

GPU-accelerated Faint Streak Detection for Uncued Surveillance of LEO

Peter C. Zimmer

Go Green Termite, Inc.

University of New Mexico

Mark R. Ackermann

Go Green Termite, Inc.

University of New Mexico

John T. McGraw

University of New Mexico

ABSTRACT

By astronomical standards, small objects ($<10\text{cm}$) in LEO illuminated by the Sun under terminator conditions are quite bright, depositing 100's to 1000's of photons per second into small telescope apertures ($<1\text{m}$ diameter). The challenge in discovering these objects with no a priori knowledge of their orbit (i.e. uncued surveillance) is that their relative motion with respect to a ground-based telescope makes them appear to have large angular rates of motion, up to and exceeding 1 degree per second. Thus in even a short exposure, the signal from the object is smeared out in a streak with low signal-to-noise per pixel.

Go Green Termite (GGT), Inc. of Gilroy, CA, in collaboration with the University of New Mexico (UNM), is building two proof-of-concept wide-field imaging systems to test, develop and prove a novel streak detection technique. The imaging systems are built from off-the-shelf optics and detectors resulting in a 350mm aperture and a 6 square degree field of view. For streak detection, field of view is of critical importance because the maximum exposure time on the object is limited by its crossing time. In this way, wider fields of view impact surveys for LEO objects both by increasing the survey volume and increasing sensitivity. Using our newly GPU-accelerated detection scheme, the proof-of-concept systems are expected to be able to detect objects fainter than 12th magnitude moving at 1 degree per second and possibly as faint as 13th magnitude for slower moving objects. Meter-class optical systems using these techniques should be able to detect objects fainter than 14th magnitude, which is roughly equivalent to a golf ball at 1000km altitude.

The goal of this work is to demonstrate a scalable system for near real time detection of fast moving objects that can be then handed off to other instruments capable of tracking and characterizing them. The two proof-of-concept systems, separated by $\sim 30\text{km}$, work together by taking simultaneous images of the same volume to constrain the orbits of detected objects using parallax measurements. These detections will then be followed-up by photometric observations taken at UNM to independently assess the objects and the quality of the derived orbits. We believe this will demonstrate the potential of small telescope arrays for detecting and cataloguing heretofore unknown LEO objects.

1. FINDING AND TRACKING LEO OBJECTS WITH OPTICAL TELESCOPES

SSA of LEO has to date been the domain of radar, especially in the case of uncued surveys for LEO objects, whereas optical surveys have been more suited to GEO. Optical techniques are not likely to supplant radar, but can supplement it, taking advantage of some of the strengths of optical techniques while accommodating shortcomings. The biggest shortcomings of optical techniques compared to radar for uncued LEO surveys are that optical telescopes must contend with weather and observe during terminator conditions. If we approximate the average terminator illumination duration as 1.5 hours after dusk civil twilight and 1.5 hours before dawn twilight, this leads to a 1/8 factor in coverage of the LEO volume accessible to radar deployed at the same place. Weather will further reduce this (conservatively) by 50%, leading to a total effective coverage of roughly 6% that of radar simply from sky and illumination conditions. The comparative strengths of optical SSA of LEO is that it is passive, using the Sun as the illuminator rather than having to project that power. Optical telescopes are also relatively inexpensive, making them easily deployed around the world, assuming that one can detect sufficiently small objects in the LEO

volume with small telescopes. We hope to be able to show over the course of this ongoing work that small telescopes (by which we mean meter-class or smaller) can contribute in exactly this way.

The work presented here was conducted as a SBIR Phase I study, so the results shown are based on models and simulations of telescope and detector performance. To the best of our ability we have tried to create these simulations so that they reproduce real data where we have it. Where there are unknowns, we make pessimistic estimates. The second phase of this project where these techniques will be demonstrated with real data has recently begun.

In the next section, we will examine the challenges of observing objects in LEO orbits with ground-based optical sensors. Section 3 details the hardware we will use to make detections along with the deployment and operational parameters that are key to making our techniques effective. Section 4 describes the detection algorithm and how implementation on GPUs makes it practical and inexpensive. Section 5 concludes with some possible ways that arrays of small optical telescopes could be deployed to address LEO SSA.

2. SIGNAL-TO-NOISE OF FAST MOVING OBJECTS

Objects in LEO under terminator conditions are actually quite bright by astronomical standards. Ackermann et al. 2003 [1] demonstrate the radiometry of LEO objects illuminated by the Sun at 90 degrees to the observer. Figure 1 shows the apparent magnitude of a golf ball (4.2cm diameter, 7 cm² projected area illuminated) at various distances above sea level. At 1000 km, the golf ball is a bit brighter than 14th magnitude in V, an intensity that would deposit hundreds if not thousands of photons per second into a small telescope. If an object this bright can be tracked, it is relatively easy to detect and measure. This brightness estimate does not take into account the typical albedo of objects in space, which is about two magnitudes fainter (13%) than the white of a typical golf ball. Making this correction, 14th magnitude corresponds to a softball-sized object at 1000 km.

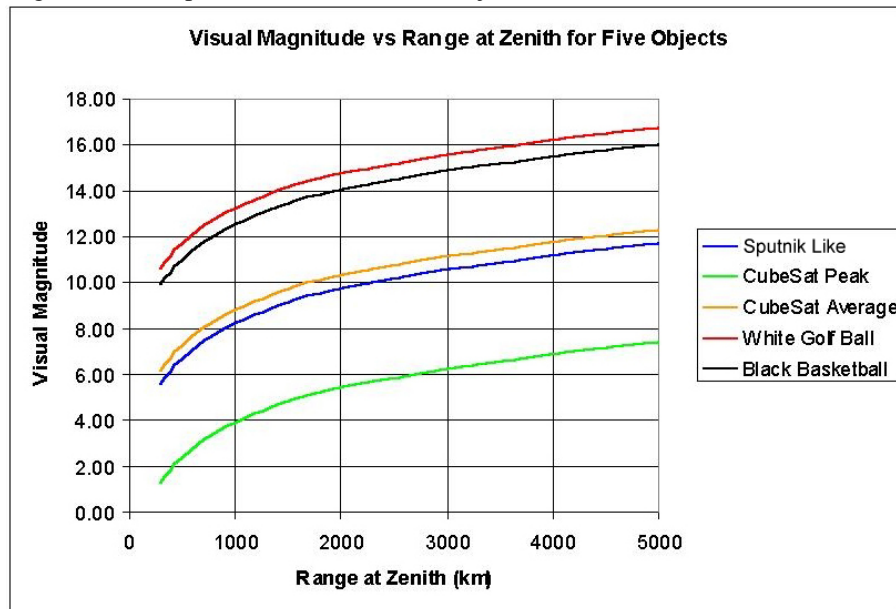


Figure 1- This figure (from Ackermann et al. 2003 [1]) shows the estimate brightness of canonical targets, one a 95% reflective golf ball (4.2cm diameter) and another a 5% reflective basketball (diameter) at various orbit altitudes under 90 degree phase (i.e. terminator) illumination.

