Automated Algorithms to Identify Geostationary Satellites and Detect Mistagging using Concurrent Spatio-Temporal and Brightness Information

Phan Dao[1]
Air Force Research Laboratory
[1]Correspondence

Elisabeth Heinrich-Josties and Todd Boroson
Las Cumbres Observatory Global Telescope Network

Abstract

Automated detection of changes of GEO satellites using photometry is fundamentally dependent on near real time association of non-resolved signatures and object identification. Non-statistical algorithms which rely on fixed positional boundaries for associating objects often results in mistags [1]. Photometry has been proposed to reduce the occurrence of mistags. In past attempts to include photometry, (1) the problem of correlation (with the catalog) has been decoupled from the photometry-based detection of change and mistagging and (2) positional information has not been considered simultaneously with photometry. The technique used in this study addresses both problems. It takes advantage of the fusion of both types of information and processes all information concurrently in a single statistics-based framework. This study demonstrates with Las Cumbres Observatory Global Telescope Network (LCOGT) data that metric information, i.e. right ascension, declination, photometry and GP element set, can be used concurrently to confidently associate (identify) GEO objects. All algorithms can easily be put into a framework to process data in near-real-time.

1. Introduction

Non-statistical algorithms which rely on fixed positional boundaries for associating objects, such as the standard Report Observation Association (ROTAS) codes, work well with the uncongested space and small space objects catalog of the past. However, the catalog, as well as the population of objects not listed in the catalog, has grown in size substantially. To complicate the situation further, satellites in allocated GEO clusters are actively kept in slots and require frequent station-keeping. As a result, orbital information often becomes stale as soon as or soon after an ephemeris is generated, which may result in cross-tagging and corrupt reported metric and photometric data.

In optical images taken by Las Cumbres Observatory Global Telescope Network (LCOGT), one telescope pointing may show multiple satellites, but identification cannot be done with certainty given only positional information derived from the most recent TLEs- two or more satellites may be closer than the uncertainties in their calculated positions.

In this paper, we test an algorithm that combines both metric and photometric data to identify satellites in a system that can be set up to process data and provide answers in near real-time. The tests were conducted on LCOGT data taken at various cadences and focused on the stressing scenario of GEO clusters and the occasional presence of non-resident objects. Our use of metric data is different from techniques which rely on improving the ephemeris— the latter typically suffer from an inherent time lag from the last maintenance maneuver.

2. Data collection and cadences

Las Cumbres Observatory Global Telescope Network (LCOGT) owns and operates a network of 0.4m optical telescopes located in Chile, Maui, Australia, and the Canary Islands (Spain) that are used for SSA applications. Each telescope [Figure 1] has an SBIG STX-6303 CCD camera that provides a FOV of 0.3 by 0.5 degrees. Pixels are 1.16

arcseconds square on the sky. All observations are made using a Pan-STARRS w filter, with a central wavelength of 665 nm and a width of 442 nm.

Data collection for this study emulates operational satellite observations. To maintain the orbital parameters of a satellite, up to 6 observations are collected per day with a cadence of at least one hour. This is to ensure that a significant portion of the orbit is sampled. In addition, observations are scheduled at different times across different days for the same reason. The scheduling algorithm utilized schedules observations on a rolling basis by first checking if a site is experiencing good weather, if there are any conflicting observations, and if at least one hour has passed since the previous observation. This process somewhat randomizes the sampling and is able to fill in the orbit.

Since geostationary satellites appear to remain over the same spot on Earth, they are always visible from a given sensor if above the sensor limits. Consequently, there are portions of the day when the satellite is not observable in optical wavelengths-when the Sun is up. This amount of time can be minimized by using multiple sensors located at different longitudes, provided the satellite is visible from those sites. Even though LCOGT has this ability, solely the 0.4m telescope in Chile is used for this study.

Figure 1. LCOGT network telescope.

During an observation, the telescope tracks the satellite making it a seeing-limited point and the background stars streaks [Figure 2]. The centers of the star streaks correspond to the position of that star during the middle of the exposure (i.e. 1 second into a 2 second exposure). To determine the coordinates of the image, each streak center is measured and matched to stars in a known catalog using astrometry.net [2]. Once a successful solution is reached, the position of the satellite is measured. This technique produces positions accurate to within 0.8 arcsec. Image time stamps are accurate to 10 milliseconds. Each exposure is processed to remove the electronic bias or pedestal and is corrected for the pixel-to-pixel sensitivity variations.

Figure 2. Frame image with satellite and star tracks.

The wide-band brightness of the satellite is measured by integrating its flux using pre-existing astronomical image processing packages [3][4]. This flux is converted into an astronomical magnitude using Eq. 1 by referencing to a zero-point determined from catalogued stars of known brightness. The zero-point of each telescope is a constant known value.

$$mag = 2.5 \log_{10}\left(\frac{flux}{exp\,time}\right) + zeropoint \qquad (1)$$

