FireOPAL: Technical Performance and First Results

Greg Madsen[1], Phil Bland[2], Matt Bold[3], Robert Howie[2], Ben Hartig[2], Trent Jansen-Sturgeon[2], James Mason[3], Liam Smith[3], Dane McCormack[1], & Rod Drury[1]

[1]Lockheed Martin Space, 8 Brisbane Avenue, Barton, Canberra, ACT 2600, Australia
[2]Space Science & Technology Centre, School of Earth and Planetary Sciences, Curtin University, GPO Box U1987, Perth, Western Australia 6845, Australia
[3]Lockheed Martin Advanced Technology Center, 3251 Hanover St., Palo Alto, CA 94304, USA

ABSTRACT

FireOPAL is a distributed network of standalone, fixed mount optical sensors that is designed to monitor a large number of RSOs at high cadence. Here we describe the technical performance and first results of our prototype network. For each 5 second image, we find that FireOPAL measures apparent angular coordinates with an accuracy of a few arcseconds. The system measures absolute time within milliseconds, absolute flux to within 10%, and detects point sources with visual magnitude of around 16 or brighter. Each unit has sufficient computational resources to completely process an image onboard within seconds, reports measurements to a central server in near real time, which are then used for orbit determination and catalogue maintenance. The data from our prototype network across Australia as used to calculate accurate orbits for hundreds of objects per clear night. When orbits for both LEO and GEO objects are propagated forward more than one day, the predictions match the observations to better than one pixel. Over the past few months, the system has recorded more than one million light curves of satellites which are used to deduce properties such as size, tumbling, composition, pattern of life, etc. Finally, we report on the capture rate of objects of different sizes and in different orbital regimes, and discuss ongoing plans to improve the performance of the network.

1. INTRODUCTION

FireOPAL (Fireball OPtical ALert) is a network of optical SSA observatories that traces its engineering heritage to the successful Desert Fireball Network, the world's largest planetary science observational facility [1]. The observatories are standalone, weather-hardened systems housing a COTS camera, with a high resolution prime lens, a GPS receiver, and a multi-core CPU on an embedded PC. The camera is a full frame 36MP cooled CMOS sensor. When coupled to a 105mm lens, the field of view is approximately 20° x 12.5° with 10 arcseconds/pixel. The observatories are solar powered with battery storage, and are capable of 3G/4G, wireless, satellite modem, or Ethernet connectivity. They are designed to function as a coherent network, to monitor a large number of artificial satellites simultaneously, and provide wide area surveillance and precision tracking in near real time. FireOPAL observatories take 5 second exposures on a fixed mount, with observations synchronized and time-stamped to millisecond precision across the entire network. Synchronised observations over a large distributed geographic network enables range estimation through triangulation as well as comparison of apparent brightness from different points on Earth. Thirty TB of hard drive space allows all imagery and other data products to be stored onboard. The entire package is designed to be low cost, and easy to deploy. The observatories are capable of operating completely autonomously for 24 months in a harsh environment without maintenance. As described in detail below, all the image processing is done onboard within seconds after each image is taken and results are reported back to a central server in near real time.

2. PIPELINE

A key feature of the FireOPAL system is the ability to process all images in near real time. Considerable effort was put into building an automated pipeline that is optimised for speed, while providing robust performance and

flexibility to employ a range of processing algorithms. Figure 1 shows an overview of the typical workflow that occurs onboard each unit during the night.

The first item is an optional step: upload a list of all recently updated TLEs publicly available from space-track[1]. All TLEs, with an epoch within 7 days of the current night, are uploaded. This can be used to correlate the observations with objects in the space-track catalogue and is useful for system diagnostics.

An image with an exposure time of 5 seconds is taken once every 10 seconds. At the moment, the images are taken with a Bayer filter (RGGB) in place. By default, the primary processing takes place on the green channel image with suitable interpolation over the red and green channel pixels. After an image is recorded, a dark correction and flat-fielding is applied. From here, the processing may be directed toward one of two directions: 'LEO/MEO' processing or 'GEO' processing.

The 'LEO/MEO' processing is aimed at finding streaks in individual images. Here, consecutive images are astrometrically aligned, smoothed, and subtracted from one another in order to facilitate the identification of streaks while minimizing artefacts. The 'GEO' processing is aimed at finding much fainter objects that do not move very much over the time span of several images; see Figure 2 for an example. A collection of 59 frames, spanning 10 minutes, are combined in such a way so as to suppress the signal from objects that move between frames (ie. stars), while enhancing the signal from objects that do not move very much between frames (i.e. objects near GEO). The decision as to which processing step to take is determined by the number of images that have been processed. The LEO/MEO processing occurs immediately after every image, until 59 frames have been processed. After that, the GEO processing occurs, followed by a return to LEO/MEO processing, and so on.

