Shadow imaging efforts at MIT Lincoln Laboratory

Jane Luu, Leaf Jiang, Bert Willard
MIT Lincoln Laboratory

ABSTRACT

The technique of shadow imaging allows one to take pictures of distant objects that are beyond the reach of the largest ground-based astronomical telescopes. The technique relies on the fact that stars are nearly ideal point sources and are conveniently located behind man-made deep space objects, such as satellites. As a space object passes in front of a star (an "occultation event"), the object casts a shadow on the ground – a diffraction pattern that can be inverted to reveal the silhouette of the object. This idea of reconstructing the image from the shadow on the ground was first proposed by Burns et al. (2005, Proc. SPIE 5896), but the idea has never been implemented nor has its feasibility been studied in detail. There is no doubt that the technique is very challenging: the ability to predict the location of and capture meter-size shadows (roughly the size of the space object) that travel at 0.5 km/s across the surface of the Earth has never been demonstrated experimentally; furthermore, effects such as turbulence, background light, non-monochromatic light, and atmospheric dispersion have not been addressed or even considered. We have addressed these issues and our simulations show that, among other capabilities, the technique offers near diffraction-limited "imaging" from the ground, with the resolution limited only by the angular size of the occulted star. The technique is also remarkably robust against atmospheric turbulence. We are thus proposing shadow imaging as an alternative method to image deep space objects that elude ground-based optical telescopes and radars, such as non-spinning geosynchronous satellites. Most man-made satellites can be imaged with centimeter resolution by existing radars, but geosatellites that display very little motion cannot be imaged by radar and there is still no good solution for this problem. Shadow imaging can fill this imaging gap, and compares very favorably to current capabilities in both capability and cost. We are setting up an experiment to capture the first geosynchronous satellite occultation in order to prove the feasibility of shadow imaging.

1. INTRODUCTION

Imaging (i.e., resolving) geosynchronous satellites is beyond the current capability of deployed radars and optical telescopes. Modern imaging radars [1], making use of a technique known as range-Doppler imaging, do not have the required pulse energy (~ 40 J), carrier frequency (~ 100 GHz, higher carrier frequencies can access geosynchronous orbits with smaller inclinations), or waveform bandwidth (~ 4 GHz) to reach geosynchronous satellites. Current optical telescopes cannot resolve geosynchronous satellites. For example, a 3.6-m telescope with adaptive optics has a diffraction-limited resolution of 13 meters at the geosynchronous orbit altitude of 36,000 km.

The ability to image a satellite is crucial as it is the most direct way to identify a satellite and its function. Understandably, the inability to image spin-stabilized satellites has attracted a lot of attention, and proposed solutions range from ground-based (e. g., a very large radar array to cover the entire geosynchronous belt) to space-based (e. g., a Hubble-like space telescope). However, the proposed methods tend to be very large scale and costly and the problem of imaging geosynchronous satellites remains largely unsolved. In a 2005 SPIE paper, Burns et al. proposed the idea of shadow imaging [2]. In this paper we explore the concept in more detail and show that its resolution could be improved by an order of magnitude using spectrally-resolving detectors.

The shadow imaging concept is shown in Fig. 1. The technique relies on the fact that stars are nearly ideal point sources and are conveniently located behind man-made satellites. It is an extension of the occultation method, which makes use of the fortuitous events when the target of interest – any space object, whether man-made or natural – passes in front of a star and blocks out the light of the star (in astronomical parlance, the star is "occulted" by the object). While the target remains in front of the star, it casts a shadow in the surface of the Earth. Due either to the Earth's rotation or the motion of the target itself, the shadow moves across the Earth's surface. If the target is a GEO satellite, it will be mostly stationary overhead while the stars move East-West behind it, so the shadow would move West-East. The shadow is simply the diffraction pattern produced by light from the star being diffracted around an obstacle, and which can be inverted to reveal the silhouette of the object. The situation is completely analogous to the familiar textbook example of light being diffracted by an aperture in an opaque screen. In an occultation, one has the opposite situation: instead of an aperture in an opaque screen, there is an obstruction in a transparent screen. The resulting diffraction pattern shows a drop in the star's intensity where the obstacle blocks out the light. The shape of the drop is roughly the same as the shape of the obstacle, just as geometric optics predicts, but superimposed on the dip are ripples which are purely a diffraction effect. The shadow can be captured by an array of telescopes aligned in a North-South direction. Each telescope serves essentially as a light bucket that captures the intensities in the shadow as the star moves across the sky. Since the shadow traverses the telescope at 0.5 km/s, an occultation event for a 10-m-long satellite lasts only 20 ms, requiring the telescopes to be equipped with a fast detector (e.g., a photon-counting detector) to temporally resolve the event.

