The BRITE Constellation Space Telescope
Design and Test of a Wide Field, High Resolution, Low Noise Optical Telescope for a Nanosatellite Constellation

Jake C.T. Cheng, Jakob Lifshits, C. Cordell Grant, Mihail Barbu, Robert E. Zee
Space Flight Laboratory, University of Toronto Institute for Aerospace Studies
4925 Dufferin Street, Toronto, ON, Canada M3H 5T6: (416) 667-7700
jcheng@utias-sfl.net, jlifshits@utias-sfl.net, cgrant@utias-sfl.net,
mbarbu@utias-sfl.net, rzee@utias-sfl.net

ABSTRACT

The *BRIght Target Explorer (*BRITE) Constellation is a nanosatellite-based astronomy mission tasked with measuring stellar variability of the brightest and most massive stars in the Earth's sky. Charged with the task of capturing the stellar cycles through photometry, the payload of the BRITE nanosatellite is a wide-field, high resolution, low noise CCD-based refracting telescope capable of observing stars with apparent magnitude of +3.5 and brighter with a signal to noise of 3,000 per 1,000s cumulative exposure. A number of payload system design aspects of particular interest to the small satellite community are discussed including a low-noise CCD-driver design that respects the power limitations of a nanosatellite, a primarily passive thermal control strategy for stabilizing the CCD temperature, a software architecture that maximizes scientific output by enabling highly flexible yet automated payload operations, and a modular mechanical design that ensures maximum flexibility during assembly, focusing, and debugging. In addition, this paper discusses key aspects of instrument integration and testing including the focusing process and the challenges associated with achieving acceptable Point Spread Functions in a wide-field scientific instrument over the entire field of view, imager characterization of bias, gain, saturation level, dark current, and readout noise over temperature.

INTRODUCTION

Scientific Objective

The primary mission objective of BRITE Constellation is to provide milli-magnitude (0.1% error) differential photometry of bright stars. It so happens that, in Earth's sky at least, the most apparently bright stars are also among the most intrinsically bright stars. These stars, and in particular the Asymptotic Giant Branch (AGB) and OB stars, most affect the ecology of the universe by creating and distributing all of the heavy elements that are necessary for life as we know it.

Despite their prominence in the sky, bright stars have not been studied to the same extent as fainter stars. As such, there remain several questions about the life cycles of these stars that BRITE Constellation hopes to help answer using photometry/asteroseismology. The ultimate goal for BRITE is to be able to provide photometric data on all 286 stars brighter than visual magnitude +3.5.

Mission Concept

It is well known that massive stars experience periodic, semi-periodic and irregular variations in intensity due to factors such as change in density, magnetic field, surface temperature and internal seismic phenomena. The periods associated with these variations can range from minutes to months. BRITE Constellation fills an observational niche by providing scientists with almost continuous precise photometric time-series measurements with very long baselines (up to six months). By identifying the modes astronomers are able to extract information regarding the internal structure and density profile of these stars. [1]

To achieve its goals, each BRITE satellite will take stellar photometric measurements of a target star field at least 15 minutes per orbit, every orbit, for up to six months at a time. To do this, each satellite in the constellation is equipped with a wide field of view (FOV) optical instrument (24° x 19°). The wide FOV ensures almost 100% coverage of the sky for differential photometry of the target stars (Figure 1).

Each BRITE instrument is tuned to either a red or blue passband (and possibly UV). As designed, each instrument is capable of achieving S/N of 3,000 per 1,000 cumulative exposure on the faintest target stars, which will likely make it possible to do photometry on even fainter stars as well.

Although a subset of the scientific goals of the mission could be met with a single satellite, the science is greatly improved by having a constellation, in this case, of six satellites. First, the time coverage for targets is improved since the visibility windows will typically be different for each satellite. Hence, targets can be observed for a greater proportion of each campaign. Second, targets can be observed over longer baselines since the launch and eventual death of the satellites is likely to be staggered in time. Third, more target fields can be observed since not all satellites need to be observing the exact same target all the time. Fourth, and finally, targets can be observed in multiple passbands without the need to add complexity to the satellite or instrument design (e.g. filter wheels).

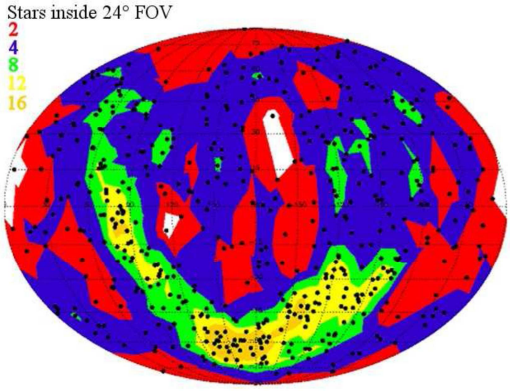

Figure 1: Number of Stars with Limiting Magnitude V=+3.5 Across the Whole Sky for a 24° FOV

Birth of BRITE Constellation

BRITE Constellation is a truly international mission. The original mission concept for BRITE was developed by a Canadian astronomer, Dr. Slavek Rucinski, as a nanosatellite follow-on to the highly successful MOST microsatellite stellar photometry mission. [2] Since BRITE was conceived, the single-satellite mission concept has grown to a six satellite constellation with science teams, engineering teams and funding sources in Canada, Austria and Poland.

BRITE Constellation began in earnest in 2005 when a single BRITE satellite (UniBRITE, see Figure 2) was funded by the University of Vienna. Although funded from Vienna this satellite would be built and tested entirely at the Space Flight Laboratory (SFL) in Canada. Funding for a second Austrian BRITE satellite (BRITE-Austria, aka TUGSat-1) was secured in January 2006, this time funded by the Austrian Space Agency (FFG). Unlike UniBRITE, BRITE-Austria would be assembled and tested in Austria by engineers at the Technical University of Graz using a kit of parts and with mentorship from SFL. While each Austrian BRITE satellite would carry a payload tuned to a different optical band the satellites were otherwise almost identical.

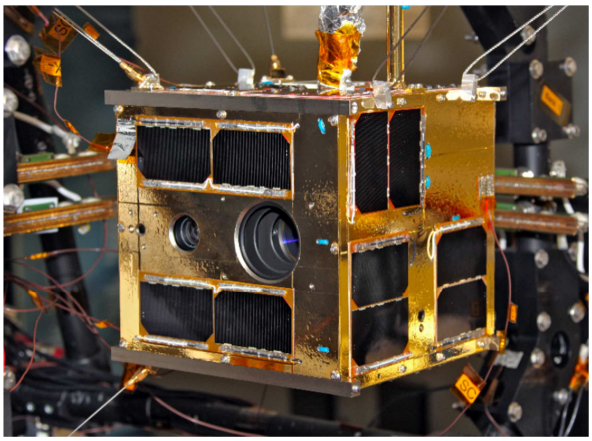

Figure 2: The UniBRITE Spacecraft (Left: Startracker Aperture, Right: Instrument Aperture)

The third and fourth BRITE satellites were funded several years later, in 2010, by Poland. Like BRITE-Austria, the Polish satellites are being assembled and tested at the Space Research Centre in Warsaw from a kit of parts delivered by SFL. The last pieces of that kit were delivered to Warsaw in June 2011. The first Polish BRITE is a copy of BRITE-Austria in that it carries the Blue version of the BRITE instrument. The second Polish satellite may carry an ultraviolet (UV) version of the BRITE instrument, pending a favourable feasibility study.

After several unfortunate delays, the fifth and sixth BRITE satellites were finally funded by the Canadian Space Agency in January of 2011 ensuring full Canadian participation in what was originally a Canadian idea. The Canadian BRITE satellites, currently under construction at SFL, will be almost identical copies of the two Austrian BRITE satellites.

Launch and Operations

The two Austrian BRITE satellites are currently scheduled for launch in December 2011 on the Indian Polar Satellite Launch Vehicle (PSLV). The first Polish BRITE is scheduled for launch in September 2012. The second Polish satellite and the two Canadian BRITEs are each expected to launch in 2013. When launched BRITE-Austria and BRITE-Poland #1 will be the first

satellite for each of their nations and, as a result, have each attracted significant public attention in both Austria and Poland.

The accelerated development schedule of the Polish and Canadian satellites with respect the Austrian ones (which, of course, bore all design and development effort) helps ensure that most or all of the BRITE satellites are likely to be functioning in orbit at the same time, albeit slightly staggered. This arrangement provides a perfect opportunity to both extend the observation baseline in time (from the time the first satellite comes online until the last one dies) while also permitting simultaneous observation of the same targets in multiple passbands.

OPTICAL DESIGN

Optical Challenges

During design the BRITE instrument had very specific requirements from the science team on the Point Spread Function (PSF) required to overcome undersampling issues and achieve milli-magnitude stellar photometry. The desired PSF was Gaussian with a Full Width at Half Maximum (FWHM) of 5-7 pixels with 99% of the energy enclosed in a 12-pixel diameter. Unfortunately, this defocused PSF is quite different from most commercial optical instruments, which require very sharp focus. As a result, optical optimization algorithms for achieving such defocused PSFs were simply not available. Further complicating the design was the very wide field of view (24º x 19º was selected) required to ensure full sky coverage for differential photometry of stars brighter than visual magnitude +3.5.

In addition, because stray light had been an issue on MOST, there was a strong desire to incorporate as much baffling on BRITE as possible. This factor further constrained the volume available onboard the spacecraft for the optical assembly and played a role in the optical design.

Lens Design

The first iteration of the BRITE optical design used a double-gauss lens system with a central aperture stop. While performance of this preliminary design on paper was quite good, the system required relatively complex lens prescriptions and would have occupied more volume and mass than was available in the spacecraft. In particular, the volume and external surface area required for the baffling would have been very extensive, even for a very short baffle.

With this desire in mind, the second iteration of the optical design made use of an external aperture stop. This design choice substantially reduced the overall size and mass of the instrument and drastically reduced the external aperture size required. After speaking with a stray light expert it was also felt that the external aperture stop would be better suited for stray light reduction. At the same time, meeting the PSF requirements over the entire FOV became more difficult, particularly given the lack of suitable optimization tools available for Gaussian PSFs and the desire to keep the number of lens elements low to keep costs down.

Also of note is the fact that BRITE Constellation was required to observe in multiple pass bands. Specifically, for the Austrian BRITE satellites (the first two built) a different optical filter would be used on each satellite, one at the blue end of the spectrum (390 to 460nm) and the second at the red end of the spectrum (550 to 700nm). These passbands were selected so that the flux from an average star in the target list (~10,000K) would be approximately equal in both instruments.