Each measurement is grouped into time, w-mag, and J2000 right ascension and declination coordinates. In this report, we use the term tracklet to denote a metric and photometric measurement associated with a single signal integration interval. For geosynchronous objects, a tracklet corresponds to a fixed point in geocentric coordinates.
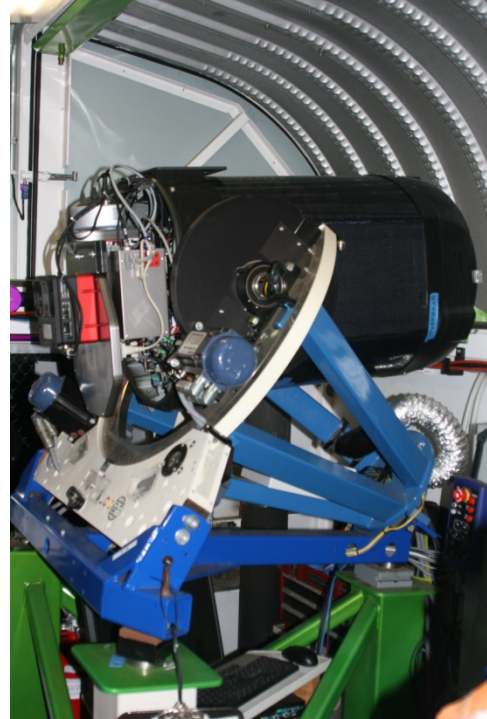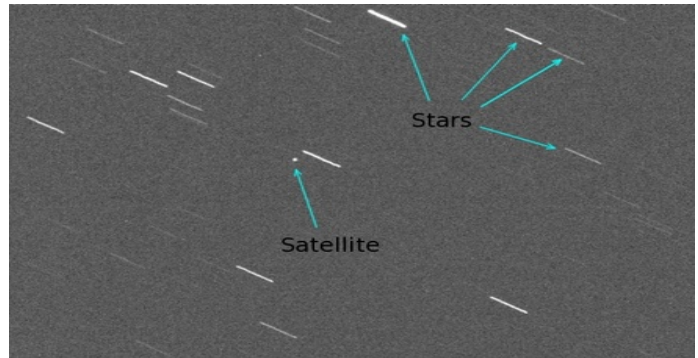
We collected three separate sets of data. Two sets are of the three ANIK satellites in the cluster located at 107W: satellites Anik F1 (NORAD #26624), Anik F1R (NORAD #28868) and Anik G1 (NORAD #39127). Observations in the first set span from August to December 2015, and depending on the weather conditions at site, they cover up to about 10 hours per night. In the second set, high cadence data of the same cluster was taken on Dec-10-2015. The last set was of the 101W cluster, the results of which are not presented in this paper.

3. Concurrent spatiotemporal and brightness technique (COSTB)

Correlating the tracklets with orbital element sets in the satellite catalog is a challenge when there are multiple objects in the sensor's field of view and orbital information is not current or accurate enough to distinguish one from another using positional information. The problem is worsened by the occasional intrusion of objects which are not resident to the cluster. We propose using both positional information and reproducible photometry to distinguish the stabilized satellites and associate the tracklets with their identities. Reproducibility of photometry is expected of three-axis stabilized satellites.

Positional information can be represented by an object's apparent longitude and declination. Apparent longitude is derived from right ascension and TLE-based range. The main idea behind the COSTB technique is the use of both sources of information to infer a satellite's identity. When the evidence presented by one source of information is inconclusive, we can check the evidence from the other source. This form of information fusion is performed via p-values, which are metrics of likelihood.

P-values are calculated based on the differences between expected and measured positions/magnitudes. The likelihood of a measured difference also depends on the empirically derived distribution of differences. COSTB does not provide better identification by improving the accuracy of the satellite's orbit. It concurrently considers positional trends, spatiotemporal information, and photometry to perform association and correlation. It also refrains from assigning an identity when the evidence is not probabilistically robust.

4. COSTB in production

4.1 Generating baseline photometry or model light curves

The flow of COSTB in production is shown in Fig. 3. Before kicking off automated satellite identification, model light curves must be generated for each satellite, and those light curves must be tagged correctly. The algorithm uses observed trends of apparent longitude to do so. At this time, trends in declination have not been used effectively to identify the satellites- occasional station-keeping maneuvers affect the trends of declination more so than those of longitude.
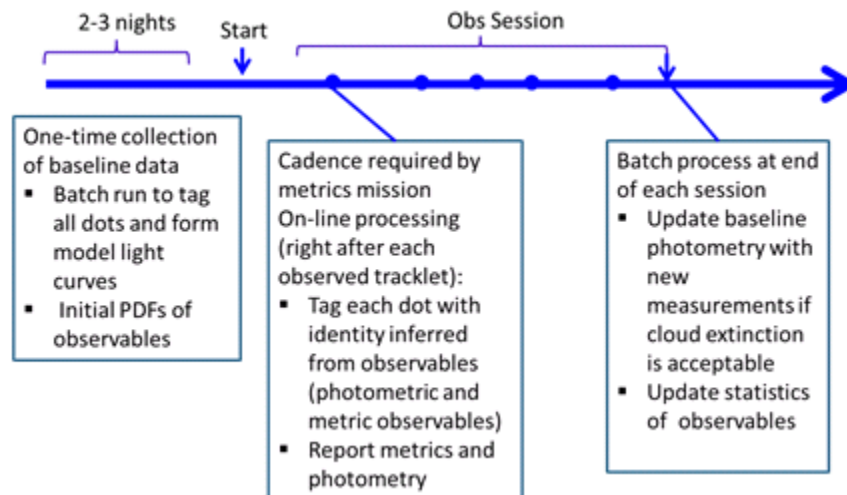

