Power of Persistence:
Persistent custody of objects through repurposed meteorite trackers and observations processing at real-time rates and volume

Joseph Diamond, Rich Osedacz, Weihan Yang, Matthew Sirocky
Peraton

Greg Madsen, Abraham Chan, Timothy Payne
Lockheed Martin Australia

Matthew Bold
Lockheed Martin

ABSTRACT

The current SDA paradigm of optimizing and prioritizing collection of a tasked system poses challenges in today's space domain where there is an ever-growing number of participants. It seems more probable that a critical space event will be missed, given the pressures to direct collection resources to "higher value" objects. Peraton and Lockheed Martin Australia are exploring an alternative approach to collaboratively demonstrate the potential of persistent, untasked optical observations that reduce the effective space object revisit time to near zero. With this collection strategy, we propose to evaluate the use of uncued/untasked, very wide field, high-cadence optical observations utilized in conjunction with a statistically driven association/fusion processing engine. We hypothesize this paradigm will satisfy custody mission requirements and information needs on both benign and high-interest objects simultaneously, in near real-time.

Our primary goal is to understand whether this collection and processing methodology can result in persistent SDA custody, specifically with regards to objects in GEO, by reducing the effective revisit time to approximately 10 seconds for observable objects. Our observation source, FireOPAL, leverages a repurposed meteorite-tracking system developed as part of an R&D partnership with Lockheed Martin Australia and Curtin University Space Science Technology Center, termed FireOPAL. The FireOPAL system is comprised of six ground-based optical sensors along the east coast of Australia covering GEO longitudes 98E – 185E, with fields of view covering 20°x13°. Each of these sensors performs uncued/untasked nighttime operations (0800-2000Z), autonomously publishing astrometric observations on all objects in all orbit regimes. The resulting dataset is comprised of 200k+ observations of 2,000+ satellites per night. Peraton will process and fuse the observations to reconstruct orbit trajectories for both maneuvering and non-maneuvering objects alike. Peraton will apply our statistically driven fusion engine that is specifically designed to handle the observation rate and volume FireOPAL produces.

Exploiting this data to achieve SDA presents several notable challenges. Foremost, the observations FireOPAL produces do not include a pre-assigned/pre-associated space object identifiers. This may present significant challenges to processing systems that rely on such pre-assignments. Further, observations collected in this way are not aggregated into tracks, rather each measurement is represented as a standalone observation. Finally, the rate and volume of this data presents an intrinsic challenge from a data processing and scalability standpoint. This study will evaluate whether Peraton's mature statistically-driven association/custody processing software and tools are able to overcome these challenges under the load of FireOPAL's near real time observation stream. Peraton built this tool to process observations strictly based on astrometric measurements, irrespective of sensor-provided assignments, appropriately model and apply errors to the process and handle large volumes of near-real-time observations.

This collaborative study aims to show the benefits of using untasked, continuous optical observations for space domain awareness. We will use Peraton software to process FireOPAL data, a system that tracks satellites and space debris with wide-field-of-view sensors. We will also address any issues with the data, assess the outcomes and implications for the broader SDA community, and compare our method with the conventional tasked approach. Our hypothesis is that by observing most geosynchronous regime objects with passive sensors, we can reduce the revisit time to almost zero and enable an SDA system to respond significantly faster to potential threats and hazards in space.

1. INTRODUCTION

Some current space object collection paradigms for Space Domain Awareness (SDA), especially for geosynchronous (GEO) space objects, are centered around the creation and execution of sensor tasking and collection schedules. The collection and tasking optimization problem itself is a challenging one, especially considering disparate sensor systems may operate under different command structures, prioritization schemas, and schedule formats. The premise of this challenge is to optimize the set of taskable sensors to ensure adequate revisit rates against high priority targets, while maintaining custody of the larger body of cataloged objects, and to provide adequate spatial coverage for the discovery of unknown objects.

The primary concern with tasked systems is the risk of missing critical, threatening events that affect other space systems. Collection resources are generally weighted toward High Interest objects (HIOs). Therefore, there may not be enough collection resources to detect an event or provide enough data to assess the “threat potential” of the event. Other events such as breakups also suffer as resources are focused on "what you know” vice “what you don’t know.” The collection optimization trades among HIOs, remaining space objects, and object discovery are exacerbated by the exponential growth of new space objects. These trades can lead to delays in determining potential threats thereby shortening or even eliminating the time necessary for an affected space system to respond to the threat. So, what can we do to monitor all objects and assess their threat potential to create a safe operating environment for operators in real-time?