The challenge with uncued observations of LEO objects is that their relative motion is fast enough in relation to their distance that they appear to move at a high angular rate. Figure 2 shows the apparent angular rate for objects in circular orbits as they appear when they pass through the zenith. For objects in the LEO regime from 500 - 2000 km altitude, the rates range from 0.2 – 1 degree per second. Thus even if those objects deposit thousands of photons into the aperture of a telescope, the light spread out across the detector in a streak that may be over a thousand pixels long, where the signal-to-noise per pixel is much lower than in the tracked case, rendering the object undetectable to most detection algorithms.

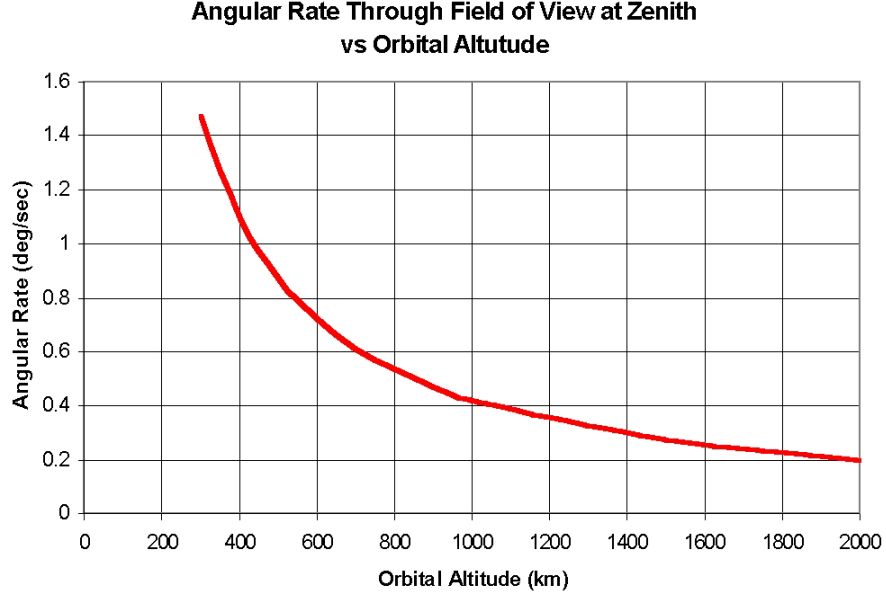


Figure 2 - This figure (from McGraw et al. 2003 [2]) shows the apparent angular rate of objects in circular orbits as viewed passing through the local zenith as a function of orbit altitude. This figure demonstrates the typical angular rates we expect for LEO objects.

In the uncued case, there is no *a priori* knowledge of streak direction or length, other than some constraints on the minimum and maximum rates for LEO objects. To render an observation useful, a few more constraints need to be applied. Simply finding a streak is not in itself very meaningful if the object traversed the entire field of view during the observation; the object could have entered and left at any time and also has an uncertainty in which way along the streak it went (the train track problem). If the object traverses half or less of the field of view during an exposure time, then it will leave endpoints in at least one image for most trajectories across the field of view. Preceding and subsequent images can break the along track direction degeneracy and measuring the endpoints constrains the velocity.

Much like traditional astronomical surveys, this puts a premium on field of view, both because it limits the useful integration time and it ultimately sets how much of the LEO volume can be surveyed. Until exposures are fully sky noise limited, meaning the number of sky counts per pixel is much greater than the variance of the read noise, the signal-to-noise of a streak continues to grow, asymptotically reaching the sky noise limit. This differs from the case of point source astronomy where the background contribution to the noise variance grows linearly with time.

$$S/N_{streak} = \frac{St_s}{\sqrt{St_s + R_{streak}t_s(Bt_{exp} + Dt_{exp} + \sigma_{RN}^2)}}$$

Equation 1 shows the signal-to-noise ratio for a streak imaged onto a detector, where S is source photon flux through the optical system onto the detector in photons per second; t_s is the time the streak is within the field of view of the optical system and t_{exp} is the exposure time, and in most of the cases we examine, the object making the streak is in the field of view both at the beginning and end of the exposure so these are the same. R_{streak} is the rate of growth of the streak in pixels per second – if the streak is undersampled this will include both the angular motion and the width of the streak. B and D are the sky background and dark current, the latter of which is negligible compared to the former for the detectors we consider. The last term, σ_{RN} is the standard deviation of the detector read noise.

For a streak, the effective size of the object is growing linearly with time compared to the signal to noise of a tracked object, so that the background variance is growing as the square of the exposure time. This time behavior in the background cancels out the linear increase in signal detected in the streak when the background rate is large compared to the read noise variance. For most telescope systems that we have examined, this crossover point between read noise and sky noise dominated growth happens after a few seconds. The consequence here is that until the exposures are background noise dominated, longer exposure times continue to yield increased detectivity.

To address these factors, we consider observing cadences such that an object moving at an apparent 1 degree per second angular rate will traverse the short side of a rectangular field of view in two exposures cycles. For the case that we will detail in Section 3, the system field of view is 3 degrees by 2 degree, thus we need a cadence of 1 second. This has important implications with respect to observation overheads, which can significantly reduce overall efficiency with shuttering and read out time, which we will also address in Section 3.

3. PROOF-OF-CONCEPT SYSTEMS

During the study phase, we investigated commercially available telescopes systems to try to find the most cost-effective telescope system to provide a real demonstration. We found that the most consistently limiting component were the detectors. Our budgetary constraints effectively eliminated backside illuminated CCDs, which are desirable for their high quantum efficiency though they do require a mechanical shutter and the readout time overheads are often of the order 100% for one second cadence images, reducing much of the photoelectric efficiency gain. The rapid cadence required drove us to examine 35mm format interline CCDs that are commercially available in cameras from several vendors, because they can be read out while the next exposure is accumulating, eliminating readout overhead. Non-interline, front illuminated sensors were eliminated as the readout time overhead reduced the allowable exposure times.