Figure 3. Flow chart of COSTB processing

While orbital information contained in two-line element sets (TLEs) is not fresh, not sufficiently accurate and sometimes not at all representative if generated before a maneuver, the measured positions still display trends that qualitatively fit those of the element sets. Measurements from Aug-13-2015 as well as corresponding TLE derived positions of the three resident satellites in the 107W cluster are shown in Fig. 4. The trends displayed by the measured longitudes show qualitative similarity to the TLE-predicted values.
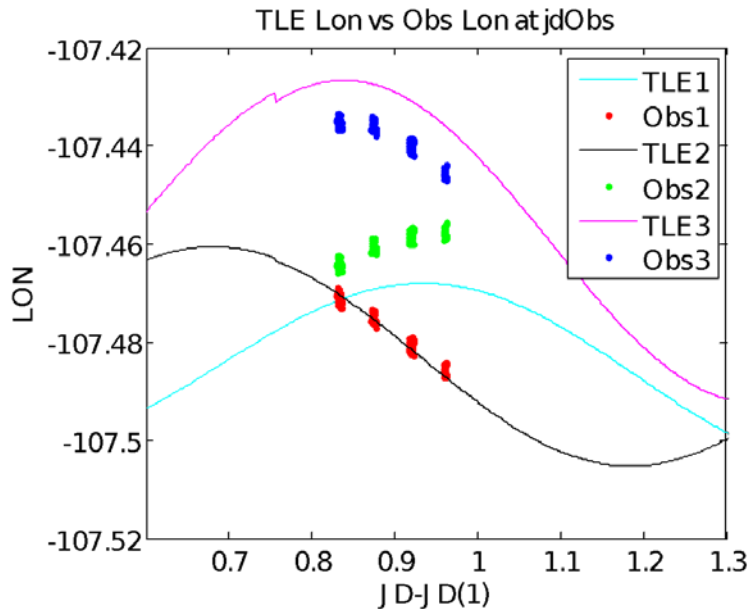

Figure 4. Measured apparent longitudes (color-coded dots) and TLE-based longitudes (colored lines) of the three Anik satellites in the 107W cluster on August 13, 2015. The measured and calculated positions show similar trends.

The trends of measured longitude approximately agree with the corresponding TLE-based trends, while the designated satellite identities do not. We use the residuals of the TLE and measured longitudes to estimate the likelihood of each identity, as shown in Fig. 5. A RANSAC-based algorithm [5] is used to find the best fit linear trend for each element set [Fig. 5].
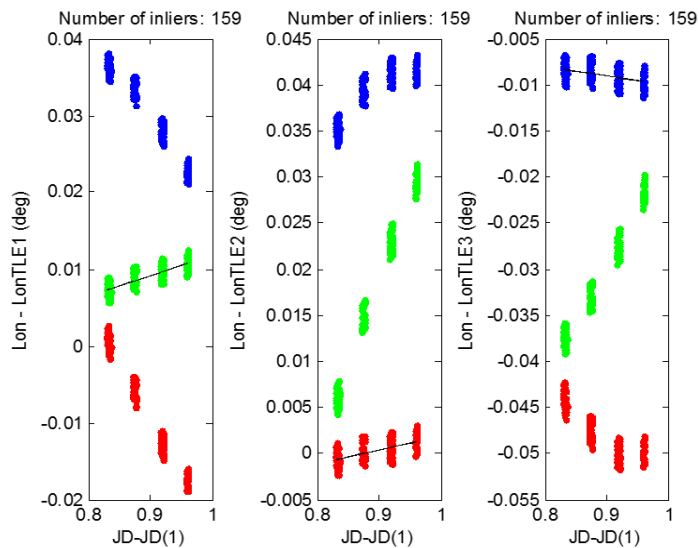

Figure 5. Detrended apparent longitudes (color-coded dots) and linear trend retrieved by RANSAC for each satellite (TLE) for Aug-13-2015. Color coding is according to the preliminary site-assigned IDs.

For each measured longitude, a p-value score is computed. These scores range from 0 (not likely) to 1 (very likely). The parameters of the distribution of measured longitudes are empirically estimated and used to calculate the p-value of each tracklet. The distribution is assumed to be normal and p is calculated using Eq. 2, where Δl is the tracklet's apparent longitude residual and σ an a-priori estimate of the standard deviation of the Δl distribution. Tracklets near the best fit line will have a high p-value.

$$p = 2\left(1 - N\left(\Delta l, 0, \sigma\right)\right) \qquad (2)$$

It is fairly straightforward to compute a composite p-value using Fisher's method [6] and estimate the likelihood of combined evidence, or multiple p-values. The tracklet p-values for Aug-13-2015 [Fig. 6] are well separated and readily distinguishable: for each TLE, the highest p-values are clearly higher than the second highest p-values. In this data set, the algorithm-assigned identities agree with the site-assigned identities.
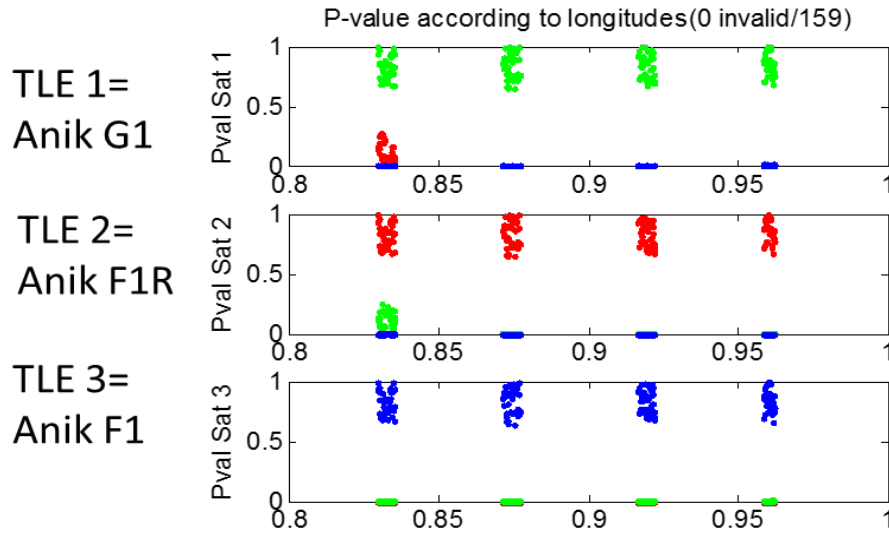

Figure 6. P values of the tracklets recorded Aug-13-2015. Note the substantial separation between the high p-valued dots from the lower p-valued dots. For this night, preliminary identification assigned by the site, shown as color code, agrees with the order of p-values.

