

The Pan-STARRS Gigapixel Cameras

John Tonry
Peter Onaka
Sidik Isani

Institute for Astronomy, University of Hawaii

ABSTRACT

The Pan-STARRS project comprises a complex, interlocking set of sub-systems, of which the Gigapixel Camera (GPC1) was deemed to be the most challenging. We were able to design and build GPC1 (including detectors) in four years without significant false steps, and it has been operating for two years now without significant failure or downtime.

1. INTRODUCTION

The Pan-STARRS project comprises a complex, interlocking set of sub-systems, and the Gigapixel Camera (GPC1) itself is a complex, interlocking set of subsystems comprising detectors, cryostat, controller, and software, with considerable supporting infrastructure and interfaces to the rest of the project. Many details have been published elsewhere [1]. Below we describe each of these subsystems in turn. Fig. 1 illustrates the camera.




Figure 1: GPC1 in solid model and actual hardware.

Pan-STARRS’s mission placed major demands on the camera’s performance. Pan-STARRS can only survey the sky to the requisite depth by producing and critically sampling a sub-arcsecond point spread function (PSF) over a large area, which translates to a need for at least a billion pixels. The rapid cadence and the short exposures required to avoid trailing of asteroids mandated that the readout time of the camera be short, well less than the ~ 30 second exposure times. We wanted high red sensitivity for many science goals as well as read noise significantly less than the Poisson statistics from the sky. Other requirements such as mass, volume, stiffness, durability were also a challenge.

2. DETECTORS

The need to read out many pixels in a short amount of time without compromising noise performance dictates that we read out through many channels. We decided that a device that is 50×50 mm, subdivided into a checkerboard of 8×8 independent CCDs, dubbed “cells” would suit our needs. The full device is called a “orthogonal transfer array”, OTA. The focal plane and the OT function is shown in Fig. 2. The full focal plane has an 8×8 array of OTAs.

Our desire was to address all cells independently and simultaneously, but time and cost mandated that we find a way to multiplex the 64 cells from a single device into 8 channels only. Mike Cooper and Barry Burke at MIT Lincoln Laboratory (MITLL) designed NMOS logic into the gaps between cells on the OTA that allow us to address the parallel gates for charge transfer or else latch in standby gate potentials to hold accumulating charge fixed, and to selectively connect the output from a cell onto a bus common to an entire column of cells. We therefore have 8 output channels per OTA and read it out row-by-row. (Continued, ensuing development has brought us close to the point of being able to address all 64 cells independently.)

The science that Pan-STARRS wanted to pursue was heavily weighted to the red end of the spectrum,

and we therefore wanted as much quantum efficiency (QE) at 1 μm as possible. Of course the devices are back illuminated, meaning the silicon is thinned and light admitted from the to the back of the detector, away from the gates and buried channel. High QE in the visible infra-red is achieved by making the device as thick as possible because the photon absorption length becomes very long as the energies approach the bandgap. A very thick device may not be fully depleted if the electric field from the gate potential can be matched by removal of free charge carriers in the silicon, unless we apply a bias electric field by putting a potential on the back side of the detector. Even so, charge created at the back side will undergo some lateral diffusion as it migrates down to the charge collection channel, with an extent that goes as square root of diffusion time or square root of electric field. On the other hand, we opted for pixels of only 10 μm in order to sample the PSF well and to minimize cost per pixel. We therefore met the trade-off between thickness and substrate bias by choosing to thin the devices to 75 μm and apply a substrate bias of about -8 V . This gives us about 30% QE at 1 μm , charge diffusion of about 4 μm RMS, and a reasonably low level of bright defects whose current goes exponentially with substrate bias.

We give up some fill factor because of the gaps between cells, and although we have designed the package to be 4-side buttable, there is inevitable loss in the gaps between devices. Our goal was to exceed 90% fill factor that we have not quite achieved because of excessive defects in the devices we were able to produce.