Initially it was desired to have an optical prescription that could be used for either passband. However, on top of the difficulties already encountered achieving an acceptable PSF across the entire field, this idea was soon abandoned in favor of creating slightly different variants of the same basic optical design. After a lengthy and iterative optical design cycle it was decided that the optics were as good as they were likely to get with the tools and resources available and it was agreed that while the optics were not perfect, they were good enough to meet mission requirements. Hence, in March of 2007 the lens designs for BRITE were frozen. Figure 3 shows the final optical designs.

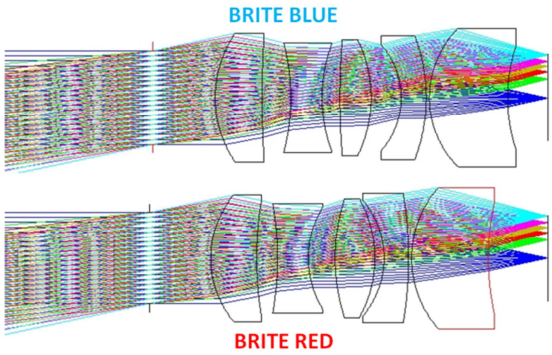

Figure 3: BRITE Blue and Red Optical Designs

The resulting optical designs each have an aperture of 30mm with a focal length of 70mm (focal ratio of 2.33). Theoretical spot diagrams for the two designs at 0º, 5º and 10º from boresight are shown in Figure 4.

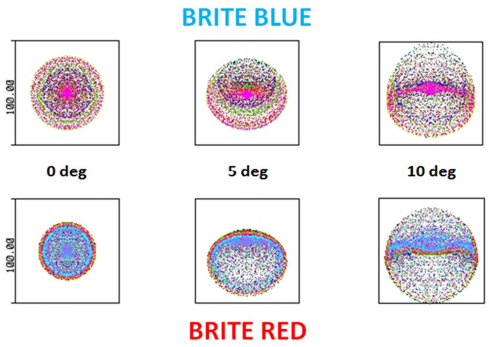

Figure 4: BRITE Blue and Red Spot Diagrams

Focal Position Exploration

Although the lens designs were frozen at this point, the science team was not entirely satisfied with the PSFs that were being achieved over the entire field. Hence, a secondary optical study was undertaken to determine if overall performance could be improved by placing the CCD outside of the baseline focal position assumed by the optical designer. That is, an additional layer of defocus was explored on top of the already defocused system.

The defocus exploration examined the structure of the PSFs using Zemax data input into a custom photometric simulation tool developed at the University of Vienna. The conclusion of that study was that an intra-focal CCD position (with respect to the sharpest possible image) produced PSFs that minimized spikes and kept the vast majority of the light contained within a reasonable number of pixels. This containment (to about 8 pixel diameter as shown in Figure 5) ensured that crowding of stars would not become a major issue.

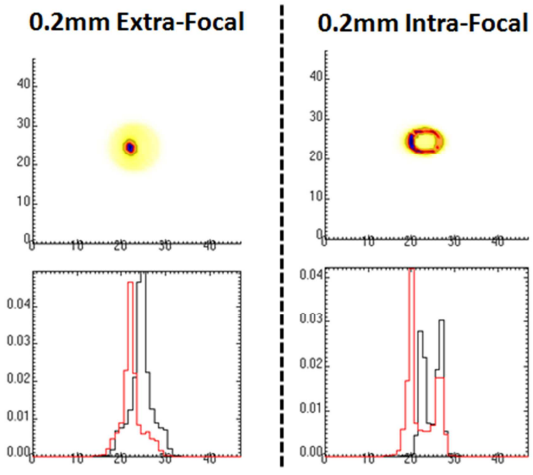

Figure 5: Sample BRITE Blue Intra and Extra Focal PSFs at 5 Degrees from Boresight

Intra-focal PSFs were largely toroidal rather than Gaussian. Extra focal positioning produced something closer to Gaussian PSFs, but the photometric performance was not significantly different from the intra-focal position and the light from each star was distributed over many more pixels. Hence it was decided to focus to an intra-focal position. The details of the focusing strategy used to achieve this PSF can be found in the *Instrument Focusing* section.

MECHANICAL DESIGN

The mechanical design of the BRITE instrument is fairly basic, which helps to ensure that maximum integration flexibility can be designed in. As shown in Figure 6, the telescope is composed of three modules, the header electronics tray, the optical cell and the baffle. The baffle includes the aperture stop as well as the filter. The optical cell houses five lenses and the spacers that position the lenses with respect to each other. The electronics tray contains the CCD header board, which includes the CCD and thermal control electronics. The associated computing and CCD driver electronics are contained separately on the instrument on-board computer (IOBC), which is stacked with the other on-board computer (OBCs) on the satellite bus to reduce payload size and limit heat dissipation within the instrument itself.

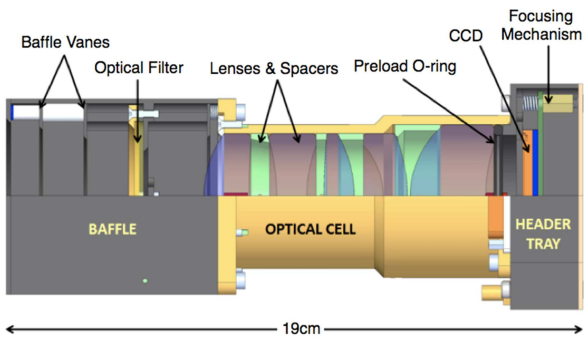

Figure 6: BRITE Instrument

The lenses are preloaded with approximately 200N (50G) of force using a rubber o-ring sandwiched between the fifth lens and the electronics tray. This preload ensures that there is no chatter of the lenses against the housing during vibration. The o-ring also adds compliance so that thermal expansion of the lens stack does not overly stress anything. In the radial axis, the clearance between the lenses and the optical cell was 0.1mm on the diameter, fairly generous by most optical standards. This generous clearance allowed for thermal expansion and contraction of the lenses/cell without having to worry about tightly matching the coefficients of thermal expansion (CTEs) of the

structure to the lenses. The small size of the BRITE instrument also allowed aluminum to be used despite its relatively high CTE.

Focusing

When the BRITE instrument was originally designed, it was envisioned that the header electronics and CCD would be finely positioned and focused using shims placed between the PCB and the electronics tray. It was felt that this solution would be more robust against vibration than having the PCB 'floating' on a set of springs. It was accepted that the shimmed design would require more time and effort to focus.

Years after this mechanical design was settled, another optical device, the S3S Startracker [3] was designed in which the PCB was floated on springs in the same manner originally conceived for the BRITE instrument. Once the S3S design had been qualified and it had been shown that focus was maintained after vibration, it was decided to retrofit the BRITE instruments (which had not yet been focused due to delays in the electronics) with the new focusing design. This modification saved a very large amount of time since it likely reduced the focusing effort per instrument by at least an order of magnitude. Due to the size of the BRITE Constellation this time savings gets further multiplied by almost another order of magnitude since eight payloads (six flight, two proto-flight) will need to be focused at least once before the mission is complete. Figure 7 shows the current focusing approach described above.

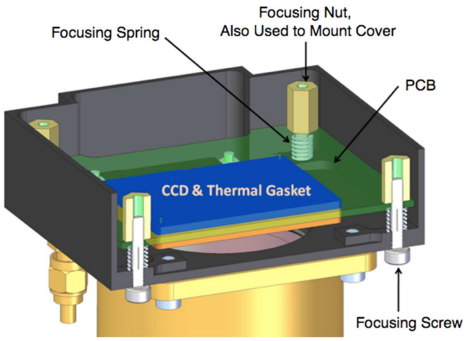

Figure 7: Focusing Mechanism Details

Once the electronics have been assembled and coarsely positioned in the header tray, the focusing nuts are finely positioned using a simple right-angled socket wrench fitted with an angular scale (Figure 8). During focusing the angular scale is used to measure rotation of the nuts by eye. While the angular scale has divisions of one degree, it is felt that slop in the socket on the nut combined with inaccuracies in monitoring angles by eye mean that each rotation is accurate to approximately ±5 degrees. Since the focusing thread has a pitch of 0.5mm this corresponds to a positioning accuracy on the order of ±0.007mm and a CCD rotation accuracy of about 1 arcminute, more than sufficient to focus the BRITE instrument, and also slightly better than the ±0.013mm accuracy that would have theoretically been achieved with the shims known to be available in the size required. The full focusing procedure is described in detail in the *Instrument Focusing* section.

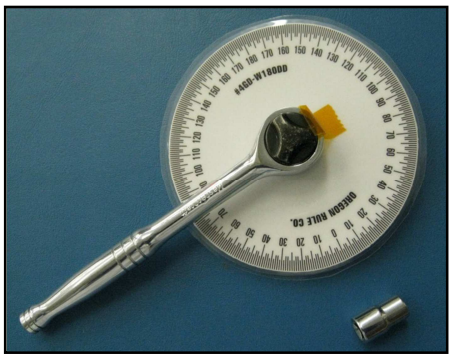

Figure 8: Socket Wrench Outfitted with Angular Scale Allows Fine Adjustments by Eye

It should be noted that the header board has four adjustment points rather than the three points nominally required to adjust tip, tilt and focal position. To avert undesirable flexing and stressing of the PCB and possible ambiguities during focusing, one of the four springs was chosen with a significantly lower spring rate than the other three (by a factor of 25). This choice ensures that the header board is still fully supported in all four corners while also avoiding the problem of over-actuation during focusing since the weaker spring is driven by, rather than driving the PCB position. Combined, the four springs provide about 150G of preload to the header PCB in the nominal focused position.

Once the CCD has been fully focused the PCB is staked to the focusing nuts to avoid in-plane translation of the PCB/CCD under vibration. Such translation would tend to change the apparent alignment between the startracker and the instrument since stars would shift to different pixels on the CCD. While it is planned to calibrate this alignment once in orbit anyway, such calibration will be easier if the starting point is known to be close to what was measured on the ground prior to launch.

Modularity

Each of the three modules of the instrument (see Figure 6) is almost completely self-contained. That is, each module can be assembled or disassembled on its own without having to integrate it with other modules. The

one small exception is that the lenses in the optical cell are preloaded by the electronics tray. So, while the optics can be integrated independently, they do not get fully preloaded until these two modules are integrated with each other.