On the night before, Aug-12-2015, the site-assigned identities disagree with the COSTB IDs [Fig. 7]. In the first panel, the RANSAC algorithm associates all tracklets in the vicinity of the nearly horizontal line as belonging to TLE #1. Within this panel, the site assigned the first group of tracklets as TLE #2 (green) and the last three groups as TLE #3 (blue). The site ID's sequence is green-blue-blue while the algorithm selects the dots near the line. The assignment of TLE #3 in the third panel shows similar issues with mistagging.

The Aug-12-2015 p-values are shown in Fig. 8. The assignment of tracklets to TLE #2 is straightforward, however the second group of tracklets in the top and bottom panels are difficult to identify as the p-values are not easily separable. The site assigned some of these tracklets as green (TLE #2) and others as blue (TLE #3). When p-values are ambiguous, the algorithm does not assign identities- the p-values must all be above a user-defined threshold.

Using longitude-based p-values, we are able to identify the tracklets [Fig. 6 and 8]. While it is not clear from Fig. 9 that the algorithm tags are correct, the corresponding tagged light curves, shown in Fig. 10, suggest that the tagging was done reasonably well. Note that color-coding is based on algorithm-assigned identities in these plots.
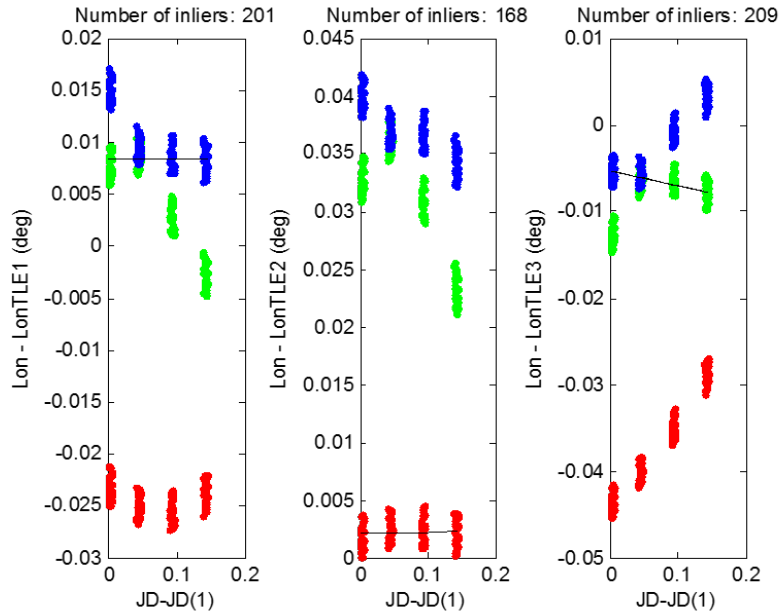
Figure 7. Detrended apparent longitudes (color-coded dots) and linear trend retrieved by RANSAC for each satellite (TLE) for Aug-12-2015. Color coding is according to the preliminary site-assigned IDs. For this session, the tracklets with different site-assigned IDs line up with the linear trends found for the first and third TLEs.
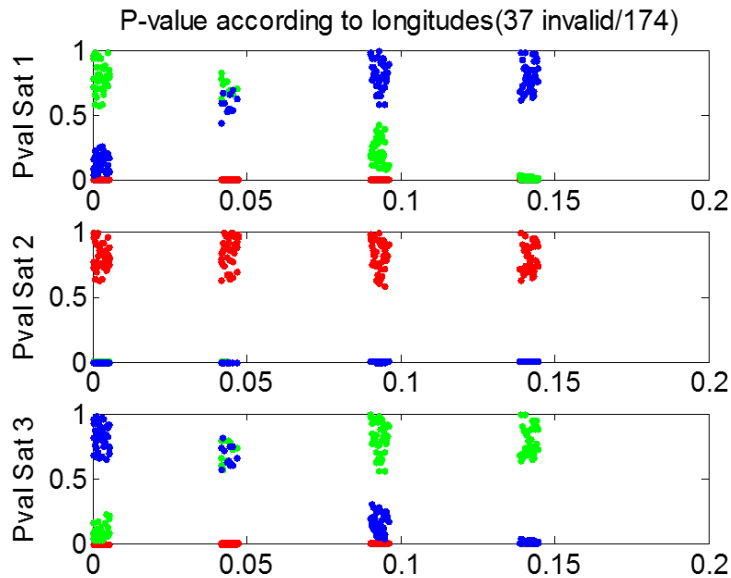