The next phase of the processing is to remove the smoothly varying background sky signal and to identify statistically significant streaks (LEO/MEO) or point sources (GEO). After applying appropriate filters based on candidate object's signal-to-noise ratio, shape, proximity to edge of image, etc., a cutout/thumbnail is extracted from the original image, centered on each candidate. In the case of the GEO processing, the original frame taken at the temporal midpoint of the collection of images is used. An example of such a cutout for LEO/MEO processing is shown in the left panel of Figure 3. The stars present in the cutouts are used to calibrate both the astrometry and photometry of the candidate objects. By doing the calibration in-frame, we minimise the complexities and uncertainties inherent in other techniques, e.g. using models for the telescope mount or irregular optical distortion in the focal plane.

For LEO/MEO processing, an extra step is taken to further refine the measurements of streaks. The position angle, streak length, and streak endpoints are estimated with higher accuracy. The streak is compared to all streaks found in the previous or next image in order to determine a time arrow for the streak (e.g. which endpoint is the beginning/end). A covariance is estimated for the x/y coordinates of the streak endpoints. A light curve, defined as the apparent flux of a streak as a function of position along the streak, is extracted. If requested, an object's apparent location is compared to predictions to recent space-track TLEs, and a correlation is made if the prediction of the object's apparent location and velocity is within some tolerance of the observations.

Finally, all the metric track measurements, meta-data, optional correlations, light curve, and diagnostic thumbnails are recorded to a local database as well as sent to a central sever. Due to the low bandwidth connectivity to the units in the field, the transmitted data rates are kept to a minimum. Full frame images are generally not downloaded from the units.

The procedure above is, on average, executed and completed with 10 seconds of an image being recorded. This enables the system to report results in near real time and ensure the system does not 'fall behind' during the night.

The entire procedure, from taking images, to processing, to publishing results, occurs without any human intervention. All the units carry out this process every night without the need for any monitoring. During the day, the previous night's performance across the network is analysed and reported to the team for monitoring and quality control purposes.