In addition to being modular, the order in which the two end modules are assembled to the optical cell is also completely free. Hence, the baffle can be added or removed from the assembly without touching the electronics and vice versa. This modularity of the design proved to be crucial at two different points in the development of the first BRITE instrument.

First, while focusing the BRITE-Austria flight instrument for the first time, once the CCD had been focused and tip/tilt removed, an aberration in the PSF was observed. The aberration manifested as a 'tail' that was always oriented in the same direction, regardless of the star's position on the CCD. It was suspected that this aberration was caused by decenter of the fifth lens. In the BRITE Blue instrument the fifth lens has a larger diameter than the other four lenses. As a result, the housing changes diameters as shown in Figure 6. It was believed that the two diameters in the lens barrel had not been machined perfectly concentrically, which was tending to misalign the fifth lens (and possibly the fourth lens as well).

To test this theory it was decided to try shimming the fifth lens on one side to force it toward one side of the lens barrel. By changing the angular position of the shim the lens center could then be offset in any direction. If a change in the aberration was observed as the angular position of the shim was varied, then it would be strong evidence that decenter was the cause of the problem.

The modular nature of the mechanical design enabled the electronics module (with the CCD still in its focused position) to be removed allowing access to the fifth lens. Once the lens was shimmed, the electronics tray could be reattached and the resulting change in PSF immediately checked without having to refocus the CCD at all. After several iterations in which lens shim thickness and angular position were varied, one at a time, a combination was found that almost completely eliminated the aberration (Figure 9).

A sensitivity analysis performed by opticians at the Space Research Centre in Warsaw, Poland later concluded that the optical system was more sensitive to decenter of the fourth element than the fifth. Hence, it is possible that shimming the fifth element actually helped correct aberration caused by the fourth element. This is certainly possible given that the spacer between the third and fourth element has one flat surface that would allow the fifth element to drive the position of the fourth.

Regardless of the source of the aberration, had this shimming process required refocusing of the CCD after each iteration, it would have taken many days. Instead, the entire process was completed in a matter of hours, making it practical to perform multiple times should the same aberration appear in other instruments in the constellation.

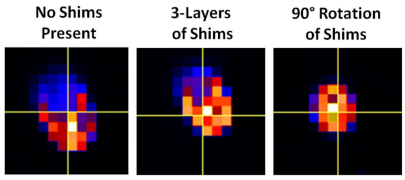

Figure 9: BRITE Blue Boresight PSF During Shimming Process

The second instance when the modular nature of the mechanical design saved considerable time and effort was when it was discovered that the epoxy used by the manufacturer to affix the optical filter glass to a supporting ring was not strong enough to withstand vibration loads. As a result, the glass could come loose allowing it to rattle slightly inside the baffle and possibly damage itself. The fix was simple; apply more epoxy to hold the filter in place. Unfortunately, at the time this problem was discovered, four BRITE instruments had already been integrated, two of which had already been focused. Fortunately, because the baffle could be removed and completely dismantled without affecting the focus or assembly in any way, all four instruments could be repaired with only a few hours of effort for each.

Stray Light Suppression and Venting

Stray light having been a problem for the MOST satellite [2], the BRITE science and engineering teams wanted to ensure that sufficient measures were taken to ensure BRITE did not experience similar problems. Volume for baffling was set aside early in the design. Using this volume a stray light analysis and baffle design was performed by a stray light specialist, the output of which was a set of vane positions and the conclusion that the same vane spacing could be used for both the red and blue instruments without a significant impact on performance.

In addition, several features were included in the mechanical design to suppress stray light. First, the instrument itself was designed to be light tight, which was considered much easier than making the satellite light tight. Second, any surface within the instrument that might possibly reflect light directly into the optics

was anodized matte black using a process known to be compatible with the space environment and that would not fade under exposure to UV radiation. Third, knife edges were machined into the apertures of the baffle vanes and into any parts that were used to hold optical elements in place. Fourth, since most black coatings are not terribly absorptive at such shallow angles, the inside wall of the optical cell was threaded with a very fine pitch thread to minimize shallow angle reflection of light travelling almost parallel to the optical axis. Such stray light stands the greatest risk of being reflected back onto the CCD rather than redirected and absorbed onto another surface.

Finally, the venting strategy for the payload was carefully designed to make it very difficult for photons to enter through the path(s) that air takes to escape. In particular, it was decided that the simplest and safest way to vent the air from all of the various cavities in the telescope was let it flow out along the outside of the optical axis. Figure 10 shows the details of how this was achieved using a combination of narrow channels containing right-angled turns, carefully positioned holes, and loose tolerances.

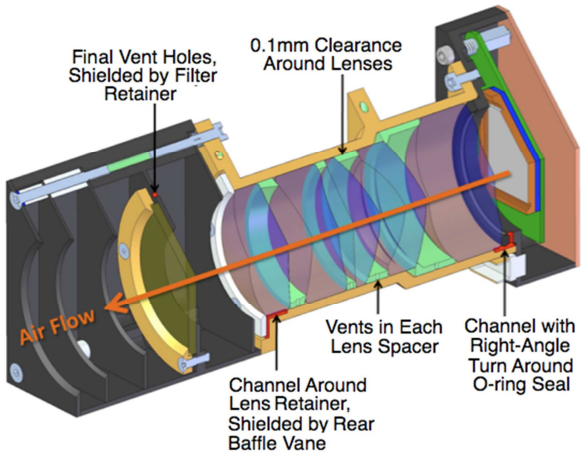

Figure 10: Payload Venting

Originally it had also been planned to use opaque RTV to seal all seams between mating components. However, after stray light testing of the first integrated instrument, it was determined that this was not necessary since metal to metal contact, even with a relatively small number of fasteners, was sufficient to keep leaked light to a virtually undetectable level.

ELECTRONICS DESIGN

The Transition to CCD Technology

The original mission concept proposed by Dr. Slavek Rucinski was to take advantage of the low cost, low power and ease of integration offered by CMOS imager technology. Unfortunately Fabry–Pérot fringe patterns were identified in the images taken by the prototype IBIS4-14000 CMOS imager during the early stage of the BRITE design phase. Tests performed with a uniform light source and spectrograph revealed intensity variations up to 40% across the imager at all visible wavelengths. The cause of the fringing effect was traced back to unevenness in the surface passivation layer on the CMOS imager chip. As a result the IBIS4 imager was not appropriately suited to achieve the high level of accuracy required for the scientific objective.

In the search for an alternative candidate, the team came across the STL-11000M astronomy research CCD camera assembly offered by Santa Barbara Instrument Group (SBIG). [4] Spectrograph tests performed with this camera revealed no evidence of Fabry–Pérot fringing and intensity variation over the imager was significantly lower at approximately 6%. The result of this test steered the choice of photometer towards the Kodak KAI-11002 CCD used inside the STL-11000M SBIG camera. This CCD had relatively low power consumption and very good noise and dark current performance even at room temperature.

The Kodak KAI-11002 CCD is an 11MPixel interline, buried channel CCD with 4008 x 2672 effective pixel dimensions. Each pixel is 9x9µm in size and BRITE uses the option fitted with micro lenses. At only 37.25mm(H) by 25.7mm(V) in size, this imager offers an impressive 66dB dynamic range with built-in electronic shutter and anti-blooming protection. [5]

CCD Driver Design Challenges

The main challenges introduced by the shift of imaging technology were the increase in power and the complexity of the electronics required to operate the CCD. At the same time, the instrument design still had to meet the primary mission requirements including:

1. Perform differential photometry with error less than 0.1% per 15-minute observation.
2. Set exposure time between 0.1 and 100s with an accuracy of 0.01%.
3. Achieve integrated signal to noise ratio (S/N) of 3000 for a cumulative exposure taken over 1000s on a V= +3.5 star.

This set of requirements can be translated into goals of providing very stable analog signal levels that operate the CCD, stable thermal control, high fidelity clocking and very low noise.

A Two-Board Photometric CCD Driver Design

The performance of CCDs is extremely sensitive to temperature. Generally speaking, lower imager temperature brings higher S/N. The original STL-11000M camera assembly contains a two stage thermal electric cooling unit with liquid assist that can bring the operating temperature down to 50°C below ambient. Unfortunately active thermal cooling is not achievable on BRITE due to volume and power constraints for the BRITE mission. These constraints drove the design towards a two-board system that would minimize the heat dissipated alongside the CCD, in the payload itself. The first board, known as the CCD header board, contains the CCD and heater control electronics. The second board, known as the instrument on-board computer (IOBC) contains all on-board computing, power regulation and CCD driver electronics. The IOBC board is an amalgamation of a complete set of standard GNB on-board computer (OBC) and the entire driver electronics required to operate the CCD imager. For the purpose of this technical paper only the driver design aspects will be discussed.

The two-board configuration not only mitigates thermal issues, it also isolates the imager from the noise generated by supporting electronics located on the IOBC. The tradeoff made here is the increased distance that the data, in analog form, must travel in order to arrive at the analog to digital converter (ADC). In order to safe guard against signal depreciation, an amplifier unit was placed on the CCD header board and a coaxial cable is used for signal transfer. These two implementations help preserve the high S/N achievable by the CCD.

CCD Bias Power Electronics Design

A Kodak KSC-1000 timing generator chip lies at the heart of the driver electronics. It is designed specifically by Kodak to control timing on a wide range of their imagers by providing all the clock signals and bias rails necessary to run the CCDs. The timing generator takes in the system clock, produced by a 20MHz crystal oscillator, and outputs the pixel clock rate at 10MHz which results in an exposure time resolution of 100ns; significantly better than the fidelity required on even the shortest exposure (i.e. 0.01% of 0.1s, or 10μs).

Switching power supplies were adopted to maximize power conversion efficiency. There are two branches of power regulator present on the IOBC for DC-to-DC conversion of the spacecraft bus voltage. The first branch is a switching regulator operating at 1MHz to regulate the 1.8V and 3.0V rails used to power all the digital logic components including the IOBC processor. The second branch of the switching power regulator (SEPIC topology) operates at 300 KHz and produces six rails of ±6V, ±12V and ±18V that are used for substrate biasing and clock voltage generation. These two branches of power regulator clocks are tied to the same 20MHz clock source with frequency division provided by a clock divider to achieve the lower switching rate and to allow the power supplies to be synchronized with the pixel clocking. It is important for the power supplies to be synchronized because this ensures that noise signatures produced by the switching power supplies, should there be any, are deterministic. Specifically, phase differences will remain constant and noise will present itself as a repetitive pattern that can be subtracted out.