Having settled on a family of detectors, we examined camera/telescope combination to find the best etendue, collecting area (A) multiplied by the field of view (Omega), per dollar. The standout system in this analysis was a 14" Celestron Schmidt-Cassegrain optical tube with a HyperStar f/1.9 prime-focus corrector. Combined with a KAF-16070 4864 x 3232 sensor with 7.4 micron pixels, this gives a $3^\circ \times 2^\circ$ field of view with 2.4 arcsecond pixels, albeit with significant vignetting in the corners of the field. The net result is an instrumented etendue of almost $0.5 \text{ m}^2 \text{ deg}^2$ that can read out at 0.7Hz with $10e^-$ of read noise, available essentially off-the-shelf for around \$20,000. The cadence is slightly lower than the desired 1 Hz, but is the best that fits in the available physical volume and budget.

We use this hardware combination as the basis of our image simulations to process with our detection techniques. Using the estimated system throughput, a fiducial 12^{th} magnitude object is expected to deposit approximately 11,000 photons per second onto the detector. If we choose the worst case rate of 1 degree per second, the streak is roughly 1600 pixels long. When combined with expected sky brightness (four days from new Moon conditions at a moderately dark site), the resulting background rate is approximately 35 photons per second per pixel, leading to a noise RMS per pixel of 11.6 e^- . Thus the streak significance per pixel is about 0.6σ , which is pretty challenging for typical objection detection techniques and invisible to the eye. The images are slightly undersampled, but this is desirable for maximum detectivity.

Figures 3 through 7 show the results of one simulated 12^{th} magnitude streak. The first figure shows the full 16 megapixel field of the simulation, which includes stars along with the 12^{th} magnitude streak. The star field chosen is fairly dense at a galactic latitude of approximately 10 degrees. This region was chosen to test star light rejection, which is discussed below. Figure 4 shows a zoom in of the region indicated in the box in Figure 3. The stars are more apparent in this view and while the streak passes through this region, its low per-pixel significance renders it invisible. Figure 5 shows the same zoomed in region as Figure 4 but with the streak intensity enhanced by a factor of 10 (2.5 magnitudes). It is now clearly visible and would be easy to detect for just about any method.

Most moving object detection systems reject background stars and galaxies before attempting to find the objects of interest, be they artificial satellites in Earth orbit, minor planets in a solar orbit, supernovae or anything that goes bump in the night. The most common approach for current and planned astronomical surveys is to subtract a model of the sky based on previously acquired data (e.g. Drake et al. 2009 [3]). This is very effective, especially at finding objects like supernovae that would otherwise be blended with the host galaxy and has the advantage of better noise characteristics than a simple difference image. The downside of this from the standpoint of detecting streaks is that the objects being modeled and subtracted also contribute Poisson noise from their signal photons. For streak detection, stars are both a source of systematic error because chance star alignments can be confused for streaks but also the starlight itself contributes random errors. To mitigate this, we mask star contaminated pixels and effectively ignore them and their associated noise.

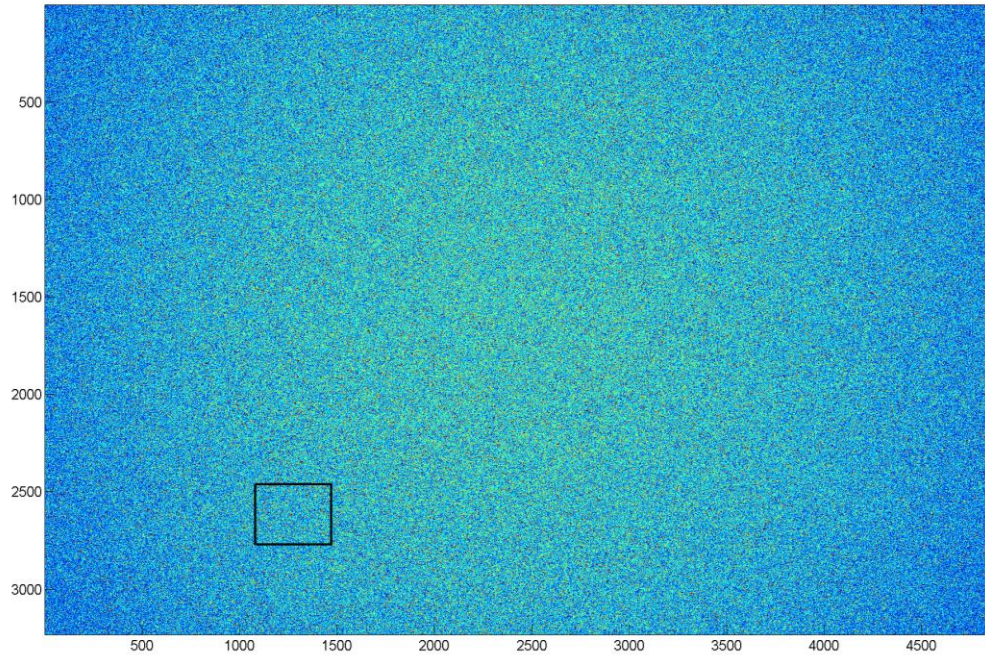


Figure 3 - Simulated image based on expected performance of the proof-of-concept telescope systems. This image shows the full 16 megapixels of the POC imager, covering 6 square degrees on the sky. This image also contains a simulate LEO object streak from a 12th magnitude object moving one degree per second. The effect of the vignetted field of view of the Hyperstar corrector is apparent in the corners of the field of view.