Figure 8. P values of the tracklets recorded Aug-12-2015. In the top and bottom panels, there is no clear separation between the high p-valued dots from the lower p-valued dots. For this night, preliminary identification assigned by the site, shown as color code, does not agree with the order of p-values.
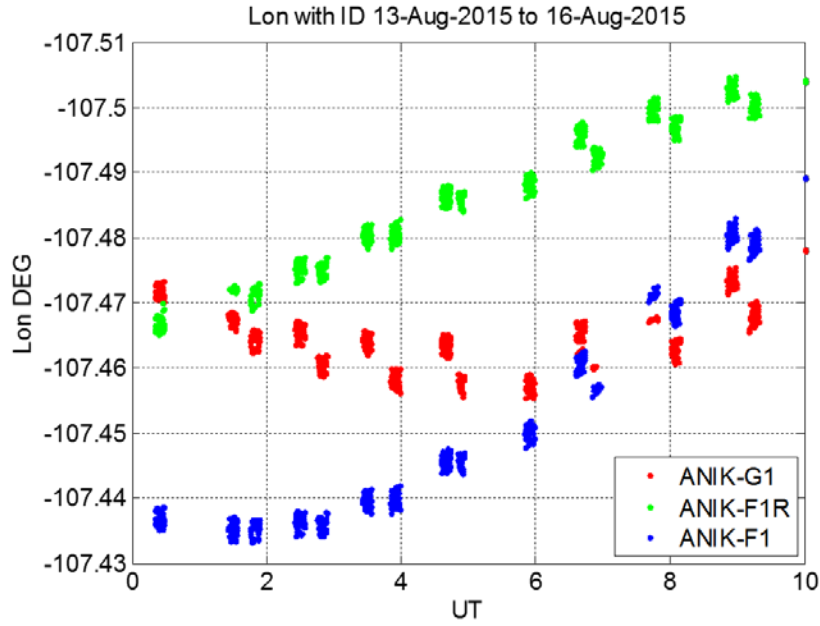
Figure 9. Longitudes of tracklets with COSTB-assigned IDs. Color coding is based on the algorithm.
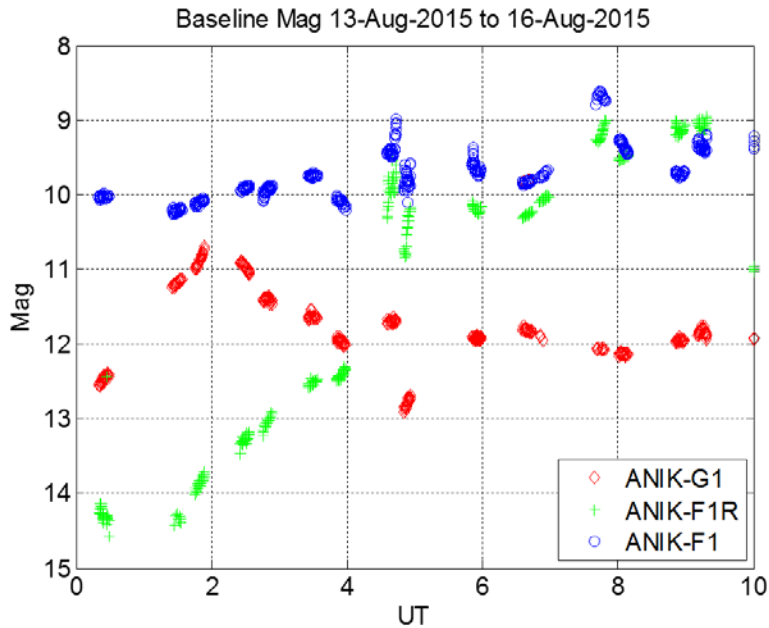

Figure 10. Magnitude of tracklets with COSTB-assigned IDs. Color coding is based on the algorithm. The magnitudes follow three distinctive traces. Color-coding is based on algorithm-assigned identities.

In general, the three light curves in Fig. 10 are well separated and visually continuous except for near 8UT. At first, the crossing of light curves near 8UT seems to be problematic because there is no clear distinction between the green and blue tracklets based on magnitude. However, one also notices that the measured longitudes of those two satellites (blue and green in Fig. 8) are well distinguishable for those times. On the other hand, while the red and blue longitudes are nearly indistinguishable near 7UT, their magnitudes are quite different and easily distinguishable. This illustrates why using both positional and photometric information is more effective than using either position or photometry alone.

4.2 On-line tagging

With the photometry baseline established, the COSTB algorithm can be applied to measured longitudes and magnitudes in near-real-time. The term "on-line" is used to mean that identification is required after each measurement and in this context, is equivalent to "near real time". The delay is due to data processing and data retrieval. The COSTB algorithm executes in less than 10 seconds with a PC running Matlab and, depending on factors such as network connection, data processing can begin on the order of minutes after being collected.

The longitude-based p-values are calculated as in 4.1. For on-line tagging, we need to add photometric evidence. To demonstrate the utility of COTSB in congested areas of the GEO belt, we studied the 107°W ANIK cluster and the 101°W cluster. Only the results on the more challenging ANIK cluster will be presented here.

For each measurement, there are typically three tracklets for the three resident satellites and additional ones for non-resident objects. For each tracklet, three p-values [Eq. 2] are calculated to represent its relative likelihood of belonging to each of three satellites. In the absence of non-resident objects, only three tracklets are measured, and 9 longitude-based p-values calculated. Those 9 longitude-based p-values are combined with the 9 magnitude-based p-values. The discussions of magnitude-based p-values will be presented later. The best combined p-value for each resident (ID) is selected based on two criteria: being the highest p-value in the cluster and its ratio to the next highest value greater than a user-defined threshold. While absolute likelihood is not available, the use of probability ratio is justified by the law of likelihood (ratio).

The magnitude-based p-value is calculated as follows. For three-axis stabilized satellites, visual magnitudes are reproducible in UTC from one night to the next. For each satellite, its expected magnitude can be inferred from previous photometry (model or baseline light curves).