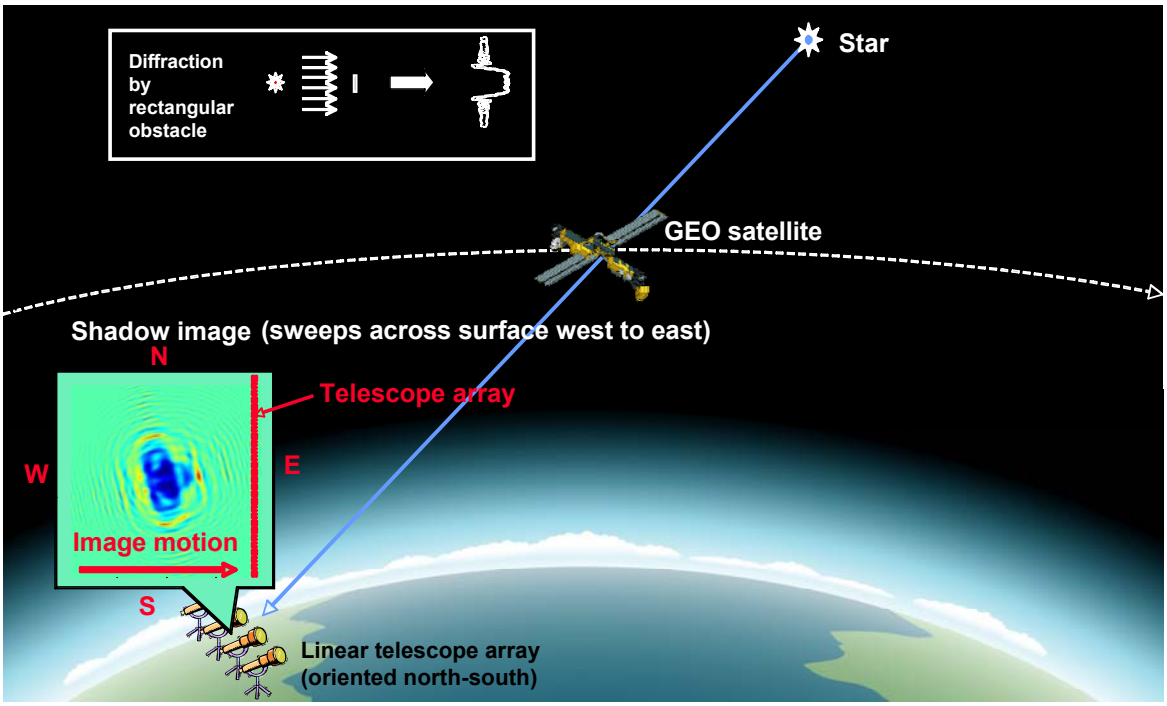

Fig. 1. Shadow imaging concept.

Several chromatic effects, such as atmospheric dispersion and non-monochromatic star light, smear out the diffraction pattern and limit the resolution of the technique. By using capturing the shadow for each wavelength and using this information to reconstruct the silhouette, we can remove degradations due to atmospheric dispersion and non-monochromatic star light. We call this technique "spectrally-resolved shadow imaging," since it is a refinement of the shadow imaging technique. It can be achieved by simply replacing the receiver with a spectrometer consisting of a grating and a Geiger-mode avalanche photodiode array.

2. RESOLUTION LIMITS

Shadow imaging as described by Burns et al. assumes ideal conditions: 1) point-source stars, 2) monochromatic light, 3) no atmosphere, and 4) noiseless detection. In reality, starlight is polychromatic, stars are not point sources and therefore do not produce perfect planar wavefronts, atmospheric effects affect all ground-based observations, and finally, the detection of light always introduces noise. These factors are the main physical phenomena that degrade the resolution of a shadow image, and they are explored in detail below. We cannot change the size of stars, but we can minimize the effects of polychromatic light and the atmosphere via a modified version of shadow imaging which we call "spectrally resolved shadow imaging." The theoretical performance of shadow imaging is summarized in Table 1.

