Please cite this paper as:
J. Aristoff, A. Ferris, A. Larson, A. Hariri, N. Singh, J. Horwood, J. Shaddix, K. Wilson, T. Lyons, and N. Dhingra, "Non-traditional data collection and exploitation for improved GEO SSA via a global network of heterogeneous sensors," in Proc. of the 2018 Advanced Maui Optical and Space Surveillance Technologies Conference, (Wailea, HI), Sept. 2018.

Non-Traditional Data Collection and Exploitation for Improved GEO SSA via a Global Network of Heterogeneous Sensors

Jeff Aristoff, Neil Dhingra, Alex Ferris, Austin Hariri, Joshua Horwood, Ari Larson, Tyler Lyons, Jeff Shaddix, Navraj Singh, Kevin Wilson

Numerica Corporation, 5042 Technology Parkway, Suite 100, Fort Collins, CO 80528

ABSTRACT

Space operations centers and satellite operators are continually being challenged with maintaining situational awareness on a growing number of resident space objects and with protecting key assets from hazards and potential threats. This trend necessitates pursuing multiple non-traditional approaches to not only free-up sensor/operator/analyst resources, but also to expand one's ability to characterize the evolving and dynamic space environment. To help address this need, Numerica Corporation has combined advanced algorithms, high-performance software, and a newly designed and deployed global network of small telescopes to demonstrate a responsive, robust, and affordable commercial deep-space tracking system. The network consists of both small-aperture, wide field-of-view sensors that provide persistent coverage of a large swath of the night sky, and medium-aperture telescopes that provide increased detectability and resolution but with a smaller field of view. This paper provides an overview of the Numerica Telescope Network (NTN) and describes how it is being used to maintain an independent catalog of objects in geostationary Earth orbit (GEO) and in other deep-space orbital regimes, as well as how it can be used to fill tactical needs through a full tasking, collection, processing, exploitation, and dissemination (TCPED) pipeline.

1. INTRODUCTION

Space-based systems play an important role in civilian life and national defense. The increasing number of satellites in space, and the many small or maneuverable objects that are difficult to track, are leading to a more congested and contested space environment. The ability to maintain an up-to-date catalog of all resident space objects (RSOs) is becoming increasingly challenging and continues to place a large burden on space operations centers. To buck this trend, Numerica has taken an active role in providing alternative solutions that free-up government resources and improve our situational awareness of the evolving space environment. In particular, for the past decade, Numerica has been developing algorithms and software to support improved space situational awareness (SSA) which involves the detection, tracking, identification, and characterization of all RSOs. More recently, Numerica has begun producing and providing high-quality observational data and information products to help inform operator action to protect military and commercial satellites.

In collaboration with the Center for Rapid Innovation at the Air Force Research Laboratory Space Vehicles Directorate (AFRL/RV), Numerica developed and demonstrated the feasibility of collecting non-traditional satellite tracking data to build and maintain a deep-space catalog, independent from the public catalog, with accuracy as good as or better than the Government-maintained space catalog. High levels of catalog completeness, accuracy, and timeliness were achieved through the use of advanced algorithms and high-performance software, together with the customization, deployment, and operation of a global network of small telescopes. Numerica also developed and demonstrated the ability to task third-party telescope networks, ingest observations collected by these networks, and fuse this data with data from its own network to update orbital state estimates.

This effort culminated in a 50-day evaluation of Numerica's prototype catalog maintenance function during which daily snapshots of its space catalog were created and later delivered for analysis by the Government. An independent analysis was also performed by Numerica, and no recourse to the public catalog was made during this evaluation period (and 30 days beforehand). In summary, using just 11 telescopes deployed and operated by Numerica (and a very small amount of data collected by contributing sensors), Numerica built and maintained a catalog of 2,272 objects (including 1,284 GEOs), viewing each object roughly 1.5 times per night during the evaluation period. For comparison, the public (deep space) catalog contained 3,367 objects (including 1,281 GEOs) at the start of the evaluation period (not including objects on the lost list). Further, Numerica achieved sub-arcsecond measurement accuracy about half-way through the evaluation period, tracked many dim objects, and collected over one million correlated observations. A visualization of the catalog Numerica built and maintained during the evaluation period is shown in Figure 1.

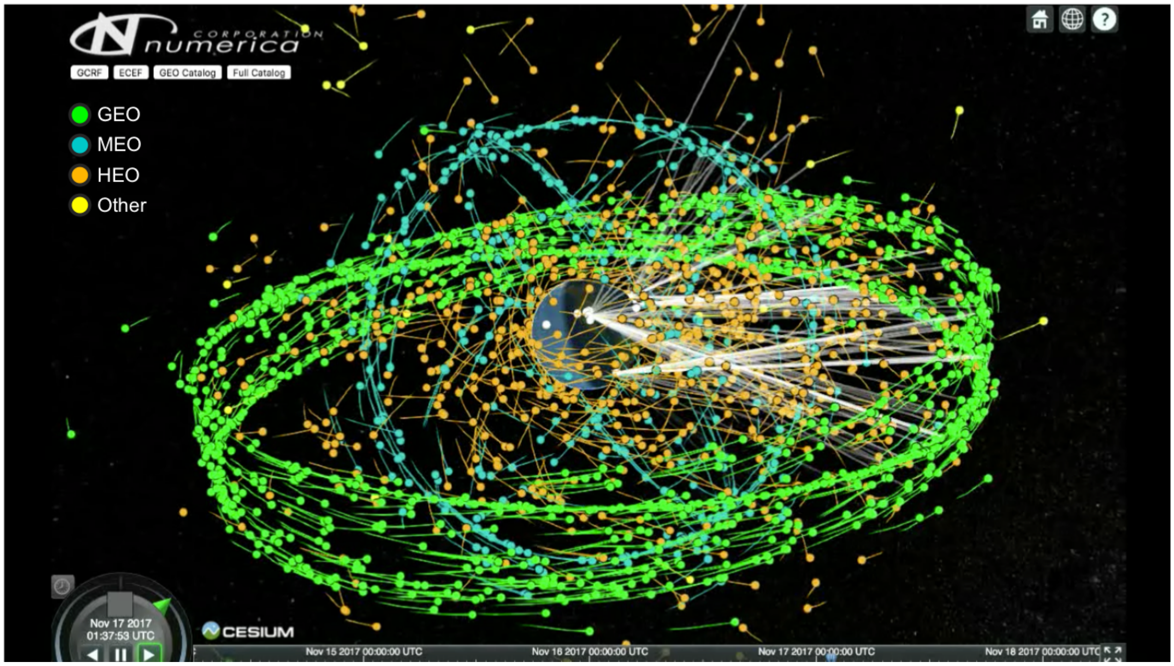

Figure 1: Visualization of Numerica-Generated Catalog at Conclusion of Evaluation Period (Nov 2017)

In the months since the conclusion of this evaluation period, Numerica continues to operate, mature, and grow the telescope network, with particular attention given to significantly increasing the network capacity and to maturing the data collection, data processing, and data dissemination pipelines. The current NTN consists of 25 total telescope systems (over 130 sensors), including of a mix of sensor arrays and medium-aperture telescopes, operating at 15 observatories world-wide to provide real-time tracking information on thousands of objects in deep space, including those in GEO, medium-Earth Orbit (MEO), and highly-elliptical orbit (HEO), as well as persistent nighttime coverage of GEO and near-GEO.

The layout of this manuscript is as follows. In Section 2, we provide an overview of the NTN and describe our technical approach. In Section 3, we present results from the 2017 catalog maintenance evaluation, focusing on completeness, accuracy, and timeliness. Improvements to theses key metrics that have been achieved since the end of this evaluation are also discussed, and we present conclusions in Section 4.

2. TECHNICAL APPROACH

Numerica is building the world's largest commercial deep space catalog, inclusive of GEO and near GEO, with accuracy as good as or better than the that of the Government-maintained catalog. Additional network capacity can and is being used for external tasking and improved quality of service on specific objects of interest based upon customer request. Our approach centers on a global telescope network, including commercial-off-the-shelf hardware and advanced algorithms, that is purpose-built from the ground up to cost-effectively support SSA missions through the provision of accurate and timely tracking data, ephemeris products, and other advanced analytics. Although this manuscript focuses on the development and demonstration of an operationally-viable deep-space catalog maintenance, search, and UCT processing function, the NTN can also support other use cases such as those on shorter timescales or involving fewer objects.

2.1 Network Overview

The NTN currently spans 15 sites across 5 continents, as depicted in Figure 2. Blue triangles denote the telescope systems used during the 2017 evaluation period (not including the prototype telescope in New Mexico), while the red circles denote systems that have been deployed since the end of the 2017 evaluation period. Deployments within the United States are located in California, Arizona, Colorado, New Mexico, and Texas. Deployments outside the United States are located in Chile, Morocco, Spain, France, South Africa, Crete, Western Australia, South Australia (two different sites), and New South Wales Australia. Together, this layout provides 100% coverage of all deep space orbital regimes, including GEO, with robustness to both regional and seasonal weather. For example, each GEO will be visible from several geographically-diverse sites and by multiple co-located sensors. Additional deployments are anticipated for 2019.