[1] https://www.space-track.org

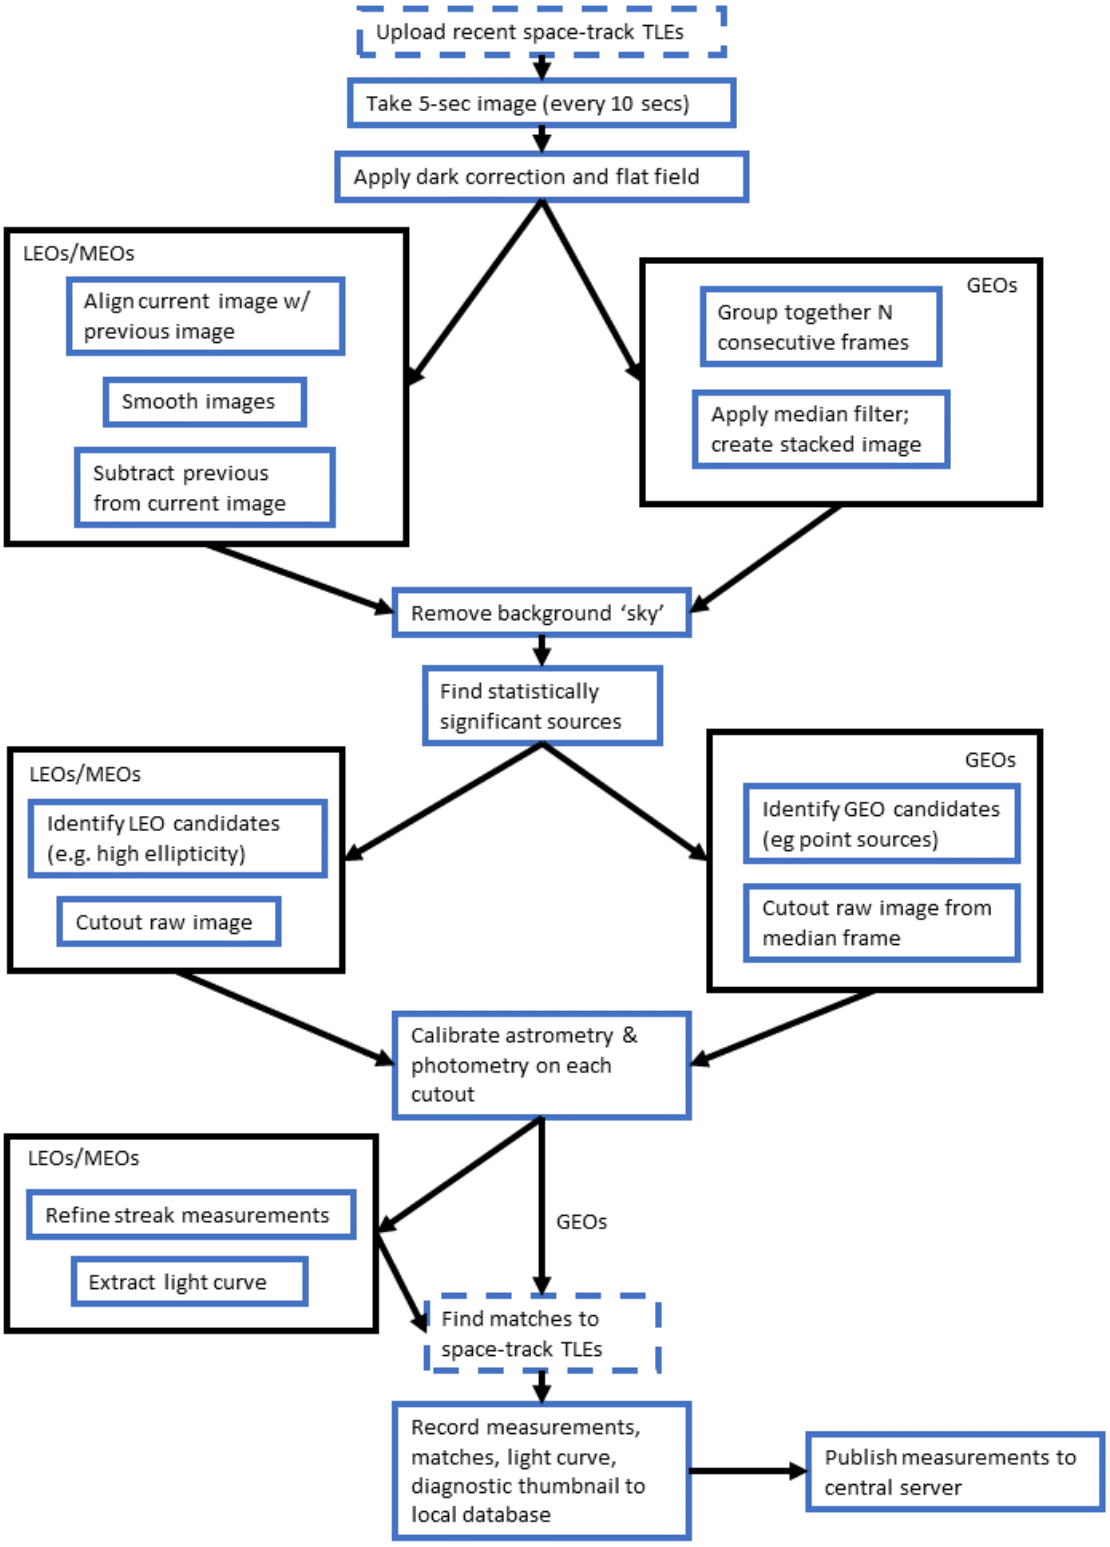

Figure 1: Overview of the FireOPAL pipeline processing, illustrated as a workflow (see text for details).

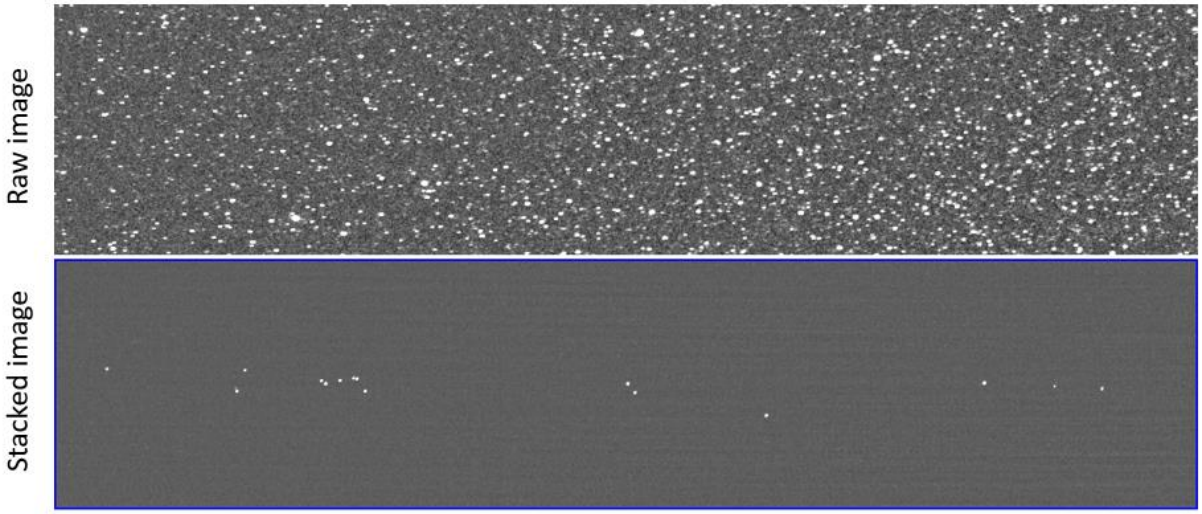

Figure 2: Comparison between raw image (top panel) and stacked image (bottom panel) used during the GEO processing. Both images show the same alt/az part of the sky (subtending 1.5° x 7.5°) and are placed on the same colour scale. A total of 59 raw frames were combined with a median filter to create the stacked image. Point sources brighter than m_V = 13.7 mag (5-sigma) are detectable in the raw image. The limiting magnitude in the stacked image is almost 3 magnitudes deeper (m_V = 16.6 mag; 5-sigma). Objects that have not moved more than a few pixels over a 10 minute time span are easily detected in the stacked image.