This paper addresses an alternative collection paradigm where wide field-of-view sensors are always on – persistent – and collect on everything within the field-of-view meeting detection criteria. In short, how does collection, processing, and actionable information generation change in the absence of tasking and scheduling since sensors are always collecting data? With this new paradigm observations are continually streaming into the processing/fusion processor on all detected objects, updating orbit states, detecting and reconstructing maneuvers (as appropriate), and assessing threat events. This aspect simultaneously meets the High Revisit Rate (HRR) collection requirements on HIOs as well as other custody requirements. Additionally, benign and threatening events can be detected and assessed nearly as soon as the event occurs as the revisit time is effectively zero thus maximizing the time for operators to respond to potential threats in an operationally relevant timeframe.

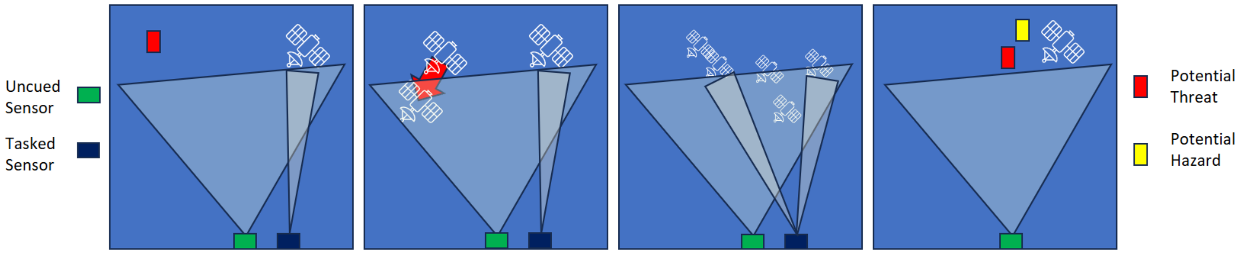

Early detection of unknown hazards & threats

Space event detection, alerting and cueing of sensor systems

Target prioritization for sensor networks

Operations support, battle damage assessment & resiliency

Fig. 1. Uncued vs Tasked Sensing for Threats and Hazards

2. COLLECTION METHODOLGY

The observational data used for this study is from the FireOPAL network of electro-optical sensors. These sensors were developed jointly by Lockheed Martin and the Curtin University Space Science and Technology Centre in Perth, Western Australia. The FireOPAL sensors are modified versions of systems from the Desert Fireball Network that is used to track and recover meteorites. A summary of the FireOPAL capabilities as they pertain to this study are provided here; further details may be found in [1,2].

FireOPAL comprises several wide-field, ground based staring sensors. For this study, data from five prototype sensors, all co-located near the East coast of Australia, were used (Fig. 2). Each sensor system is primarily composed of commercial off-the-shelf components. A 40MPix commercial DSLR camera is combined with a 105mm/f1.4 lens for a field of view of approximately 20° by 13° per sensor. The five sensors are pointed at different locations along the GEO belt covering GEO longitudes from 98E - 185E (see Fig. 3).

Fig. 2. FireOPAL sensors deployed near Uralla, NSW Australia

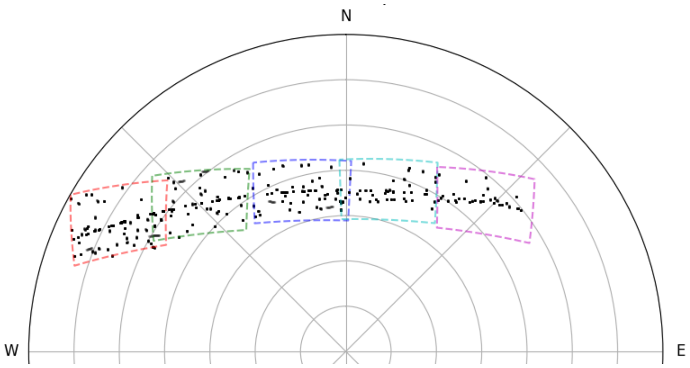

Fig. 3. FireOPAL sensor coverage for this study.

The above Fig. 3. FireOPAL sensor coverage for this study. The zenith is at the center of the bullseye, with North up and East to the right. Concentric circles are spaced by 10○ in elevation and extend to 20○ above the horizon. The boundaries of the fields of view of five sensors are shown as dashed colored lines. The predicted locations of GEO objects with the fields of view using the public space-track.org catalogue are shown as dark dots.

