP which runs the whole operation. We currently run with a CDS with 1usec integration on signal and reset for a net pixel rate of 125 kpix/sec at 4 e- read noise.

2.1. Orthogonal Transfer

The unique feature of the CCID28 is the structure of the pixels. These are laid out in as an “orthogonal transfer” design, as described by Tonry et al². Figure 1 illustrates how the gate layout permits charge movement in all directions. Therefore the developing electronic image within the silicon can be clocked during exposure to follow




Figure 1. The layout of the orthogonal transfer gates is symmetric to 90° rotations, and so the charge can be 3-phase parallel clocked in the vertical direction or the horizontal direction by setting the appropriate vertical or horizontal gates negative to act as a barrier.

an optical image which may be dancing around on the detector.

This process can cause blurring, but since the offset between photons and electrons never needs to be larger than 0.5 pixel, this blurring is equivalent to convolution with a 1-pixel PSF. Non-perfect charge transfer efficiency can also cause blurring, but with a CTE of more than 0.99999 and typical exposures of a few thousand shifts this creates a very small skirt (few percent at most) of charge around a given charge packet and this is lost in the optical PSF. OT clocking can also create spurious charge if clock voltages are excessive, but since a normal readout involves 2K shifts anyway, it is not significant. The advantages conferred by OT clocking include some amount of clocked anti-blooming (saturated stars tend to create a mesa of charge which never overflows the pixel boundaries), and automatic dithering of CCD defects or obscuring particles.




Figure 2. The OPTIC camera consists of two CCID28s next to each other in a cryostat. These two devices therefore offer 4K x 4K pixels, and the plate scale at the UH 2.2-m is 0.138" per pixel so the camera subtends about 10' on the sky. The four "lower" regions of 0.5K x 2K are designated as "guide regions", and the four "upper" regions of 1.5K x 2K are designated as "science regions".

2.2. Camera and Operation

Figure 2 shows how two CCID28 devices have been deployed in a cryostat. Operation of OPTIC first requires a quick snapshot of the four guide regions in order to locate guide stars. The observer is presented with an image and selects zero or one guide star in each of the guide regions for a total of 1–4 guide stars. The shutter is then reopened, a few video frames are taken of the stars, the precise centroid of each is determined, and then the shutter is closed and the arrays are cleared. The observation then commences when the shutter is opened, and OPTIC goes into its basic video loop.

The first stage of the video loop consists of a "shutterless video" readout of a small region around the 1–4 guide stars, typically 32x32 pixels binned 2x2. This is done simultaneously for up to four stars, with a horizontal and vertical parallel transfer to quickly bring each of the 1–4 regions down to the amplifier, and then a standard subarray readout of 1–4 guide regions. The science regions and any guide areas which are not currently being used for guiding are left quiescent. Although the shutter is open throughout this process, the speed of the parallel transfer and readout, and the preeminence of the guide stars virtually guarantees that the guide star images will be uncontaminated by other stars in the field. This process is reasonably fast, requiring up to 1–10 msec for parallel transfer depending on whether guide stars are close to or on the opposite side of the CCD from an amplifier, and 2–5 msec for the readout.

The guide star images are then analyzed by the host computer for centroid, width, and flux. The time sequence of each guide star's position is analyzed to predict where it will be at the next iteration. We have found that an algorithm which finds coefficients which predict observation N in terms of a weighted sum of observation N-1, N-2, etc. works well. The coefficients are periodically determined from the earlier observations, and a three coefficient algorithm does a good job even on periodic motions with only a few samples per period.

Once the host computer has determined the predictions of the guide stars, it rejects bad predictions based on discordant offsets from multiple guide stars (if there is more than one) or on recent positions of a single guide star. The offsets are then used to update the location of the guide box for each guide star.

The offsets of the guide stars are then used to make a prediction of how the next iteration's optical image will differ from the accumulating image in each of the science (or unused guide) regions. We use various algorithms here: use closest guide star, use various averages of the guide stars, based on their distance from the center of the science region, or use time-lagged averages of the guide stars. These predicted shifts are then applied to the science regions by performing orthogonal transfer shifts to bring the accumulating image into registration with the optical image.