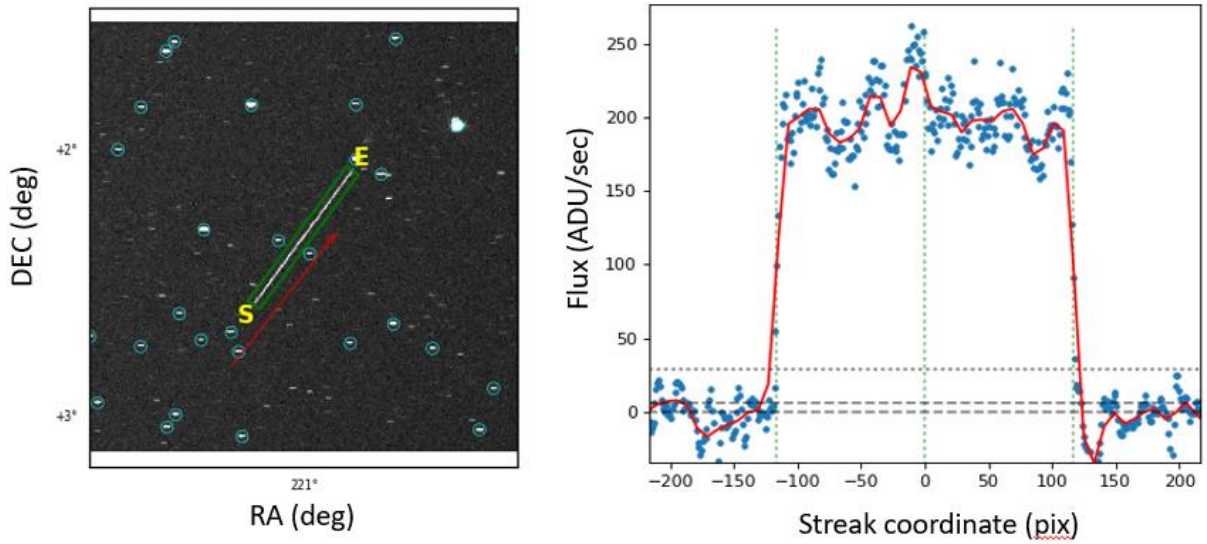

Figure 3: Sample diagnostic thumbnail from LEO processing. The left panel is a cutout of a raw image centred on the detected streak and subtends 1.5° on a side; North is up and East is to the left. The green box shows the estimated location of the streak; the yellow 'S' and 'E' indicated which streak end point is the temporal Start and End, respectively. The cyan circles show the predicted location of Tycho-2 reference stars using the in-frame astrometric calibration. The red arrow shows the predicted position and motion of BREEZE-M DEB (TANK) (SSN# 39477) using the most recent TLE from space-track. The integrated visual magnitude of this streak is 9.9mag. The blue dots in the right panel shows the extracted light curve for this streak. A periodic variation in the light curve with a period of about 50 pixels is evident, highlighted by the smoothed data shown by the red line.

3. CALIBRATION

The location of the observatories, time stamps of observations, the apparent location of candidate satellites, and their apparent brightness are all carefully calibrated as part of the processing pipeline. The observatories are equipped with commercial GPS receivers that continuously report the geodetic coordinates with an accuracy of a few meters; those coordinates are part of the meta-data for every image. The cameras are triggered to open the shutter every 10 seconds, starting at the nearest 10-second mark after local sunset (eg 9:03:50, 9:04:00, etc.). Experiments have shown an average delay of 56 milliseconds between triggering and actual opening of the camera shutter, with a scatter of a few milliseconds. We are in the process of independently confirming the timing accuracy using observations of objects with very high accuracy ephemerides (eg. ILRS targets). We believe the time stamps are accurate to milliseconds; this is supported by agreement between obseravtions and orbit predictions described in section 5.

As described in Section 2, the astrometric calibration is performed in-frame and within a few degrees of a detected satellite streak or point source. First, the x/y coordinates of objects in the cutout with shapes consistent with the expected motion of stars are measured. An approximate World Coordinate System (WCS) conversion [2], using a simple gnomic projection, is calculated using the known timestamp of the observation, location of the sensor, and the altitude/azimuth of the sensor. The coordinates of stars in the Tycho-2 catalogue [3] are then used to further refine the calibration. (The sensitivity and accuracy of the Tycho-2 catalogue is well matched to the FireOPAL detectors; larger catalogues are not needed and introduce unnecessary complexities in crowded fields.) The measured x/y coordinates of the stars are matched to Tycho-2 stars after a series of iterations with outlier rejection. The end result is a WCS solution with a typical residual scatter of a few arcseconds. We find that on the relatively small angular scale of a cutout (~1°), the relative distortion across the narrow field of view is negligible; there is no need to use a more complicated WCS projection, e.g. high order polynomials. If an astrometric solution cannot be found, e.g. thin cloud reduces the detectability of a sufficient number of reference stars, the object and cutout frame are discarded. If successful, the astrometric calibration is adjusted so that the apparent angular coordinates (RA/DEC) of satellites are reported in the ICRF inertial frame, but precessed to the epoch of the timestamp of the observation and corrected for stellar aberration, enabling a simpler conversion to ECI/ECEF as needed. For streak measurements, care is taken to account for the fact that the WCS solution is valid at the mid-point of the image exposure, but the streak end points must be astrometrically calibrated at the beginning or end of the exposure.