Each camera is triggered by GPS to acquire a 5 second exposure image every 10 seconds. Before the next image is taken, the image processing system identifies and measures all objects in each image that appear to move like satellites, i.e. long streaks (LEO), short streaks (MEO), and quasi-stationary point sources (GEO). The calibrated results are aggregated onto a system hosted by Amazon Web Services (AWS). The fully calibrated observations are available to be published to customers within 10-20 seconds after each image is taken (e.g. through the Unified Data Library). This cycle is repeated all night between local nautical dusk and dawn. Because the sensors are pointed toward the GEO belt, this provides near-continuous observational coverage of objects in GEO within the approximately 1,300deg^2 field of view of the system.

The widefield, staring nature of FireOPAL provides several significant, unique advantages for optical SDA sensing, including:

- Observing hundreds of satellites simultaneously, in all orbit regimes;
- Delivering large volumes of data (approximately 60,000 metric observations per sensor per clear night) at low cost with high accuracy, good sensitivity, and very low latency;
- Operations do not require any tasking; each image is processed independently and all objects that appear to move like satellites are reported, regardless of their presence in any catalogs; and
- Operations are not affected by changes in tasking requests, i.e. there is no need for schedule de-confliction or to prioritize observations of some targets over others.

One principal limitation of the system is that coverage is limited to local nighttime and only under clear or partly clear skies. The image processing system is able to identify satellite candidates in small areas within the field of view when other parts of the field of view are obscured by clouds. Another limitation of the current prototype/demonstration system is geography; the sensors are all co-located at one site. However, the system architecture and software is designed to be scaled to a large number of sites and large number of cameras per site.

FireOPAL data have been used in a number of demonstrations and trials, e.g., Sprint Advanced Concept Training (SACT) run by the Joint Commercial Operations (JCO). Independent analysis of the results by data curation providers have shown the astrometric accuracy of FireOPAL observations to be a few arcseconds with timing accuracy of approximately one millisecond. The sensitivity of the system is approximately 14th magnitude for single images that are used to find LEO and MEO objects. Several images are combined to identify objects in GEO, resulting in a limiting magnitude of about 16. The photometric accuracy is typically less than 0.1 magnitude but may be greater for some objects in GEO on inclined orbits.

Although this study is primarily focused on GEO, it is worth highlighting the capabilities of the FireOPAL system for LEO. A single site comprised of 12-16 cameras would be sufficient to detect LEO objects as they rise and set over the site, providing sufficient coverage to refine their orbits. In addition, the staring nature of the system means that objects in LEO appear as long streaks. These streaks can be processed to measure the change in brightness along the streak that, in turn, could be used as a ‘fingerprint’ or to infer properties such as spin rate or stability of the object. Because each image is a 5 second exposure, and depending on the apparent speed of the object, the change in brightness could be sampled as often as once every few milliseconds. An example of this capability is shown in Fig. 4.

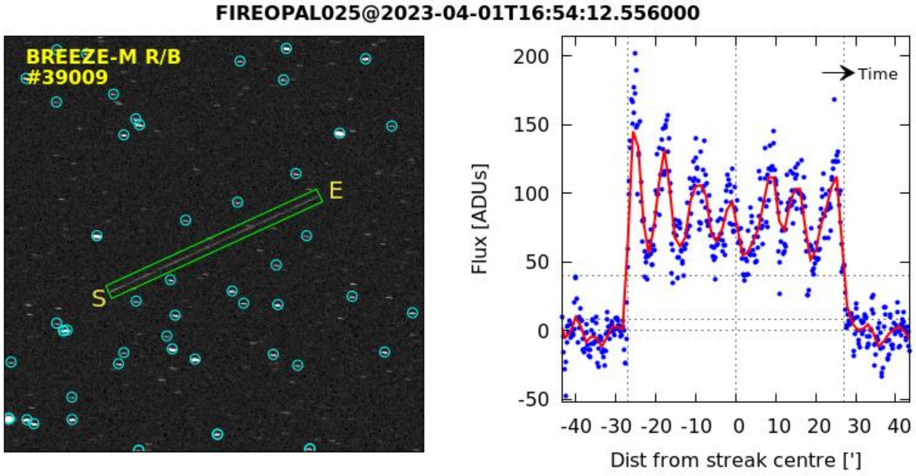

Fig. 4. Sample showing high time resolution light curve from FireOPAL

The panel on the right shows the change in apparent brightness of the streak over time, suggesting a regular periodic change in brightness. With an exposure time of 5 seconds, the figure suggests there is a rotation period for this rocket body of 1.2Hz.

Because of the uncued nature of FireOPAL, processing the metric observations into tracks and orbits can be a challenge. The sensors are not tasked to deliver observations of particular objects; observations of all objects that appear to move like satellites are reported for each image simultaneously. The image processing system is optimized to deliver high rates of metric observations with low latency; it is not designed to reliably associate observations with each other over time. Consequently, there can be ambiguity as to which satellites corresponds to which metric observation as published by the image processing system. This is a challenge for many orbit determination software tools that assume each observation is uniquely associated (or 'tagged') with a particular object.