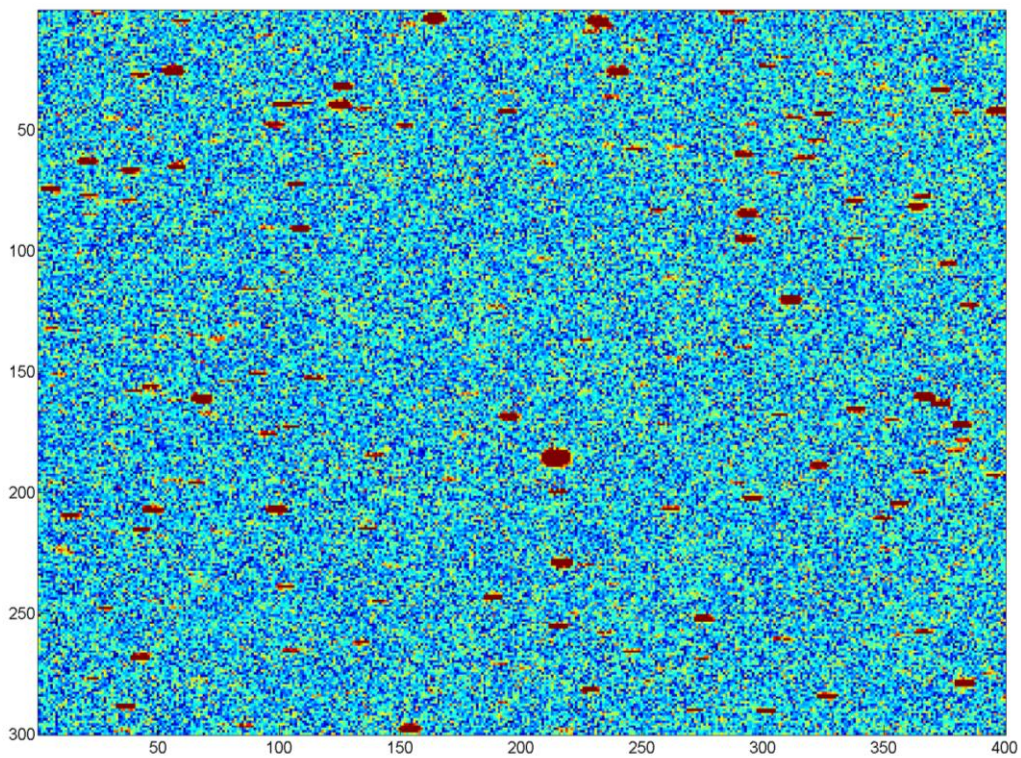


Figure 4 - This is a subframe of the image shown in Figure 3. The red spots are stars which are themselves streaked by a few pixels left-to-right (West-East) because the POC systems do not track sidereal motion. The 12th magnitude streak passes through this subframe, but is difficult if not impossible to detect by eye.

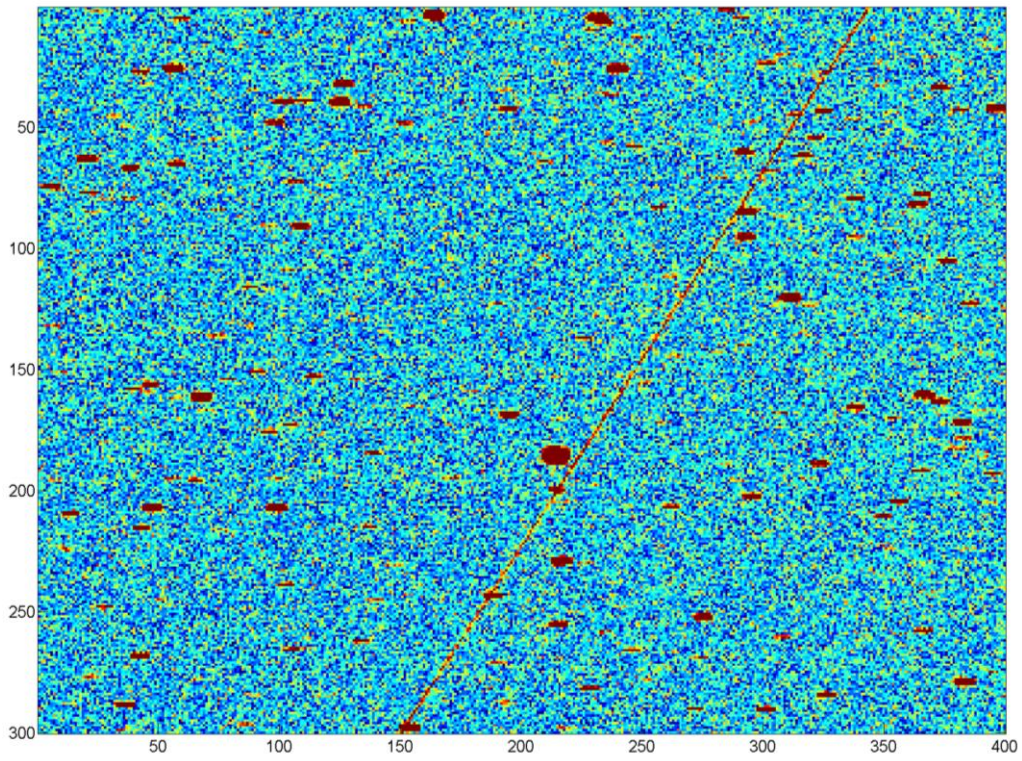


Figure 5 - The same as Figure 4, except that the streak has been enhanced by a factor of 10 (2.5 magnitudes) for visibility.

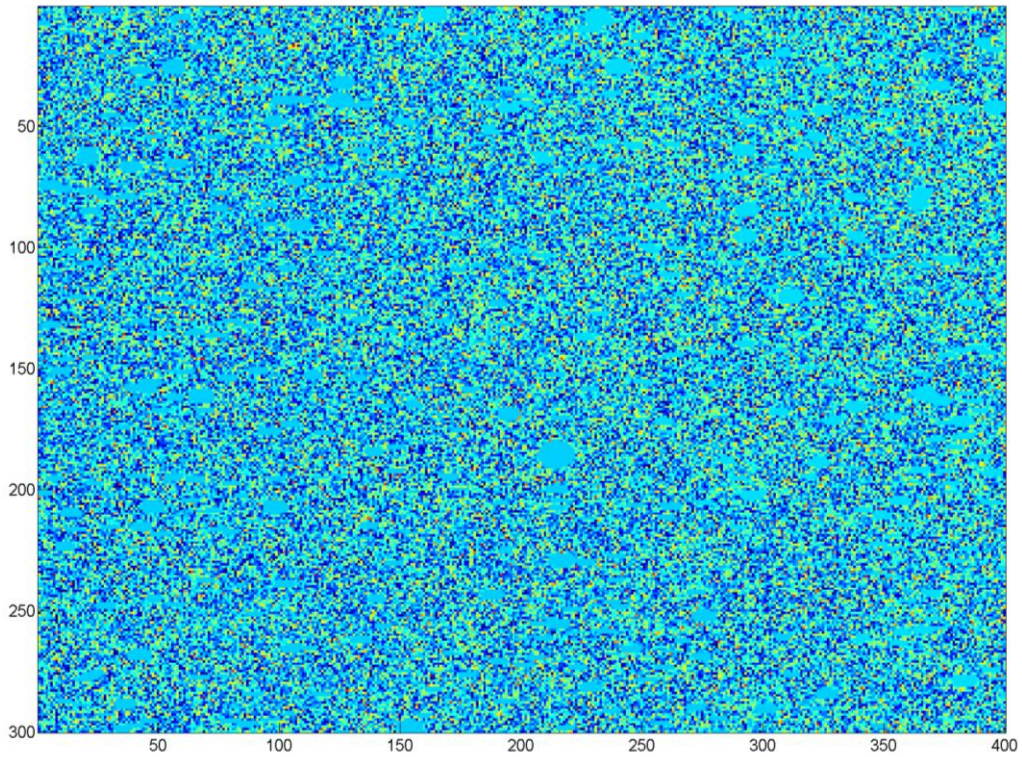


Figure 6 – The same subframe as Figure 4 but with the stars masked.