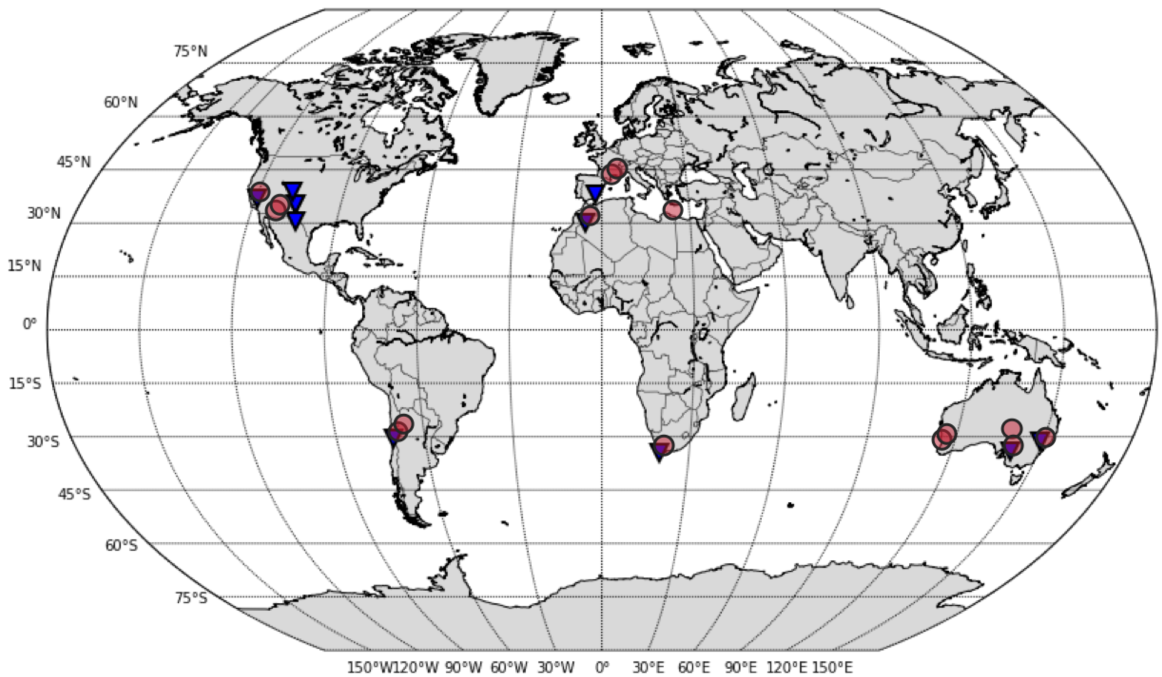

Figure 2: Locations of Dedicated Network Assets (Deployments in 2017 and 2018)

The NTN consists of multiple complementary telescope systems, each of which is comprised of several hundred carefully-selected commercial-off-the-shelf components together with components that have been designed and manufactured by Numerica, to provide a responsive, robust, and affordable commercial deep space tracking system. This includes 15 medium-aperture (0.3-0.4m) robotic telescopes providing precise astrometry and photometry, good detectability, and the ability to collect data on objects in all regimes of deep space. Further, the telescopes use high-speed professional-grade robotic mounts that enable rate tracking for objects moving at high angular speeds. The NTN currently includes 10 robotic sensor arrays that exchange breadth for depth and collect continuous observations on all objects in GEO (or across a given swath of sky). Like the medium-aperture telescopes, the sensor arrays provide high-quality astrometric and photometric observations, thus enabling object characterization and persistent monitoring beyond the typical scope of catalog maintenance missions. The capabilities of these two systems complement one another, and given their global distribution, provide a powerful, innovative, and flexible system for deep-space situational awareness and space control.

Specifically, augmentation of the medium-aperture telescopes with the sensor array systems enables persistent surveillance with the ability to tip-and-cue the larger systems when events are detected. It also facilitates photometric characterization and rapid maneuver detection, ultimately resulting in a more accurate, timely, and complete catalog of GEOs and near-GEOs. The sensor arrays are capable of sweeping up and down in elevation, and arrays from different observatory locations can be combined to widen the "optical space fence" or to provide more frequent and geographically diverse observations on the same set of GEOs. While Numerica is still collecting periodic updates on these objects with the larger telescopes, a substantial amount of work is being offloaded to the sensor arrays, allowing the larger systems to focus on objects outside of GEO, or those too dim to be detected by the arrays in a reasonable amount of time. In addition, freeing the larger telescopes from routine GEO catalog maintenance allows more prompt and thorough response to time-sensitive events such as the detection of UCTs or of unexpected maneuvers. Finally, this more complete and persistent monitoring of the geostationary belt allows Numerica's sensor tasking software to quickly respond to events of interest and to obtain precise data within minutes to aid in real-time decision making as well as post-event forensic analysis.

2.2 Supporting Algorithms and Software

To enable a relatively lean network to perform at a level comparable to networks several times its size, Numerica's solution utilizes a suite of advanced algorithms and a suite of six high-TRL software components, as depicted in Figure 3 and as described below. Some of these algorithmic components, most notably MFAST, SLATE, and NITRO, use rigorous statistics-based principles and include many of the recommendations made by the Astrodynamics Innovations Committee (AIC) Covariance Realism Working Group.[1]

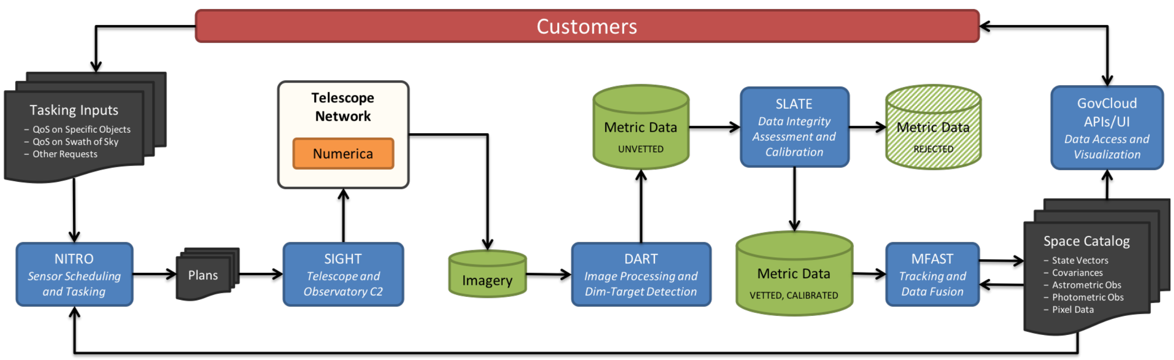

Figure 3: High-Level System Architecture

2.2.1 NITRO (Numerica Intelligent Tasking and Resource Optimization)

NITRO determines which telescope should look at which object(s) and when by optimally and dynamically scheduling sensor dwell times and durations based on tasking requests, viewing conditions, and the state/type of each object in the catalog, among other factors. This sensor resource management (SRM) task is posed as a generalized version of the canonical Traveling Salesperson Problem (TSP) in which a salesperson must find an optimal route between several cities. In NITRO, the "salesperson" represents a telescope, each "city" corresponds to an object to observe, and "traveling between cities" signifies slewing to point at different objects.

The full formulation of the SRM problem is as a Time Dependent Team Orienteering Problem with Time Windows (TDTOPTW). The Team Orienteering Problem (TOP) is a generalization of the traveling salesperson problem that allows for multiple "salespeople" (i.e., telescopes) and does not require visiting every "city" (i.e., observing every object). The Time Dependent (TD) component of the problem is necessary because slew times between objects change due to the motion of the observer with respect to that of the satellite. The Time Windows (TW) express times at which observation is feasible and ensure that telescopes are not tasked to dwell on objects near the moon, below the horizon, or when solar geometry is not favorable. Moreover, they enable the imposition of rapid revisit constraints for frequent viewing of specific objects.

The TDTOPTW is a nondeterministic polynomial (NP) hard problem in general and even determining feasibility is NP complete. To approach it, we employ a two-stage heuristic algorithm that first constructs telescope schedules using a generalized nearest neighbor approach and then further refines the schedules via local refinement techniques. Several instances of the algorithm are run in parallel to generate a set of candidate dwell schedules. The resulting schedules are graded on a number of objectives and a final dwell schedule is chosen based on strategic (or tactical) priorities. Compounding the complexity of the TDTOPTW problem are hard constraints that require certain tasks to be executed. These tasks may involve mandatory viewing of high priority objects at certain times or persistent monitoring of specific objects.