The photometric calibration is also performed in-frame. The fluxes of objects consistent with stars are calculated on each cutout using aperture photometric techniques. Aperture corrections are estimated and applied in an attempt to recover the total flux from trailed stars; on a 5 second image a star near the GEO belt has a length of about 7-8 pixels with current optical setup. Stars that do not suffer from crowding are matched to stars in the Tycho-2 catalogue that have a reported B and V band magnitude; Tycho-2 magnitudes are in the Vega system. An instrumental magnitude is calculated for the observed stars and corrected for attenuation by the atmosphere. The median difference between the instrumental magnitude and Tycho-2 V band magnitude defines the magnitude zero point for the image. We find a root-mean-squared scatter in the difference between instrumental and Tycho-2 V band magnitudes of about 0.1 mag. The scatter appears random; there is no evidence that the inclusion of a (B-V) colour term or other terms reduce this scatter. This suggests that the Bayer green channel is an approximate match to the Tycho-2 V band system. This relatively large uncertainty is due to the combination of the challenge in accurately measuring the flux from trailed stars on a de-Bayered image, as well as the different spectral response between the Bayer green channel and the Tycho-2 V band system. We are pursuing efforts to understand and minimize the source of this uncertainty but are limited at the moment to uncertainty in the absolute photometric calibration of 10%.

If the photometric calibration fails, usually due to an insufficient number of matched reference stars, a default value for the magnitude zero point is used and the object is reported as potentially having unusable photometry. We have monitored the network over enough nights to build reference maps of the magnitude zero points on clear/photometric nights for each sensor. The 'photometric' zero point is reported along with the measured magnitude zero point and enables an assessment of the photometric quality of the data. In addition, the point source sensitivity for each cutout, taking into account the measured noise in the background sky, is reported. On a clear night, the limiting magnitude for a 5-second image taken near sea level around the Australian desert, with a 105mm lens, is typically around $m_V = 14$ mag (5-sigma). For a group of 60 images combined for GEO processing, the sensitivity is typically around $m_V = 16$ mag.

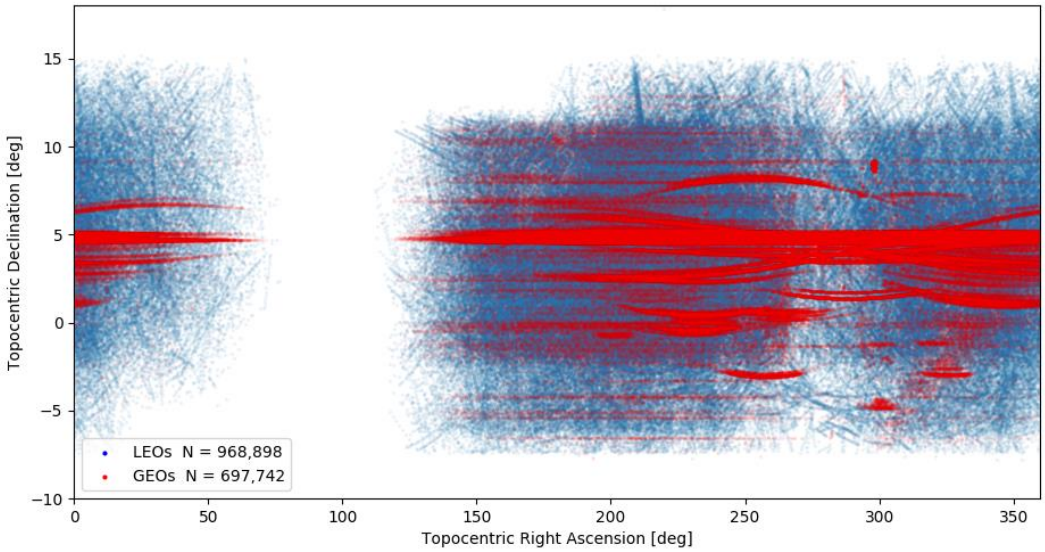

Figure 4: Apparent angular coordinates for LEOs and GEOs observed from all sites in the prototype network between 1 June – 15 Aug 2018. The gap near RA = 100° is middle of the day. The gap in LEOs near RA=275° is near midnight (LEOs in shadow). The diminished number of GEOs near RA=275° is due to noisy background from Galactic plane.

4. PERFORMANCE

Since June 2018, eleven FireOPAL observatories are deployed around Australia. Six units are located at a cluster of five sites in central South Australia. Another five units are located in a cluster near the centre of Western Australia. These form our prototype network and are used to estimate the performance and capabilities of a much larger network of sensors.

