Report on 2020 Mega-Constellation Deployments and Impacts to Space Domain Awareness

Ryan Hiles
Omitron, Inc.
Amos Alexander
18th Space Control Squadron
1st Lt Nicholas Herzer
18th Space Control Squadron
Diana McKissock
18th Space Control Squadron
Aaron Mitchell
18th Space Control Squadron
1st Lt Jonathan Sapinoso
18th Space Control Squadron

ABSTRACT

The rapid proliferation of low-Earth orbit satellite constellations came into full-force in 2020. The primary difference in these launches compared to historical launches involved the number of simultaneous deployments, the frequency of deployments, and the scaled use of electronic propulsion for orbit-raising. We examine the impacts of this emerging methodology on the space surveillance mission and the improvements made to date to meet the challenges of this new environment. Starting with pre-launch conjunction assessment, new techniques have been adopted to blend risk mitigation practices, system capabilities, and screening responsiveness. During the launch phase, existing sensor management and tasking processes have evolved to ensure custody of all newly launched objects as well as the existing space catalog. This also drove changes during the object separation phase which required new orbital modeling techniques and analyst expertise to distinguish the clustered objects in a short period of time. Novel approaches towards satellite operator-provided ephemerides, in addition to rapid software upgrades, enabled a new field of orbital analysis which will soon dominate the efforts of resident space object custody. The increase of payloads and data also increased the volume of orbital conjunction assessment data, which drove the need for increased collaboration between data providers and satellite operators to ensure safety of operations in the space domain. Finally, the increase in satellites has resulted in an increase in reporting as satellites re-enter the atmosphere prompting a more efficient approach on how these events are managed and reported.

DISCLAIMER

The views expressed in this article are solely those of the authors in their personal capacity and may not necessarily represent the views of the U.S. Department of Defense and its components, Omitron, Inc., or the Advanced Maui Optical and Space Surveillance Technologies Conference.

1. INTRODUCTION

The number of satellites operating in Low-Earth Orbit (LEO) in recent years has shown a fast and steady increase. This expansion, referred to as LEO proliferation, has induced changes across all phases of Space Domain Awareness (SDA) operations. Starting in 2019, LEO proliferation was further accelerated by multiple commercial entities beginning launch operations of their mega-constellations. Unlike previous launches of large numbers of satellites, these mega-constellations are identified by the large number of payloads combined with their orchestrated movements following orbital insertion. The mega-constellation footprint was primarily a result of the SpaceX Starlink fleet which launched twice in 2019 and 14 times in 2020. OneWeb also contributed to the environment with a single launch in 2019 and three launches in 2020. In the near future, the number of mega-constellation launches is expected to grow with 17 mega-constellation launches occurring within the first five months 2021[1]. These are the first of a number of entities expected to field mega-constellations in the coming decade. As such, a concerted effort has been placed on understanding how to manage observation-based tracking of these satellites in an automated manner to maintain accurate and timely SDA despite the changing landscape.

The primary organization charged with executing United States Space Command's (USSPACECOM) SDA mission is the United States Space Forces' (USSF) 18th Space Control Squadron (18 SPCS), whose mission is to provide and advance a continuous, comprehensive, and combat-relevant understanding of the space situation. In 2020, the first full year of mega-constellation activity, the unremitting stream of launches paired with the unprecedented manner of satellite orbit-raising posed an interesting challenge to the traditional approach utilized by 18 SPCS as it pertains to maintaining the USSPACECOM's resident space object (RSO) catalog. From launch through reentry, the legacy SDA architecture was built around the expectation of small numbers of satellites per launch, large inter-satellite separation distances, and rapid orbit-raising and lowering via chemical thrusters. Mega-constellations have turned these traditional concepts on their head, which have affected every aspect of space operations. Ultimately, the 18 SPCS and the US Space Surveillance Network (SSN) have met and overcome the most pressing challenges with a mixture of invention and innovation, and, in turn, postured USSPACECOM to successfully overcome the challenges that remain for the US space enterprise and the space operations community writ large.

2. PRE-LAUNCH CONJUNCTION ASSESSMENT

In accordance with Air Force and USSPACECOM instructions, every launch provider (LP) launching from a USSF installation will receive pre-launch conjunction assessment (LCA) screening from the 18 SPCS. This process requires the LP to provide trajectory data for each piece separating from a launch vehicle including rocket boosters, fairings, satellite deployment vehicles, and satellites. Along with the trajectory data, the LP must also provide 18 SPCS with a screening request and other supplemental information. The process relies solely on LP data for the launch vehicle which 18 SPCS screens against the most current version of the RSO catalog. The results of these screenings are closure reports defining which time-based launch opportunities yield potential conjunctions.