Figure 4 illustrates the complexity of a NITRO-generated tasking schedule that was executed using four telescopes located in South Australia, California, Spain, and Chile, respectively. These telescopes were tasked to routinely observe 849 different objects over a 24-hour period, and to persistently monitor 31 of these objects (17 GEOs, 7 HEOs, and 7 MEOs), 3 of which are highlighted in the schedules and denoted by the red circles (●), black squares (■), and blue stars (⋆).

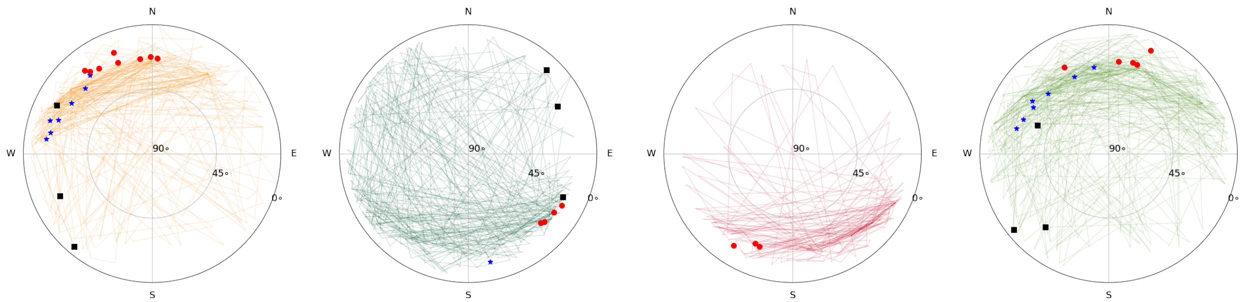

Figure 4: Example NITRO-Generated Tasking Schedule Executed Over a 24-Hour Period

2.2.2 SIGHT (Satellite Image Gathering Hub, Terrestrial)

SIGHT is a sophisticated observatory-level command-and-control software suite built from the ground up for the NTN. Numerica's SIGHT automation architecture is modular and extensible. It is composed of independent services, running on multiple PCs throughout each observatory. Properly authorized hardware commands may be propagated throughout the observatory network, so requests for hardware action may be made without a direct interface to the hardware. This distributed architecture allows the SIGHT system to be incrementally expanded and to respond flexibly to unexpected events.

As the NTN addresses ever more critical data collection tasks, sensor uptime is key to success. To support this requirement, the NTN automation system is highly reliable. Each observatory control network can recover from individual component failure, and component services will automatically restart upon failure where possible. Autonomous watchdog utilities constantly monitor hardware safety, and the observatory will respond immediately to threats such as inclement weather, observatory power loss, server health issues, unexpected or unsafe hardware positions, hardware malfunction, network communication loss, and even low disk space. Moreover, NTN sensors all carry small position sensors as a safety feature, to provide verification of hardware position independent of manufacturer drivers. Reliance on third-party hardware control utilities have been mostly eliminated, further improving sensor reliability. By communicating directly with ASCOM and manufacturer drivers, SIGHT automation software cuts unnecessary latency and avoids restrictions on hardware resources often seen in other automation tools. For example, PCs controlling NTN telescope systems can communicate directly with tens of sensors at a time, and can do so with low download latency.

The majority of sensor tasking is handled by our centralized tasking and scheduling software, NITRO, which delivers secure, encrypted tasking plans to each observatory. However, the SIGHT automation system does have some autonomy. In the case of extended network communication loss, for example, NTN sensors will execute default scan patterns (or use a plan that was provided well in advance), and deliver the resulting data once a connection is restored. In the event of UCT discovery, NTN sensors can bypass the full NITRO scheduling algorithm and respond flexibly before the new object is lost. Any co-located sensors can rapidly follow-up on the UCT until enough data has been collected to form an orbital state estimate, at which time the full NTN can be tasked by NITRO to refine the orbit on the newly-discovered or or recently-maneuvered object.

The 2018 expansion of the global telescope network and inclusion of heterogeneous hardware configurations posed new challenges for SIGHT. For example, the new sensor array systems produce massive amounts of data. A single sensor array generates over one gigabyte of raw imagery per minute under routine tasking conditions. This volume of data necessitates on-site image processing even at observatories with high bandwidth internet

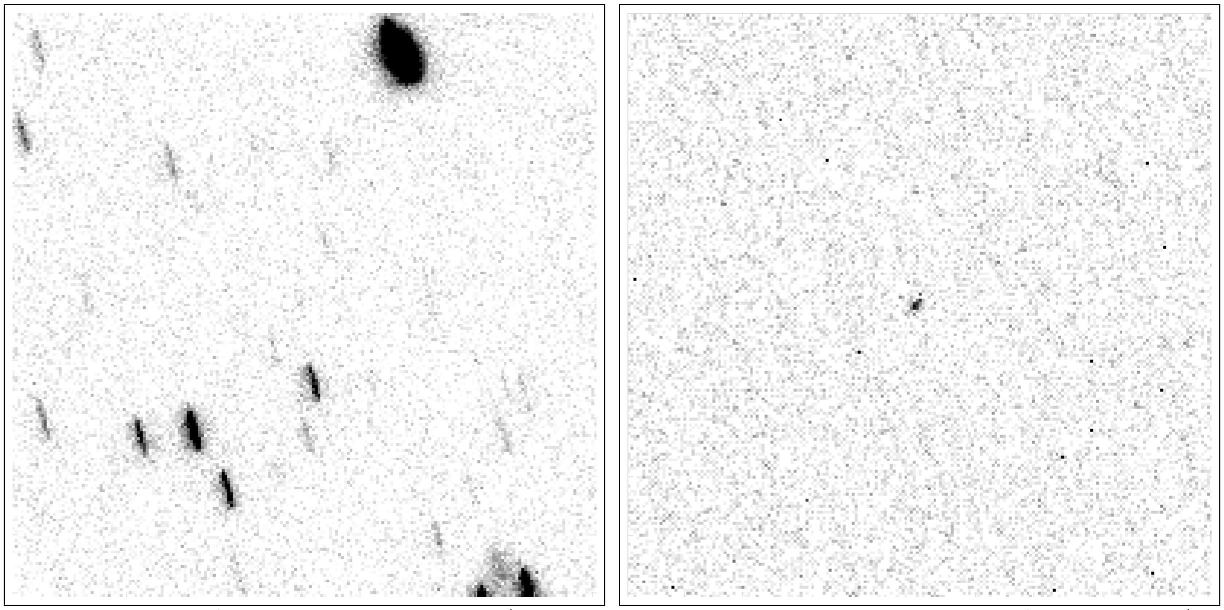

Figure 5: Sensor Array Stacking Example (Left: Single Image Before Stacking, Right: After Stacking)

connections. To meet this challenge, most Numerica image-processing capabilities have been modified to run on observatory computers, without the full computing power of Numerica (or cloud-based) servers. Due to the distributed, service-oriented design of SIGHT, new image-processing PCs may simply be plugged into an observatory network to bolster the image-processing capabilities of that observatory. As part of the deployment of the expanded telescope network, Numerica has delivered enough computing power to every NTN observatory to fully accommodate local image processing, which will not only improve robustness, but also reduce latencies relative to those observed during the 2017 evaluation period.

2.2.3 DART (Detection and Astrometric Reduction of Space Tracks)

DART processes the raw imagery, pulling astrometric and photometric observations, and enabling the detection of relatively dim objects without the need for much larger, more expensive telescopes. The DART image processing suite was developed in-house at Numerica to facilitate rapid and accurate automated processing of raw imagery from the NTN or other data sources. At a high level, the algorithm takes as input a set of consecutive images (i.e., a single dwell) from a single sensor and generates astrometric and photometric detections as well as image statistics and metrics. The routine is designed to operate on rate tracking collects in which the optical system follows an object of interest, and will detect both stationary and mobile objects in these images. In addition, it is also capable of ingesting sidereal collects in which the system follows the stars. Multiple detection routines are implemented; a default single-image detection routine produces the highest possible data cadence on relatively bright objects while image-stacking techniques, including a prototype track-before-detect algorithm, enables the detection of significantly dimmer objects at a reduced observation cadence. Figure 5 shows an example of results from a single image and after an image-stacking technique is applied to sensor array data.

Identified stars in each image are cross-referenced with a star catalog to allow quantification of the current imaging performance and viewing conditions, as derived from the astrometric and photometric offsets of these individual matches. This process includes photometric color correction of catalog magnitudes to match that of our sensor passbands. Finally, a set of filtering routines operate jointly on all detections produced within a dwell to reduce or eliminate false positives while preserving objects of interest, independent of any state estimates or other external information.