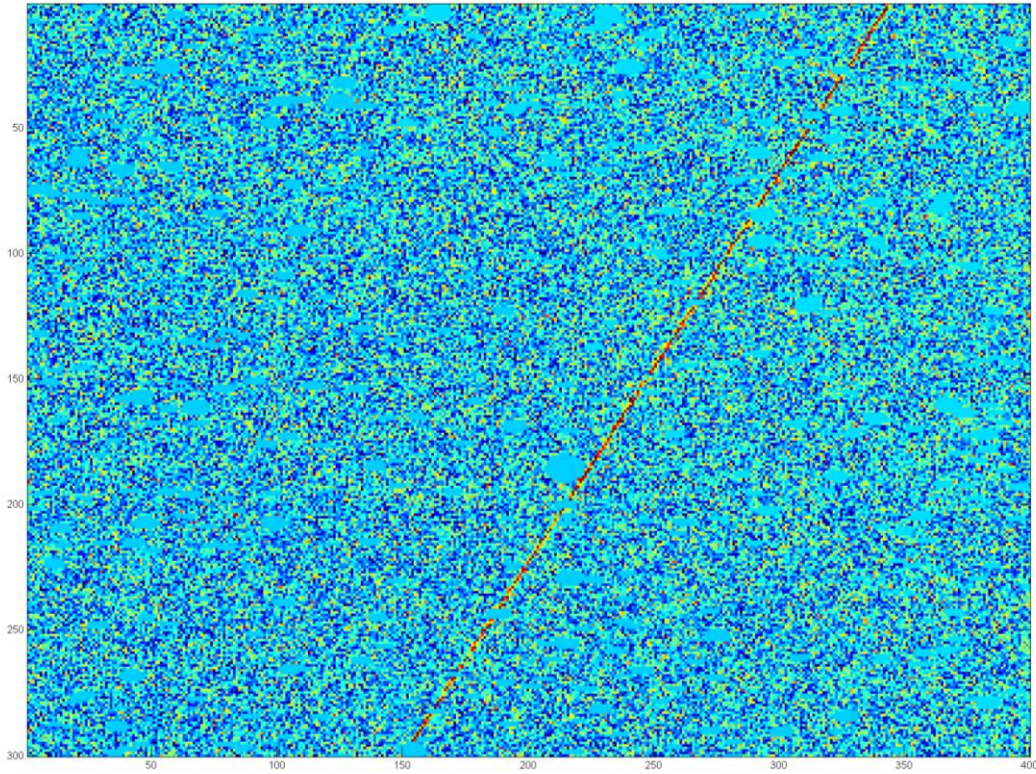


Figure 7- Masked subframe as in Figure 6 but with the streak enhanced for visibility. This shows that the vast majority of pixels in the streak are unaffected by the masking.

For small telescopes and the short exposure we are using, stars and galaxies that are bright enough to cause problems are well-catalogued already. Thus we can create simple masks at known positions for all stars expected to contribute more than one quarter of the expected noise RMS in photons. For typical star fields away from the galactic plane, this amounts to roughly 0.1% of the pixels being masked. At 10 degree from the plane, this number rises closer to around 1% and is likely feasible down to 5 degrees, but that will be confirmed with observations. A masked subframe of the same image as in Figure 4 is shown in Figure 6 and Figure 7 is the same but with the streak enhanced to show that the number of pixels in the streak that are also masked is small. With stars masked we can proceed to detect linear features in the image and the full frame.

4. GPU-ACCELERATED STREAK DETECTION

To detect streaks in images, we simply project line integrals across the image at various angles using a modified form of a Radon Transform (Radon 1917 [4], Dean 2000 [5]). We then compare those sums to what would be expected from just noise, which in our case is just the read noise and sky noise.

Dark current contributions for short exposures using modern detectors are negligible except for hot pixels or columns, which are masked and constitute a small fraction of the field of view even for low grade detectors. Detector bias levels are removed as they are in typical observations. Flat field type corrections are not applied to the image but are included in the sky background component of the noise model.

This approach has the benefit of covering all possible trajectory angles but does so at cost in detectivity compared to an optimal detection technique. An optimal technique would use only the pixels in the streak and none from outside it. Our detection transform projects across the entire image, which for most trajectories is at least twice as long as the streak and therefore distinctly suboptimal. This is a bigger effect for shorter streaks which an optimal detection

scheme would find as being a more significant detection for a given target magnitude -- the slower object leaves its light over fewer pixels and thus has a higher signal-to-noise. One could imagine performing matched filters using kernels with all possible lengths and orientations, but this would come at a very large computational burden. This Radon-based approach essentially does a matched filter except that the length is simply the dimension of the image along that path.

Based on simulated images with faint streaks, we have found that performing the detection transform over an angular grid of 0-180 degrees at 0.02 degree increments is sufficient to detect streaks expected to have a projected significance 6 times greater than the pure noise model would predict. Regions where a significant detection is found are re-run on a finer grid of 0.001 degrees over a restricted region, which improves precision of the measured angle for brighter streaks. Based on multiple simulations the precision of the recovered angles is 0.02 degrees at a signal-to-noise ratio of 6 and decreases proportionally to the signal-to-noise. Figure 8 shows the full detection transform for the image shown in Figures 3-7. Figure 9 zooms in on the detection of the 12th magnitude streak.