Our production consisted of a test lot of wafers, the CCID45, two production lots of wafers, the CCID58, and access to a third production lot. The devices work quite well, although there are some shortcomings. The read noise is about 6 e^- at 0.5 Mpix/sec where we expected 4, the incidence of bad cells, bright defects, and the yield were worse than we had hoped, and Lots 1 and 2 of the CCI58 had a problem with radiation damage during etching that badly degraded the charge transfer efficiency and dark current at the outer parts of the wafers, manifesting itself as a “bad corner” in each device. The make-up, third, production lot permitted us to replace the 20 worst devices in the focal plane. Despite the imperfections we have an excellent focal plane for science.

The CCD package is a relatively thin piece of molybdenum that carries an alumina ceramic circuit tucked underneath between three mounting pads. The wire bonds from the OTA attach to the ceramic on top and a pin grid array on the bottom attaches to a flexprint circuit. We have found that the molybdenum is too thin and too poor a CTE match to silicon, so the packaged OTA bows upward when it cools. We therefore dish each molybdenum surface before silicon attachment so that it will be flat at the operating temperature. We require an overall flatness of $\pm 20\text{ }\mu\text{m}$ in order not to compromise the quality of the focus. The thermal problems and the difficulty in adjusting the three mounting pads require that we improve the package for GPC2.

3. CRYOSTAT

The dark current of a CCD is far higher than the night sky at room temperature, so the CCDs must




Figure 2: Carbon Fiber focal plane (left), fully populated GPC1 focal plane (middle). The four gates in an orthogonal transfer CCD permit image correction in all four directions (right).

be cooled to about 200 K. This, of course, necessitates installation of the CCDs in a vacuum cryostat. In principle such a cryostat is straightforward. In practice GPC1 had some unusual challenges because of the very large focal plane (400 mm), the need to make the cryostat both light and stiff, the need to get thousands of electrical connections out of the cryostat, and the use of in-situ diagnostic apparatus with which to test and maintain the health of the system. The optical design of Pan-STARRS placed the last lens about 200 mm in front of the detector, and this is used as our cryostat window. It sits in a lens cell on a seat that has an oring as a vacuum seal.

Below the lens cell is the “Antechamber” that serves as a spacer as well as the attachment to the telescope. The telescope interface is via the “lower cass core” (LCC) which is a welded steel structure with four feet. Since the cryostat is aluminum we needed a flexure between the LCC and the Antechamber, and the LCC also required us to provide a stiffening tie between the feet. This is the function of the “iron ring” (FeRIT) that is guided into connection with the LCC via a pin and slotted pin, and that has four titanium flexures and aluminum posts in each corner to the Antechamber.

The third large component is the GPC1 body, that carries the devices on a carbon fiber, honeycomb plate, again attached to the aluminum body via flexures in order not to create thermal stress. The GPC1 body has 16 rigid-flex boards, potted through opposite side plates, that provide the electrical attachments for each device. On the other two sides we mounted a pair of CTI-1050 helium cryocooler heads to cool the devices.

Unlike most cryostats, GPC1’s focal plane is so large that it is impractical to mount the devices on a plate that can be considered isothermal — the radiative heat loads are uneven and many heat paths would have to be provided to obtain anything near isothermality for the devices. Instead, we chose to build the carbon fiber support (Fig. 2) so as to insulate the devices from each other as much as possible, and then use an independent thermal path for each to the CTI heads. The heat load on each device is about 2 W (approximately evenly split between radiative load and power dissipation by the amplifiers), and we therefore need a considerable thermal resistance between device at 200 K and CTI head at 70K. We regulate the temperature on each device by changing its power dissipation by altering drain voltage when the amplifier is not called on for readout. The coldest parts of the thermal path are behind a single layer of radiation shield.

The vacuum integrity of the cryostat is quite good — it will hold below 1 mtorr (i.e. negligible conductive heat load) for many months. Steady vacuum degradation is inevitable because of permeation through the many orings (we have about 400 inches of oring in total), but since the primary gas that diffuses is water, we use a room temperature zeolite getter, as well as 30 g of activated charcoal at 80 K.