The FireOPAL image processing system does attempt to provide a crude association/tag for each observation. If an individual observation is consistent with the predicted location and motion of an object in the current public space-track.org catalogue then the observation is tagged with that SSN ID. However, this assignment may not be accurate, TLEs from space-track.org may be based on stale observations and not be accurate, the public space-track.org catalogue is not complete and FireOPAL regularly observes satellites that are not in that catalogue (UCTs) and in regions of space with a high density of satellites, e.g. GEO, cross-tagging may occur. To better match observations to objects it is necessary to implement a comprehensive near real-time, statistically driven data fusion and association process, discussed in the following section.

3. DATA PROCESSING APPROACH

The fundamental objective of any SDA support system is to assist decision makers with a timely, relevant, and actionable understanding of the operational space environment. This requires real-time foundational and accurate knowledge of space objects, an ability to produce continual real-time updates and hazard/threat assessments on not just high interest objects, but on all space objects. Without timely, accurate, and geometrically diverse observation data that lends insight into event detection, it is possible there will not be enough time for affected operators to respond resulting in mission failure.

The Peraton fusion system used to process the FireOPAL data performs the fundamental functions: associate observations, update orbit states, and maneuver detection and reconstruction. So what is this system doing differently? The primary difference is the system is designed to process streaming data vice observations that come in the form of "tracks". Thus Peraton's fusion engine is a natural fit with FireOPAL's collection paradigm. All observations are processed as they come in with no checks performed to determine whether an update is necessary (e.g., Length of Update Interval (LUPI) or Vector Magnitude (VMAG)). This critical processing ability enables timely event detection and threat assessment through continual and timely state updates and event processing. Because the primary mission of our fusion engine is real-time threat assessment, our fusion engine is statistically driven in all its processing. Processing observations in this manner promotes accuracy (provided the proper observations and types are processed) and maximizes the time for affected systems to respond to threats in operationally relevant timeframes. This real-time processing approach is not without its challenges, primarily weather effects and resource costs.

Being an optical system at approximately 800 meters altitude, FireOPAL suffers from cloud cover and other weather conditions affecting operation. Thus, the fusion engine will not get the anticipated streaming data on all objects in the field-of-view. This occasionally leads to inaccurate states and change detections on the uncollected objects whose effects could ripple through the system. However, being a statistical system, the orbit state errors commensurately grow allowing the system to correctly associate observations and produce state updates. One way to mitigate this issue is to fuse observations from non-affected sites either within network or from other data providers. While this processing system is specifically designed for such multi-source fusion, for the purposes of this paper the FireOPAL data was processed independently.

The other challenge is managing cloud compute costs. Cloud infrastructures bring the ability to process large data volumes within actionable timeframes via a wide array of software services and hardware configuration options. The key cloud processing attribute for processing FireOPAL observations is the ability to automatically scale the processing based on demand.

Though automated scaling based on demand is a key feature of cloud computing, it also necessitates optimally designing the fusion engine to be horizontally scalable with processing load. Understanding the impacts of continuous, voluminous, streaming observation data processing on compute resource can be a challenge, and likely requires characterizing baseline performance and costs relative to the demand for those resources.

4. PROCESSING RESULTS

For GEO observations processed from a 25 day period from May 17 – June 12, 2024, there were 1,102,277 observations that were associated to known objects. This processing was initialized from TLEs from spacetrack.org from which observations initially associated. As these observations are associated, new state vectors are continually generated with covariance, against which future observations are associated. The total associated observations over this time period are seen in Figure 5 below, separated by color. Importantly, this data set only includes observations provided from the FireOPAL sensors; however, this data would ideally be fused alongside other observation sources to mitigate local weather and lighting outages.