Angular size of stars. Although stars can usually be considered point sources for most purposes, they have finite angular sizes and therefore do produce perfectly planar wavefronts on the ground. The star diameter affects the occultation intensity patterns in that (1) they broaden the occultation width, (2) reduce the occultation depth, and (3) they smooth out the diffraction ripples. All these factors not only render the occultation signature more difficult to detect, they also reduce the resolution of the reconstructed satellite image. Recent advances in optical interferometry have increased the number of measurements of stellar diameters, and Fig. 2 plots the stellar diameters from the largest sample measured to date [3]. It should be noted that the 85 stars plotted in Fig. 2 correspond to some of the brightest and closest stars, and therefore, the data is biased towards larger angular extents than the average star. The smallest stars in Fig. 2 have diameters ~ 10 nrad, limiting the shadow imaging resolution to ~ 4 cm.

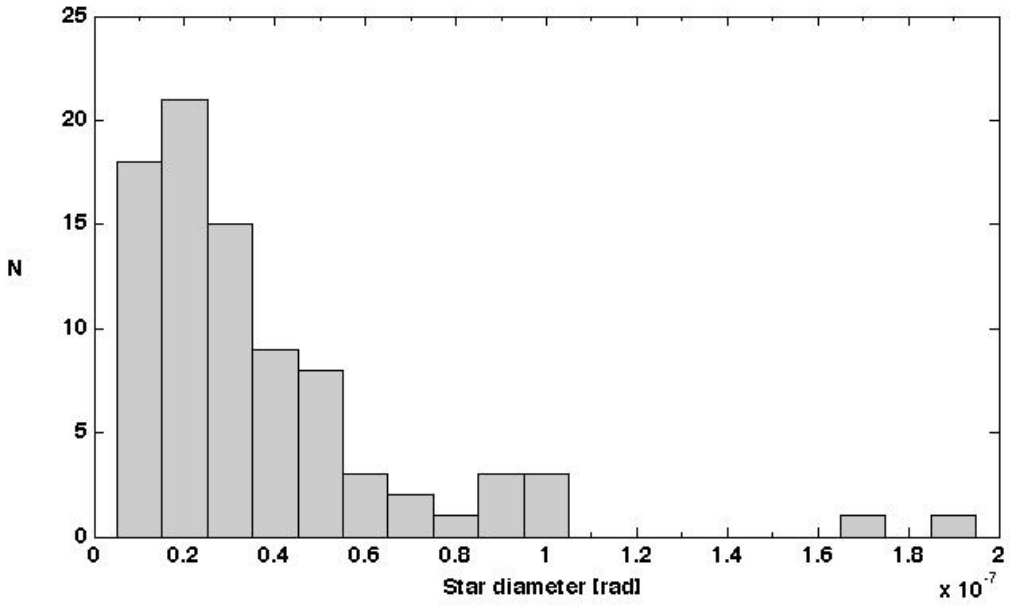

Fig. 2. Distribution of 85 measured star diameters (Mazurkewich et al. 2003).

Non-monochromatic star light. Starlight is highly polychromatic, and each wavelength produces diffraction ripples at its own distinct spatial frequency. The 2D intensity pattern of the shadow on the ground is a summation of all the ripples of the different wavelengths, and the net effect is that the diffraction ripples, which contain all the high resolution features in a shadow image, will start washing out at some distance away from the origin, resulting in blurred images. (There is thus, for polychromatic illumination, a maximum array length over which no additional telescopes improve resolution.) The blurring due to polychromatic light is illustrated in Fig. 3, which shows that as the spectral width of the star increases, the image becomes more blurred.