We have equipped GPC1 with a number of important auxiliary features. We have a deployable Shack-Hartmann head that can be brought out from the side into the field of view, casting a Shack-Hartmann set of spots onto the focal plane when a star is guided into its entrance aperture. On the head we have put 9 colored LEDs and a photodiode, and we have provided a means of deploying an Fe-55 xray source if desired. These are extremely helpful for measuring and optimizing amplifier gain and CCD performance within the cryostat, on the mountain. We also use a novel optical arrangement of lens and calcite in two corners of the focal plane that creates above- and below-focus adjacent donuts out of each star falling within their field of view. This instantly tells us the quality of focus when an image has been read out. Lastly, we have placed a few pinholes on the Shack-Hartmann head that can image the optical train above the camera and detect glints and stray light paths.

4. CONTROLLER

Because we read out 512 channels simultaneously, and because of some of the unusual features of the OTAs (four gates for parallel transfer, logic signals) we needed to develop a custom controller for running the CCDs. Our goal was to build a controller that could “live the the shadow” of the devices it services, permitting infinite tiling. As it happened, the telescope clearance was small enough that we needed to bend the controllers to the side instead of packing them underneath, but in principle we could make an arbitrarily large focal plane. The layout of the focal plane and the connection to the controllers is illustrated in Fig. 3.

Each controller chassis is about $4 \times 7 \times 12$ inches, and includes 4 board sets, each of that can run 2 OTAs, as seen in Fig. 4. Thus a single chassis runs a 2×4 block of OTAs. The signals emerging from a 1×4 block

of OTA flexprints pass through a rigid-flex board, through the cryostat wall, and then to a pair of board sets.

The board sets comprise a preamp board, a “DAQ3U board” that provides biases and clocks as well as AD9826 A/D converters, and a “FPGA board” that has a Virtex 2-PRO FPGA, enough laptop DDR memory to store a readout, and a gigabit ethernet interface. These are shown schematically in Fig. 5 and the boards themselves in Fig. 6. A power distribution board below the rigid-flex provides power to the chassis, and the entire assembly of four board sets is guided by a pair of rods. Some effort was expended to ensure that the precision and tolerances were adequate that the entire chassis could be safely plugged and unplugged.

The controllers on the side of GPC1 are cased in a boxes on each side that draw in cooling air from the bottom and eject it through a glycol-cooled heat exchanger on the top, thereby delivering all the heat from the controllers into the glycol stream instead of the dome air.

Our AD9826 A/D converters are 16 bit, 10 MHz, three color devices. At present we use just one channel and sample each of pedestal and signal four times, as seen in Fig. 5. The averaging and subtraction of these data values are performed in the FPGA, and a single, correlated double sample value is staged and passed via ethernet to our “pixelservers”. We regard the use of commodity memory and ethernet to be a significant strength of our “STARGRASP” controller, since it creates a lot of flexibility at low cost, and uses a standard,




Figure 3: Focal plane, rigidflex, and controller layout.




Figure 4: Controller and OTA form factors (dimensions in inches).

non-real-time interface protocol for communications. We commonly run a controller chassis with essentially any, cheap laptop computer.

The ethernet traffic from the controllers is passed down individual 1GE fibers to a 10GE switch. Each controller has its own IP address, and the switch allows us to make an essentially arbitrary mapping of pixel servers to controller boardsets. Fig. 7 shows the basic layout

5. SOFTWARE

GPC1 is an integrated system, and the software is at least as important and difficult as the hardware. There are five layers of interfaces to the GPC1 Pan-STARRS camera system, each built on the next layer below:

- A. Script Level

The highest level of orchestration is currently handled by scripts, primarily for ease of development. The scripts use toolkit level commands (often, many in parallel) to control a group of controller boardsets.
- B. Toolkit Level

The STARGRASP toolkit consists of C programs. Toolkit programs can be invoked on any host with the ability to connect to the controller using IP over Ethernet. The most important tools in the toolkit are those which receive log messages, boot the controller, send arbitrary commands, and retrieve FITS images. Additional engineering tools exist to visualize and optimize noise performance and generate clocking patterns.
- C. STARGRASP C Libraries