The increased frequency of multi-payload launches has added a new challenge to LCA. For SpaceX Starlink launches alone in 2020, the average launch frequency reached one launch every three weeks with some launch requests happening just days apart to meet the desired cadence. Screening requirements for LCA paired with the high launch cadence stressed 18 SPCS resources in supporting all of the requested, and often customized, LCA services. In addition to launch frequency, the increased payload-to-launch ratio has compounded this load due to the increased number of trajectories that require clearing for each launch. In the initial stages of configuring LCA for multi-payload launches, 18 SPCS attempted to execute standard procedures where each trajectory is screened separately in a linear fashion. In many cases, this resulted in 120 separate screenings per launch, each taking several minutes to complete culminating in hours of processing time for a single launch event. To reduce this strain, 18 SPCS reconfigured LCA screenings for multi-payload launches by having a single trajectory with an over-estimated hard body radius represent a subset of the satellites on board the rocket. Through this method, groups of satellites would be screened as a single trajectory with the hard body radius representing the maximum distance between all

[1] 2021 launch count includes the SpaceX Transporter-1 launch which carried 10 mega-constellation Starlink satellites as well as over 120 additional non-Starlink payloads.

satellites in the cluster. This reduced the number of trajectory files and the overall screening time, allowing accelerated delivery of results to the LP that represented an effective keep-out zone.

As more low-Earth orbital shells populated throughout 2020, LPs reported reduced launch opportunities within their planned launch windows. Due to this, there is growing interest in LCA based solely on probability of collision rather than the historical methodology of miss distance based on stand-off separation. Both methodologies are permitted by current instructions, but the lack of covariance standardization and validation of trajectories precludes full reliance on this approach. Additionally, the hard-body radius modifications to accommodate multi-payload launches mentioned above compromise a probabilistic screening as the results would impart risk assumption for the individual satellites at varying rates while the screening represented the risk of penetrating the satellite cluster. Overall, 18 SPCS has made great strides in meeting LCA challenges, they continue to work with LPs, mission system experts, and higher headquarters to identify the best way to represent a proposed trajectory's path through the RSO catalog. The data, software, resource, and policy limitations have created a need for rapid updating of the entire pre-launch CA segment to continue providing meaningful and actionable results to launch agencies. The details of those updates and solutions are yet to be determined with efforts ongoing in each realm.

3. LAUNCH PROCESSING

For the 18 SPCS, launch processing consists of tracking and identifying new RSOs and establishing orbital state data for them. The current methodology generally includes sensor sites executing search patterns along the projected path of the objects to acquire observational data. For domestic launches or launches following previous launch patterns, 18 SPCS creates nominal element sets to model the initial orbit of the new satellites. For non-domestic launches not following previous launch patterns, initial orbit determination is executed against the collected observations to create the first orbital state. In either case, the launch processing generally ends when the first element set is generated and represents trackable data for subsequent sensor passes.

With mega-constellation launches, the impacts on launch processing are directly related to the number of payloads deployed and the method used for their deployment. Without discrete element sets for each object, the sensor sites must continue to execute a search pattern and collect observations to account for all RSOs on subsequent passes. These observation periods are extended by default when there is no corollary element set in the RSO catalog, causing a larger strain on the sensor site to track the entire deployment on each pass. Even when a single representative element set is produced, the sensor sites continue to treat additional unknown pieces in this manner which extends the tracking load. Providing pre-launch predictive element sets for all objects does not alleviate this issue as the accuracy of predicted insertion element sets is greater than the inter-payload separation when a payload deployer is not used. Finally, software limitations on the initial piece separation process drove the 18 SPCS to prefer that sensors tag all initial observations to a single placeholder analyst number to ensure data consolidation regardless of the sensor site's correlation results. A software upgrade in mid-2020 further improved this approach, but the problems of providing complete, accurate element sets for tracking and correlation persist.

While the initial element set is generally attainable within any of the observed mega-constellation satellite deployment models, attaining the initial element set for all payloads becomes impossible for days after the launch when passive deployment models are employed. The lack of element sets has downstream effects on sensor tasking and tracking which is driven by element sets and on-site correlation. Further, poor on-site correlation degrades the ability to process piece separation quickly or automatically. The technological improvements made to enhance other areas of mega-constellation support offer no major improvement on resolving these issues, but the impacts of the proliferation are far less dramatic.