Telescope guiding is also carried out by periodically applying to the telescope drives a leaky average of the average offsets of the guide stars. The host and DSP then pause for some number of msec to fill out the guide

star exposure time and the process iterates. The rates at which we guide depend on the brightness of the guide stars and the signal to noise we want for the guide stars, but typically range between 30 and 5 Hz. It is possible to run as fast as 100 Hz with guide stars near the amplifier.

When the desired exposure time has elapsed, the shutter is closed and the CCD read out. There are four basic files saved: the science observation, all of the video frames of the guide stars as a 3-D FITS file, a table of all the guide star analysis, and a table of all the OT shifts applied to the science regions. The only complication in subsequent reduction is that each observation has undergone a unique pattern of OT shifts. We have a utility which will convolve an unshifted flatfield with the shift pattern from the shift table of a given observation. This does a very good job of flattening, but obviously causes problems for creating a “superflat” out of the sky from all of the observations during a night. OPTIC requires that high signal to noise flats (e.g. dome flats) be used to remove the pixel-to-pixel and OT shift pattern from all images. They can then be shifted into pixel registration, a superflat process applied, and a second round of flattening performed if necessary.

We have not yet pressed OPTIC to see how faint it can stably carry out guide operations, but scaling from present performance we expect to be able to work usefully to $m = 18$ on the UH 2.2-m telescope. At these magnitudes, the probability of finding a guide star within one of the four guide regions is very high, even at the galactic pole.

During the exposure, the observer is presented with a video display which shows the image of each guide star in its box, a “radar” display which shows how the images are moving around on the CCD (even though the readout boxes are adaptively following the stars), a leaky average, a strip chart of the image FWHM, and a strip chart of the magnitude of the stars. The latter two are particularly useful for monitoring the quality of the seeing, transparency, and telescope focus.

3. IMAGE MOTION

We expected that tracking image motion at modest rates (5-10 Hz) would lead to a significant improvement in image quality, and indeed found that the motion compensated images were approximately 80% the width of images with only telescope guiding at 0.5 Hz or so. It was not uncommon to approach or hit the optical floor of the UH 2.2-m telescope ($0.4''$, mostly triangular astigmatism). We were surprised, however, to find that the image quality did *not* vary appreciably with distance from the guide star. Previous observations² at MDM Observatory had hinted at this, but had only demonstrated it to a modest separation of $2'$. This camera is capable of covering a full $10'$ separation between guide stars, and we never saw the characteristic radial elongation of image motion which heralds the edge of the isokinetic patch.

In order to investigate the degree to which motions of distant stars are correlated, we searched the USNO-A2 catalog for all constellations of bright stars where all four guide regions would have a star. At $m < 12$ these constellations are rare, but for $m < 13$ they are plentiful. A good example is f1121-0852, which has two pairs of stars separated by about $2'$, and about $8'$ separation between the pairs.

Over the course of nine nights throughout Spring semester 2002 we imaged these constellations and recorded the image motion at typical rates of 15–30 Hz. We analyzed these videos to look for correlations between the image motion of the various pairs, the angle at which the motions decorrelate (isokinetic angle), and for correlations with non-zero time lag.

It was immediately obvious that the images *always* show a striking degree of correlation in their motion, regardless of separation up to $12'$ or so. Figure 3 shows typical motion of the stars in the f1121-0852 field. Of course, any flaw in the telescope pointing from bad tracking, wind or dome shake will give us a correlated motion in all stars imaged, so this result is not very significant of itself. However, the *differential* image motion comes purely from atmospheric distortion. Figure 4 illustrates this differential motion of the stars relative to the first one. There is a modest amount of decorrelation at $8'$ relative to $2'$, but the overall amount of image motion is quite small. Even at $8'$ the differential image motion is less than $0.18''$ FWHM, which will degrade $0.4''$ seeing by only 10%, whereas with no correction the seeing would degrade by about 20%. At $2'$ the degradation is only 4%.




Figure 3. Image motion of $0.32''$ and $0.38''$ FWHM in x and y seen between stars at separations of $8'$, $2'$, and $8'$ (top to bottom with respect to star on the bottom). The overall image size during this exposure was $0.5''$, or $0.4''$ with image motion removed.