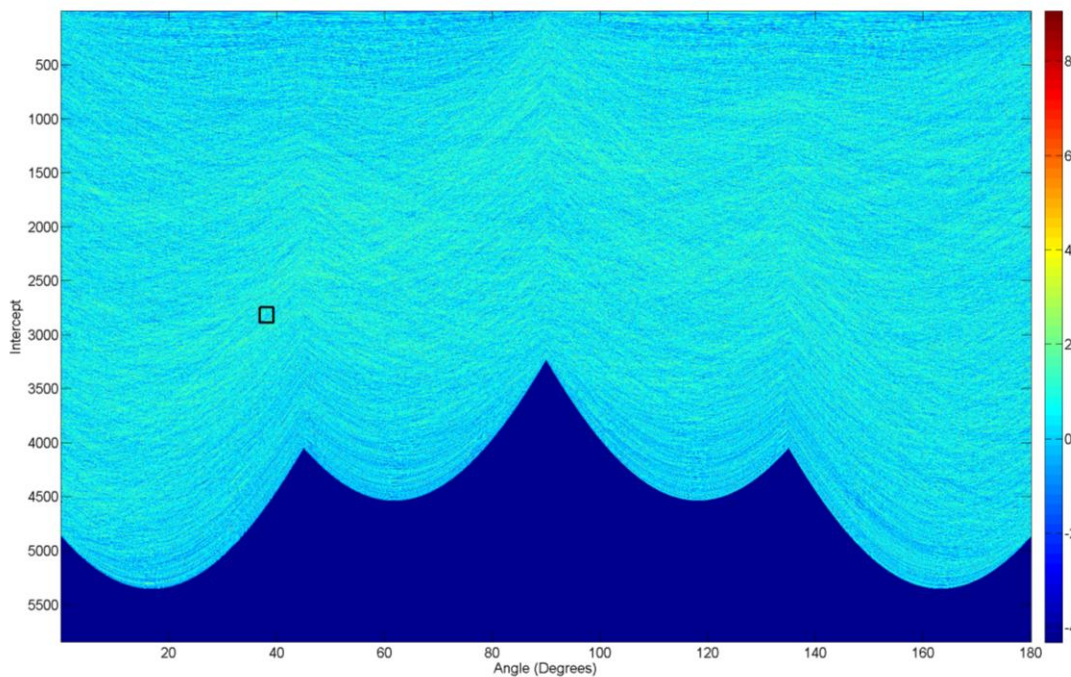


Figure 8 - This figure shows the entire 0.02 degree grid detection transform for the image shown in Figure 3. The color scale depicts the significance of the projected sum at that angle and intercept. The only pixels above 6 sigma of the noise model are inside the indicated black rectangle.

To enable rapid follow-up of detected objects, the process needs to be fast - essentially real time. In this case, for one second cadence image, the entire process must be similarly fast. Prior to parallelizing the detection transform for operation on Graphical Processing Units (GPU), the time to run the detections on a fairly powerful (for 2012) desktop PC running the built-in Matlab `radon()` function was hours for the 16 and 30 megapixel images run on 0.01 or 0.01 degree angular grids. This was slow enough to keep us from running them except when needed as a not-particularly-practical comparison to some other technique. These transforms require a large number of computations per image and scales as the number of pixels times the number of angles. For a 16 megapixel image and 0.02 degree grid, this is 150 billion calculations, each of which is a few dozen floating point operations.

It is no use working to use small, inexpensive telescope and detector hardware only to have to invest far more in computing to make them effective. Luckily, our detection transform is highly parallelizable and has a large number of symmetries which can be exploited. We have implemented it with the CUDA architecture (NVIDIA Corporation) for GPU programming and so far achieved a roughly 400x increase in processing speed. On a system with a pair of NVIDIA 680 GTX graphics cards (which presently can be bought for \$400 each) the detection transform runs in 7 seconds. While this is not quite fast enough to meet our real-time goals, we continue to develop it and the GPU

marketplace continually improves, often faster than Moore's Law, especially when looking at cost per floating point operation. Our best GPU implementation at present only uses 16% of the theoretical computational power because it is memory bound. GPU cards that have 3x the theoretical performance of our current cards are expected on the market within weeks of this writing with expected cost around \$1000 each.

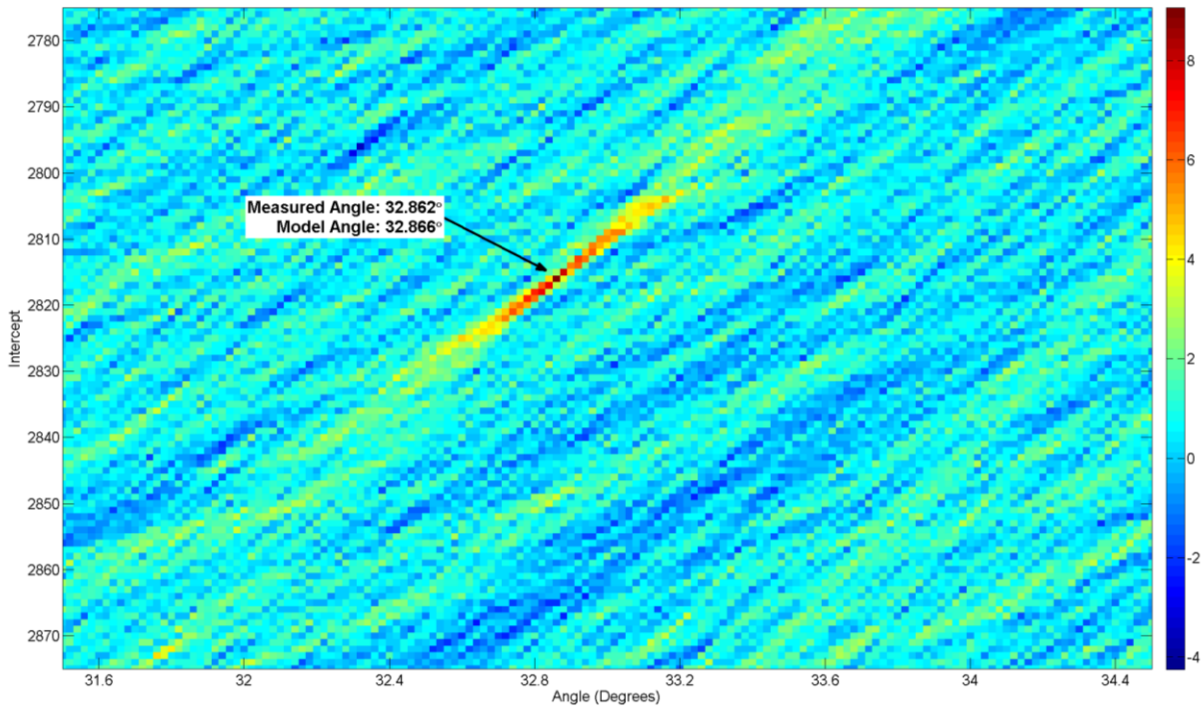


Figure 9 - This image shows the same transform as Figure 8, but zoomed into the streak detection. The measured angle is shown to match the input model streak angle very well. For a streak of this significance, the angular precision is usually a little worse than the example shown, approximately 0.015 degrees.