Active development continues on these algorithms to improve runtime efficiency and further refine performance and robustness. Recent advances soon to be incorporated include a high-dimensional optimization routine to more accurately map optical distortions through reference star catalogs, significantly improving astrometric

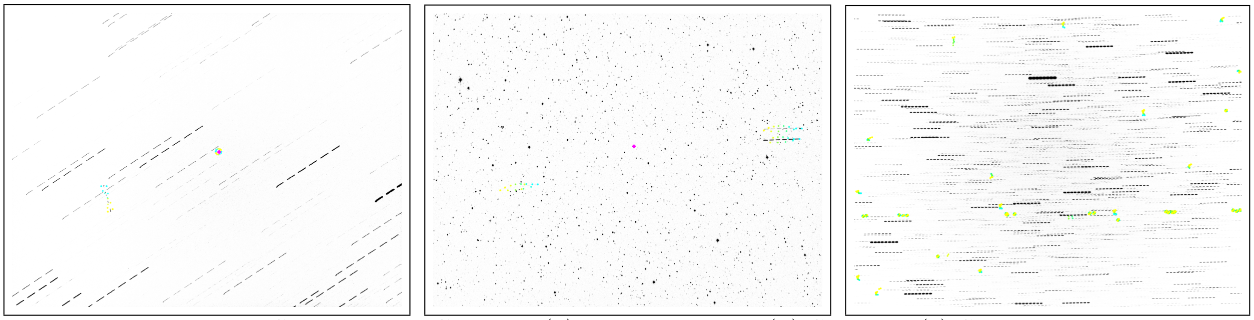

Figure 6: Sample Data Collects: (a) Rate Tracking, (b) Sidereal, (c) Stationary

accuracy of observations. Nevertheless, DART as it currently stands has been extensively vetted on real data and deployed operationally for over a year, automatically processing all imagery produced by the NTN. Figure 6 shows some example diagnostic snapshots, generated by DART for each dwell processed.

2.2.4 SLATE (Sensor Level Assessor of Track Errors)

SLATE is a suite of data processing tools which acts as both a data preprocessor and as a gatekeeper. Specifically, it ingests sensor observations from a variety of sources and data formats and through robust sensor calibration and tracklet generation routines ensures that only well-calibrated, properly-tagged, high-quality data is passed to MFAST and any other downstream components. Thus, SLATE enforces the assumptions of any downstream algorithms on its output.

SLATE performs three main tasks. First, it performs sensor calibration, estimating the characteristics, including weights and biases, of each sensor in the network. Second, the track partitioner component of SLATE solves the observation-to-track assignment problem, ingesting line-of-sight observations and performing sensor level tracking to determine which observations emanate from a single track on a single object. Finally, given the calibrated sensor parameters from the first task and the correctly partitioned observations in the second, SLATE forms bias-aware tracklets with realistic covariances that are ready for processing in MFAST and other downstream functions. For optical data, a tracklet is a data type consisting of a single four-dimensional angle/angle-rate state vector and covariance matrix produced by processing a track via the tracklet generation component of SLATE. The high-level system architecture of SLATE is shown in Figure 7.

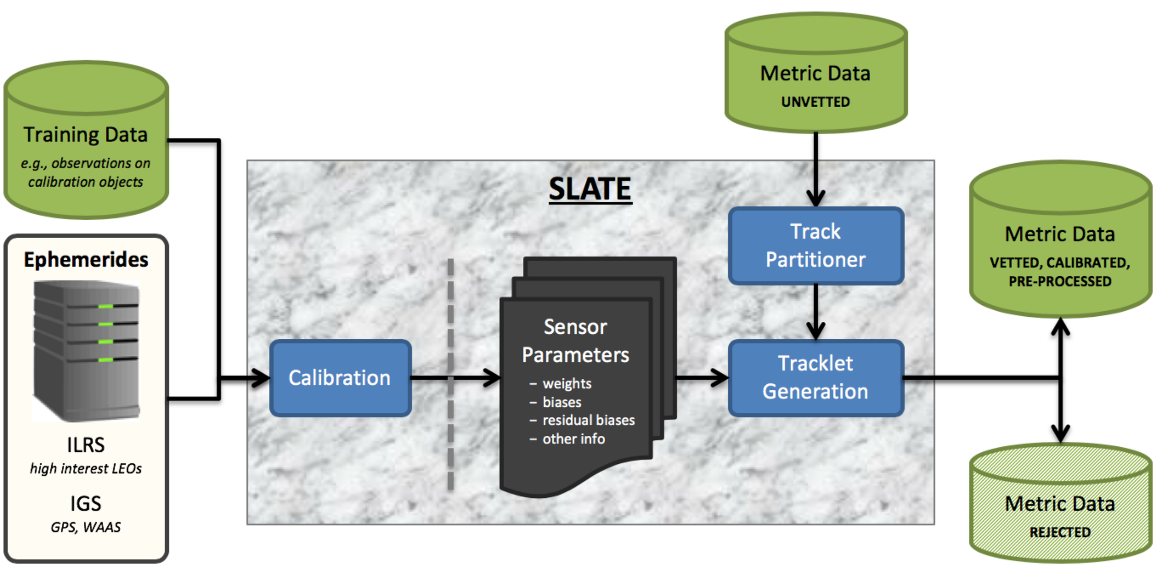

Figure 7: High-Level System Architecture of SLATE

Elaborating on the track partitioner component of SLATE, it ingests a stream of sensor observations with non-existent or unreliable sensor tags (i.e., RSO identifiers from which the sensor believes the observations emanate) and robustly partitions the observations into pure tracks (i.e., free from incorrect sensor tags, also known as cross-tags) and outliers. This functionality is of particular benefit when processing data from closely-spaced objects such as GEO clusters, or when the provided observations for what is claimed to be a single track do in fact come from multiple objects, or when a sensor does not produce track-level associations. The track partitioner can correct these mistakes by ensuring that pure tracks are passed to downstream SSA functions. The algorithm can also be used to re-tag observations at the sensor level.

2.2.5 MFAST (Multiple Frame Assignment Space Tracker)

MFAST fuses astrometric data into high-quality orbits with realistic covariances through joint orbit initiation and orbit determination. MFAST uses Numerica's special optimization-based formulation of the data association problem, namely the MFA formulation, that is well-suited for large-scale tracking problems. MFAST includes customized algorithms for non-linear filtering, orbit determination, orbit and uncertainty propagation (including Numerica's implicit Runge-Kutta orbital and uncertainty propagator[2,3]), and advanced physics-based complexity reduction (or hypothesis gating) techniques that are used to control runtime without sacrificing accuracy. Such methods along with the portability of MFAST allow it to be run on most platforms and hardware setups, including laptops and both serial and parallel computing environments. Some results from MFAST, including assessments of its scalability, obtained by processing real-world historical radar and optical data from the Space Surveillance Network (SSN) in a "UCT processing mode" are presented in Aristoff et al.[4] MFAST excels in deep-space (optical) UCT processing in terms of the number of high-quality orbits that it is able to establish, and the software has helped personnel at space operations centers to more quickly identify breakups, discover lost objects, and substantially reduce the size of the attention list.

In the current operational pipeline, MFAST is largely run in its joint UCT resolution and catalog maintenance mode, so that existing objects in the catalog database can be updated with new estimates while simultaneously allowing new objects to be discovered and added to the catalog. The catalog maintenance mode is implemented within the MFAST C-level application programming interface (C-API). Figure 8 provides an overview of the MFAST data flow pipeline C-API.

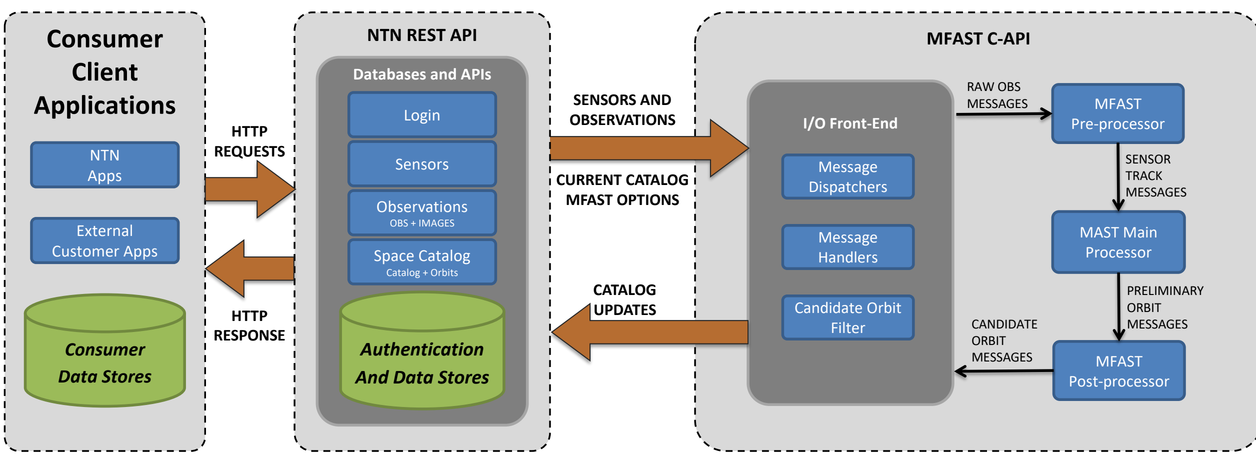

Figure 8: Architecture of the MFAST Data Flow Pipeline

MFAST also provides some orbit post-processing components that perform maneuver detection, ballistic coefficient estimation (for drag and solar radiation pressure), and duplicate orbit resolution. For maneuver detection, an algorithm has been developed specifically for detecting routine station-keeping maneuvers performed by active GEO satellites, utilizing the typical objectives of these maneuvers (inclination reduction, compensation for longitudinal drift) to better characterize the maneuver with potentially limited data. This post-processing algorithm complements internal MFAST-provided maneuver detection and rapid processing capabilities.