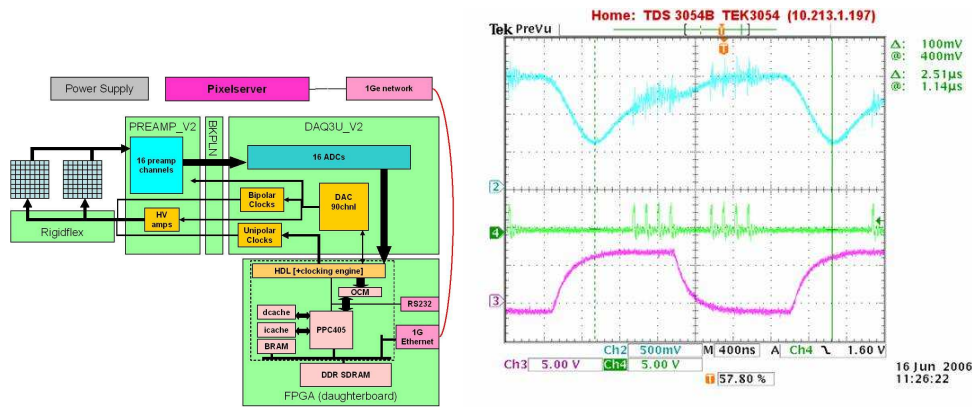


Figure 5: STARGRASP board set block diagram (left), and a multiple sampling example (right).




Figure 6: Preamplifier, DAQ3U, and FPGA boards.

Each of the toolkit tools are built on a set of C API Libraries. If in application is adversely affected by processes starting and stopping, a monolithic or persistent server process can be easily implemented using these APIs.

- D. Controller Socket Commands

The STARGRASP C libraries must accomplish all controller communication through a TCP socket on the controller’s port 915 (by default.) The majority of these commands are sent by using `grasp_cmd()` (the C function, or the toolkit tool.) `grasp_cmd` is a pass-through to the extensive list of commands that the controller supports at the socket level. For diagnostics, many of these commands can also be typed manually if a serial console is connected to the controller, and most can be sent interactively with “telnet” or “nc”. A command exists to control each analog and digital function of the controller.

- E. Registers and Clocking Instructions

Registers and Clocking Instructions are the lowest level API. All controller socket level commands are implemented in C code which modifies these registers, and/or writes instructions to the clocking engine processor.

Within the FPGA we run HDL which supports constructs called the “math engine” and “clocking engine” (Fig. 8). The Virtex2Pro FPGA operates at 2 levels. The PowerPC405 is running a kernel that allows for high speed control and data communication over gigabit Ethernet. This level of software also interfaces and controls a lower level of embedded functionality that provides high speed programmable clocking of the OTAs. This hardware description language based “Clocking Engine” design is a command FIFO structure that executes the following command set and provides repeatable, precise timing:

- Delay for precise amount of time
- Wait for external trigger (sync)
- Set voltages (biases, clock levels)
- Do 1 or more parallel clocking sequence
- Do 1 or more serial+video clocking sequence