Further, all bias rails are filtered with double RC filters to eliminate any DC-DC conversion or external sources of noise that would contribute to photometric inaccuracy. Additionally, a latching mechanism is employed with the "readout start" signal feeding through the mechanism and setting the latch. The latch release is triggered by the rising edge of the bias clock ensuring the readout always begins at the same point in the pixel clock period since all the clocks are synchronized. This ensures the start of pixel readout occurs at the same phase point of the pixel clock in every image and avoid an undesired stripe pattern. [6]

A dedicated power up sequence circuit was adopted to divide the entire bias rail power-up process into three stages to be compatible with the built-in electrostatic discharge (ESD) protection circuit in the CCD. When the bias is turned on the ±6V, ±12V and ±18V rails become active first. The ESD circuit comes up 60ms later and a regulated precision rail is powered up at 120ms. The precision rail is used as references for a total of 18 high-speed voltage follower circuits that produce all the clock biases used to readout the pixels on the imager. It is crucial to use a stable, precise reference source since any variations at the input of the voltage follower circuits will show up as fluctuations at the output and may produce noticeable effects on the quality of the resulting image.

The substrate voltage level determines the charge capacity of the photodiodes and is therefore an important factor to consider in achieving high S/N. The choice of substrate voltage is also a tradeoff between saturation level and anti-blooming protection (prevents charges from overflowing to neighboring pixels). Each KAI-11002 CCD comes with a prescribed substrate voltage level in the range of 8V to 17V that has been tested to provide a charge capacity of 60,000 e^- and anti-blooming protections on the order of 100 times the saturation level. In addition, the photodiodes have zero charge capacity at a substrate voltage level of greater

than 40V. Hence, sending a short pulse of 40V on the substrate voltage line effectively provides electronic shutter functionality. An adjustable voltage regulator is used to bring the substrate voltage, powered from the +18V rail, down to the prescribed level specified by the manufacturer. The electronic shutter functionality is achieved by a separate signal-controlled shutter rail feeding off the +18V and -18V rails to provide a combined potential difference of 36V which is superimposed on the nominal substrate voltage (typically around 10V) to achieve an approximately 46V pulse signal that clears off all the charges held by the photodiode region of the CCD.

CCD Pixel Clock Driver Design

Two sets of pixel clock drivers were used to separately control the VCCDs and the HCCD on the imager chip. The VCCDs are responsible for both the charge transfer from the photoactive region to the VCCD and the charge transfer from one row to the next. This dual responsibility requires two clock drivers V1 and V2 to control the VCCD gates and three signal states, low, mid, and high on V2. The V2 clock signal employs two signal inputs: V2A and V2B from the timing generator. V2 remains in a low state when V2A is low (-9V). V2A high signal corresponds to V2B selection, with V2B high giving V2 the high state (+8V) and V2B low the mid state (0V). The high and low state voltages are generated from the precision rail to ensure a stable signal level.

The HCCD is responsible for shifting an entire row of pixel charges out of the CCD for analog amplification after every VCCD cycle. The HCCD clock does not require the tri-state configuration and therefore only requires two signal inputs, H1 and H2, from the timing generator. The H clock high/low states are also derived from the precision rail for the same purpose of signal level stability. Figure 11 illustrates the VCCD clocks in conjunction with the HCCD clock to illustrate the relative clocking associated with the three fundamental stages of pixel readout: photodiode transfer, VCCD transfer, and HCC transfer.

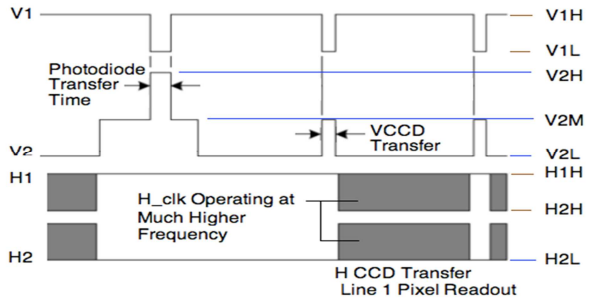

Figure 11: VCCD Clock Signal [7]

Analog Front End (AFE) Chain Design

The AFE chain starts at the analog data output of the CCD on the header board. The signal is first stored in a set of buffers and amplified by an ultra low noise amplifier to provide a gain such that the full scale range of the CCD signal corresponds to just under the signal processor's full scale input range at peak CCD saturation of 60,000e^-. The data is then transferred through the coaxial cable to the signal processor input on the IOBC to become digitized. The signal processor is a 14-bit analog to digital converter (ADC) set to 10 mega-sample per second with correlated double sampler (CDS) functionality. CDS allows the ADC to sample each pixel data stream twice, once to obtain a reference level (sample and hold positive, SHP, signal) and the second time to obtain the data level (sample and hold data, SHD, signal) as illustrated in Figure 12. The difference of the two readings gives the actual pixel value to be passed on to the next stage.

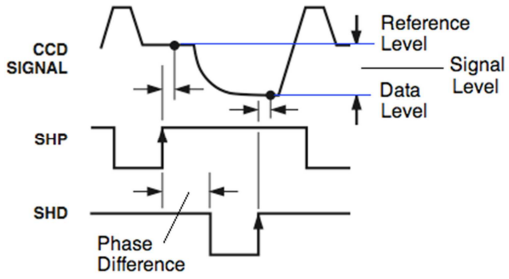

Figure 12: CDS on CCD Signal Output [8]

The CDS is an important feature since any fluctuations in the analog signal level can be effectively eliminated through the differencing. The timing offset (i.e. phase difference) between the reference and data sampling points is programmable through the timing generator and proper fine-tuning is required so that sampling does not occur during the transition region of the CCD signal, which would contribute to overall noise.

The next stage of signal processing involves the variable gain amplifier (VGA) that offers a 10-bit programmable gain setting, providing amplification in the range of 0 to 34dB. The gain value should be chosen to maximize the resolution available in the resulting digital signal, but should not be so great such that signals at saturation exceed the 14-bit output capability of the ADC, which would mean the full dynamic range of the CCD imager was not being utilized.

The final stage of signal processing before the signal becomes digitized is the application of a programmable offset to the analog signal. Noise manifests itself in the signal as both positive and negative variations on the

mean and it is desired to capture the full range of signal fluctuation in order to determine the full effect of noise (and enable its measurement). This offset is achieved by biasing the signal such that the zero reference is offset to a level high enough to ensure all negative noise fluctuations remain positive. This offset produces the bias pattern present in the resulting images (illustrated in Figure 20) and will be discussed further in the *Bias Level and Stability Test* section.

CCD THERMAL CONTROL

Perhaps one of the largest challenges in CCD-based astronomical missions is thermal control of the imager itself. CCD detectors are susceptible to two main detrimental effects: noise, and dark current. Both can significantly reduce the detector's ability to image faint stars, and great measures are taken to reduce these effects by lowering the CCD temperature, sometimes to as low as -110° C.

Due to the small size, extremely limited power and full sky coverage requirement of the BRITE satellite, neither active nor passive (radiative) cooling schemes were feasible. Fortunately, since BRITE will image only the very brightest objects in the sky and because the chosen CCD has relatively low noise, even at room temperature, reducing the absolute temperature of the detector is a secondary concern. Of higher importance is the thermal stability of the CCD over the course of an observation window. Since an observation consists of a number of co-added exposures, it is essential that all exposures be taken with the detector at the same temperature. The BRITE requirements state that temperature of the CCD must be stable to ±2.5° C throughout an observation, with a desire of ±0.5° C.

As controlled cooling of the BRITE instrument is not practical, thermal stability is achieved by way of heating, which is much simpler to implement. Located towards the center of the spacecraft, the thermal fluctuations of the instrument are significantly lower than those of components located closer to the surfaces. Trim heaters are used to bring the CCD temperature to just above the highest temperature it normally experiences during the orbital cycle. Figure 13 shows a theoretical uncontrolled CCD temperature fluctuation throughout the orbit, as well as the temperature stabilized with heaters.

Figure 14 shows the CCD mounted to the header board with four resistive heaters placed near the corners of the CCD. Four temperature sensors are located underneath the CCD and are coupled to it with a thermal gasket. A microcontroller is located on the opposite side of the header board. The microcontroller collects temperatures from the four temperature sensors, controls power to the heaters, and communicates with the IOBC.

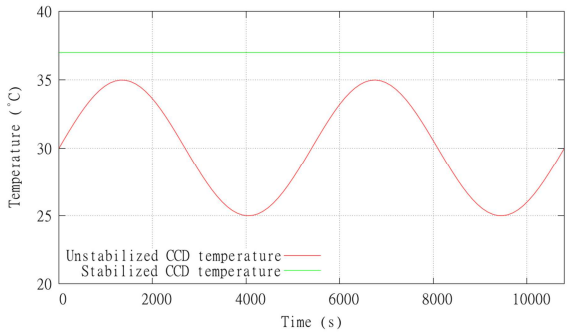

Figure 13: CCD Thermal Control Strategy

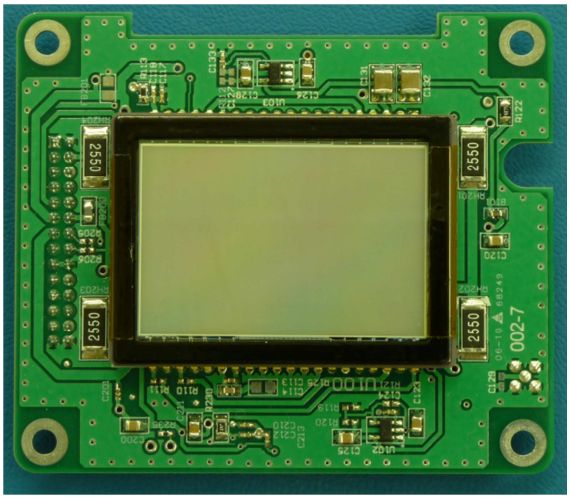

Figure 14: CCD Header Board Containing the CCD, Heaters, and Supporting Electronics

As the heaters are resistive, the heat produced is basically equal to the power dissipated. This suggests that the power output of the heaters can be controlled by either varying the voltage across it, or the current through it. However, analog circuitry to accomplish such control would add complexity and increase the part count. Instead it was decided to use a pulse-width-modulation (PWM) scheme, in which the heater is either fully on or fully off, with a variable duty cycle. The duty cycle would then correspond to the average power produced by the heaters.