2.2.6 GovCloud-Based Application Programming Interfaces (APIs) and User Interface (UI)

In addition to the back-end algorithmic components discussed above, Numerica has deployed a cloud-based (Amazon Web Services GovCloud) web service to provide access to some of the key data products generated by the pipeline shown earlier in Figure 3. The web framework utilizes a modern container-orientated virtualization paradigm: different types of data products, such as orbital state vectors, astrometric observations, light curves, associated uncertainty information, sensor health information, data analytics, etc., are exposed via "microservices" deployed in Docker containers (see `www.docker.com`) in a scalable compute infrastructure. A high-level architecture of the NTN web service is shown in Figure 9.

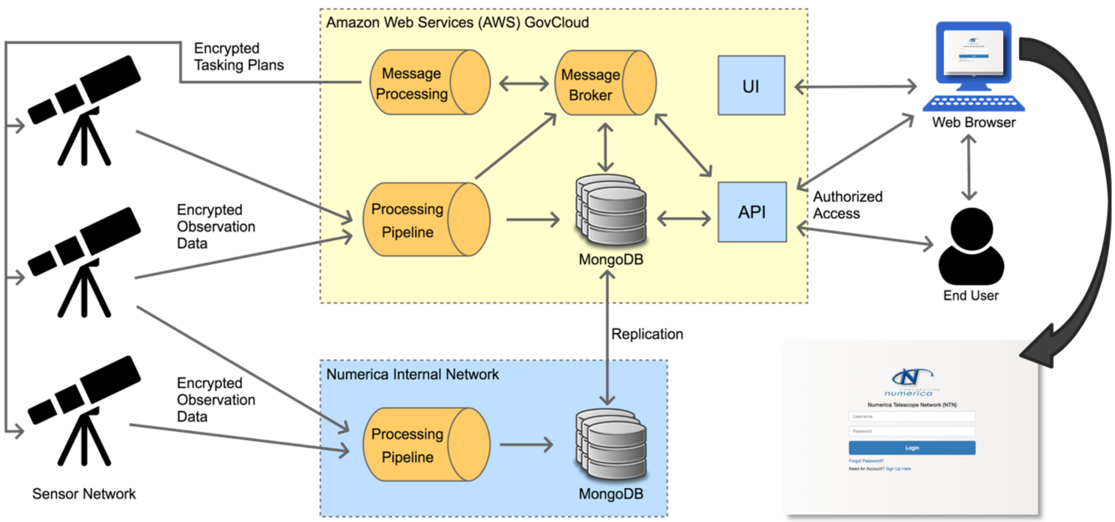

Figure 9: High-Level Architecture of the NTN Web Service

In addition, a secure front-end UI has been built with login and authentication utilities, allowing users to browse, search, and filter objects of interest within the Numerica data stores through secure connections to the back-end APIs. Authorized users can export orbital state vectors, astrometric and photometric observations in various formats, and tag objects of interest for quick viewing or data export. The UI also provides some basic real-time visual analytics tools, such as light curves visualizations, observation collection statistics, etc.

The latency between an object being observed by one of the NTN sensors and the corresponding fully processed astrometric and photometric observation being made available via the NTN APIs is currently about two minutes (on average) and dropping. Thus, near real-time data can be pushed to any external clients subscribing to the NTN data APIs, allowing support for a variety of integration experiments with third party sensor networks or data processing nodes.

The combined NITRO-SIGHT-DART-SLATE-MFAST software suite currently supports a number of capabilities (e.g., catalog maintenance, search, UCT processing, breakup processing, maneuver processing, attention-list processing, routine and dynamic sensor tasking). Additionally, Numerica's high-TRL KRATOS (Kollision Risk Assessment Tool in Orbital Element Spaces) software,[5] performs conjunction assessment screening and efficiently computes probabilities of collision (PC) to support space flight safety. KRATOS has been demonstrated using data collected from the NTN, and its PCs rival the accuracy of Monte-Carlo methods but with little added computational cost relative to the traditional 2D integral method (also known as the Foster method). Active research and development efforts are underway to mature and transition additional SSA capabilities including change detection given both dense and sparse light-curve data using Numerica's Athena software.[6]

3. RESULTS AND DISCUSSION

In this Section, we present the results from the 2017 catalog evaluation and discuss improvements that have been made since then, focusing on completeness, accuracy, and timeliness. Recall that only 11 telescopes (located at

10 different sites) were used during the 2017 catalog evaluation period, whereas the 2018 NTN consists of 25 telescope systems (located at 15 different sites) and over 130 sensors.

3.1 Completeness

A vitally important element of any deep-space catalog is its completeness: i.e., how many objects are successfully tracked. Catalog completeness requires (i) sufficient geographic diversity of sensor assets to render all objects viewable; (ii) sufficient capacity to view all objects of interest with a sufficient update rate; and (iii) sufficient detectability to ensure the sensors can detect the objects of interest.

At a high level, the Numerica data pipeline during the 50-day evaluation period in 2017 correlated approximately 160,000 tracks (each consisting of multiple line-of-sight observations) to new or existing objects in its deep-space catalog. This resulted in a catalog which, at the conclusion of the evaluation period, contained 2,272 objects receiving regular updates. A summary of the tracks and orbits produced is provided in Figure 10. Across this deep space catalog, an average object update rate of 1.5 updates per day was achieved. We note that the relative production of a given site is a function of both weather, season, and the time at which the telescope became operational (e.g., the telescope in South Africa became operational near the very end of the evaluation period, hence its low contribution relative to the other sites).

Regime	Updated in Final Week	Mean Rate (tracks/day)	Median Rate (tracks/day)
GEO	1284	1.77	1.34
MEO	337	1.3	1.14
HEO	638	0.88	0.78
OTHER	13	0.48	0.4

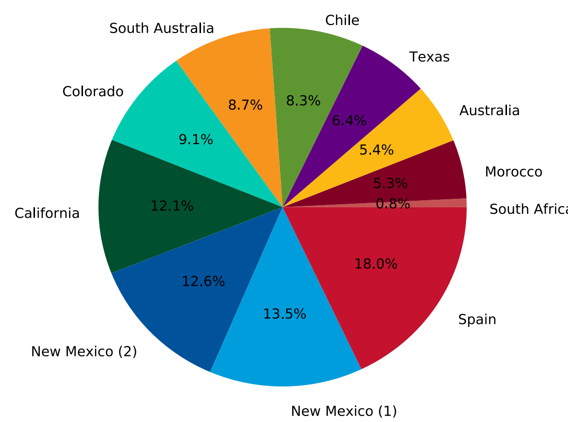

Figure 10: Summary of Evaluation Catalog by Regime and Site Production Statistics

Since the conclusion of the evaluation period we have continued to collect observational data with the NTN while simultaneously upgrading hardware and software components; the observation database now contains over six million correlated observations comprising more than 850,000 tracks on 2913 unique objects. We anticipate reaching 100 million observations in late 2018 or early 2019 based on current projections.

The distribution of GEO objects in the evaluation catalog by longitude and with non-public objects removed, contrasted with the public catalog, is given in Figure 11. Here, all objects from each catalog updated within the last week, as of the end of the evaluation period, are propagated to a common epoch and their longitudes plotted. At a high level, the catalogs appear to correspond well and similar trends can be seen in each, such as the relatively high frequency of geostationary objects over the CONUS and the lower density over the Pacific Ocean. One notable difference between the two catalogs is the smaller number of objects in the 50 to 80 E region in the Numerica catalog. This area had minimal coverage during the evaluation period; the current location in South Africa, now fully operational, and new deployments to Crete and Western Australia (and an additional system in South Africa), will increase coverage for this longitude band.