Figure 4. Differential motion of $0.17''$ FWHM (stars at $8'$ separation) and $0.11''$ (stars at $2'$).

3.1. Image Motion and Image Size

If we look at all pairs of stars in all constellations observed throughout the nine nights spread through semester we find that the median seeing was about $0.7''$, as is commonly reported at Mauna Kea. Figure 5 shows the seeing and differential motion from all the observations this Spring at the UH2.2-m telescope.

However, the differential image motion was far less than what we would have expected if the star pairs were undergoing independent Kolmogorov turbulence. In that case we would expect the differential image motion to be something like $3/4 \sqrt{2}$ times the intrinsic image size. In fact, given the amount of differential motion, if we could somehow remove the image size which occurs because of correlated wavefront distortion, the median seeing at the UH 2.2-m would be about $0.3''$, not $0.7''$.

What is going on? We expect that the wavefronts of guide stars to become uncorrelated when the telescope pupil to each passes through different regions of turbulence, i.e. in the simple case that a single layer of turbulence exists at a height h , the star motions should decorrelate at an angle D/h , where D is the telescope diameter. The high degree of correlation at greater than $10'$ means that the height of the majority of the turbulence is very much less than 500m altitude. It is possible that we are seeing a boundary layer at a few 100m altitude, but it seems more likely the problem is local to the telescope and is related to the thermal environment in the dome. Presently the 2.2-m mirror typically has a temperature of $T \sim +1^\circ\text{C}$, the closed tube $\sim -2^\circ\text{C}$, and the ambient air $\sim -1^\circ\text{C}$, which is a certain recipe for generating plumes of rising air. Ventilating the tube with a fan, and the short period at the start of a night when the tube has not cooled below the mirror

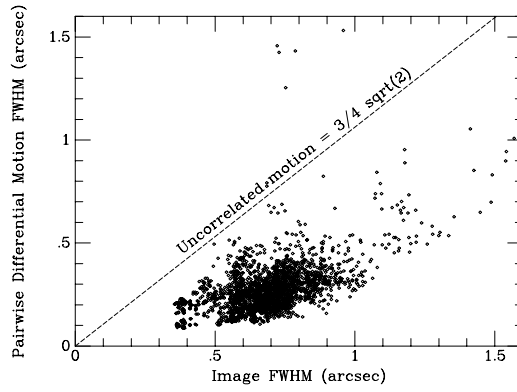


Figure 5. Sample of seeing and pairwise differential image motion over nine nights at the UH2.2-m.

help the seeing; the seeing does seem to show some sensitivity to the orientation of the dome with respect to the wind. These suggest the UH 2.2-m has a very bad thermal environment and that the seeing is primarily generated in the dome or the telescope tube. The newer generation of telescopes on Mauna Kea such as Subaru and Gemini are starting to report much better median seeing than has been supposed to be possible on the older telescopes.

The Kolmogorov turbulence model and the supposition of thermal problems at the UH2.2-m may or may not be correct. What is undeniable, however, is that images become much sharper when their motion is corrected according to a guide star as distant as $12'$. Figure 6 shows the image improvement from motion compensation as a function of separation between guide stars. Typically we get a 15–20% improvement in image size by

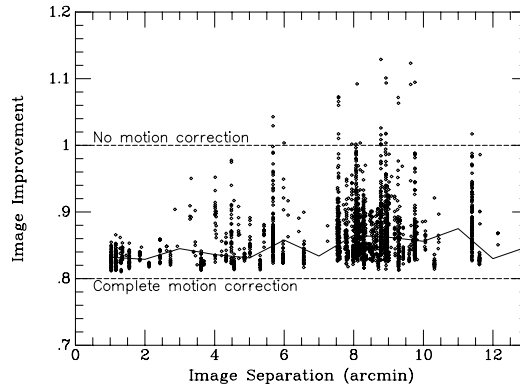


Figure 6. The image improvement as a function of guide star separation.

application of reasonably rapid image motion correction. This may seem like a modest gain, but making a psf 20% smaller corresponds in signal to noise to increasing the quantum efficiency or throughput of the system by 50%, and the gains are even larger in crowded fields. Image motion compensation is very much worth doing when the cost does not exceed 50% of not doing it.