A proportional-integral-derivative (PID) controller was selected to control the PWM duty cycle of the heaters. This scheme was selected due to its simplicity and ubiquity in digital control applications. A PID controller is also very intuitive to tune, and does not require substantial computational resources to execute.

Bench-top tests conducted at room temperature and atmospheric pressure have shown that the controller can achieve a control accuracy of ±0.1° C, which is 25 times better than the required control accuracy. Figure 15 shows these results. The green line represents the average CCD header temperature. The time interval from 0s to 4200s shows an uncontrolled state, with temperature fluctuations due to changes in ambient temperature. Once turned on at 4200s, the controller quickly achieves the target of 28° C and successfully maintains it for the remainder of the test. The blue line represents the duty cycle of the PWM signal applied to the heaters. Due to the naturally slow response time of the system, it was found that a mostly proportional controller can achieve good results. Comparable performance was achieved while performing thermal vacuum testing on the integrated spacecraft, whereby the spacecraft was subjected to the thermal oscillations expected in orbit.

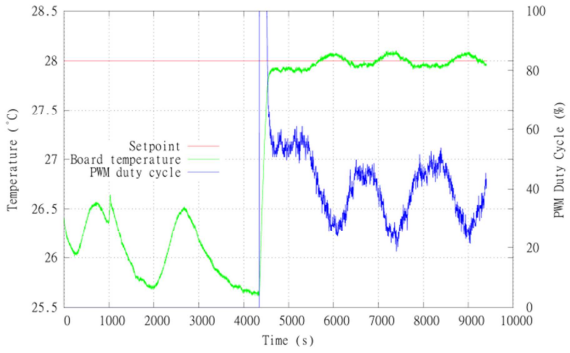

Figure 15: CCD Thermal Control – Bench-Top Performance [10]

MISSION OPERATIONS CONCEPT AND SOFTWARE

Mission Operations Sequence

As previously stated, the main objective of each satellite in the BRITE Constellation is to observe the brightest stars in the sky and measure variability in their brightness over time. As target fields can be anywhere in the sky and there will be times during which the target fields will be obstructed by the earth, the sun, or the moon, a minimum 15 minute interval of view of the target is expected for each orbit. This window will henceforth be referred to as the observation window. During an observation window the satellite will point at its target and perform as many observations as possible. An observation consists of a number of individual exposures which are co-added together to produce a single data point on the brightness curve.

To enable differential photometry, target fields will be selected such that more than one star of interest falls into the field of view of the instrument. The region around a star of interest is called a raster, and up to 16 rasters could be imaged during any given observation. To accommodate the satellite's limited downlink capacity, most of the data processing will occur on-board the satellite, with only minimal information being stored and forwarded to the ground. Full or even partial images will be downloaded only for commissioning and debugging purposes. All data sent to the ground will be compressed on-board first.

To accommodate the observation plan described above, an operational concept was developed and a high-level operations cycle defined for each of the satellites in the BRITE Constellation. This concept was then flowed down to define the required ground and flight software. The top-level operational cycle is shown in Figure 16.

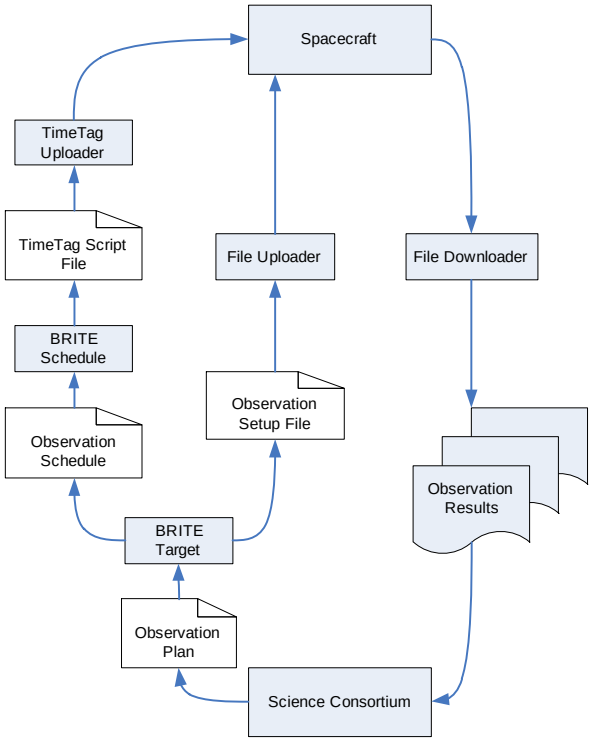

Figure 16: Mission Operations Cycle [10]

The operational cycle begins with the science consortium, which consists of 12 lead scientists and investigators for the mission. The science consortium collects observation requests from the astronomy community, selects targets and times for observations and releases this information through the observation plan. The observation plan is entered into the BRITE Target ground application, which takes into account the

time of year, positions of stars in the sky, and the orbit of the satellite to determine feasible observation times as well as the required spacecraft attitude. BRITE Target produces two files: an Observation Setup File that is uploaded directly onto the satellite's instrument computer, as well as an Observation Schedule File that is passed to the BRITE Schedule ground application.

The Observation Schedule File contains the times at which the target will be visible to the satellite, and the attitude to which the instrument must point. BRITE Schedule converts this information into a series of timed commands in a format understandable by the spacecraft. It also adds other commands that are necessary to initiate an observation. BRITE Schedule outputs these commands to a TimeTag Script file, which is then uploaded to the housekeeping computer on the spacecraft through the TimeTag ground application.

The Observation Setup File contains instructions for the instrument computer that specify the number and duration of exposures to be taken for each observation, the location of the targets in the imager's field of view, and the type of processing that must be done on the raw data.

At the time specified by the TimeTag commands, the satellite's housekeeping computer turns on the instrument computer, performs initialization routines, and begins the observation. Observations continue until the time at which the target is no longer visible, is polluted by stray light or fine pointing lock is lost, at which point the instrument computer is shut down.

During every pass over the ground station observation results are downloaded through a File Downloader application. The observation results are then returned to the science consortium for processing, concluding the observation cycle.

Spacecraft-Level Operations Sequence

Each BRITE satellite is designed around SFL's Generic Nanosatellite Bus (GNB). [9] Each BRITE satellite contains three onboard computers: the housekeeping computer (HKC), the attitude determination and control computer (ADCC), and the instrument onboard computer (IOBC). Additionally, as previously described, a dedicated microcontroller is responsible for the thermal control of the CCD. In nominal operations the HKC receives all ground commands and distributes them to the appropriate destinations on the spacecraft. Likewise, all outgoing traffic is forwarded to the ground through the HKC. Since the CCD heater controller is connected only to the IOBC, all commands to it must be forwarded by the IOBC. Figure 17 shows a simplified interconnect diagram for the onboard computers.

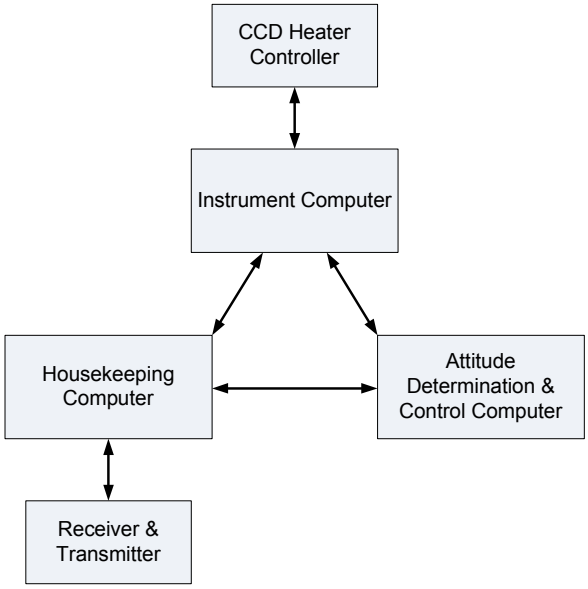

Figure 17: Command & Data Handling Interconnect Diagram

The HKC houses the TimeTag command interpreter. Each entry in the TimeTag queue contains a command as well as the time at which it should be executed. When the command time elapses, the HKC forwards the command to the appropriate computer for processing. TimeTag provides a necessary means of automation as observations will most often occur outside of contact periods with the ground.

During a ground pass, a TimeTag Command File and an Observation Setup File are uploaded to the HKC and IOBC respectively. At a predetermined time before the observation is scheduled to begin, the HKC begins configuring all of the satellite subsystems. First, the attitude determination and control system is enabled and sensors are given a warm-up period. If necessary, a slew is performed to point the instrument at the target. Following the slew, commands are issued to the power distribution module directing it to turn on the instrument, IOBC, and CCD thermal control subsystem. The CCD heater controller is enabled and allowed to settle to the desired temperature. With the satellite pointing in the right direction and the CCD thermally stable, the instrument is configured and the observation cycle initiated.

Observations continue autonomously until the HKC issues a time-tagged command to halt observations. While in an exposure the IOBC continuously polls telemetry from the ADCC and CCD Heater Controller

to ensure that the spacecraft attitude and CCD temperatures are within an acceptable range for imaging. Once observations are finalized, to conserve power the power distribution module is commanded to power down all instrument electronics and the satellite returns to its pre-observation state. This cycle is repeated once per orbit until new targets are selected and new observation scripts are uploaded.

Instrument Computer Operations Sequence

Once an observation start command is received by the IOBC, it begins an autonomous cycle of observations with the parameters defined in the observation setup file. As described previously, an observation window spans the time between an observation start command and an observation stop command. It may consist of multiple observations, and each observation may be composed of multiple co-added exposures. The number of exposures, the locations and sizes of rasters within the field of view, and the type of processing to be performed on them are all defined within the observation setup file. Figure 18 shows the state transitions occurring on the instrument computer during an observation window.

When first turned on, the IOBC application software performs hardware initialization routines and enters an IDLE state where it awaits command from the ground or from the other onboard computers. Upon receiving an observation start command, the observation setup file is read from permanent memory and observation parameters are extracted. In order to maintain a constant time interval between exposures, an onboard timer is populated with the desired time for the next exposure and the state machine transitions into the EXP_START_TIMER state. When the exposure timer elapses, the CCD sensor is cleared of any built-up charge and an exposure is started. The state machine transitions into the EXP_IN_PROGRESS state to await its completion. While an exposure is underway, the attitude of the spacecraft is continuously verified to ensure the validity of the exposure. If the attitude falls outside of the allowable range, the exposure will not be used in the co-added stack.