Another key requirement for achieving catalog completeness is detectability. Maintaining an object requires not merely occasional but consistent detection to ensure a timely and high-quality orbital state estimate is produced to enable successful follow-up sensor tasking.

Figure 12 gives the computed apparent magnitudes for the observations collected during the evaluation period by the dedicated network (higher magnitudes mean dimmer objects). These magnitudes are relative to

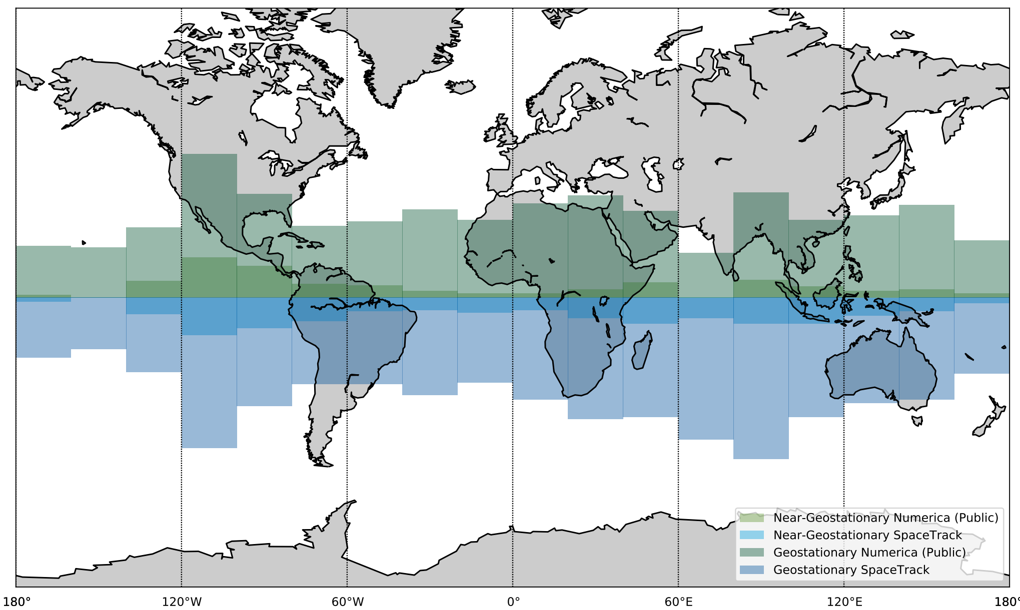

Figure 11: Longitude Distribution of Recently-Updated GEOs at Conclusion of Evaluation Period

our sensors' passbands, which have relative spectral responses similar in shape to the ESA Gaia green passband. Since our sensors are more responsive to longer wavelengths than the Johnson V band, objects reflecting solar radiation (a cooler stellar spectrum than Vega), will have slightly brighter magnitudes in our Vega-calibrated scale than they would in the standard visual magnitude scale. If a satellite has uniform spectral reflectivity across our sensor passband, the visual magnitudes would be approximately 0.13 larger (dimmer) than our apparent magnitudes. Since we do not know the spectral reflectivity of the objects we observe, we keep our apparent magnitudes in their native scale to be self-consistent.

As can be seen, the majority of the objects in our catalog, at least in the GEO and MEO regimes, reside in the 10-15 range of apparent magnitudes. False detections as determined by SLATE are omitted. The HEO objects are significantly more diverse in magnitude (observations on HEO objects comprise only 19.7% of those below 15.5 magnitude but 51.1% of those above). Note that because NITRO generally attempted to schedule dwells on objects at high solar phase angles, the reported apparent magnitudes may have been lower (brighter) than the true average apparent magnitude for the observed objects.

Maintenance of particularly dim objects was attempted but was not a focus during the evaluation period; consequently, the dimmest objects regularly tracked during that time had average magnitudes (as observed from our produced observations) in the upper 16 range. Subsequent work on DART and the overall catalog maintenance pipeline has substantially improved the detectability thresholds and allowed for more consistent maintenance of dim objects as well as those with highly-variable apparent magnitude. Figure 13 shows the apparent magnitudes of recent observations on thirty dim objects in our current space catalog, indicating that the image processing pipeline is capable of detecting objects down to about 18.5 apparent magnitude (and that the catalog maintenance pipeline is capable of tracking said objects).

3.2 Accuracy

A fundamental goal for the NTN from its beginning has been to produce robust high-quality orbital state estimates through the fusion of astrometric observations. Therefore, we have devoted significant time toward the development of image processing and observation reduction algorithms.

Our primary means of astrometric calibration/reduction has been through comparisons with publicly-available high-accuracy ephemerides, particularly the GPS constellation. All sensors in the NTN regularly collect images on GPS satellites to allow for continual calibration and detection of any issues in the processing pipeline. These

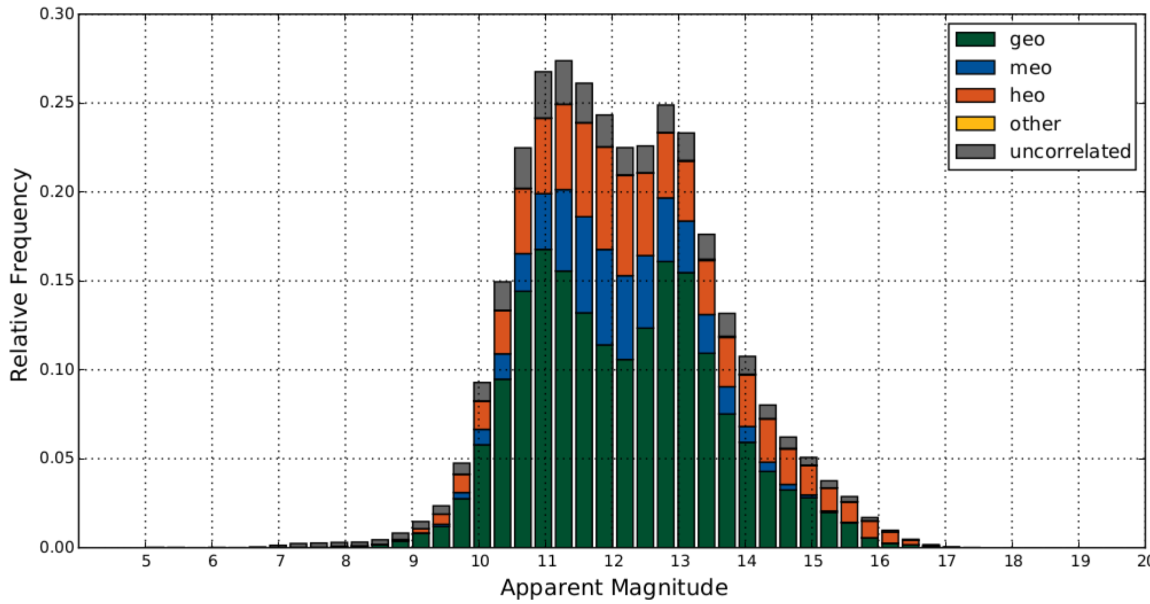

Figure 12: Distribution of Magnitudes for Correlated Observations During 2017 Evaluation Period

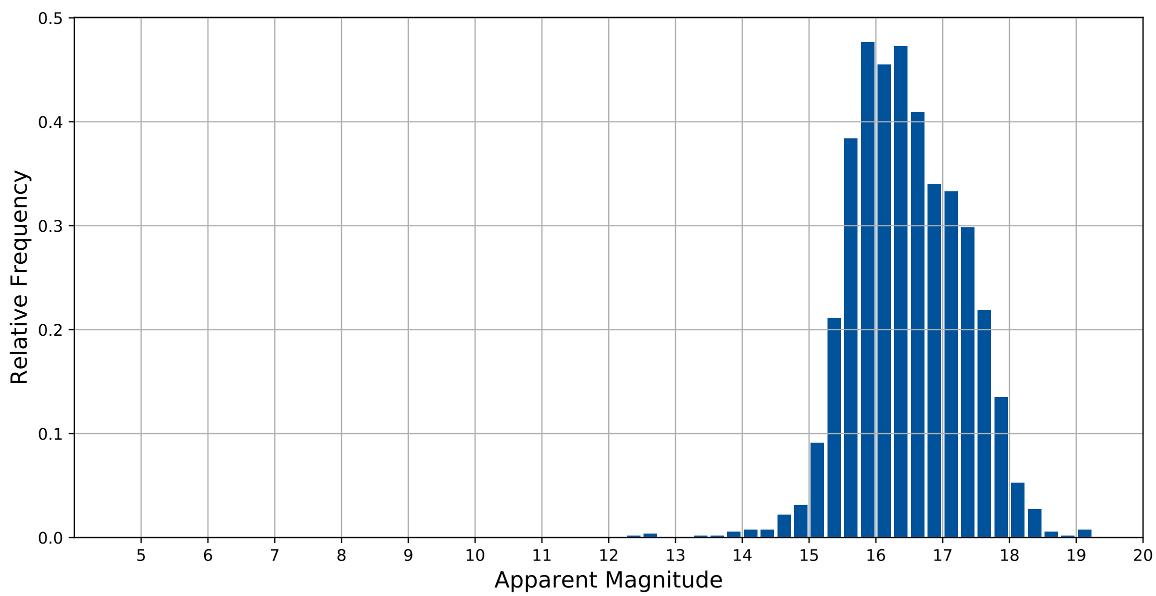

Figure 13: Distribution of Magnitudes for Recent Observations on 30 Dim Objects

calibrations along with dozens of other sensor and system-level metrics are monitored to ensure any problems are quickly identified and resolved.