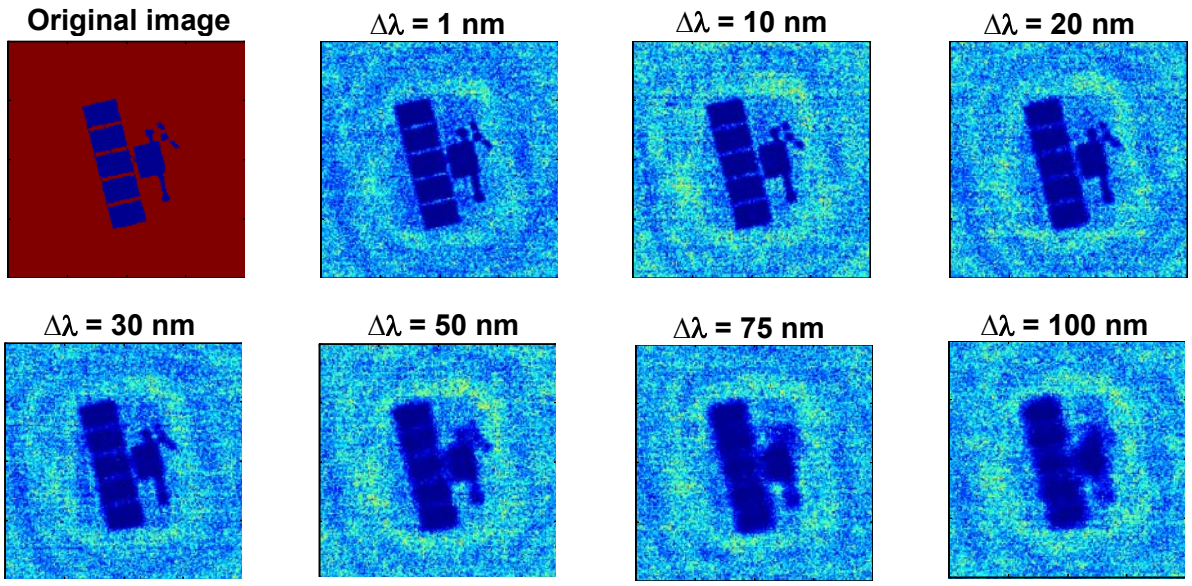

Fig. 3. Shadow image of the satellites in 7 different bandpass filters, each with the same central wavelength, but different bandwidth. The bandwidth of each filter is indicated above each image.

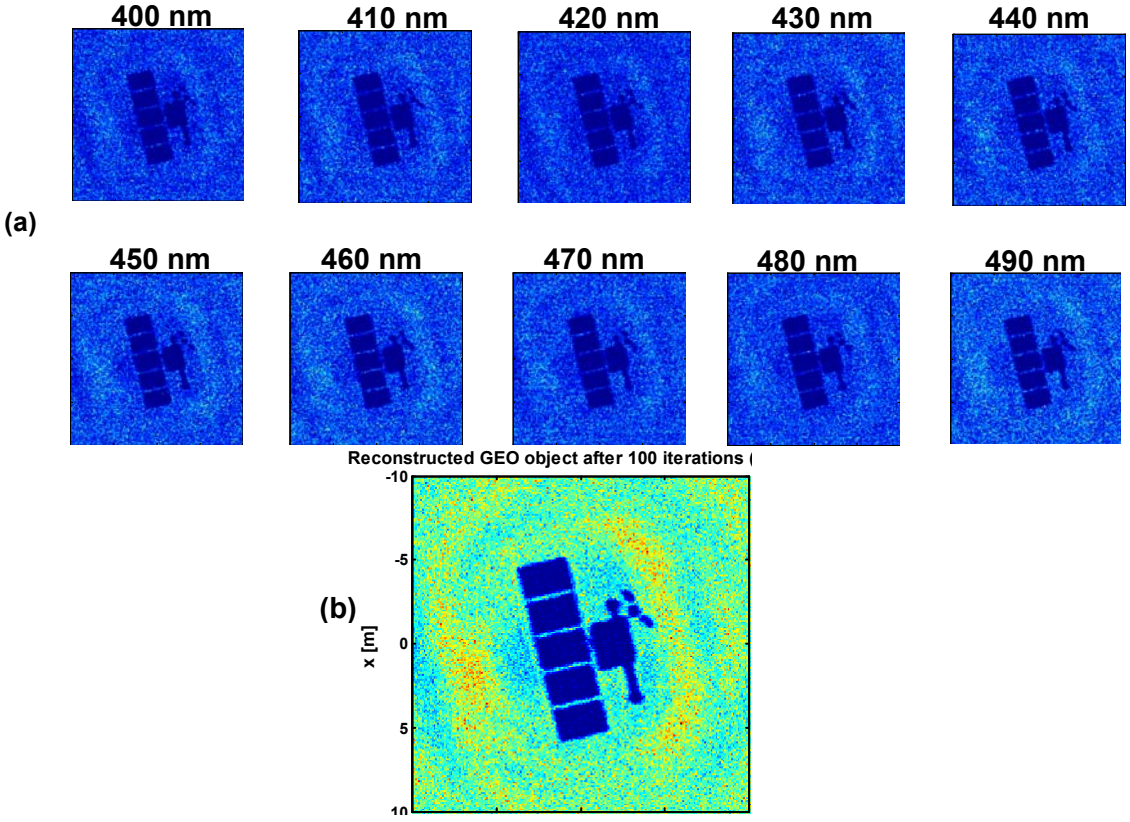

Fig. 4. (a) Shadow image of the satellites in 10 different bandpass filters, each with 10 nm bandwidth. The center wavelength of each filter is indicated above each image. (b) The sum of the narrowband images shown in (a). Although the image is effectively a broadband image (wavelength range = 400 – 500 nm), it retains the resolution of a narrowband (Δλ = 10 nm) image.