The completed exposure is downloaded from the CCD into memory and processed in the EXP_POST state. Individual rasters are extracted and co-added to those from previous exposures within the observation. If the number of exposures taken matches the number of exposures per observation as specified in the observation setup file, the state machine transitions into the OBS_FINALIZE state. Observation results are compressed and written to permanent memory. A new observation is started with a transition into the OBS_START state. Otherwise, if the number of exposures per observation has not yet been reached, the time of the next exposure is calculated, and the exposure timer configured. Observations continue until an observation stop command is received, at which point unsaved results are stored into permanent memory and the software returns to the IDLE state.

The presented software scheme provides ease of implementation and testability due to its relative simplicity, while at the same time allowing for a high level of configurability and autonomy in order to maximize the scientific throughput of the satellite.

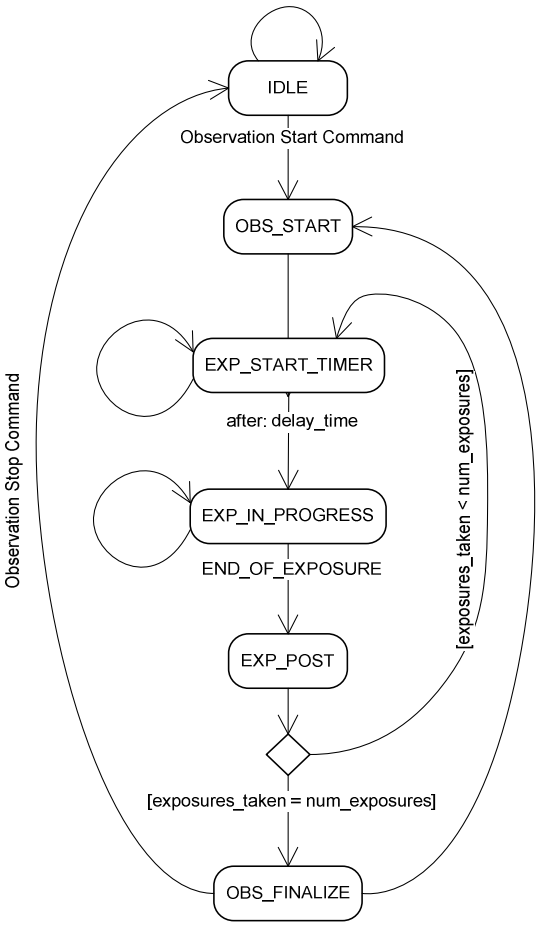

Figure 18: IOBC Level Observation Sequence [10]

IMAGER PERFORMANCE TESTING

The BRITE science team specified a set of acceptance criteria for the performance of the payload, targeting the imagers' bias level and stability, gain, saturation, dark current, and readout noise levels. Each payload must undergo these tests in order to confirm

performance meets requirements and to obtain a baseline characterization of the key parameters affecting photometry.

The acceptance test process involves assessing the imager performances at temperatures of -20ºC, 0ºC, 10ºC, 20ºC, 30ºC, and 60ºC. The two temperature extremes at -20ºC and 60ºC test the performance limits of the hardware and do not need to meet the requirements, while tests done at other temperatures capture performances within the expected on orbit operating temperature range.

Visualization, analysis and processing of the gathered data images, stored in FITS format, were performed through the image reduction and analysis facility (IRAF) in combination with SAOImage DS9 software.

Data Types

The data gathered during acceptance testing at each temperature can be broadly characterized into two types: bias/dark images and gradient images.

Bias/dark images are exposures taken without an external light source and provide information regarding the intrinsic electronic and thermal signatures of the payload. The distinction between bias and dark images is that bias images are taken with zero exposure time. The measured analog-to-digital unit (ADU) levels in a bias image are strictly due to noise and the software offset programmed into the ADC. A dark frame, on the other hand, can be composed with any exposure time. The ADU levels observed in the resulting data images increase proportionally with exposure time due to thermal effects within the semiconductor (dark current).

Gradient images (Figure 19) are taken with an external light source and exhibit a very smooth transition from saturation to the base bias-level. The profile of these images is designed to extract photon transfer curves that detail the signal-to-noise relationship over the entire dynamic range of the imager system.

Figure 19: Raw Gradient Image Sample

Bias Level & Stability

Bias level provides an offset from zero to prevent negative analog signal caused by noise fluctuation about zero mean. Negative signals are misinterpreted by the ADC and result in artificially large values. An appropriate bias level is considered just high enough to eliminate any negative signals out to a four-sigma level. In addition, the bias pattern must be sufficiently stable in order to support accurate differential photometry over long observation periods. The acceptance criteria for bias are provided below:

Table 1: Bias Level & Stability Requirements

1	The mean bias level should be set to 100±30ADU.
2	Bias level should remain stable to a level of ±2ADU averaged over 1000 pixels at constant temperature (±0.5ºC) over a period of 1 hour.
3	Bias level must not drop below 40ADU at any operating temperature.
4	The deviation of each bias frame from the average bias frame is fitted to a Gaussian. The fitted curve and the data should be consistent to within 10% within 3 sigma of the peak, there should be one obvious peak.

Requirement-1 specifies the minimal bias level chosen to safely capture the full range of noise variations centered about the mean and requirement-2 states the required stability level in the bias pattern. Requirement-3 and 4 are in place as checks to ensure the bias is sufficiently high to offset any noise variations. Finally requirement-5 is posed to ensure that the noise signature is of Gaussian nature.

Figure 20 shows an example of a single raw bias frame taken with the BRITE imager at room temperature using the proto-flight payload software. The noise fluctuations riding on top of the bias level are obvious. Furthermore, the presence of a gradient in the vertical scale is evident and is a direct byproduct of the particular pixel clocking strategy adopted. The timing table programmed into the final flight payload software eliminates this gradient.

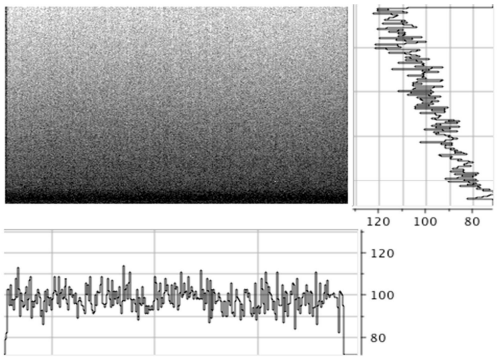

Figure 20: Bias Image Profile (Z-Scale)

To perform the bias validation, five full-frame bias images are obtained followed by a series of partial-frame images. The partial frames capture the bottom sections of the full-frame bias and are taken at a cadence of three minutes throughout a one hour period of continuous operation. The same process is repeated for the top section of the bias and finally, another set of five full-frame bias images are taken to examine the characteristics after continuous operation.

The statistics of the full-frame bias images were examined to ensure consistency in addition to verifying the mean bias level remained within 100±30 ADU. The partial frame images are strategically chosen to capture the bottom and top sections of a full image. Because the CCD readout process occurs in a top-down manner, the upper rows spend more time traversing through transistor gates and are generally noisier than the bottom rows. As a result the top and bottom segments provide an envelope of bias fluctuations. Smoothing is performed by placing each partial-frame image through a 1000 x 1 pixel size boxcar moving-average process. Each pixel in the resulting image is then compared against its relatives from the series to determine the mean, maximum, and minimum values and thus the range of variation from the mean during the one-hour data-gathering period.

Requirements-3 was readily verified by examining the statistics of all the data images to make sure the mean bias level never dropped below 40ADU.

The mean bias frame was subtracted from each of the bias images to remove the constant features and reveal the noise signals. The histogram of the noise signature (Figure 21) is then examined for a Gaussian distribution.

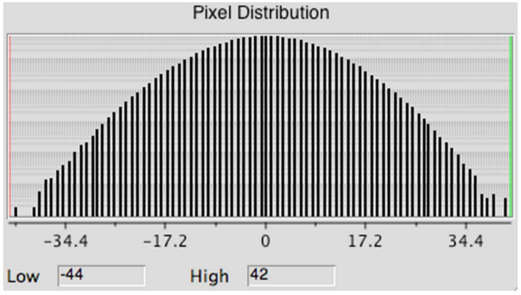

Figure 21: Gaussian Noise Distribution

Table 2 summarizes the results of the bias level and stability investigation performed on the UniBRITE flight instrument.

Table 2: UniBRITE Bias Level and Stability Results

1	Mean bias levels are within 100±30ADU from -20ºC up to and including 30ºC. Mean bias level exceeds 500ADU at 60ºC. This is acceptable as 60ºC is a hardware survival test.
2	High degree of bias stability up to and including 30ºC. Less than 0.5% of the pixels exhibit fluctuations greater than ±2ADU, and are bounded within ±3ADU.
3	Bias level of all the pixels never dropped below 40ADU at any of the tested temperature.
4	All frames with the mean-bias removed reveal a Gaussian-shaped histogram.

The Gain Test

The gain of the CCD system specifies the conversion from the number of electrons collected at the photodiode of each pixel to the ADU count level at the output of the ADC and has the unit of [e^-/ADU]. Direct approach to gain measurements is typically made using a precision photon source that is beyond the resources available for the mission. The methodology adopted for gain measurement for BRITE is a variation of the popular variance method that can be performed with a simple bench-top setup.

Mathematically, the Poisson process describes the relationship between photon signal and the associated shot noise variance as follows:

$$S_E = \sigma_E^2 \tag{1}$$

where S_E and σ_E are the signal and the associated natural noise respectively, both measured in electrons at each pixel. The total noise present, N_E, consists of the readout noise (electronic noise) R_E, shot noise σ_E, and noise due to flat field variations [11] σ_{fE}. Hence the total variance is:

$$N_E^2 = R_E^2 + \sigma_E^2 + \sigma_{fE}^2 \tag{2}$$

Through the definition of gain (ie. $S_E = GS_{ADU}$) and equation (1), equation (2) can be translated and rearranged to reveal that signal and total variance are related linearly through gain:

$$N_{ADU}^2 = \frac{1}{G} S_{ADU} + \left(R_{ADU}^2 + \sigma_{fADU}^2 \right) \tag{3}$$

This linear form provides the foundation for gain measurement using the variance method. The typical procedure involves collecting a series of images at different exposure times to construct a plot of signal versus variance. The modified approach taken herein utilizes the gradient image to do away with the need to take images and deal with dark current calibration at different exposure times. The gradient image contains all the necessary information on the dynamic range of the imager if the illumination is set up properly with a

sufficiently gradual transition of photon intensity from apparent darkness to saturation.