During the evaluation period, a constraint was added to NITRO to require it to assign at least one GPS satellite for viewing from each dedicated telescope each hour (i.e., a revisit rate constrain across a class of satellites was enforced). This produced a rich dataset of GPS objects (containing 24,265 observations during the evaluation period) that allowed Numerica to perform regular, automated sensor calibrations by comparing the resulting observations to the public GPS ephemerides.

Specifically, 304 full calibrations were performed using SLATE during the evaluation period on Numerica-operated sensors (an average of one calibration per sensor every 1.8 days). Each of these calibrations produced an estimate of (i) sensor biases in right-ascension, declination, and time; (ii) sensor weights (uncertainties) parame-

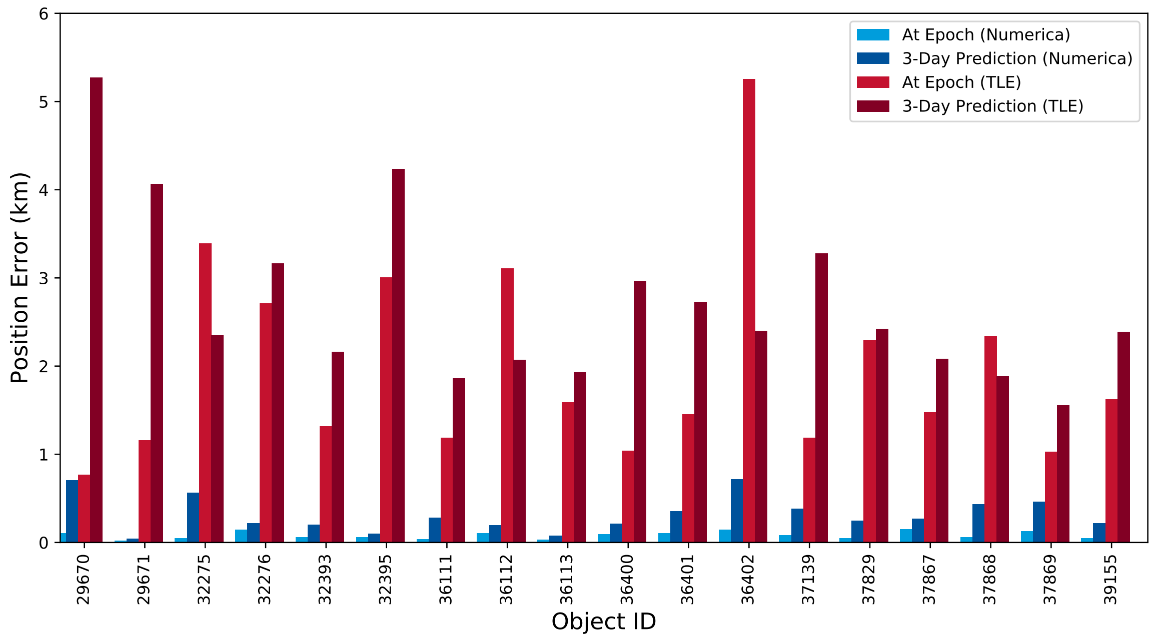

Figure 14: Position Errors for GLONASS Satellites

terized in the in-track and cross-track direction of the object on the focal plane; and (iii) the amount of temporal correlation between the errors in the observations produced by the sensor. Processing chain improvements made during the evaluation period reduced the average observation error on GPS objects to 1.1” in-track and 0.7” cross-track, i.e., subarcsecond. Additional improvements are underway to further improve astrometric accuracy.

Both GPS and GLONASS ephemerides were used to evaluate the quality of our orbital state estimates on these objects throughout the evaluation period. As the GPS objects received preferential tasking for calibration purposes and were used to inform the estimated sensor biases, one would expect to see good performance on these objects and that was generally the case, with median epoch and three-day-prediction errors of 80 and 180 meters at the end of the evaluation period. Encouragingly, performance on the GLONASS objects, which were not treated differently by the pipeline in any way, was similar with median epoch and three-day-prediction errors of 70 and 260 meters, respectively. Figure 14 gives the GLONASS errors per object, as compared to the closest public TLE. Present-day performance is slightly improved with current GPS median epoch and three-day-prediction errors of 50 and 130 meters, respectively.

In addition to astrometric line-of-sight measurements, the NTN also produces instrumental photometry as part of each metric observation. These instrumental measurements are then calibrated using the measured brightness of identified stars in the source image, taking into account both the emission spectrum of the star and the response spectrum of the individual sensor. Various techniques are also in-place to handle data corruption due to transiting stars and atmospheric conditions including haze and partial cloud cover. Work is currently underway to improve this calibration process, but the results thus far are promising. Within a typical image on a clear night, the errors between measured and cataloged star brightnesses across a wide variation of emission spectra deviate, when adjusted, by less than 0.1 magnitudes from a single photometric zeropoint.

A sample set of light-curves are displayed in Figure 15. The figure contains several nights of data in the first two weeks of 2018 on the SGDC satellite (NORAD ID 42692) from an NTN telescope located in Chile. Under good viewing conditions, noise levels are low and a high level of night-to-night consistency is maintained, allowing for the identification of interesting changes such as the peak which gradually develops at around 28 degrees.

3.3 Timeliness

The utility of a real-time space catalog greatly expands as both the age of the orbits cataloged, and the operational latency, are reduced. Reducing the operational latency allows data to more quickly make the transition from the

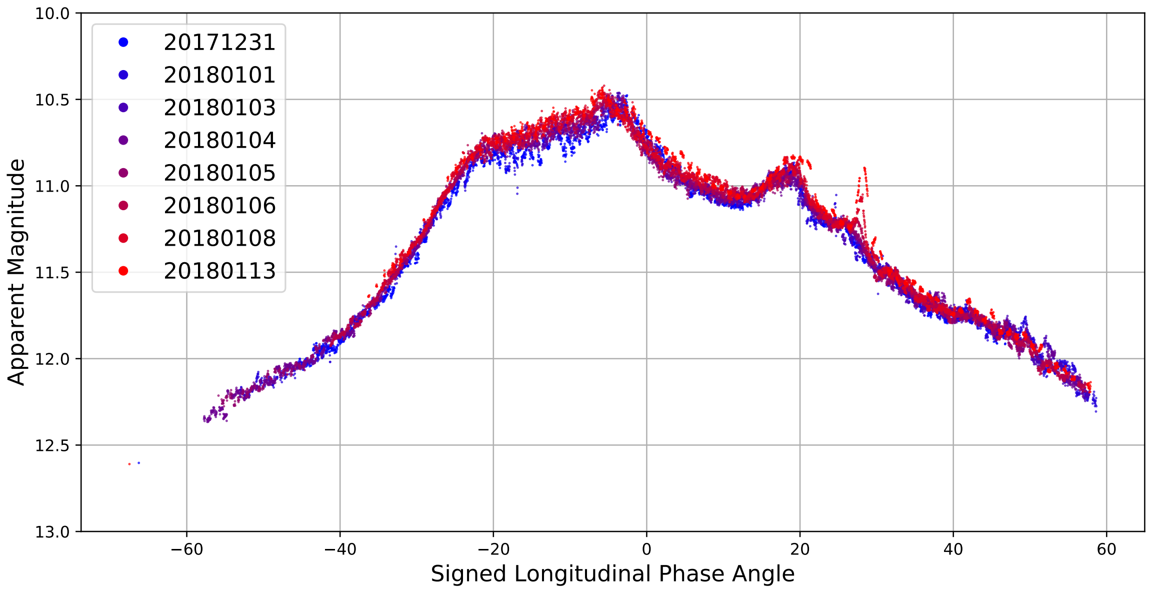

Figure 15: Sample Multi-Night Photometric Collect on SGDC

optical system to the catalog, resulting in more up-to-date orbital states and supporting additional capabilities such as tactical surveillance and automatic detection and follow-up on interesting events. This then improves the "best-case" latency on objects and events which are currently of interest. Complementary to this, increasing network capacity and the efficiency of sensor tasking reduces the average age of objects in the catalog by enabling more frequent follow-ups, ensuring that at all times the catalog as a whole contains up-to-date, actionable state estimates.