In order to recover the ripples (and thus recover the full resolution capability of the shadow imaging technique), it is crucial to use narrow bandpass filters (~10 nm) to prevent ripple washout. Unfortunately, the narrow bandpass filters also mean that less light is incident on the detector, and the shot noise limits shadow imaging to only the brightest stars (mag 7.5 or brighter). One way around this challenge is to use an array of photon-counting detectors to capture light from different wavelengths – much like a spectrometer. Now, instead of a single shadow image, there are N shadow images each corresponding to a different wavelength, where N is the number of spectrometer elements, say each corresponding to a 10 nm chunk of spectrum. The spectrally-resolved shadow imaging reconstruction algorithm then proceeds as follows: (1) generate initial guess for the 2-dimensional E-field for all wavelengths, $E_1(\lambda_i)$, where $i = 1 \ldots N$ denotes a particular wavelength; (2) propagate the N $E_1(\lambda_i)$ from the satellite to ground plane, resulting in $E_2(\lambda_i)$; (3) apply the N amplitude constraints at the ground plane (from the N detectors) to $E_2(\lambda_i)$, and store result in $E_3(\lambda_i)$; (4) back propagate $E_3(\lambda_i)$ from the ground to the satellite plane; (5) apply flat phase constraint to the $E_3(\lambda_i)$; (5) repeat steps (1) – (4) until the shadow image converges. An illustration of this technique is illustrated in Fig. 4, which shows (a) 10 narrowband images ($\Delta\lambda$ = 10nm) covering the 400nm – 500nm range, and (b) the sum of these images. Although the effective bandwidth of the final image (Fig. 4b) is 400nm – 500nm, the resolution is that of a 10 nm bandwidth image. Note that the dark counts increase by a factor of N since there are now N detectors for the same measurement, and the computation time increases by a factor of N. Nevertheless, this is a small trade for enhanced resolution.

Atmospheric refraction. Atmospheric refraction refers to the fact that the atmosphere acts like a prism and bends any light that enters at an angle away from the zenith (Fig. 5). Assuming that the angle of refraction is small, atmospheric refraction R depends only on the refractive index of air n at the telescope, and on the zenith angle z, defined as the angle between the zenith and the source of the light: $R = (n-1) \tan z$. For wavelength λ = 0.5 μm, a zenith angle z = 45º (typical for geosatellites observed from 42º, the latitude of MIT Lincoln Laboratories), and 20 km for the effective thickness of the atmosphere, the refraction angle is R = 279 μrad, i.e., the satellite shadow will be displaced by 5.6m on the ground. This offset can easily be accounted for when calculating the location of the shadow. However, the refractive index n also depends on the wavelength, so the shadow on the ground will also be smeared due to the spread in refraction angle caused by different wavelengths – an effect called "differential refraction" or "atmospheric dispersion". The resulting shadow on the ground is the sum of all the shadows created by each wavelength, each shadow slightly shifted from the next one (in the North-South direction) due to refraction. If left uncorrected, the shadows from different wavelengths will add incoherently, and as with polychromatic light, the diffraction ripples will be smeared out, resulting in a blurred image. But in spectrally-resolved shadow imaging, at each telescope the light is dispersed into several narrowband components, with each component being detected by its own photodetector. It then is straightforward to run the image reconstruction algorithm for each separate wavelength band (as in dealing with polychromatic light, see previous section), then add them together after applying the proper refraction shift. In addition, given the latitude of the observer site, temperature, and pressure, the atmospheric dispersion can be calculated and used in the reconstruction algorithm.

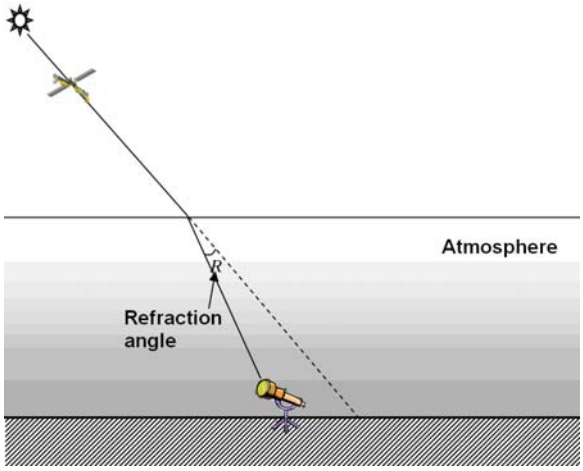

Fig. 5. Atmospheric refraction.