Gradient images (Figure 19) are constructed by illuminating the CCD with a single LED through a neutral density filter and a single lens element as shown in Figure 22. Creating a working gradient illumination test setup involves iterations to achieve a working combination between the characteristics of the LED, the geometric positioning of the LED, the voltage applied to the LED, the optics, and the exposure time. Once a suitable gradient is achieved, a set of eight to ten gradient and dark image pairs are sufficient to construct a photon transfer curve.

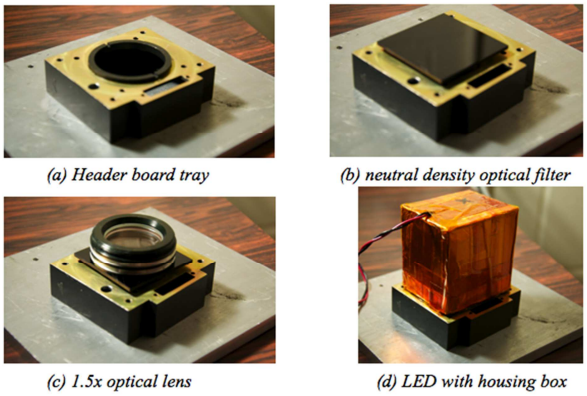

Figure 22: Gradient Image Scenery Setup

A standard dark frame removal is performed on all raw gradient images by subtracting away the mean dark frame. A section of the gradient image with roughly constant vertical intensity profile and the full dynamic range in its horizontal span is selected as the region of interest from each raw image. Figure 23 is an example of such a region of interest.

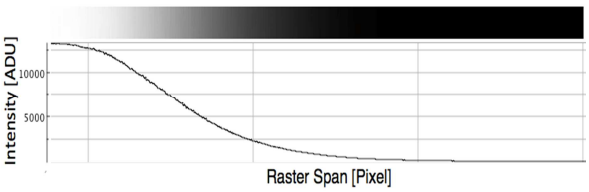

Figure 23: Gradient Image Raster Intensity Profile

Crucial to the data processing is normalizing the mean signal value of each raster region to eliminate any signal fluctuations or potential flat field effects that will contribute to artificial variance in the photon transfer curve; especially at high signal regions. To illustrate this point, consider a pixel in image A measures 10ADU nominally. Due to noise, the same pixel may measure 12 ADU in image B and 8ADU in image C. In this case the artificial contribution to variance is only 4ADU. Now consider the signal level is amplified by a factor of 100 (and hence noise is increased by a factor of 10) and the three measurements are now 1000, 1020, and 980 ADU respectively. The variance contribution in this case is 400ADU. Without the normalization step, the artificial variance would drastically distort the photon transfer curve into an upward trending line (Figure 24). Normalization is done on the mean signal of each gradient image raster by multiplying each image by an adjustment factor so that the mean signal levels are uniform.

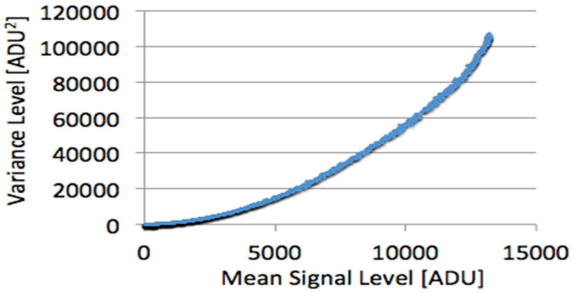

Figure 24: Photon Transfer Curve Constructed without Normalization Step

A set of mean and sigma image is generated from the normalized rasters. Each pixel within the mean and sigma image corresponds to the mean and sigma value calculated based on the set of corresponding pixels from the raster inputs. This means the number of input images correspond to the sample number from which mean and sigma are calculated. Arguably eight to ten raster images is a small number of samples but nevertheless sufficient to extract the general trend of the photon transfer curve. Finally, in order to reduce the degree of variance scattering in the resulting plot, a vertical boxcar smoothing average of 1 by 100 pixels is applied on the sigma image. This step does not affect the accuracy of the resulting signal-variance relationship if the chosen raster has minimal intensity variation along the vertical axis.

Next, the mean and sigma images are exported as a set of data points and the sigma values are squared to obtain the variance. Finally the mean and variance values are plotted to obtain the photon transfer curve. In order to extract the slope of the signal-variance relationship, a linear regression is performed over the low signal region as shown in Figure 25, which tends to be the most linear. Equation (3) reveals that the gain of the system is simply the inverse of the slope of the resulting linear equation.

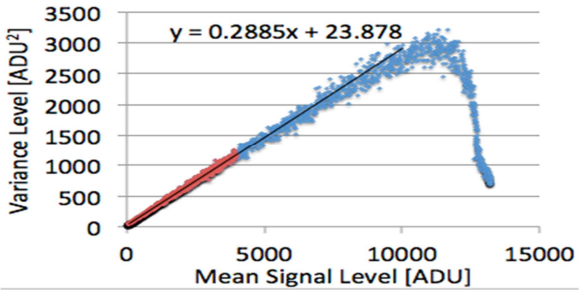

Figure 25: Photon Transfer Curve Constructed with Normalization Step

An important point to take note is that the gain value obtained through this procedure gives the effective system gain. This value accounts for both the amplifier gain and the VGA of the ADC, as well as any signal depreciation throughout the electrons to ADU conversion process. Because the ADC only expresses outputs up to 16383 (i.e. 14 bits), an appropriate gain value would optimally capture the linear response region of the photodiode while maximally utilizing the 14 bits to express the dynamic range of data output. The BRITE Science team has decided an effective system gain value within the range of 3.0 to 3.5[e^-/ADU] is appropriate for this purpose.

Table 3 below summarizes the UniBRITE imager's effective system gain at the different test temperatures.

Table 3: UniBRITE Gain Results in [e-/ADU]

-20°C	0°C	10°C	20°C	25°C	30°C	60°C
3.43	3.42	3.49	3.49	3.44	3.40	3.60

Dark Current Test

Valent electrons in the silicon substrate of each pixel of the CCD may be freed by absorbing the thermal energy alone. This effect contributes to additional electron counts introducing artificial signal level and additional noise. Dark current refers to the rate at which these "dark" electrons are generated at each pixel. The requirements for the dark current test are:

Table 4: Dark Current Test Requirements

1	Dark current should not exceed 5 e-/sec/pixel overall for exposures up to 10 seconds at +20°C.
2	Not more than 0.5% of pixels should exceed a dark current of 50 e^-/sec.
3	Not more than 0.05% of pixels should exceed a dark current of 100 e^-/sec in a 10 seconds exposure.
4	At +20°C, the CCD must make 60 seconds dark exposures with no saturated rows or columns.
5	At +20°C, the CCD must produce exposures with no more than 5% of the detector area saturated at 90 seconds.

In order to quantify the dark current present in the CCD, multiple dark frames are gathered at different exposure times (1, 3, 10 and 30s) in a completely dark environment. Standard bias removal from each dark image is performed using the mean bias frame created in the Bias Level and Stability Test. Next; in order to reduce the random noise level, dark images taken with the same exposure setting are averaged. ADU counts for each pixel are tabulated into bins with a width of 1ADU. This binning provides information on the number of pixels at each discrete dark count measurement. The top 0.05%, 0.1%, 0.5%, 1%, and 5% most dark current sensitive pixels are examined and the dark count level plotted against exposure time for each group of interest. A second order polynomial was found to be sufficiently accurate to describe the dark count development over exposure time in each case (Figure 26). The first derivative of the quadratic formula is evaluated at 1, 3, 10, and 30s to obtain the dark current level at these exposure times.

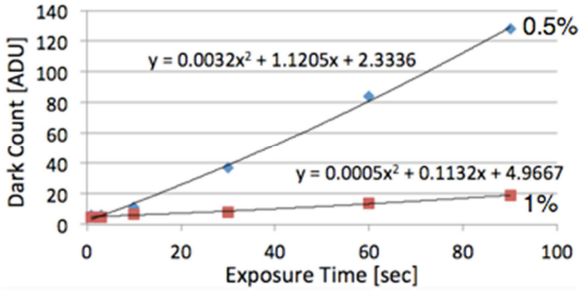

Figure 26: Dark Count vs. Exposure Time for 0.5% and 1% Most Dark Current Sensitive Pixels

Table 5 summarizes the dark current test results for the UniBRITE flight instrument.

Table 5: UniBRITE Dark Current Results

1	99.5% of pixels have dark current levels below 5e-/sec for exposures up to 10 seconds at +20°C.
2	Pass at all operational temperatures. In fact no more than 0.05% of the pixels exhibit dark current level of 50 e^-/sec.
3	Pass at all temperatures investigated, except at 60°C. 99.95% of the pixels exhibit dark current level less than 40e^-/sec in a 10 seconds exposure.
4	Hot pixels are apparent, but no saturated rows or columns present.
5	Considering 10,000ADU as the saturation level (see *Saturation Test* section), only 0.01% of the imager pixels are saturated in a 90 seconds exposure at +20°C.

Saturation Test

As shown in equation (1), the magnitude of the shot noise is the square root of the signal. This relationship is true in the low-signal, shot-noise-dominated region of the dynamic range. However, the positive portion of the

random noise imposed on top of the measured signal will be "clipped" off as full well capacity is approached. [12] As a result, saturation information comes free from the gain test as it manifests itself as the rapid dip in variance evident in the photon transfer curve (Figure 25). The signal level at which this happens provides a measure of the full well saturation level. The minimum saturation level requirement is that full-well capacity must exceed 10,000ADU and results indicated this requirement is satisfied at all the temperatures tested.

Readout Noise Level Test

The readout noise represents the total electronic noise present in the signal chain of the imager system. Because the readout noise is intrinsic to the system, one must eliminate noise contributions from external light signal and dark current in order to quantify it. Eliminating these factors is accomplished by placing the imager in a dark surrounding and setting the exposure time to zero.

The bias frames gathered in the *Bias Level and Stability Test* are used for analysis in this case. Any two bias frames are co-subtracted and the resulting measurement histogram is checked for a Gaussian distribution (Figure 21). The standard deviation of the resulting Gaussian frame is then calculated and scaled down by a factor of $\sqrt{2}$ to account for the fact that the variance level is doubled when two bias images are combined into one. The scaled standard deviation provides the readout noise level of the system.