The expected magnitude of each satellite can be inferred from baseline photometry. The abrupt effects of atmospheric extinction are effectively removed by applying the technique of differential magnitude [7]. To calculate the expected magnitude, we use a third order polynomial fit to model data in the relevant time interval and find the expected value by interpolating. The polynomial fit is weighted to de-emphasize old photometry and emphasize more recent data. This step is necessary because while light curves are reproducible, they also exhibit a noticeable seasonal dependence.

A high cadence continuous light curve is not necessary as long as there is enough magnitude data sampled around the time of estimation. The likelihood of each tracklet relative to each resident satellite is represented by a magnitude-based p-value. Deviations of the measurements from the expected magnitudes are used to calculate the p-values.

The final selection of the most likely identity for a tracklet is based on a combination of the best longitude and magnitude p-values. The selection logic has been designed (1) to minimize the number of false identities and (2) to make use of both positional and brightness information. By setting the selection criterion stringent, we can manage the occurrence of false identification.

4.3 Updating baseline photometry

After each session, baseline photometry is updated to reflect seasonal changes that continuously transform the light curve. The magnitudes of all identified tracklets are added to the database with time stamps. The stamps are used to track the age of photometry data and to calculate the weights of the polynomial fit mentioned in section 4.2.

Fig. 12 shows all the magnitudes of the baseline database and new measurements from the test. The baseline was originally formed with measurements described in section 4.1 (Aug 13-16). The color-coded dots represent the Aug 13-16 database. The color-coded plus symbols represent new data contributed by Aug 18 measurements. As an example, the new data points (+ symbols) near 0 UTC are circled in Fig. 13. Other clusters of new magnitudes can be found at 4, 6 and 8 UTC.
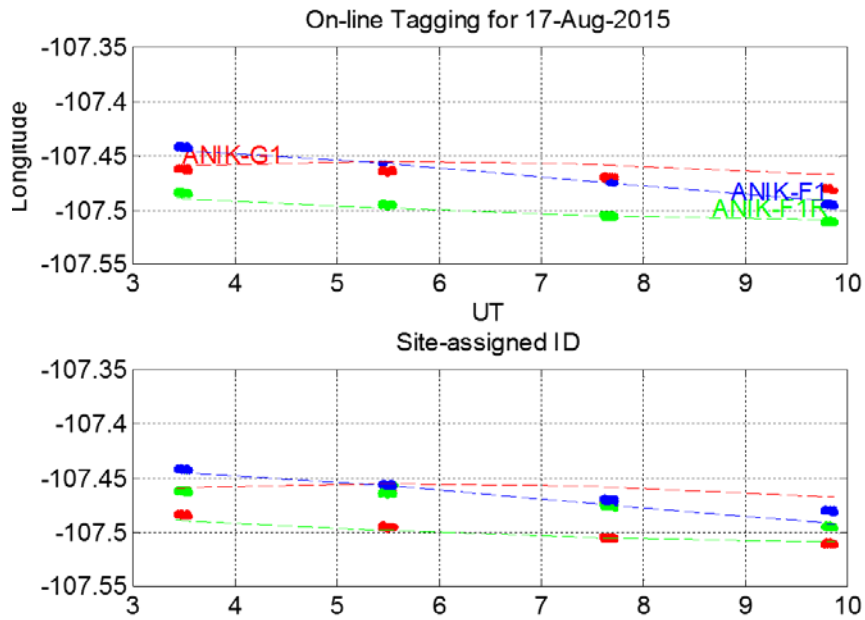
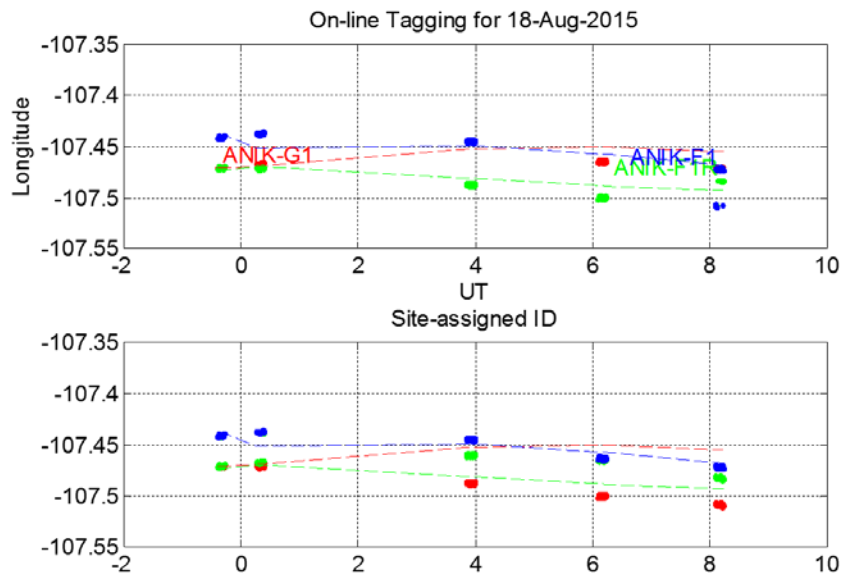*Figure 11. On-line or near real time tagging of the tracklets.*