In the clusters, each sensor is separated by about 100km and are pointed at overlapping regions of the GEO belt. This enables imaging of the same objects from multiple sites at exactly the same time. With a 100km baseline, the geocentric parallax effect causes an apparent shift in position of the same objects of several arcminutes at GEO, and much larger shifts for closer objects. If the objects seen by different sensors can be associated correctly in each image, a direct estimate of the range to the object can be calculated. Taking into account the uncertainty in angle and timing, we find that two units, observing the same object, can estimate the *instantaneous* range to an object with an accuracy of about 0.2%. Note that this accuracy is greatly improved after fitting larger sets of observations with high fidelity numerical propagation models (see Section 5 below).

On a typical night, each unit will take about 4,000 images. Images are taken after the sky is dark enough to avoid saturating the images; this occurs when the sun is about 10 degrees below the horizon. At the moment, images are taken without regard to local weather conditions. If the weather is partly cloudy, we are able to detect objects through thin cloud (or places on an image that are cloud-free). A future development is to implement a computationally inexpensive method to check for adverse weather before taking images.

Figure 4 shows the number and distribution of all objects detected by the network over the period 1 June – 15 August 2018. Over those 76 days, almost one million objects were detected with ‘LEO/MEO’ processing (streaks), and almost 700,000 objects detected with the ‘GEO’ processing (slowly moving point sources). This includes a significant number of nights during affected by adverse winter weather.

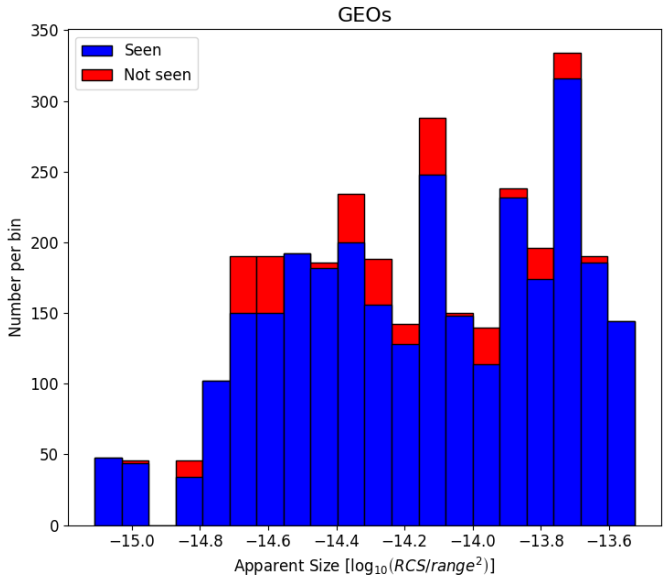

Figure 5: Histogram of the number of GEO objects in the fields of view of four sensors on the clear night of 24 June 2018; objects are binned by a proxy of their apparent size. The blue bars show the number of objects detected by the GEO processing; the red bars shows those objects predicted to be in the field of view but not detected by the processing. More than 90% (1,474 out of 1,622) of objects are detected, with no evidence that objects of smaller apparent size are missed more often.

The number of LEO/MEO streaks detected per clear night per camera varies according to the population going overhead. On average, on a clear night each unit reports measurements for about 1,500 streaks. When those streaks are correlated to objects in space-track, the number of unique SSN objects observed by each camera each night is about 200. There are a large number of streaks that do not appear to be associated with objects in space-track.

When stacking 10 minutes (59 frames) for GEO processing, the number of GEO measurements reported each clear night by each camera is about 1,000. Note that this number strongly depends on the number of frames stacked. If fewer frames are stacked, the rate at which measurements are reported goes up. However, the sensitivity of the stacked image is reduced. If more frames are stacked, than the reporting rate goes down, and the number of GEOs that don't move much over a longer time from goes down as well. We find that a 10-minute stacking cadence is a good match for our system that optimises the frequency and number of reported GEO measurements.

Figure 5 shows a measure of the 'capture rate' of objects in GEO. We used data from a night that was clear all night at four sensors (24 June 2018). For each stacked image, we propagated the orbits of all objects with recent TLEs from space-track and noted those objects that are a) predicted to be in the field of view of the sensor, b) predicted to have an apparent angular motion of less than a few pixels over a 10 minute period centered on the stacked image, and c) have a known radar cross section (RCS). We then matched those predictions to the observations and recorded how many 'known' objects were seen. A 'match' here means that the apparent angular coordinates of the prediction are within an arcminute of the observation (many matches are better than this). The objects were further binned according to a measure of their apparent size (RCS divided by the square of the slant range). This binning allows us to see if there is any correlation between 'missed' objects and apparent size.

Figure 5 shows the number of objects per bin matched to space-track objects (blue bars), and the number per bin that were not matched (red bars). A total of 1,472 out of 1,622 (90%) of the predictions were well matched to the observations. The 10% that were 'missed' are likely due to a combination of unfavourable viewing angle of the satellite (or the satellite is intrinsically faint) or its position was coincident with detector artefacts or other sources of

noise. We also see no correlation between apparent size and capture rate, suggesting that the choice of a 10 minute cadence is ideal for measuring objects in GEO.