The BRITE Science team has determined that the readout noise should be below 30 electrons/pixel at a readout clock rate of 10MHz, with the desire to keep the noise level as low as possible. Table 6 summarizes the readout noise results for the UniBRITE flight instrument.

Table 6: UniBRITE Noise Test Results in [e^-]

-20°C	0°C	10°C	20°C	25°C	30°C	60°C
13	13	14	16	18	19	63

INSTRUMENT FOCUSING

The goal of the focusing process is to finely position the CCD with respect to the optics to achieve PSFs that match those selected as being best for photometry (Figure 4). As described in the *Optical Design* section, the payload mechanical design provides a convenient way to adjust the focal position, tip and tilt of the CCD. The tricky and delicate part of this process is to create the artificial star field to be used as reference scenery.

Creating the Artificial Star Field

To assess the focus of the instrument, point sources are required to generate suitable PSFs. Because the BRITE instrument must be kept in a clean environment at all times, it is difficult to use the real stars of the night sky for focusing. Hence, a bench-top focusing setup is required. Adding to the appeal of the bench-top focusing setup is the fact that an artificial star field eliminates inconveniences caused by unfavourable weather conditions.

Unfortunately, the combination of point sources, focus at infinity and wide FOV make the bench-top focusing setup quite challenging. In general, dealing with any two of these challenges is manageable, but all three is not. Hence, it was decided that a wide FOV was not necessary since the FOV could be explored in a piecewise fashion by rotating the instrument. With this compromise a simple collimator with point sources at the source became practical.

With limited funds available and no collimating system in house, some ingenuity was required. Dr. Stefan Mochnacki, a member of the BRITE science team, realized that a commercial telescope focused at infinity is basically a collimator in reverse. That is, if point sources were placed at the focus of a telescope light at the other end would emerge collimated.

Unfortunately, since a human eye and the point sources are not analogous, it is not sufficient to focus a telescope to infinity by eye and then place the point sources at the eyepiece. Instead, an ancillary telescope, with an off-the-shelf CCD mounted to it, was taken outside, mounted on a tracking platform (Figure 27) and focused to infinity using the moons of Jupiter.

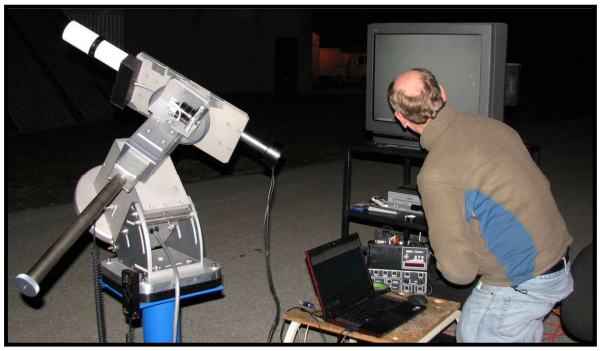

Figure 27: Focusing the Ancillary Telescope

Once the ancillary telescope was focused, it was brought back inside the lab, and was positioned so that it stared straight down the length of the primary telescope which had the point sources mounted to the eyepiece. The focus of the primary telescope was then

adjusted until the point sources appeared sharply focused on the CCD attached to the ancillary telescope. At that point, the primary telescope was correctly collimating the point sources and could be used to focus the BRITE instruments.

The point source was created by drilling several very small holes (∅0.1mm) into a thin sheet of aluminum. To ensure the sources were as point-like as possible, a layer of diffusing material was added and the entire sheet was placed at the end of a shallow tube. The tube was white on the inside so that light was reflected inside the tube providing fairly even illumination. Finally, an LED flashlight with a further layer of diffusion in front of it was used as the light source.

Figure 28: Collimating System in Cleanroom with UniBRITE Flight Instrument at the Aperture

Focusing the Instrument

Focusing the payload is an iterative process. After the header board has been integrated with the rest of the payload and roughly positioned, the first step in the focusing process is to adjust the focal position until PSFs of a suitable size are achieved at the boresight. This adjustment is made by incrementally adjusting all four focusing nuts by the same amount. Typically, angular adjustments no finer than 30° (0.04mm) are sufficient for this step.

At this point the boresight PSFs are assessed for asymmetry caused by lens decentering which, as described in the Mechanical Design section, has been observed in the Blue instruments. If the boresight PSFs appear asymmetric then the header tray is removed from the assembly and the fifth lens is shimmed and rotated until symmetric images are achieved. Figure 29 shows the fifth lens of the Blue protoflight instrument with one layer of Kapton shim applied. Up to three layers of shim are sometimes required.

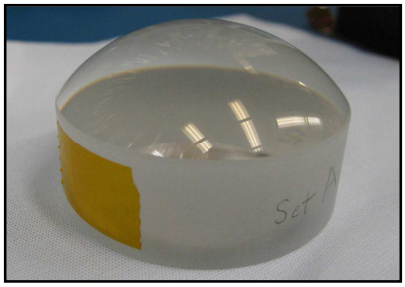

Figure 29: Shimming the Fifth Blue Lens

Once symmetric PSFs have been achieved, tip and tilt are assessed by collecting PSFs at ±5° and ±10º from boresight in both axes. At this point, because aberrations also distort off-axis images, it is difficult to figure out exactly how much tip or tilt is present from the PSFs alone. However, the direction of the offset is usually very obvious and so correcting for it simply becomes an iterative process. The angle of the CCD is adjusted one axis at a time by counter-rotating opposing pairs of focusing nuts. After each adjustment the effect on the PSFs is checked and the process is repeated. Figure 30 provides an example of PSFs from the UniBRITE flight instrument before and after tilt was removed.

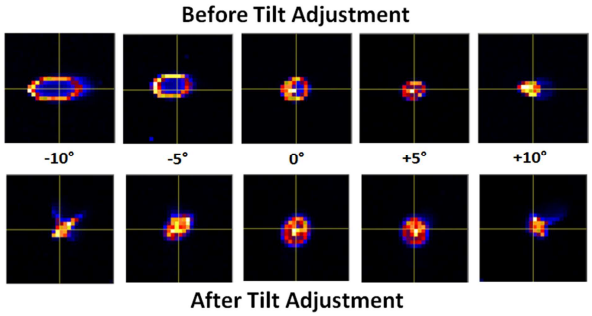

Figure 30: PSFs Before and After Tilt Adjustment

After tip and tilt have been adjusted a final map of the PSFs over the entire FOV is produced and is sent to the science team for approval. Once approved the focusing nuts are staked in place and then the instrument is sealed up. Once sealed, testing is performed to confirm that the instrument is light tight and to verify that readout noise and bias have not changed. At this point the instrument is ready to be integrated with the spacecraft.

CONCLUSIONS

Major challenges in the areas of optics, mechanics, electronics, and software were encountered when designing the payload for BRITE Constellation. Although the resulting solutions were not always ideal

due to time and budgetary constraints, they were always sufficient to ensure that BRITE meets its scientific requirements.

With its half dozen high-performance, multi-spectral payloads, BRITE Constellation will soon begin to study bright and massive stars from space. If the payloads perform as well in orbit as in ground tests, for several years to come BRITE Constellation should contribute significantly to improving our understanding of how the universe has evolved to the life bearing status it currently holds.

ACKNOWLEDGMENTS

The funding and development of the BRITE payload has been a large collaborative effort between many people and organizations. The authors wish to thank the Canadian Space Agency for its financial support of two concept studies and the two Canadian BRITE satellites. Appreciation is also extended to the Ontario Centers of Excellence Inc. (ETech Division) for their support of the development of enabling nanosatellite technologies. The authors also thank the University of Vienna, the Technical University of Graz, and FFG/ALR for financing and leading the development of two BRITE satellites. The Technical University of Vienna is acknowledged for its contributions to the ground station and frequency coordination at the University of Vienna. Walter Besenmatter (U Vienna) is acknowledged for his assistance with the optics design. Polish participation in the BRITE project was made possible by a generous FNiTP grant by the Polish Ministry of Science and Academic Education to the participating institutes of the Polish Academy of Science.

REFERENCES

[1] Norman C. Deschamps, C. Cordell Grant, Dan G. Foisy, Robert E. Zee, Anthony F. J. Moffat, Werner W. Weiss, "The BRITE Space Telescope: A Nanosatellite Constellation for High-Precision Photometry of the Brightest Stars", *Proc. 20th Annual AIAA/USU Conference on Small Satellites*, Logan, Utah, 2006.

[2] S.C. Grocott, R.E. Zee, and J.M. Matthews, "The MOST microsatellite mission: one year in orbit," *Proc. 18th Annual AIAA/USU Conference on Small Satellites*, Logan, Utah, 2004.

[3] J. Enright et.al., "Towards Star Tracker Only Attitude Estimation", *Proc. 24th Annual AIAA/USU Conference on Small Satellites*, Logan, Utah, 2010.

[4] Santa Barbara Instrument Group Inc., "SBIG Research Series STL-11000M Large Format CCD Camera," Product Specification.

[5] Kodak Inc., "KAI-11002 Image Sensor," Device Performance Specification, Rev.1.0, MTD/PS-0938, Jan. 30, 2006.

[6] Kodak Inc., "Vertical Timing Optimization for Vertical Interline Image Sensors," Application Note, Rev.3.0, MTD/PS-0686, Oct. 4, 2010.

[7] Kodak Inc., "KSC-1000 Programming Examples," Application Note, rev.2.0, MTD/PS-0736, Oct. 25, 2007.

[8] Analog Devices Inc, "AD9824 – Complete 14-Bit 30MSPS CCD Signal Processor," Datasheet, Rev.0, 2002.

[9] F. M. Pranajaya, Robert E. Zee, "Generic Nanosatellite Bus for Responsive Mission," *5th Responsive Space Conference*, Los Angeles, California, 2007.

[10] J. Lifshits, "Embedded Systems Development for SFL Satellites", Master's thesis, University of Toronto, 2010.

[11] M. Newberry, "Tech Note: Pixel Response Effects on CCD Camera Gain Calibration", Mirametrics Inc., Copy right 1998.

[12] Thomas J. Fellers, Michael W. Davidson, "Concepts in Digital Imaging Technology – CCD Saturation and Blooming", *Hamamatsu Learning Center, National High Magnetic Field Laboratory, The Florida State University,*