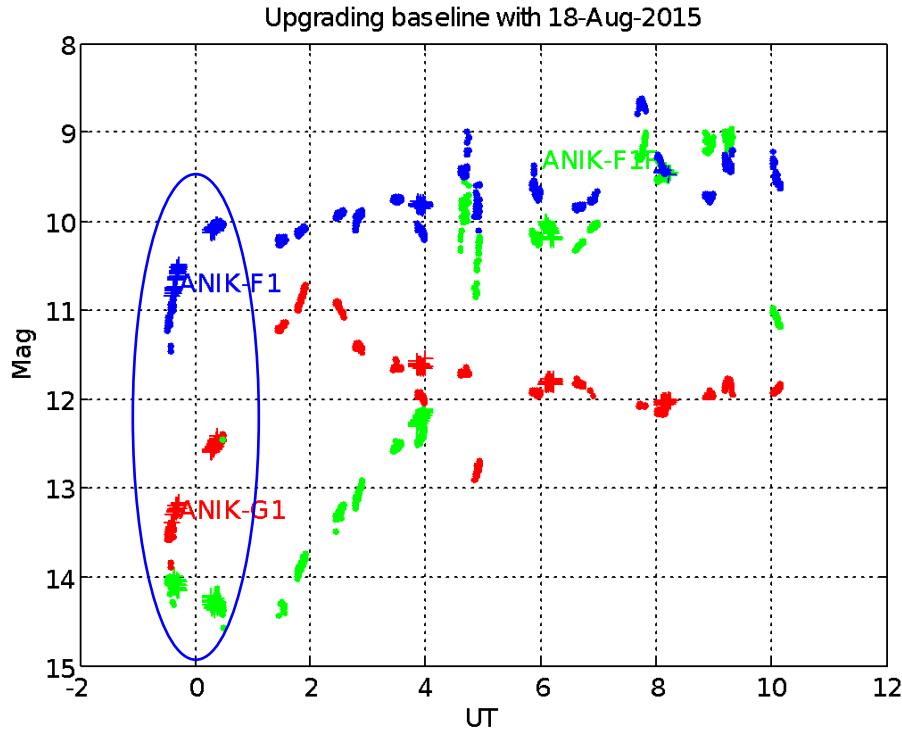*Figure 12. On-line or near real time tagging of the tracklets.*

Figure 13. Magnitude of tracklets with COSTB-assigned IDs. Color coding is based on the algorithm. The dots are magnitudes which have been identified in the baseline period. The color-coded plus symbols are the magnitudes which have been identified on-line in the Aug-18-2015 test. The magnitudes follow three distinctive traces.

5. Detection of non-resident objects

Data taken Dec-10-2015, between 1 and 9 UTC, was collected with a dense cadence in order to (a) construct a set of continuous light curves to be used as reference for the 107 W cluster and (b) demonstrate the ability to detect in near real-time the presence of a non-resident object approaching or entering the cluster. The second objective is quite important because uncorrelated objects are known to cause incorrect tags and compromise metric data. The ability to flag intruding objects is also crucial for SSA. We find in Fig. 14, a display of apparent longitudes and declinations as functions of time, a triplet of objects with similar trajectories (resident satellites), a streak of tracklets with high inclination (red) between 6:40 and 7:24 UTC and a single tracklet (black) in the vicinity of the cluster at 2:34 UTC. The small longitudinal drift of the highly inclined track (red) suggests that this object is near geosynchronous. The object responsible for the single tracklet (black) is probably in a low orbit.

The analysis shows that there were several mis-tags of the resident satellites as indicated by the flip-flopping identities (color-coding in Fig. 14 is based on initial site-assigned tags) of the tracklets in each of the three tracks in the middle of the data cube. Applying position-only COSTB techniques, the same used to generate baseline light curves in section 4.1, to this data reveals the identities of the three satellites associated with the triplet of continuous tracks. Fig. 15 shows the retagged data points for the resident satellites.
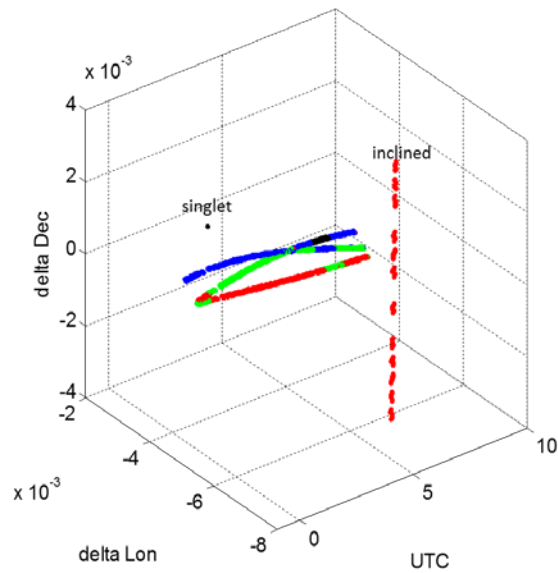
Figure 14. The tracklets were recorded on Dec-10-2015 and shown in a scatter plot in a detrended DEC-detrended longitude and UTC coordinate system. The dots are color-coded based on initial site-assigned IDs.
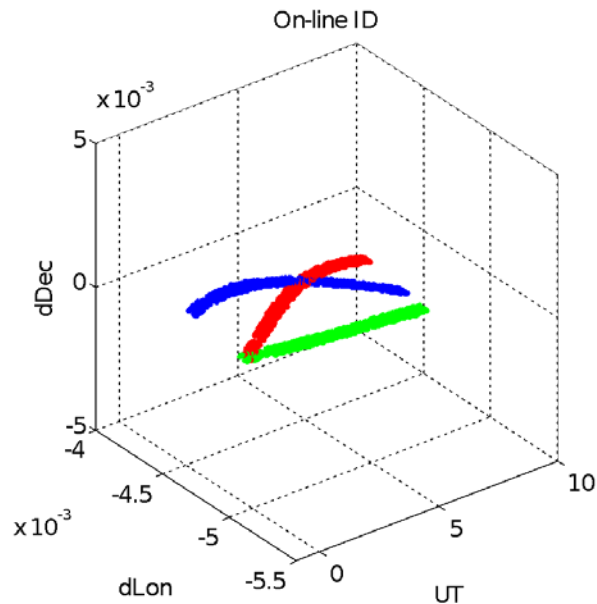

Figure 15. The tracklets were recorded on Dec-10-2015 and shown in a scatter plot in a detrended DEC-detrended longitude and UTC coordinate system. The dots of the resident satellites are shown and color-coded based on COSTB identification.