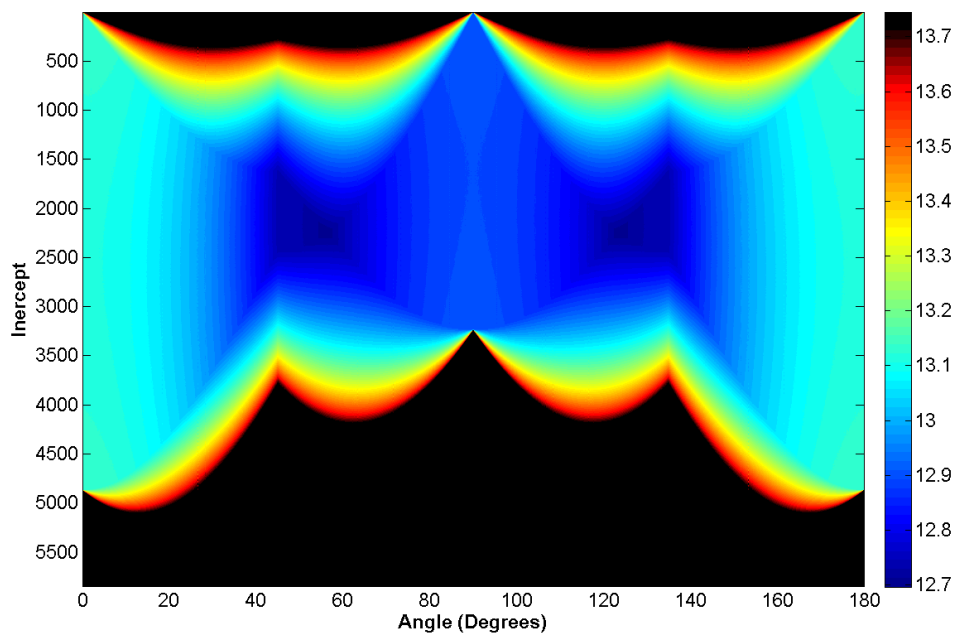


Figure 10 - This figure shows a map of the expected sensitivity of detection transform by transforming the noise model. At each angle and intercept combination, the 6 sigma noise threshold is shown, converted to a magnitude as observed by the Proof-of-Concept system.

The detection transform only determines the angle and intercept of the streak. Because we choose an exposure time to ensure we capture streak endpoints for most trajectories through the field of view, we must extract the pixels corresponding to the most significant angle and intercept. Then from that one dimensional set of values, we determine the most likely beginning and end of the streak.

This work is still under development. The initial indications are that for faint streaks near our detection threshold, the endpoint determination scales as the streak length divided by the optimal signal-to-noise. For the streak shown in Figure 3 and those following, the along-track end point error is ± 30 pixels at each end. The cross track position scales as the streak width, which is related to the sampled point spread function width, divided by the detection signal-to-noise. These errors tend to be less than half a pixel.

5. POTENTIAL APPLICATIONS AND DEPLOYMENT

We've chosen to deploy the two systems at sites separated by roughly 27km pointing in such a way to have almost completely overlapping fields of view in the LEO volume. Thus every detectable object is seen by both system simultaneously. The exposures are triggered off GPS clock signals so that the images are nearly simultaneous at each site. By comparing the detected positions at each site with respect to star positions, the range to the objects can be determined to a few hundred meters using parallax techniques (Earl 2005 [6] & McGraw et al. 2008 [7]). This means in one pair of exposures we can obtain a rough initial orbit determination, sufficient for follow-up observations on the same pass, if the detection and measurement can occur fast enough. We expect to demonstrate that the value of this extra information outweighs the cost of the extra system. The initial deployment setup is shown in Figure 11, which is a simple visualization built in the STK software package by Analytical Graphics Inc.

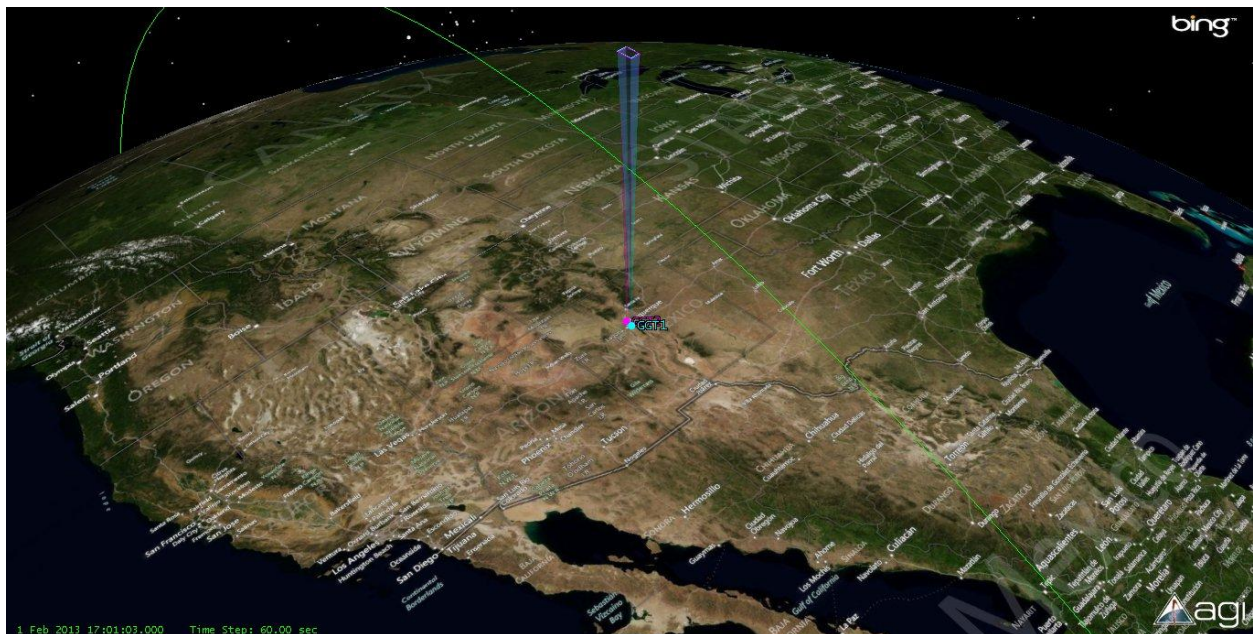


Figure 11- STK (TM AGI, Inc.) model showing the deployment of the proof-of-concept systems.

Ultimately an array of system like this should be able to cover significant volume of LEO down to 500km. So far, we've only considered observing at the local zenith. If the field is pointed to larger zenith angles, the effective distance to the object increases like the secant of the zenith angle (for small angles) and thus the brightness drops by the secant squared. If we limit the viewing angle to 30 degree from the zenith, the effect is less than 0.25 magnitudes. In this case, in order to have overlapping coverage down to 500km altitude orbits, neighboring array sites need to be roughly 450km apart accounting for the curvature of the Earth. This sort of arrangement is shown in Figure 12, another STK visualization which depicts four sites across the desert Southwest along the I-40 corridor: near Barstow, CA; near Flagstaff, AZ; near Albuquerque, NM; near Amarillo, TX – outside of the sky pollution halo of each city. Each of those four sites would host two arrays of 10 small telescopes, each of which has a $4^\circ \times 6^\circ$ FOV (close to the limit of the systems we have examined), where the array are separated by roughly 50km for parallax

observations. Above 500km, the arrays have significant overlap, and depending on the chosen goals, one could optimize deployment for either complete coverage to some altitude or to maximize independent volumes.