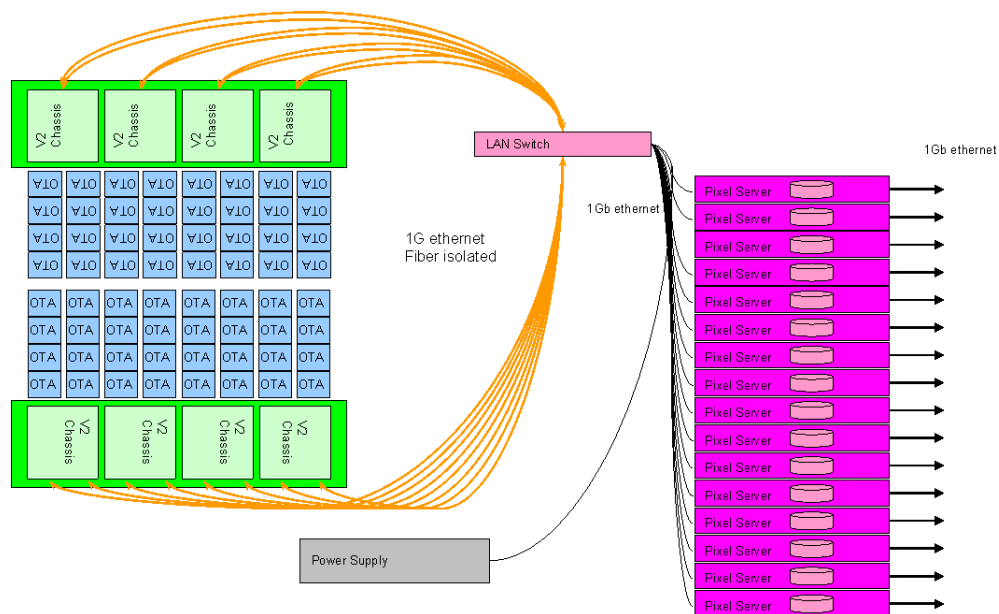


Figure 7: System Block Diagram.

- Select cells (by manipulating. digital output lines)
- Shift cells during integration (OT correction)

A software GUI has been developed to aid in the optimization and test of the different types of CCDs. Named Cestlavie (Fig. 9), it allows an operator to graphically load, edit, save and real time download a CCD clocking pattern to the STARGRASP controller.

Data communication throughput across the Gigabit Ethernet link have been excellent. 400 to 600Mbps per second speeds are achieved with a UDP protocol through commercial network switches. The large SDRAM memory on board is capable of buffering multiple device images and removes the need for real time data transport.

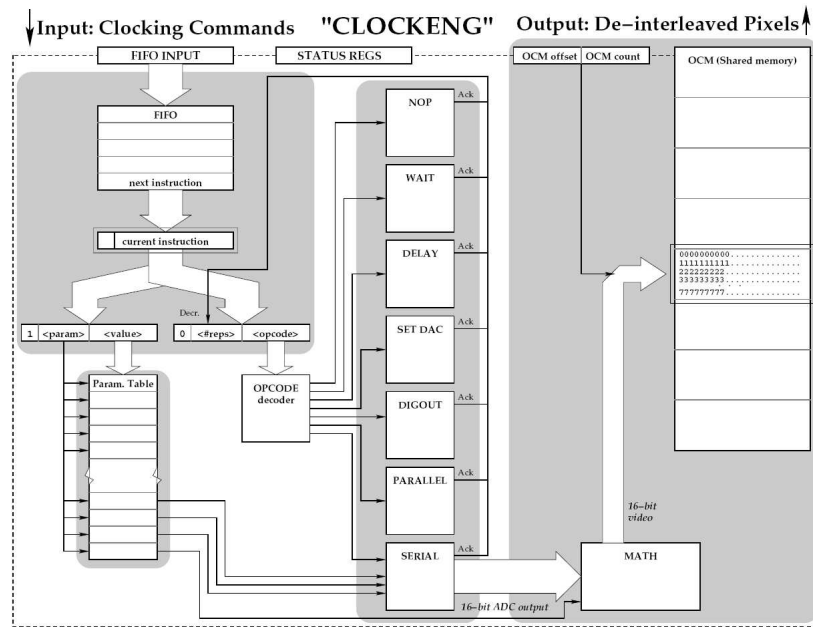


Figure 8: Clocking Engine conceptual diagram.