Minimizing latency has not been the highest priority of development thus far; nevertheless, the network was designed from the outset to support rapid data processing from image collection at the sensor, through image processing and observation formation, to association and the updating of orbital state information, as well as higher-level alerting functions. In addition, current developments including the transition to a new communication protocol between processing nodes and the remote installations and the deployment of additional on-site processing power, promise to further reduce the current operational latency.

Figure 16 shows the operational latency of the primary catalog maintenance pipeline. Both the delay between image capture and insertion of observations from that image into the observation database, and the delay between image capture and the production of an updated orbital state estimate are shown for a typical day during the 2017 evaluation period and for a typical day from the summer of 2018.

The observation generation latency is comparable across the two days, peaking at around two minutes. The heavy tail seen during the evaluation period is somewhat reduced. This is not surprising as development of the image processing pipeline thus far have been focused on improving data robustness and accuracy rather than latency. A more substantial improvement of about 50% is seen in the catalog update latency, indicative of the recent restructuring of this portion of the pipeline to improve efficiency and robustness. Additional development is underway to further reduce both types of latencies.

Outside of tactical situations, the age of any given object in the catalog is more a function of the revisit rate on that object than the latency of the data produced on it (although with the newly-deployed sensor arrays, the age of many geostationary objects is reduced to the operational latency at night across the viewing sites). It is therefore important when evaluating timeliness to also look at the age of the catalog as a whole. Here we define the "average age" of an object over some time interval to be the integral of the object's age evaluated continuously as the time since the last update at any given time, over the time interval, normalized by the length of the time interval. It is a closely related metric to the update rate previously discussed but differs in that long coverage gaps are penalized even if the total number of updates over the time interval remains constant.

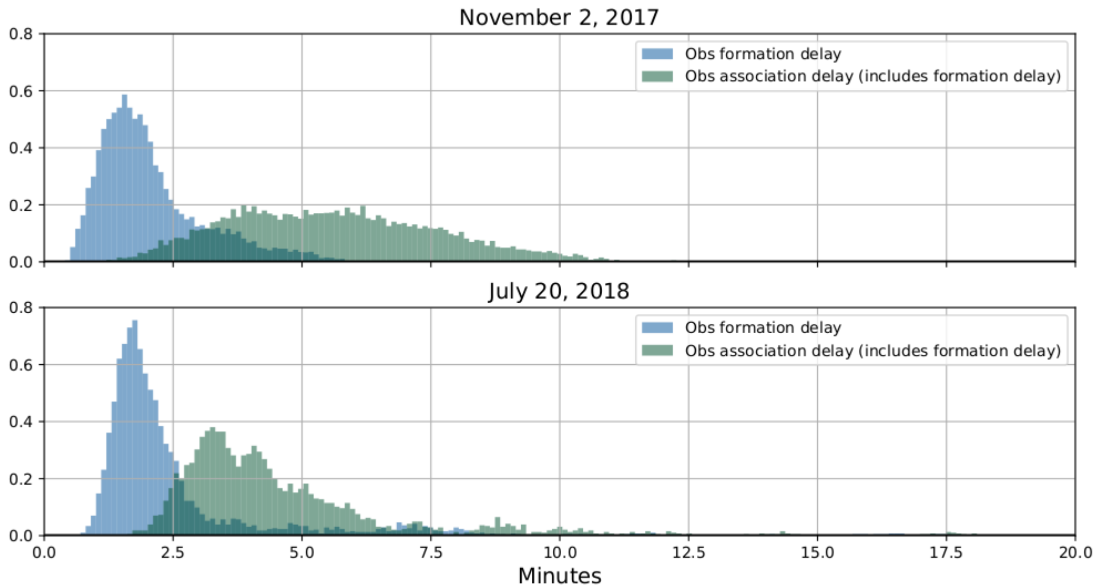

Figure 16: Operational Latency Examples

During the evaluation period in 2017, the median of the average ages of all maintained objects, by regime, was as follows: 1.3 days (GEO); 1.1 days (MEO); 1.7 days (HEO); 2.8 days (OTHER). We expect these values to significantly decrease with the expansion of the network, the deployment of sensor array systems, and the improvements to NITRO that are currently underway. In particular, the age of relatively bright, low-inclination objects in GEO should be held near-zero throughout the majority of nights due to the sensor arrays, resulting in an average age significantly less than one day (closer to a few hours). This in turn will allow the expanded telescope network to devote more resources to the HEO, MEO and inclined/dim GEO regimes, if needed. This, along with other efficiency improvements could reduce the average age in these regimes by more than a factor of two on a similarly-sized catalog. Alternately, a greatly expanded catalog including more than double the number of HEO objects could be maintained with a similar update rate.

4. CONCLUSIONS

SSA data collection (e.g., from ground-based optical sensors) is no longer limited to expensive hardware and legacy software. Pairing "prosumer-grade" hardware with modern algorithms and software yields a system that is capable of supporting mission objectives with precise data and information. Over time, improving algorithms will be able to do even more with available data, hardware will yield higher performance at lower price points, and further automation will result in systems that produce even-higher-quality information in near-real time with very little human intervention required.

Numerica's SSA data collection and processing system has been externally validated and is advancing rapidly, as evidenced by the results presented in this manuscript and summarized in Table 1, where we compare the status of the NTN in November of 2017 to its status at the time of writing. Note that the accuracy and detection results refer to the medium-aperture telescopes. Analysis of the small-aperture systems is underway.

This manuscript serves as an introduction to the NTN and describes how a global network of non-traditional heterogeneous sensors, when coupled with advanced algorithms and high-performance software, can be used to collect and exploit non-traditional data for improving situational awareness in GEO and in other deep space orbital regimes. We hope that these results help strengthen the case for pursing non-traditional and/or commercial solutions for SSA and space control.

Table 1: Summary of NTN Status and Results

	Late 2017	Late 2018
# Sites	10	15
# Systems	11	25
# Sensors	11	>130
GEO Coverage	100%	100%
Deep Space Coverage	75%	100%
Obs Accuracy (Mean)	≈1.0"	<1.0"
Obs Formation Latency (Median)	2.5 min	2 min
Detection Threshold (Apparent Mag)	16.5	18.5

ACKNOWLEDGMENTS

This work was funded, in part, by two Phase II SBIRs from the Air Force Research Laboratory Space Vehicles Directorate (FA8650-17-C-9201 and FA8650-18-C-9209). We also thank the GEO Odyssey team for their scientific inputs and for performing an independent evaluation of Numerica's space catalog. Lastly, we thank Andrew Yee and the Business Operations team at Numerica for their support in deploying the telescope network.

DISTRIBUTION

Cleared for Public Release (AFMC-2018-0311, 17 Aug 2018).

REFERENCES

[1] A. B. Poore, J. M. Aristoff, and J. T. Horwood, "Covariance and uncertainty realism in space surveillance and tracking," Tech. Rep. AD1020892, Astrodynamics Innovations Committee (AIC) Working Group on Covariance Realism, June 2016.

[2] J. M. Aristoff, J. T. Horwood, and A. B. Poore, "Orbit and uncertainty propagation: a comparison of Gauss-Legendre-, Dormand-Prince-, and Chebyshev-Picard-based approaches," *Celestial Mechanics and Dynamical Astronomy*, vol. 117, pp. 13–28, 2013.

[3] J. M. Aristoff, J. T. Horwood, and A. B. Poore, "Implicit Runge-Kutta-based methods for fast, precise, and scalable uncertainty propagation," *Celestial Mechanics and Dynamical Astronomy*, vol. 122, no. 2, pp. 169–182, 2015.

[4] J. M. Aristoff, D. J. C. Beach, P. A. Ferris, J. T. Horwood, A. D. Mont, N. Singh, and A. B. Poore, "Multiple Frame Assignment Space Tracker (MFAST): results on UCT processing," in *Proceedings of the 2015 AIAA/AAS Astrodynamics Specialist Conference*, (Vail, CO), August 2015. Paper AAS 15-675.

[5] J. T. Horwood, N. Singh, J. M. Aristoff, and A. Bhopale, "KRATOS: Kollision Risk Assessment Tool in Orbital Element Spaces," in *Proceedings of the 2016 Advanced Maui Optical and Space Surveillance Technologies Conference*, (Wailea, HI), September 2016.

[6] N. Singh, J. T. Horwood, J. M. Aristoff, and J. Murray-Krezan, "Athena: a data-driven anomaly detection and space object classification tool for SSA," in *Proceedings of the 26th AAS/AIAA Space Flight Mechanics Meeting*, (Napa, CA), February 2016. Paper AAS-16-447.