Shot noise. As the target becomes smaller, the contrast of the shadow decreases. At some point the contrast becomes smaller than the shot noise of the detectors, and this defines the resolution of the image. Shot noise causes random fluctuations in a measurement due to the random arrival time of the photons. The contrast C is defined as the peak-to-peak excursion of the diffraction pattern divided by the mean return signal without an obscuring target. Using this definition, the shot-noise-limited resolution is reached when $C \geq 1/\mathrm{SNR}$, where SNR is the signal-to-noise ratio. Now, for a random process like the detection of photons, $\mathrm{SNR} = \sqrt{N_{pe}}$, where N_{pe} is the number of collected photoelectrons per pixel and we have assumed no dark counts. For an ideal small square target, $C = 2w'^2$, where w' is the width of the target (in units of FSUs). Combining the two expressions for C, the shot-noise-limited resolution is given by $w(m) \geq \frac{\sqrt{\lambda D/2}}{4N_{pe}}$. For a wavelength of 532 nm and altitude of $z = 36 \times 10^6$ m, the shot noise-limited resolution is shown in Table 1. Other effects (such as angular extent of the star, pixel size, and atmospheric dispersion) will limit the resolution when the number of received photoelectrons per pixel is greater than ~10.

Table 1. Theoretical performance of shadow imaging.

Effect	Resolution Limit *	Solution
Angular size of stars (Section 2.1)	30 cm (50% of stars)	Use stars with small angular extent
	4 cm (10% of stars)	
Non-monochromatic light (Section 2.2)	146 cm (50 nm bandwidth)	Spectrally-resolved shadow imaging
	49 cm (5 nm bandwidth)	
	15 cm (0.5 nm bandwidth)	
Atmospheric refraction (Section 2.3)	5.6 m (at 42 ^{o} latitude)	Spectrally-resolved shadow imaging
	0 m (at 0^0 latitude)	
Shot noise (Section 2.4)	1.0 m ($m_v = 9.4$)	Brighter stars
	55 cm ($m_v = 6.9$)	
	31 cm ($m_v = 4.4$)	

* The reported resolution limit assumes a 14-inch telescope aperture and is a function of the star's visual magnitude (m_v) and the corresponding number of photoelectrons per pixel (N_{pe}).

3. PROPOSED TELESCOPE CONFIGURATION

To implement the shadow imaging technique, we propose a linear array consisting of 150 telescopes, each with a diameter of 30 cm (~ 12"), spaced 30 cm apart, yielding a baseline 45m long in the North-South direction. The North-South orientation is needed because the shadow moves East-West. The telescopes should be on a railroad track, so that all 150 telescopes can be moved as one single unit. The need to move the telescopes stems from one simple reason: to increase the occultation rate of geosynchronous satellites. The rate at which geosynchronous satellites occult stars depend on how fast the stars traverse the sky and the star density. The former (15 arcsec/sec) is set by the Earth's rotation rate, while the latter depends on the star brightness and the time of the year (i.e., whether the Milky Way is overhead or not). Bright stars are more scarce than faint stars. For a 10^{th} magnitude star (visual magnitude), the average star density (averaged over all sky and over a year) is 1.1 star deg^{-2}; the chance that the shadow of a given geosynchronous satellite happens to fall on a single telescope is thus 0.035 per day (assuming a 12-hr night), i. e., once every 287 days. This unacceptably low rate can be improved if the telescope could be moved: this is tantamount to having more telescopes available to catch shadows. If we want, on average, 1 occultation per day (per given satellite), the telescope should be able to move ~ 2 km. This illustrates the need for the shadow imaging telescopes to be on a railroad track so that they can be moved easily to the shadow's location. It is clear that the further the telescope can move, the higher the occultation rate becomes.

Each of the 150 telescope should be able to operate independently, with its own detectors and accompanying detector electronics. The detectors on each telescope include a CCD camera for acquiring the occultation star, and a linear array of occultation detectors. Since the occultation event lasts on the order of a few milliseconds, the occultation detectors need to be able of integration times on the order of microseconds. A good choice would be photon-counting detectors such as Geiger-mode avalanche photodiodes (APDs) (currently used in our proof-of-concept experiment).