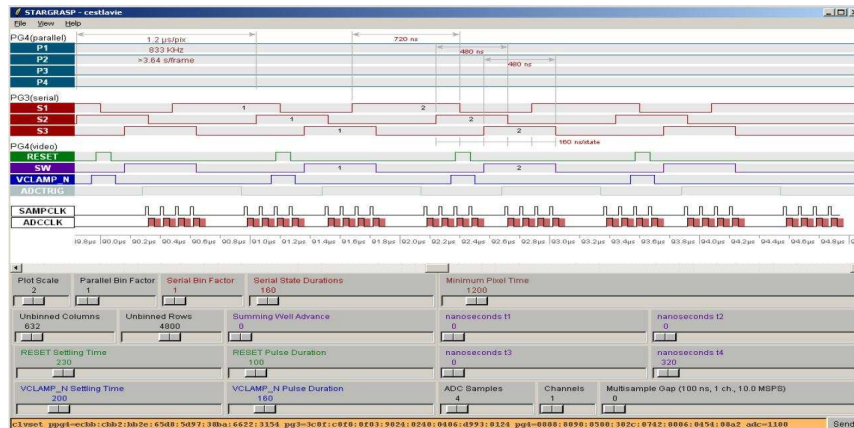


Figure 9: Cestlavie GUI

We use what we call “video mode” on designated cells that does a rapid readout of a small area containing a bright star without any shuttering. Of course we incur slight trailing, to the extent that the parallel clocking is slower than the integration time, and we are subject to other bright stars contributing charge in the readout or erase region of the small area. By choosing appropriately bright and isolated stars this is not usually a problem. This mode gives us the ability to monitor a star’s location at up to 30 Hz, and we use this signal for guiding. We typically designate 20–30 star cells for video during an exposure, and we analyze the mean position of all the stars relative to the start of exposure for offset and rotation. The offset signal we low-pass filter and post to the telescope for corrective guide offsets.

We calculate a number of “value-added products” that make efficient observing possible, for example we display binned-by-64 jpeg thumbnail images, noise images, guiding strip-charts, etc. We identify all saturated stars and bright defects in each frame that is recorded, also collecting a set of unsaturated stars. From these we calculate image quality jpegs (image size, image elongation). There are so many cells and amplifiers in the focal plane (~ 3500) that it is often the case that the patterns that appear are more diagnostic of image quality and camera health than just the mean values themselves.

6. AUXILIARY

Of course an effort of this scale required development of an extensive infrastructure. Apart from equipping a lab, we built a sequence of three test cameras for device and controller testing. The largest test camera, TC3, is half the size of GPC1 and gave us the opportunity to design and verify many of the enabling technologies. We built an extensive detector “Test Bench” that provided a completely automated, thorough set of signals for device evaluation and optimization.

A gantry microscope gave us a $30 \times 20 \times 8$ inch volume we can measure to 1 μm using a long focus microscope objective and digital camera. This was essential for determining the as-built thicknesses and profiles of the detectors, and adjusting their pads to bring their surface to the desired location. Large adjustments were required, both to bring the detector surface to where we thought the optical focal surface would be, and then to adjust it to where we measured the optical focal surface to be once we had assembled and collimated the entire telescope.

7. PERFORMANCE

The performance of GPC1 has been quite satisfactory, if not perfect. We initially installed GPC1 on the telescope in August 2007 and has operated without any significant problem for two years. We brought it back to Manoa briefly in Feb 2008 while the primary mirror supports were being reworked, and used the time to improve the cooling housing around the controllers and power supplies. During this two year operation we have had many power failures, a melted power connector, loss of cooling on many occasions. The fact that no detector suffered any damage and the controllers were unaffected is a testament to the thought and effort put into basic, good engineering practices not detailed here.

The radio frequency interference (RFI) on Haleakala is quite serious. Prior to removal of some broadcast antennas the field strength was as much as 1 V/m, near-field and unimpeded by Pan-STARRS’s fiberglass dome. This caused as much as $50 e^-$ interference noise in TC3, but careful design and phasing reduced this to about $2.5 e^-$ for GPC1. The worst offending antennas have now been relocated, but it is still a very high RFI zone, and enduring RFI seems to be an unavoidable cost of an observatory on Haleakala.

The vacuum performance of the cryostat has been as designed and satisfactory. We can go many months with a good enough vacuum that we have negligible conductive heat load and negligible condensation on the device surface. There have been enough opportunities created by bad weather or telescope work to warm up the cryostat and repump.