4. GIGAPIXEL FOCAL PLANES

We anticipate funding to build a telescope facility known as the Panoramic Optical Imager (POI) (also called the Panoramic Survey Telescope and Rapid Response System PanSTARRS). This will be an array of four telescopes, each of aperture 1.5–2m and equipped with a Gpixel focal plane, which will survey the entire visible sky approximately every week. We will build a software pipeline which will monitor the data in real time

for everything which moves (near Earth “killer” asteroids (NEOs), comets, other asteroids, KBOs) or blinks (supernovae, variable stars, gravitational lenses). We will also build an image of the sky of unprecedented quality which can be used for weak lensing, counts of stars and galaxies, discovery of exotic classes of objects, etc. The first year will be spent carrying out design studies and building prototypes, the second and third year will be telescope and detector construction, the fourth year will be integration, and we expect to be doing science by the fifth year.

Our scientific interests along these lines have been endorsed by the Astronomy and Astrophysics Decadal Survey which calls for a “Large Synoptic Survey Telescope” (LSST) which has a 6–8m effective aperture, 7–9 sq. deg. field of view, and regardless of optical design, a 1–3 Gpixel detector. Our POI project will build a system of 3–4m effective aperture, and as such is *not* intended to be the full LSST. However, should we prove successful, it is possible to build more telescope/detector units to achieve the full aperture desired at a very predictable cost and schedule.

4.1. Detector Economics

It is a very sobering prospect to build a focal plane with ~ 4 Gpixels. Currently CCD detectors cost approximately 1 cent per pixel (thinned, science grade 2K x 4K devices can be bought for \sim \$80k). Infrared arrays hybridized to a silicon readout multiplexer are at least a factor of 10 more expensive per pixel (the Rockwell 2K x 2K HgCdTe devices are sold for \sim \$500k). In addition, the total cost of the current crop of “megacams” which offer 0.3 Gpixels is more like 2-3 cents per pixel partly because of the cost of engineering and controller electronics. Even at 1 cent per pixel for a detector, the cost of 4 Gpixel is \$40M and this is prohibitively high for the system we envision.

Our proposed detector should permit us to reduce the detector cost to approximately 0.1 cents per pixel, offering 4K x 4K pixels at a cost of \sim \$20k. The total cost of each of our Gpixel focal planes would then add up to \sim \$2M once cryostat, shutter, filters, controller, and assembly are added in, with an additional \sim \$2M of NRE cost.

4.2. Proposed Improvements

We detail below a number of areas where we expect to improve the detector cost to performance. First and foremost, we expect to improve the yield of workable detectors, since the cost per detector goes nearly inversely as the yield. We will decrease the pixel size as much as possible, since the cost of CCDs is roughly proportional to the area of silicon, not to the number of pixels contained within. We will remove image motion, since a 20% improvement in PSF width improves signal to noise as much as a 50% improvement in quantum efficiency. We will reduce the readout time to approximately 2 seconds, so that the duty cycle of the detector will be nearly unity. Finally we have the luxury of designing the telescope, CCD, the package, the cryostat, the readout electronics, and the host computer all at the same time, so we will pay close attention to the costs and scalability of our total system.

4.2.1. Yield

The probability that a CCD from a given silicon wafer will turn out to be operable depends in large part on the absence of very localized, catastrophic defects such as gate to substrate shorts or bad amplifiers. For example, the engineering grade CCID28's currently in the OPTIC camera are poor detectors for science because of three tiny gate-to-gate shorts which cause massive bright defects. In order to operate these devices one cell of each array has to be “turned off” by reducing the gate-to-gate voltages below the point where the device can be clocked effectively.

It is possible, of course, to make very small devices (e.g. 1K x 1K) and then reject the ones with defects, but this leads to very rapidly rising packaging and focal plane costs, and the metrology of the focal plane also becomes highly unstable. Current CCD foundries do a very good job on yield of 2K x 4K parts, typically achieving yields of 25% which means ~ 0.1 defect cm^{-2} . To increase yield we need a way to isolate defects without losing the entire device.