The aperture size of each individual telescope also influences the ultimate resolution of the shadow imaging technique. One may naively think that larger telescope apertures would improve the performance of the system, as they would allow the system to make use of fainter stars. Although this is true, larger telescope apertures can actually degrade the quality of the resulting image since they would sample multiple diffraction ripples. For a small square target with width w such that its width is much smaller than the Fresnel scale unit (FSU) = $\sqrt{\lambda z/2}$, where z is the altitude of the object, the intensity of the shadow on the ground for monochromatic illumination is given by [4]

$$I_{mono}(x,y,\lambda)=1+\frac{w^4}{(\lambda z)^2}-\frac{2w^2}{\lambda z}\sin\frac{\pi\left(x^2+y^2\right)}{\lambda z}. \tag{1}$$

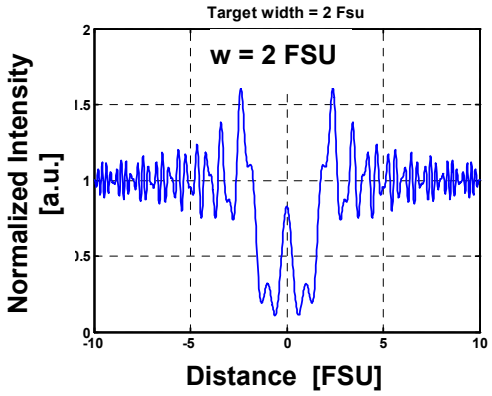

Fig. 6. Occultation intensity profile measured across the center of the shadow for a square satellite with width w = 2 FSU, where 1 FSU = $\sqrt{\lambda z/2}$ = 3m (for z = 36000 km and λ = 0.5μm) .

For the geosynchronous altitude of 36000 km and λ = 0.5μm, 1 FSU = 3m. Fig. 6 is a plot of Equation (1) with y = 0, λ = 0.5μm, w = 2 FSU. As the Figure shows, the period of the diffraction ripples decreases away from the origin, so given the requirement that each telescope samples only one ripple, the aperture of the telescopes ultimately limit the useful baseline length of the telescope array. The period of the ripples as a function of distance from the origin can be obtained from period of the sine term in Equation (1) and is equal to $2\lambda z/x$. Setting the ripple period equals the telescope diameter, the telescope aperture equals this period at an offset x_0 determined by the expression $D=2\lambda z/x_0$. The telescope array baseline length is thus limited to $2x_0=4\lambda z/D$. Hence, the aperture-limited resolution for a baseline equal to $2x_o$ is

$$\frac{2.44\lambda z}{2x_0}=\frac{2.44\lambda z}{4\lambda z/D}=0.61D. \tag{2}$$

For example, 14-inch telescope apertures would limit the resolution to 22 cm. Decreasing aperture size can yield better resolutions at the cost of decreasing signal.

Given the proposed 45m baseline of the shadow imaging telescopes, it is appropriate to compare its resolution with that of a conventional 45m optical telescope. Although no 45m optical telescope currently exists, there are plans to build 50m – 100m telescopes in the future [5]. Fig. 7a shows a simulated image of the same satellites used in our previous simulations, this time as would be obtained by a 45m ground-based telescope, with no atmosphere present. This can be compared with the shadow image in Fig. 7b, assuming a 7.5-mag star, 150 30cm telescopes and

monochromatic light. The two images have comparable resolution, as expected. In Fig. 7c, we show a simulated image from the same 45m telescope, but this time under an atmospheric turbulence characterized by r_0 = 3.9cm. (The parameter r_0 is used to describe the level of atmospheric turbulence, the smaller the r_0, the more turbulent the atmosphere. The very best observing sites on Earth have r_0 ~ 40cm, whereas r_0 ~ 3 cm for most other places). The turbulence was simulated by four phase screens located at 10, 100, 1000, 10000 meters. The constants for the turbulence model was obtained using the HV 5-7 model [6]. Fig. 7c shows a completely washed out image, with no recognizable feature, whereas the shadow image (Fig. 7d) shows little degradation under the same turbulence. This robustness against atmospheric turbulence is not surprising: turbulence corrupts mostly the phase of the wavefront, but the shadow imaging technique does not make use of phase information and hence is relatively unaffected by atmospheric turbulence.