The tagged magnitudes shown in Fig. 16 suggest that the identification of the Anik satellites is satisfactory. Only three tracklets were discernibly incorrect. Two of those occurred at the beginning of the measurement session when longitudinal trends had not been established. The third incorrect tag occurred at 4:10. It's encouraging to note that out of the 1089 tracklets recorded that night, only three seem to be incorrect in this visual inspection. The benefit of determining where the resident objects are at any given time helps us uickly identify non-resident objects. By the process of elimination, it is straightforward to single out non-residents or uncorrelated objects.

While this analysis was done after the full sets of data had already been collected and we have not yet shown that the algorithm can be used in near real-time, we are confident that on-line identification and flagging of uncorrelated objects can be achieved using concurrent spatiotemporal and photometric information.
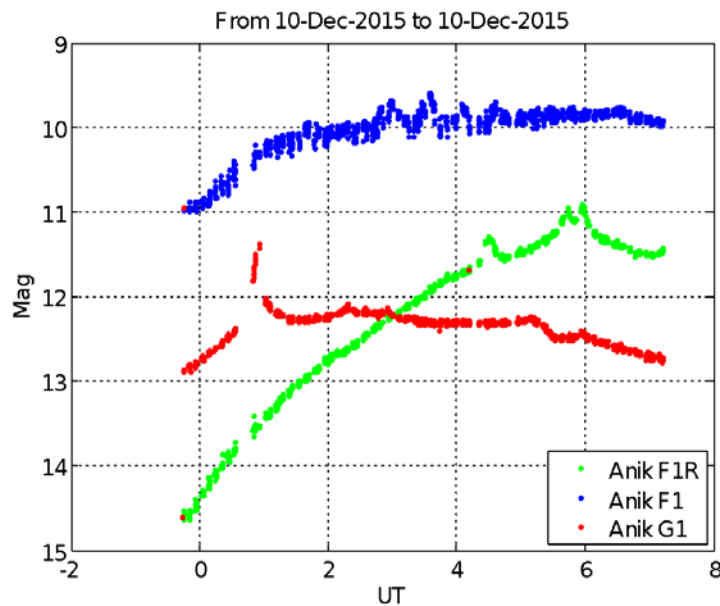

Figure 14. Magnitudes of the resident satellites with color coding based on COSTB identification.

6. Discussions and Conclusion

The objective of this study was to demonstrate that the tracklets of three-axis stabilized geosynchronous satellites can be identified and correlated with the element sets in near real time. To generate near real time tags for measurements, we developed an algorithm that combines astrometric and photometric information that computes the likelihood (p-values) of each measurement relative to each resident satellite. For each tracklet and each element set, we obtained two p-values, one representing how well the longitude fits the linear trends of the the respective longitudes, and the other showing how well the magnitudes agree with model photometry. To calculate the latter, we collected "baseline" light curves and we continue to update them with newly identified magnitudes. For a cluster of three resident satellites, there are 6 p-values for each tracklet. Selecting the best fit ID is based on all those 6 p-values. While the logic of selection is still preliminary, the results have been satisfactory.

Further development of the selection logic is warranted to reduce the number of mis-tags and non-results. An essential step for this process is to update the light curves or model brightness values to accommodate seasonal variations. Once the magnitudes are identified with high combined p-values, they can be added to the set of magnitudes. Using time stamps, the more recent magnitudes values can displace the older values.

Currently, the decision to add new magnitudes has not been automated because we have not developed an algorithm that validates the new data. Validation is needed to reject magnitudes that have been corrupted by thin clouds. While the computation of magnitude-based p-values is not affected by cloud extinction because we use differences in magnitudes, the baseline magnitudes can still be corrupted by cloud-affected data. We will work on automating this step.

Rapid identification of resident satellites will greatly reduce the adverse effect of mistagging metric and photometric information. When the combined probabilistic indicator is marginal, the site has the option of not reporting metric data. Because identification of resident satellites can be achieved as soon data is available, the ability to detect non resident objects is greatly enhanced.

Acknowledgements

References

[1] *Continuing Kepler's Quest: Assessing Air Force Space Command's Astrodynamics Standards*, Report by Committee for the Assessment of the U.S. Air Force's Astrodynamics Standards, Aeronautics and Space Engineering Board, Division on Engineering and Physical Sciences, National Research Council, The National Academies Press (2012).

[2] Lang, D. and Hogg, D.W. and Mierle, K. and Blanton, M. and Roweis, S., "Astrometry.net: Blind Astrometric Calibration of Arbitrary Astronomical Images", 2010.

[3] Kyle Barbary, Kyle Boone and Christoph Deil 2015 sep: v0.3.0 Online: http://dx.doi.org/10.5281/zenodo.15669

[4] E. Bertin and S. Arnouts, Astronomy and Astrophysics Supplement, v.117 (1996)

[5] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," Comm. Assoc. Comp. Mach, vol. 24, no. 6, pp. 381–395, 1981.

[6] Fisher, R.A. Statistical Methods for Research Workers, 1925

[7] Dao, P., Automated algorithm to detect changes in geostationary satellite's configuration and cross-tagging, AMOSTECH proceeding paper, Sept. 2015.

[8] https://en.wikipedia.org/wiki/P-value

Paper ID: 3248258