The stacking technique described in Section 2 is not restricted to objects near GEO. Consecutive images can be shifted relative to one another before stacking, thereby enhancing the detectability of object that move across images with arbitrary velocities. We have demonstrated that this works by using the predicted apparent velocity of some GPS satellites. A series of 39 consecutive frames shifted to match their predicted velocity and combined together. The GPS satellites that were not detectable in individual images were easily identified in this 'digitally tracked' image. The challenge in implementing this in the pipeline is the computational expense. We are currently exploring other high-performance computing techniques, eg using GPUs, that would enable real time reporting of blind searches of combined images over a range of velocities. This would significantly enhance the number and capture rate of objects in all orbital regimes. We are also exploring a range of other machine learning and computer vision algorithms to maximise the number of objects that can be tracked in each image.

5. TRACK FORMATION, ORBIT DETERMINATION AND HANDOFF EXPERIMENTS

Fitting orbits to FireOPAL data and comparing predictions from those orbits to observations is a strong measure of how accurately the system performs. The near real time reporting of satellite positions from across the network enables us to form tracks from sets of measurements, to fit orbits to those tracks, to identify anomalies, and to hand off orbits to larger more sensitive detectors for followup observations.

We have conducted two sets of experiments to assess this powerful capability. One experiment demonstrated that FireOPAL can detect a satellite, form tracks as seen across the sensor network, fit an orbit to all the tracks, and then perform a 'live' handoff to another sensors - all occurring within minutes of the initial detection. This experiment was conducted in coordination with the Australia Defence Science & Technology Group (DSTG). Over three nights in March 2018, FireOPAL LEO data were screened to identify objects that were rising in the sky, were moving slower than 5 arcmin/sec, were well correlated to an object in space-track and were seen by more than one FireOPAL sensor. A total of 10 sets of observations were randomly selected from among the hundreds of suitable targets. For each set of observations (grouped together by their match to a space-track object), an initial orbit was determined using a simplistic, analytical two-body propagation model. The IOD was converted into a TLE and fed into the control system of a large DSTG optical sensor that was co-located with one of the FireOPAL sensors. The DSTG sensor was tasked, in an open loop, to track the TLEs within about 10 minutes of the initial FireOPAL observation. Out of the 10 objects handed off to DSTG, four of the objects appeared right in the middle of the small (<1°) field of view. The other six objects either went into shadow or set before they could be observed by DSTG. This experiment demonstrates the accuracy of FireOPAL data and its capacity to rapidly determine orbits and hand them off for followup characterisation.

The other experiment was run at a slightly slower pace but included a much larger set of data for handoff. A 'live' handoff experiment is not very sensitive to the quality of the fitted orbit – the orbit only has to be propagated for a few minutes and for a site that is near sites that initially observed the object. A more difficult challenge is to build orbits, propagate them for 24 hours or more, and then task a larger sensor on a later date. This was the focus of the second experiment, also conducted with DSTG. Each day, the central FireOPAL server grouped together sets of observations of both LEO/MEO and GEO measurements from the previous two nights. At the moment this grouping/association is done using the proximity of the observations to predictions from space-track. In the future, we will use more robust and flexible techniques to associate measurements with one another, eg. multi-hypothesis tracking. For each group, an initial orbit is fit to the data using the OREKit[2] high fidelity numerical propagator. The fitting procedure is an iterative batch least squares technique with outlier rejection. On a typical clear night, orbits for more than 100 LEO objects and 100 GEO objects are calculated. Over a two-week period in June 2018, the FireOPAL orbits from the previous night were handed off to the DSTG sensor for automatic scheduling for observation for the current night.