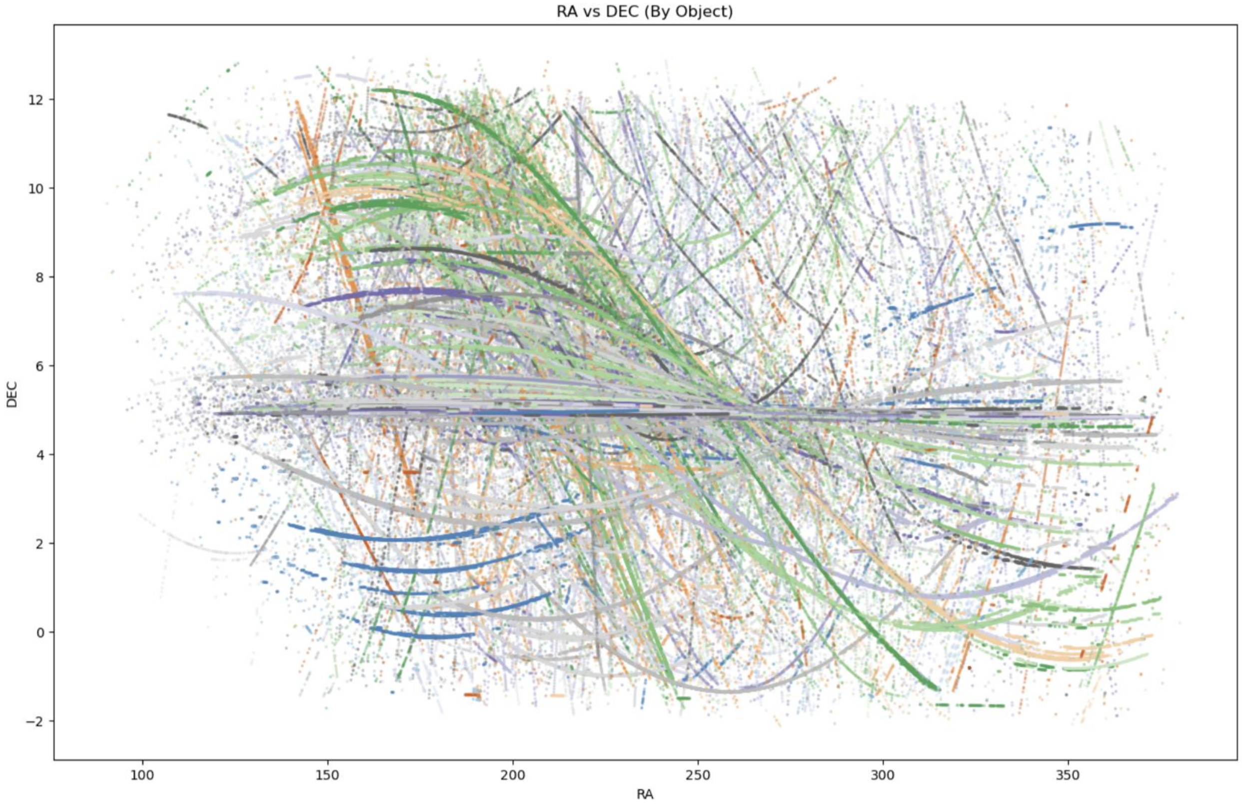

Fig. 5. Statistically Associated FireOPAL Observations

The observations were ingested and processed utilizing the observation measurements, irrespective of any sensor-provided object ids, incorporating a sensor model calibration methodology into the observation statistics. The revisit rate for observable GEO space objects was seen as frequently as 10s under clear, dark conditions. The resultant frequency of updates to space object orbit states varied object-to-object, but under ideal observation conditions, state updates were produced every 1-2 minutes. The availability of ideal viewing conditions, however, was highly variable with weather and daylight conditions. For most objects, there are far fewer observations available during dawn and dusk. Further, the effects of weather obscurations appear particularly impactful in the GEO data, due to the means by which the FireOPAL processing stacks frames to achieve the visual magnitude sensitivities required for GEO observations. The data processed for this paper covers the May and June timeframe, which are typically among the most cloudy months of the year in this part of Eastern Australia. It understood that these sensors are repurposed meteorite trackers, and as such, were neither purpose-built nor ideally located for this mission.

Heatmaps of observation frequency for several frequency observed space objects below, with the number of observations (per hour) seen in darker shades:

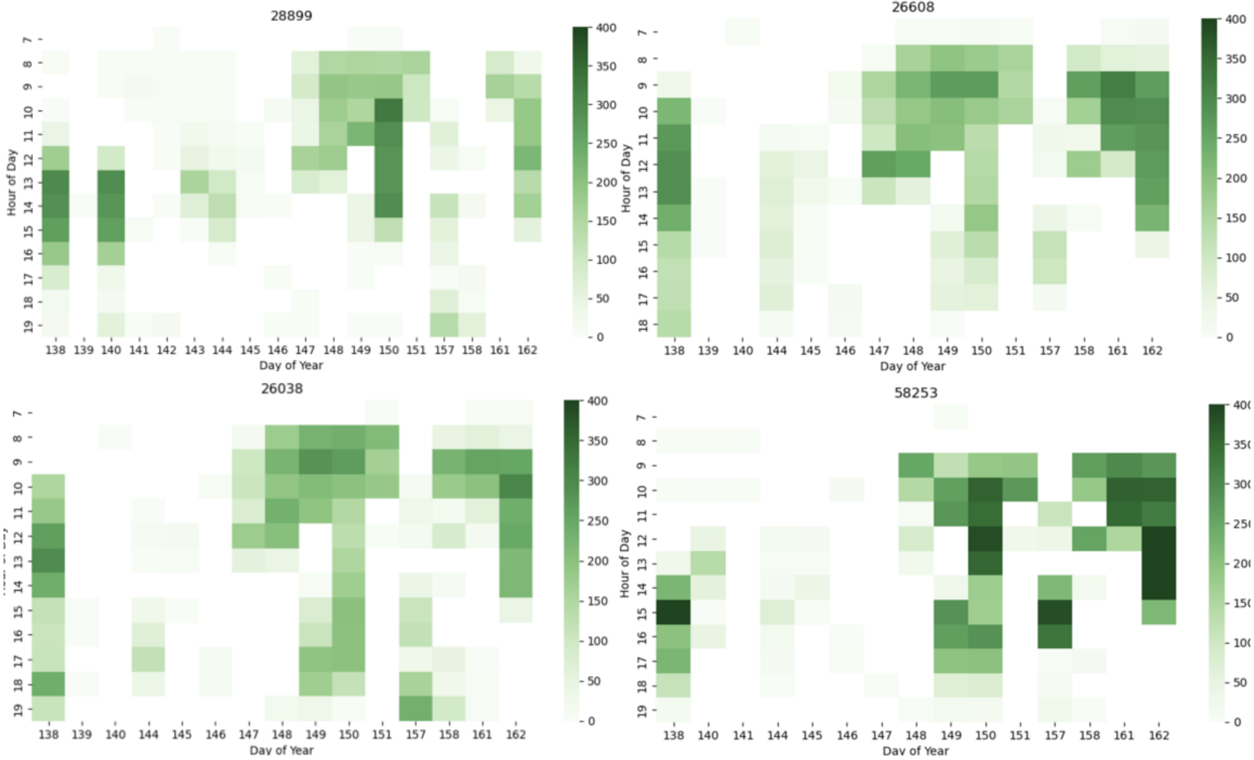

Fig. 6.FireOPAL Observation Heatmaps

These observations would ideally be fused together with measurements from other sources of EO observations, as well as radar and passive RF to achieve the most comprehensive and timely coverage of space objects. While this processing system has been able to leverage that kind of fusion with other data sets, for the purposes of this paper, the FireOPAL sensors were evaluated on their own. Despite the aforementioned environmental gaps in coverage, the sensors' observations nominally associated to its observation-derived state vectors with minimal post-association residuals, typically within 2-3 arcseconds of angular error, as exemplified below:

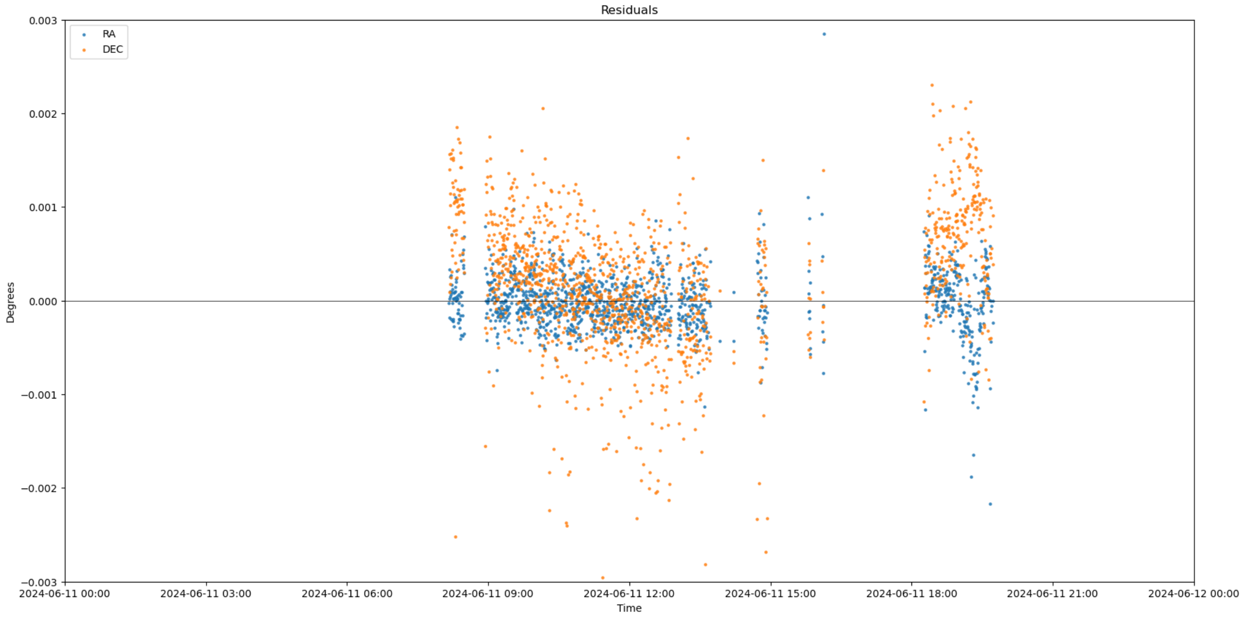

Fig. 7. Post-Association Observation Residuals for Object 54230 in RA and DEC

As previously stated, a key value of persistent observations such as FireOPAL is the ability to provide near real-time state updates and maneuver detections. The amount of time needed to complete data processing and produce results is a factor of the processing resources allocated to the processing system. Designed to be modular and horizontally scalable, the processing speed under load is able to increase as additional resources are added; however, realizing these performance gains incurs the additional financial cost associated with the necessary cloud computing resources. For the purposes of this paper, these costs were minimized, and as such the speed of output generation varied with the data load. This is especially true for maneuver detection processing.

Due to these constraints, there were cases where it took upwards of 10 minutes from the receipt of an observation to produce a maneuver detection, as the associated observation needed to wait for available maneuver detection workers to become available. In cases where the processing resources were not constrained, the time between the initial receipt of an observation and the generation of a maneuver detection message and corresponding post-maneuver state vector was between 80-90 seconds. Properly resourced, it would be expected that the time between photon-on-detector to the detection of a space object maneuver would be consistently less than 2 minutes in time.

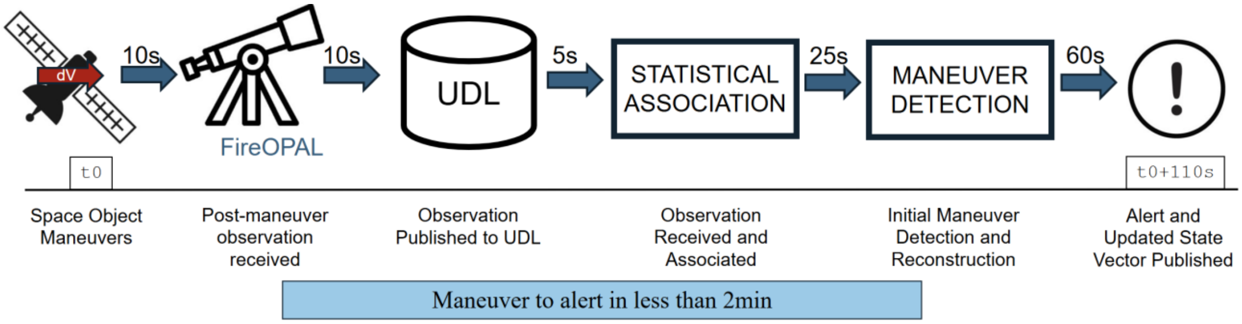

Fig. 8. Maneuver Detection Alert Timeline

As an example, the following maneuver for space object 58253 (CHINASAT 6E) was automatically detected in FireOPAL observations seen below. In this case, while there were no observations during the actual maneuver event, an initial maneuver solution was created less than 60s after the observations were received from the Statistical Association processing.