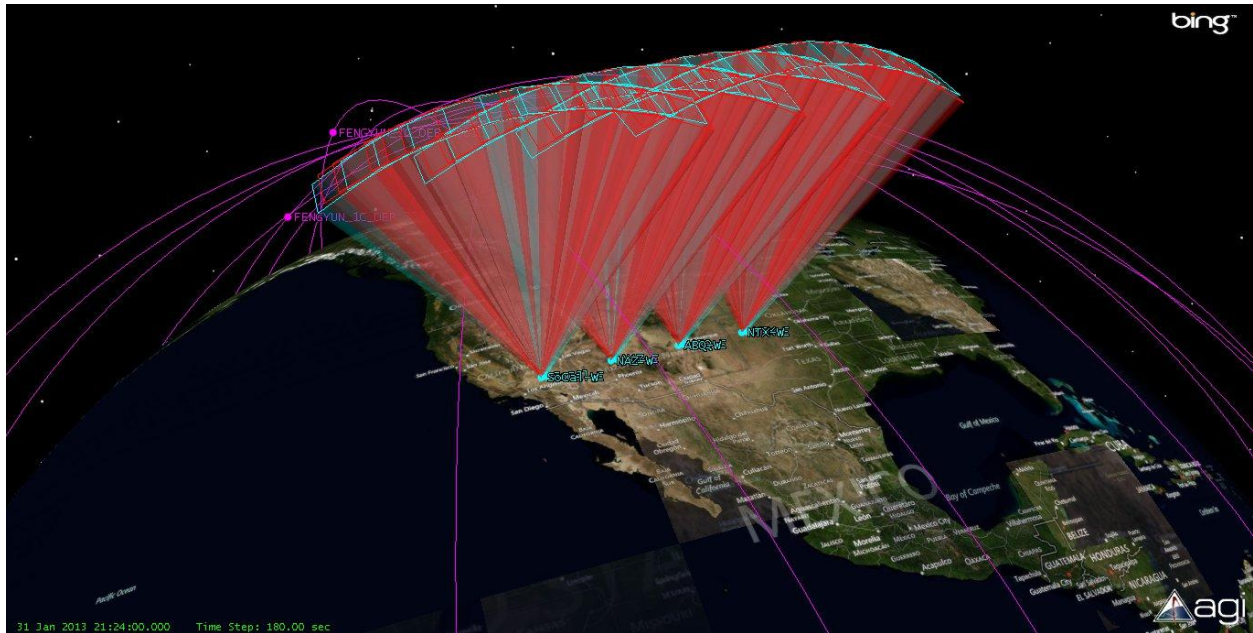


Figure 12 – Notional model of a desert southwest optical fence designed to cover the LEO volume over the western US down to 500km altitude. Each of the four sites (Barstow, CA; near Flagstaff, AZ; near Albuquerque, NM; near Amarillo, TX) consists of two arrays of 10 telescopes separated by 50km to provide parallax range and overlapping coverage.

Note that the arrays of telescopes we contemplate would all fit into a common enclosure and have very few moving parts because of they do not need tracking mounts. This arrangement means that these arrays could be very mobile, easy to deploy and maintain while keeping costs to a minimum.

The fraction of the total LEO volume instantaneously covered by an optical telescope depends on the altitude of the orbital parameter space. As altitude increases, a one square degree field of view covers fractionally more of the total parameter space at that altitude. On average throughout the volume, 1 square degree covers 8×10^{-7} of the total LEO volume. However, because the objects are orbiting, each square degree covers completely independent portions of the orbital parameter space every few seconds. Coupled to a large telescope field of view, it is possible for a single system to cover significant fractions of the LEO volume each clear day. However, some portions of the parameter space will always be unavailable to certain sites because of the site latitude or because of resonant orbits that never pass over a particular site during terminator conditions. Wide geographic distribution of telescopes is therefore crucial and some cleverness with regards to siting of telescopes to break resonant orbit aliases will be needed if optical telescopes are to contribute appreciably to LEO SSA.

6. ACKNOWLEDGEMENTS

This work has been supported by the Air Force SBIR program under contract number FA9451-13-C-0092. We'd like to thank our program manager for his support throughout this project. We'd also like to thank Mark Weeks of Go Green Termite, Inc. for providing the infrastructure to allow this work to proceed.

7. REFERENCES

- [1] Ackermann, Mark R., McGraw, John T., Martin, Jeffrey B., Zimmer, Peter C. 2003, *Blind Search for Micro Satellites in LEO: Optical Signatures and Search Strategies*, Proceedings of the 2003 AMOS Technical Conference, also published as Report No.: SAND2003-3225C, Sandia National Laboratories, Albuquerque, NM (USA), September 2003.
- [2] McGraw, John T., Ackermann, Mark R., Martin, Jeffrey B., Zimmer, Peter C. 2003, *The Air Force Space Surveillance Telescope*, Proceedings of the 2003 AMOS Technical Conference, also published as Report No.: SAND2003-3226C, Sandia National Laboratories, Albuquerque, NM (USA), September 2003.
- [3] A. J. Drake, S. G. Djorgovski, A. Mahabal, E. Beshore, S. Larson, M. J. Graham, R. Williams, E. Christensen, M. Catelan, et al., "First Results from the Catalina Real-Time Transient Survey," *The Astrophysical Journal* 696, 870–884 (2009) [doi:10.1088/0004-637X/696/1/870;].
- [4] J. Radon, "On the determination of functions from their integral values along certain manifolds," *IEEE Transactions on Medical Imaging* 5(4), 170–176 (1986) [doi:10.1109/TMI.1986.4307775].
- [5] Stanley R. Deans, "Radon and Abel Transforms," in *The Transforms and Applications Handbook*, 2nd ed., A. D. Poularikas, ed., CRC Press, Inc., Boca Raton, FL (2000). Chapter 8, pp. 1-95.
- [6] M. A. Earl, "Determining the Range of an Artificial Satellite Using its Observed Trigonometric Parallax," *Journal of the Royal Astronomical Society of Canada* 99, 50 (2005).
- [7] J. McGraw, M. Ackermann, P. Zimmer, S. Taylor, J. Pier, and B. Smith, "Angles and Range: Initial Orbital Determination with the Air Force Space Surveillance Telescope (AFSST)," presented at Advanced Maui Optical and Space Surveillance Technologies Conference, 2008, 48.