[2] https://www.orekit.org

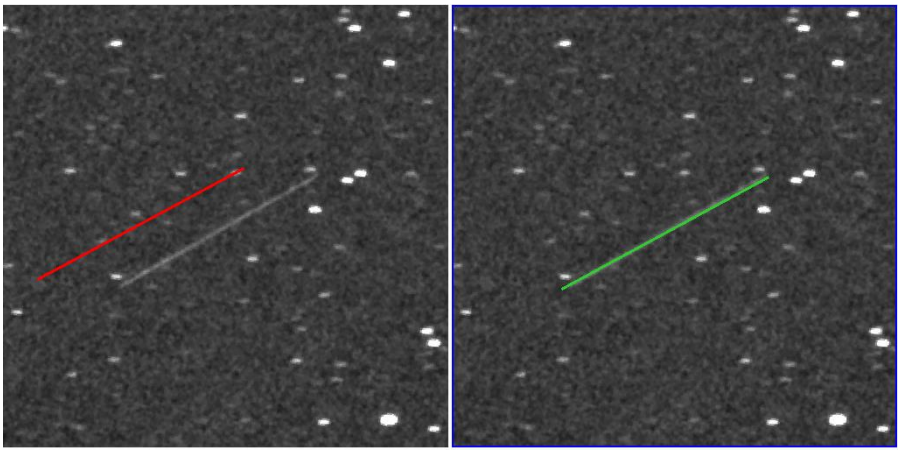

Figure 6: Comparison of orbit predictions from space-track and FireOPAL observations. The greyscale images are identical in each panel and subtend 1° by 1°. The image was taken at 5 June 2018 20:04:10UTC from a South Australian site. The images are centered on a LEO streak that is well correlated with ASIASAT 2 PKM (SSN# 23725), a rocket body in a highly eccentric orbit. The left panel shows the predicted location of this object (red) using the most recently available TLE (at the time) from space-track; the TLE epoch was 22.4 hours older than the image epoch. The space-track prediction is offset from our observation by about 9.8arcmin in the along-track direction (~60 pix) and 5.3arcmin in the cross-track direction (~33 pix). The panel on the right shows the predicted location of the object (green) using an orbit to FireOPAL observations that were taken across the network on the previous night: 70 observations spanning 20.9 hours, latest observation was 26.3 hours older than the image. Our prediction is offset from the observation by 9.2arcseconds (<1 pix) in the along-track direction, 22.5arcsec (~2 pix) in the cross-track direction.

More than 90% of the orbits handed off were successfully imaged by the DSTG sensor and appeared in the middle of their field of view. The targets that were missed could be attributed to two primary issues. One issue was that although we used a numerical propagator to fit an orbit, we had to convert this to a less accurate TLE in order to drive the DSTG sensor control system. The limitations of using TLEs to represent an orbit are well known. The other issue was that occasionally we would fit orbits to objects on highly eccentric orbits, where FireOPAL only observed it for a small part of the orbit near apogee. These objects had very low perigees (<400km) and subject to significant atmospheric drag – but that drag could not be reliably estimated from the data. Both of these are known issues that can be addressed by either avoiding the use of TLEs or by extending the time span of observations of objects on highly eccentric orbits.

The FireOPAL server continues to automatically fit orbits to all data after every night of observing. Over the period 1 June – 15 August 2018, more than 6,400 orbits have been calculated for more than 1,000 objects in the space-track catalogue. These orbits have been propagated to the next night and compared to observations across the FireOPAL network. An example of this is shown in Figure 6. Both panels show the identical cutout image around a streak associated with ASIASAT 2 PKM (SSN# 23725). This is a rocket body in a highly eccentric orbit (apogee ~18,000km; perigee ~300km) and suffers from significant atmospheric drag. The left panel shows the predicted location of this object using the most recently available TLE from space-track; the TLE epoch was about 22.4 hours older than the image timestamp. The space-track prediction is offset from our observation by about 9.8arcmin in the along-track direction (~60 pix) and 5.3arcmin in the cross-track direction (~33 pix). The offset from space-track is not the result of systematic biases in the calibration of FireOPAL data – a lot of predictions from space-track are matched to our observations within a few pixels.

The green line on the panel on the right shows the predicted location of the object using an orbit fit to FireOPAL observations that were taken across the network on the previous night. 70 FireOPAL observations spanning 21 hours were used to calculate the orbit; the latest data point was about 26.3 hours older than the image timestamp. Our prediction is offset from the observation by 9.2arcseconds (<1 pix) in the along-track direction, 22.5arcsec (~2

pix) in the cross-track direction. FireOPAL data can be used to determine or refine orbits for large numbers of objects, with orbits that are more accurate than those provided by space-track, and those orbits are reported more frequently than space-track.

6. SUMMARY

We have presented an overview of the current capabilities and performance of the FireOPAL observatories. We continue to explore new developments to further improve the network. Some examples include: employing lower noise cooled detectors, pan-tilt movable camera heads, embedded higher performance computing, advanced faint streak detection algorithms, and improved orbit determination.

The FireOPAL system is designed to be a cost effective optical SSA solution for objects in LEO, MEO, and GEO. Our observatories can be mass produced, are simple to deploy and maintain, and can operate autonomously in remote environments for extended periods. Although the imaging systems in an observatory are comparatively low cost and low resolution (compared to larger telescope solutions), multiple, synchronized, triangulated observations deliver highly accurate orbits – comparable to results from exquisite (single look) optical sensors. The additional benefit is that a distributed network is disruption tolerant, unaffected by weather, images a large fraction of the sky, and extends the optimal terminator observation period. A FireOPAL optical fence, covering the entire GEO belt, with triangulation and redundancy, can be built from a network of a handful suitably placed observatories. Several fences, distributed at different latitudes, can be built for a fraction of the cost of a single one-metre class telescope. It is our goal to build a global system that can follow objects multiple times per day, is capable of catalogue maintenance for a large fraction of all satellites and is able to detect anomalous events in space as they occur.

7. REFERENCES

1. Bland, P.A., Madsen, G.J., Bold., M. et al. FireOPAL: Toward a Low-Cost, Global, Coordinated Network of Optical Sensors for SSA. Proceedings of AMOS, 2018.
2. Calabretta, M. R. and Greisen, E.W. Representations of celestial coordinates in FITS. Astronomy & Astrophysics, 395, 1077-1122, 2002.
3. Høg, E., Fabricius, C., Makarov, V.V. et al. The Tycho-2 catalogue of the 2.5 million brightest stars. Astronomy & Astrophysics, 355, L27-L30, 2000.