4.2.2. Pixel Size

Given the freedom to design the optical telescope, we can push the size of a pixel which samples the PSF quite small, perhaps as small as 5 μm . However, we want to have quite thick devices for good performance in the near IR (800-1000nm). Ideally we would like to see 40–50 μm thick, fully depleted, high-resistivity silicon. On such devices charge diffusion may negate the value of small pixels, however. Experiments by UH, ESO, MIT-LL, and others find charge diffusion of ~ 5 μm rms. In addition, small pixels and standard NMOS lithography increase the probability of design traps or yield problems. At present we know that 12–15 μm pixels yield well, we think that 8–10 μm pixels may be possible, and we believe that less than 8 μm pixels are unlikely in a thick, deep depletion CCD. Other technologies such as CMOS have no difficulty with smaller pixels, and use of a vertical electric field can improve charge diffusion, but we do not believe that these technologies are quite mature enough to give us detectors on the 1–2 year time scale necessary for POI.

4.2.3. Motion Compensation

Image motion compensation is extremely important, since approximately 20% of an uncompensated image comes from motions which are coherent over as much as $10'$. Various schemes have been used to remove image motion optically, including small field of view cameras with tip-tilt mirrors, but the most ambitious to date is the transmissive tip-tilt plate for the 1-degree FOV CFHT Megacam. It remains to be seen how well it works but one manifest limitation is that it can only correct the entire field of view at once. No one expects image motion from the atmosphere to be coherent across 1 degree, much less the 3 degrees planned for POI. Instead what we need is a “rubber focal plane” which can effectively stretch and move to follow differential image motion across an arbitrary field of view.

We believe that the orthogonal transfer technology can be used to create a rubber focal plane. It is a very fast way to shift charge to follow a moving optical image, and if the focal plane is divided into individual cells which are smaller than the isokinetic angle we can follow image motion over an arbitrarily large field of view.

4.2.4. Readout Speed

Fast readout is essential. With the SUPRIME camera on Subaru, for example, the total readout time of 100 seconds is nearly equal to the time it takes the detector to saturate in the Z band, so the 8.2-m telescope is reduced to less than 6-m effective aperture because of readout overheads. Telescope time is simply too expensive to waste it with a closed shutter. In addition, a large part of the rationale for LSST and POI is to find NEOs which move quite rapidly.

Once a moving object has moved more than one psf width the detection signal to noise does not improve with exposure time. For NEOs this can happen in as little as 20 seconds. We therefore feel that a readout in ~ 2 seconds is essential. This requires us to have a lot of signal chains. We are willing to tolerate $4e^-$ of read noise, but cannot accept much more without becoming dominated by read noise instead of sky noise in a 20 second exposure. It is possible to run a single amplifier at 1 Mpix/sec with $4e^-$ noise, but it is unlikely to be able to run much faster. Therefore reading out 1 Gpixel in 2 sec at 1 Mpix/sec requires 500 amplifiers and signal chains, or one amplifier per 2 Mpixel. As a benchmark, the CFHT Megacam reads out in about 20 seconds using 80 amplifiers, so we need nearly an order of magnitude more readout channels and faster readout.

5. THE ORTHOGONAL TRANSFER ARRAY

We believe we can achieve these goals by building monolithic devices which are broken into an 8x8 chess board of independent orthogonal transfer CCDs (OTCCDs). A given device might have 4K x 4K pixels total, and be comprised of 64 OTCCDs which are each 512 x 512 pixels.

We call this device the Orthogonal Transfer Array (OTA). Figure 7 illustrates the composition of an OTA in terms of cells, and cells in terms of OT pixels. Each cell in this device shares common serial clocks and biases (appropriately buffered for isolation), but there are two independent sets of parallel clock lines – active and standby. Row and column addressing lines determine which set of parallel clocks a cell is connected to. When connected to standby, the charge does not move in a cell. When connected to active, the charge can be clocked for readout or for OT shifting to follow image motion. Likewise, the cell addressing also determines whether a