The detector performance has been adequate. The QE is about what we expected and exceeds our requirements. Most of the cells have a full well in excess of $50 ke^-$, but cells affected by the corner problem are considerably worse. This does not quite meet our requirement, but is acceptable. Similarly, the noise at 0.5 Mpix/sec is $6.5 e^-$ or less for most cells. This does not meet our requirements, and causes the signal to noise of our shortest, bluest exposures to be 25% worse than expected. We have not been able to unravel why the amplifiers are noisier than expected — unfortunately progress on that front tends to run in quanta of lot runs and \$1M.

Our readout and save time is 8 seconds at 0.5 Mpix/sec. While this is considerably slower than we had hoped for, it is a compromise between read time and noise, and appears to be an acceptable overhead loss for our survey programs.

We have a larger loss of area from bad cells and bright defects than we had hoped. While within our specifications and only amounting to a few percent in fill factor, it is disappointing. The OTA bad corners have been alluded to. The “ion implant laser anneal” process by which the thinned devices are annealed passivated appears to be the cause of most of the bright defects. It is as yet unknown what the exact damage mechanism is, nor what aspect of that process at MITLL causes the problems.

We have some detector artifacts that are self-inflicted by our controller. We see line-by-line bias offsets, sometimes dramatically different in some OTAs, that appear to be caused by an imperfect zero point restore. We must reapply this zero point in each row of pixels read out because the cross talk from many hundreds of clock signals and amplifiers is manageable until it drifts in time. This “streaky” effect can largely be removed by a line-by-line bias correction, that for some OTAs requires a rescaling of the pedestal signal by a few percent and acquisition of a pre-scan bias level as well as post-scan. Although annoying, these horizontal streaks can be mitigated almost perfectly by a combination of data collection method and a bit of post-processing.

Another artifact arises because the parallel gate voltages must be more positive than the NMOS logic signals, and we designed the controller to provide only positive logic signals. Since we cannot therefore invert the surface of the CCD, the charge from a star that is grossly saturated gets into long-lived traps and leaves a “persistence burn”. This manifests itself as a streak upwards from the saturated star, as the charge from the star’s location leaks into the pixels being clocked by, but also in ensuing frames as a streak *downwards* from the star, because the huge charge packet was dragged across those pixels during readout. We are preparing to effect the proper, long term solution of modifying the controller to provide negative logic signals, but in the meantime we have implemented a routine in software that detects, fits, and corrects these persistence features. The quality of the correction is mostly spectacularly good, but the sky is a big place, and not knowing whether a feature is a line of stars or galaxies, the routine has to be conservative and sometimes misses trails or leaves behind artifacts.

There are some glints and excess scattered light that we are trying to track down and fix. GPC1 has three baffles, one above the entrance window, one in the middle, and one just above the focal plane, but there are many surfaces left polished to minimize radiative load that apparently can be illuminated by light scattered above into unexpected angles. The most egregious problems have been fixed by improving the telescope baffling, but some yet remain.

In December 2008 we brought the camera back to Manoa for two months in order to replace the 20 devices with the worst “bad corners” with ones from the Lot 3 newly delivered to the IFA. By this time we had also measured very significant (~ 100 um), unanticipated differences between the focal surface and detector surface. We did not know at the time whether this was primarily an optics effect or movement of the detectors or carbon fiber focal plane, but subsequent metrology of the detector surface and understanding of the optics make it clear that the difference was primarily an optics effect. Fortunately it was a relatively simple (albeit extremely painstaking and dangerous) task to disassemble the entire focal plane, adjust all 60 device’s tilt and piston, and reassemble to follow the as-measured focal surface. GPC1 went back on the telescope in Feb 2009 and has operated without significant problem since.

Some concerns have been raised about flexures and hysteresis effects that were thought perhaps to arise from the LCC/GPC1 interface. As yet we do not have a complete picture of whether there are motions in excess of the FEA design, but the M2 support is now thought to be the proximate cause of these effects. This issue highlights how complex the Pan-STARRS system really is, and how hard it can be to track down and fix problems.