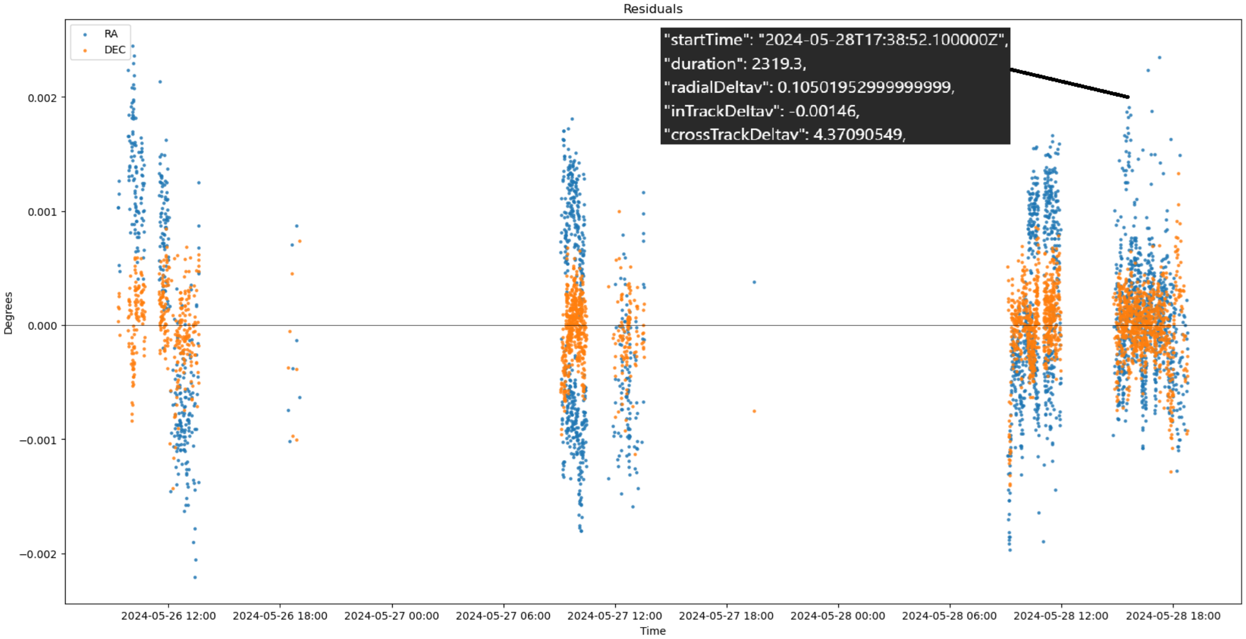

```
"startTime": "2024-05-28T17:38:52.100000Z",
"duration": 2319.3,
"radialDeltav": 0.10501952999999999,
"inTrackDeltav": -0.00146,
"crossTrackDeltav": 4.37090549,
```

Fig. 9. Rapid Maneuver Detection w/ Initial Estimate for Object 58253 (CHINASAT 6E)

5. SUMMARY OF FINDINGS AND POTENTIAL APPLCIATIONS

Persistent, streaming data processing on all detected objects provides numerous advantages over today's tasked systems. Constant object observation and data processing minimizes the time between the event time and actionable information supporting SDA, Space Traffic Management (STM), and other space system operator activities. This reduction allows operators to respond to threatening events on operationally relevant timelines. Shifting the paradigm from tasking and scheduling to a persistent "everything, always on" approach reduces the likelihood of missing a critical event as resources are no longer concentrated on high-interest objects thereby sacrificing collection on resulting unobserved threat events.

Combining wide field staring sensors like Lockheed-Martin Australia's FireOPAL with real-time processing systems like Peraton's SDA engine realizes the benefits of persistent data processing. Enabled by modern sensor, compute, and communications technologies this combination of collector and processing systems has significant advantages. First, persistence drives the revisit time to effectively zero so all events on all detected objects are collected and processed. Second, low observation delivery latency and cloud-enabled processing to dynamically scale to handle streaming data processing produces actionable information as close as possible in time relative to the event. These two advantages provide continuous state updates and statistically based predictions that enable operators to assess potentially threatening and hazardous events with sufficient response time. Optimally architecting the processing software for cloud operations is key to producing a cost-effective solution with this processing paradigm.

Follow-on studies to this paper seek to close some of the gaps this study identified. Leveraging previous architecture work [3], we wish to study how a proliferated FireOPAL-like sensor network impacts the observations' real-time processing and the mission benefits of such an architecture, especially as the capability continues to improve with advances in both hardware and software. FireOPAL is modular by design and is easily scalable to many sites and sensors per site. Fusion of data from other sources generally improves orbit solutions' accuracy but we would like to investigate how better and on what timelines can actionable information be derived; this study would include multiple phenomenologies as part of the trades to close collection gaps. Persistent collection enabling new object discovery is another subject to study as more data will be collected on such objects potentially leading to more timely and accurate unknown object determination. Lastly, we plan to study various cueing strategies resulting from processing persistent data. These strategies can bring more costly but more accurate and weather invariant sensors such as tracking radars into the architecture to determine a cost-effective sensor mix to include wide field of view daytime collection to fill SDA collection and information gaps.

REFERENCES

[1] Bland, P.A., Madsen, G.J., Bold., M. et al. FireOPAL: Toward a Low-Cost, Global, Coordinated Network of Optical Sensors for SSA. Proceedings of AMOS, 2018.

[2] Madsen, G.J., Bland, P.A., Bold, M., et al. FireOPAL: Technical Performance and First Results. Proceedings of AMOS, 2018.

[3] Bold, M., Mason, J., Renshaw, J. et al. Proliferated Sensor Network Performance Study and Architecture Design Optimization. Proceedings of AMOS, 2023.