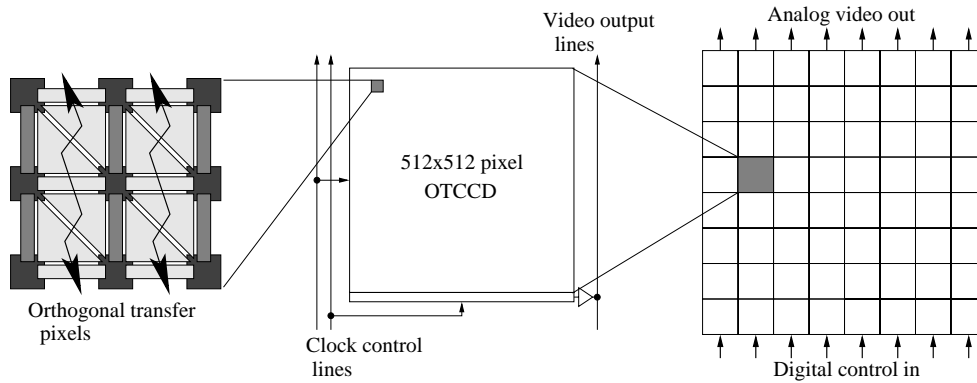


Figure 7. Layout of the OTA, showing OT pixels which can clock charge horizontally or vertically, a single CCD cell with its addressing and clock lines, and a full OTA of 64 OTCCD cells.

cell is connected to an output bus line or not. Readout of the entire array would normally proceed row by row, with each cell in a row connected to the column video bus line for readout. The connections to these cells are via metal lines which run in 100 um gaps between the devices or over the gates.

As with any device of this size, we expect to have catastrophic defects in these OTAs, but now we can isolate the defects to individual cells. With a typical defect rate of 0.1 cm^{-2} , a $5 \times 5 \text{ cm}$ OTA would have 2–3 dead cells on average. These dead cells (and the gaps between the cells) would simply be accepted as dead area, and dithered away by the four telescopes of POI and/or multiple observations with small shifts.

We therefore expect the yield of these devices to be significantly greater than 50%, and we expect to be able to get 30–40 per lot run, which is how we hope to achieve our goal of less than \$20k per device. We note also that bad columns are confined to a cell, that the OT pixel structure does not bloom for bright stars nearly as badly as three phase, and that the OT shifting which is a normal component of observation does quite a bit of clocked antiblooming.

Nevertheless, there will be very bright stars on the array (a given POI field will have a magnitude 6 star and many magnitude 10 stars). Our plan is to write off these cells for science and instead rapidly read out a little patch around these bright stars and use them for guiding. We therefore simultaneously dispense with the enormous bleeding of charge from bright stars and acquire a distributed guide signal over the focal plane. These cells will subtend 1–3' on the sky, depending on our precise choices of pixel size and focal length, and independent OT shifting of each cell gives us a rubber focal plane. Figure 8 sketches out how one of these OTAs might look and suggests a typical dispensation of cells for various uses.




Figure 8. Typical use of an OTA, showing dead cells, cells devoted to guide stars, cells devoted to science, and bad columns.

The readout of this device should meet our goal of 2 seconds, since we can use 8 amplifiers simultaneously (one per column). We can wrap the serial register around underneath each device to the output amplifier structure, and add a follower JFET for buffering onto the output bus and addressing logic. We expect to keep

the gaps between the cells to $< 100 \text{ um}$, which means that we will have a fill factor in the focal plane of close to 97%. We have experimented with some NMOS structures to carry out NAND and NOT functions, and it appears to be practicable.

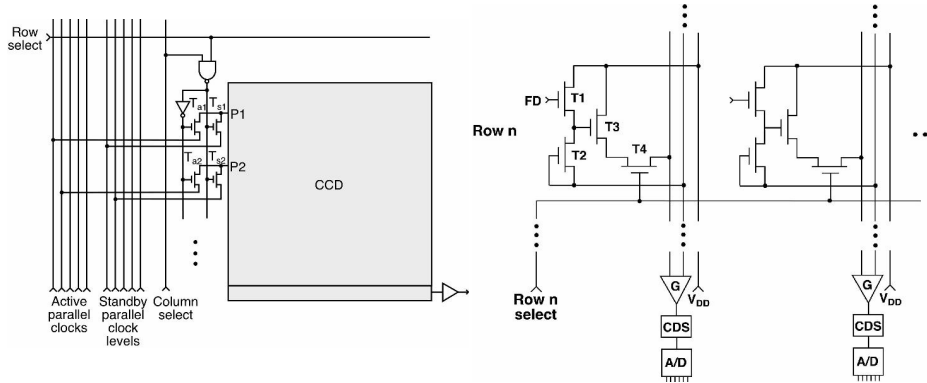


Figure 9. Sketch of how the OTA cells can be connected to active or standby clocks and output buffer lines.