8. GPC2

We are in the process of designing GPC2 for the Pan-STARRS2 telescope, to be installed initially in the adjacent dome to Pan-STARRS1. We expect to build a near-copy of the telescope, and GPC1 has proved successful enough that we anticipate only minor changes for GPC2. Some of the changes include

slight stiffening of a few of the flexures and mechanical components, level shifters for the logic, replacement of the rigid-flex boards (the primary source of cross-talk) with flexprint, more internal blackening (at the cost of increased radiative load), and better controller cooling. We have found that the ability to image up the optical train is extremely valuable and will replace the pinholes with a lens for better resolution and brightness.

With MITLL we are in the process of designing a new OTA that will incorporate the best features of the CCID58s in GPC1 while making some minor improvements. We are also in the process of designing a new package for the GPC2 devices that features AlN and invar for an excellent CTE thermal match to silicon, and different connectorization. Although the new package design was based on building to a perfect detector surface, we do believe that we must allow for three point adjustability of tilt and piston because we cannot risk the possibility that the optics will provide an out of spec focal surface.

The total development and construction time for GPC1 was about four years and consumed about 30 person-years of labor, \$2M in parts and infrastructure, and about \$3M for detectors. The total development and production time for the detectors was 5 years. Fig. 10 shows what the system cost in terms of money and time once the NRE is removed.

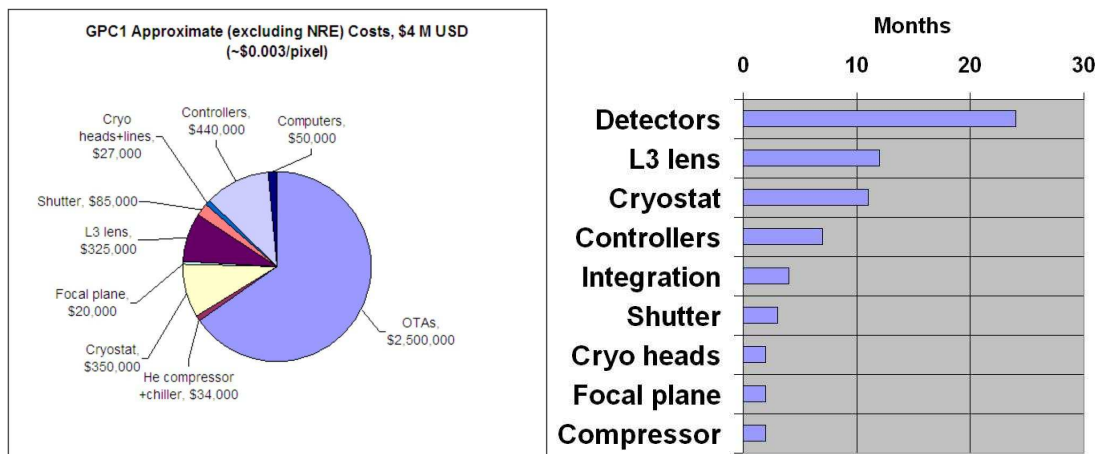


Figure 10: Cost breakdown for GPC1 (left) and approximate production times for various major subsystems. Note that this excludes non-recurring engineering cost.

We hope for GPC2 to be able to obtain the detectors required in less than 2 years for less than \$2M, but this is highly dependent on production yield and fabrication difficulties. We expect to reproduce the cryostat and controllers for about \$1M in hardware and approximately 8 person-years in labor, we hope the net cost of GPC2 will be approximately \$4M and that we can produce it in 2 years.

Almost all of the uncertainty in budget and schedule arises from the detector procurement. Large, astronomy-quality CCDs remain something of a black art, and we have been fortunate in having the expertise of our partners at MITLL. While we expect to maintain that relationship in the foreseeable future, we are developing a second source for detectors in order to augment our capability to produce many more GPC cameras as demand for them arises from the Pan-STARRS project or other needs.

REFERENCES

1. P. Onaka et al., Proceedings of the SPIE, 7021, 2008.