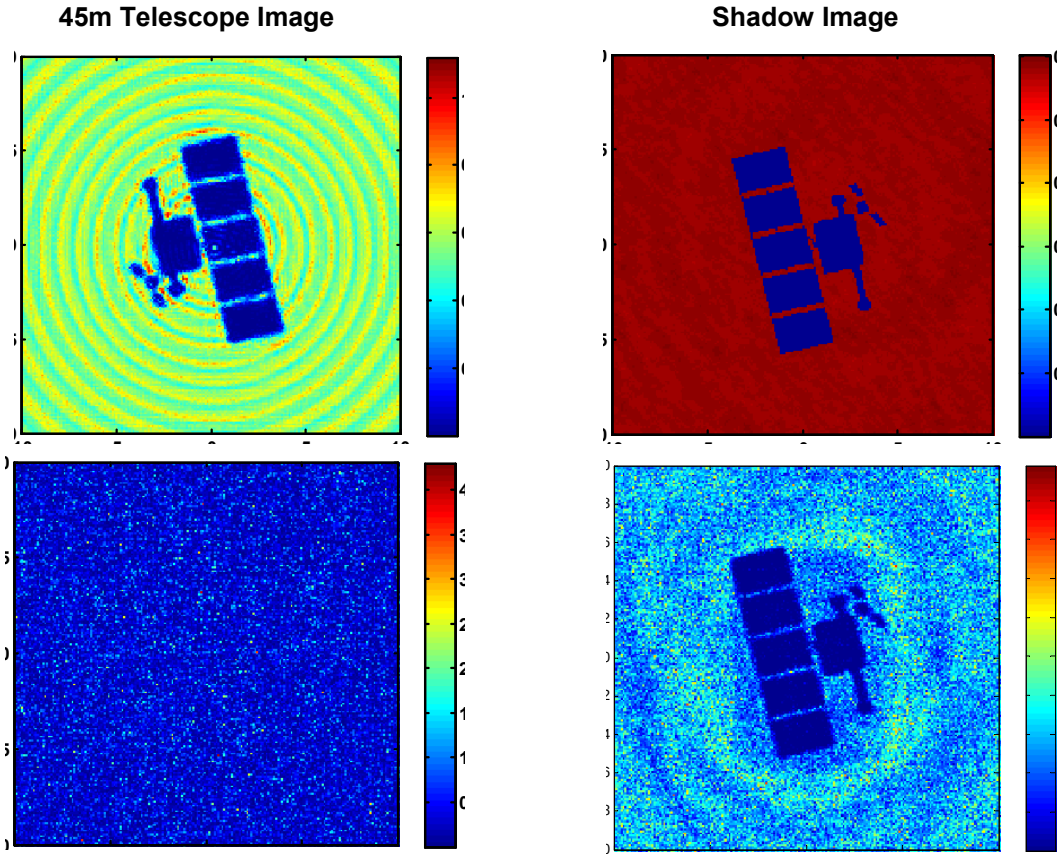

Fig. 7. Simulated image of the satellites from a 45m ground-based telescope (left column), compared with (their shadow image (right column). (a) and (b): The upper images assume ideal conditions: on atmosphere and no shot noise. (c) and (d): The lower images include turbulence represented by r_0 = 3.9 cm and shot noise. The telescope image is washed out by the turbulence, but the shadow image is not.

4. CURRENT SHADOW IMAGING EFFORTS AT MIT LINCOLN LABORATORY

Shadow imaging has never been attempted before -- in fact, no one has ever tried making observations of an occultation by a satellite. This is not surprising, given the considerable challenges. We are currently trying to prove the feasibility of shadow imaging by attempting to observe an occultation event with a single telescope. Our telescope is a 35cm telescope (the Celestron C-14), driven by a Paramount ME mount made by Software Bisque. The detectors consist of a 4096 x 4096 pixel CCD (the Apogee Alta U9000) and a 170μm-diameter silicon avalanche photodiode (APD) (the Perkin-Elmer SPCM-AQRH-15). The observations will yield a single chord of

the diffraction pattern; if successful, it is a (relatively) simple matter to extend the technique to multiple telescopes. It will also be the first observations of an occultation event by a geosynchronous satellite.

This work is sponsored by MIT Lincoln Laboratory, under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, recommendations and conclusions are those of the authors and are not necessarily endorsed by the United States Government.

5. REFERENCES

1. Camp, W. W. et al, Wideband Radar for Ballistic Missile Defense and Range-Doppler Imaging of Satellites, MIT LINCOLN LABORATORY JOURNAL, Volume 12, Number 2, 2000.
2. Burns, R. H. et al, Shadow imaging of GEO satellites, PROCEEDINGS OF THE SPIE, Vol. 5896, 2005.
3. Mozurkewich, D. et al, Angular diameters of stars from the Mark III optical interferometer, THE ASTRONOMICAL JOURNAL, Vol. 126, 2502-2520, 2003.
4. Roques, F. et al, Stellar occultations by small bodies: diffraction effects, THE ASTRONOMICAL JOURNAL, Vol. 93, No. 6, 1549-1558, 1987.
5. Ardeberg, A. et al, From Euro50 toward a European ELT, PROCEEDINGS OF THE SPIE, Vol. 6267, 626725, 2006.
6. Hardy, J. W., *Adaptive Optics for Astronomical Telescopes*, New York, Oxford University Press, 1998.