Use of orthogonal transfer shifting to remove image motion is a very attractive option, since it allows us to remove image motion from atmosphere and telescope, and substantially reduces the tolerances on telescope tracking. Even if the OT pixel topology has a lower yield than three-phase (our most current CCID28 lot suggests that this is not the case), the defects would not cause a problem with an OTA since they would simply lower the number of operable cells slightly.

The OTA package needs to be four side buttable, and we have a design using a custom multilayer AlN substrate bringing the bond pads to a pair of Nanonics connectors or a pin grid array underneath. The package sits on a 3-point pseudo-kinematic mount which can accept precision spacers to bring the device to a desired focal plane. We will use flex circuits to attach to the package and provide thermal isolation. We have also had success in passing the flex circuits through the cryostat and potting them while maintaining vacuum integrity.

The POI focal planes have so many video channels we will be forced to a new generation of controller, but fortunately the electronics and computer industry has advanced well beyond the state of the art of current astronomical controllers. For example, the SDSU-2 video board has two channels, is capable of up to 500 kpix/sec with a correlated double sample and a 16-bit A/D, and it consumes about 15W of power. In contrast, the Analog Devices 9826 part is designed for HDTV and has three channels (nominally RGB), manages 15 Mpix/sec with a CDS and 16 bit A/D. The power consumption is 250 mW, and it is more than two orders of magnitude smaller than an SDSU-2 video board. The possibility of using an ASIC for some or all of the device control and interface is extremely attractive, and we recognize that this is probably going to be the way things are done in the future, but we do not believe that ASICs will be available early enough or at a low enough development cost for POI.

Barry Starr has spoken during this conference about his plans for a “Monsteroid” controller³ which has an open and scalable architecture. Under this scenario, four OTA’s (an 8K x 8K portion of the focal plane) would be run by a quad interface unit which in turn communicate with a host computer via a Gbit fiber. Any number of these Quad OTA (QUOTA) units can be chained together, with a synchronization connection keeping readouts in lockstep.

The final Gpixel imager we envision (POICam? PanCam?) has 64 OTAs, making a 32K x 32K focal plane spanning 16 inches (40 cm) on a side and covering 8 sq. deg. It will have 512 signal chains, a cluster of host computers, a software pipeline, etc. Each of four telescopes will have a POICam, and the data from each will be brought together, flattened, registered onto a common coordinate system, combined with cosmic ray rejection, and then analyzed to find objects with parallax (things in Earth orbit), objects which move, objects which blink

relative to a growing master image of the sky, etc. We hope to see a first deployment of one of these cameras and telescopes in 2005, and routine science operations in 2006-7.

The OTA is a realizable and cost effective solution for very large focal planes. Other advantages include the fact that bright stars become guide stars, a rubber focal plane compensates for image motion, and very fast readout. We do not know of any other technology which is likely to be able to fill our need, including large monolithic CCDs (too expensive and readout too slow), small individual CCDs (packaging and electronics become overwhelming and no image motion compensation), monolithic CMOS (poor red QE and non-filled pixels are problematic for marginally sampled PSFs), and hybrid CMOS (still too expensive). We are confident that the next decade will see great strides in imaging technology, and we believe that our proposed device has the right combination of conservatism and audacity to carry out its science mission.

ACKNOWLEDGMENTS

We acknowledge support for this work from NSF grant AST97-96150.

REFERENCES

1. Tatarski, V.I. 1961, "Wave Propagation in a Turbulent Medium", Dover:New York.
2. Tonry, J.L., Burke, B., & Schechter, P.L. 1997, PASP, 109, 1154
3. Starr, B.L., et al., 2002, proceedings of the "Scientific Detector Workshop", Kluwer.