4. INITIAL OBJECT SEPARATION

The primary SDA function of the post-launch phase is to attain custody of all newly launched RSOs. This begins once an initial element set is produced for tracking purposes and concludes with the ability to reacquire an object at subsequent revolutions and enable accurate observation association and state updates for those objects. The mega-constellation deployments posed a challenge in this regard due to their large numbers, close proximity, and separation velocities. A major delineation emerged here between mega-constellation operators. SpaceX opted to apply angular momentum on a vertically stacked set of satellites to achieve inter-payload separation, which will be referred to here as the passive deployment model. This approach yielded four tension rods used to retain the satellites during flight and 60 satellites which drifted apart at a slow, regular rate until orbit-raising maneuvers began. OneWeb, on the other hand, launched its satellites upon an active deployer with timed increments between grouped payload separations, referred to as the active deployment model[2]. Under this approach, there was a single deployer and 34-36 satellites inserted in groups of four with a high rate of separation.

In the passive deployment model, observation-based satellite acquisition was delayed until the expected satellite separation was sufficient to observe unique objects via the USSF's Space Surveillance Network (SSN). The time between initial deployment and full launch custody required consistent observations of all unique objects. When an SSN sensor reported anything short of 60 objects in a pass, it was unclear to the 18 SPCS analysts which objects were missed making it difficult to clearly associate the observations to consistently enable early-orbit state updates. In 2020, it was determined that analysts would not begin to attempt separation of the objects for the first 48 hours to allow for sufficient inter-satellite spacing, though shorter or longer times were noted depending on availability of analysts as well as when orbit-raising maneuver was initiated. The controlled orbit-raising maneuver phase of the deployments greatly increased inter-payload separation, alleviating the issues caused by satellite spacing. Following this 48-hour waiting period, analysts began manually assigning observations to satellites and creating initial states in the High Accuracy Catalog (HAC). This manual process allowed analysts to begin identifying objects that had best separated from their neighbors and enable automatic maintenance of the objects' orbits. Over the period of a few days, two to four analysts would acquire custody of each object individually as they worked from least to most ambiguous sections of the deployment and acquired custody of all payloads. The initial separation phase continued for these objects until observation association errors, termed cross-tagging, were decreased to such a level that the objects could migrate to the catalog maintenance phase. Transition into the catalog maintenance phase did not indicate complete automatic maintenance, but rather a reduced level of required manual intervention to sustain orbital updates. This transition also generally coincided with entering the payloads into the public space catalog which was generally one week after launch. The total 18 SPCS analyst effort required for this initial separation phase in the passive model was approximately 160 hours of labor.

Under the active deployment model, satellites deployed in groups of four with minutes of time between groups. From an observation point of view, this reduced the requirements for complete tracking per pass and greatly reduced the time between initial deployment and full custody of all objects launched. When a sensor pass failed to collect every independent object, it only impacted the group of four satellites where the observations were missed rather than spreading the impact to the entire constellation. Aside from the inter-group spacing, this self-recovering observation association is attributed to the reliability of the object deployments where each group was known to contain four payloads. Within the partially-tracked group, 18 SPCS analysts were able to readily determine which of the four satellites was untracked due to the greater inter-payload separation paired with the much greater separation between neighboring deployment groups. This effectively negated the impact of the partial tracking in terms of the timeline required to attain and maintain custody of all objects. Additionally, the payload deployer imparted separation velocity between the payloads which accelerated the timeline to initial custody. Due to these two key features of the active deployment model, initial custody was achieved just hours after orbital insertion depending on SSN and analyst availability, though final cataloging and processing often occurred 24 hours after launch to allow for greater data density. The greater satellite separation achieved with the active deployment model was more compatible with traditional SDA processes in terms of system and analyst expectations and responses, reinforcing the systemic success for the 18 SPCS. While there were more payloads per launch than traditionally expected, the deployment strategy and timing allowed for unambiguous association of observations and thus automatic state updates. The total labor cost of initial separation of an active deployment was approximately two hours, after which the satellites could be cataloged and transitioned to the routine catalog maintenance phase with minimal subsequent manual intervention. It is unclear if these successes in SDA would be true if the active deployment model were accelerated to the launch rates and orbit-raising timelines seen in the passive deployment model. It is expected that the gains in initial custody and payload separation would allow for scalability while maintaining full custody with minimal manual intervention on the part of the 18 SPCS analysts.

The philosophies of these two deployment strategies extracted a very different cost on comprehensive SDA. All of the SDA-complicating factors of the passive deployments were mitigated in the active model. Namely, greater

[2] Active and passive deployments signify where additional mechanisms and devices are used to intentionally deploy payloads in a specific sequence beyond separation from the booster and does not imply that either method fails to attain positive control of the satellites at any point.

separation rates between constellation satellites were directly responsible for an increased reaction rate by the 18 SPCS systems and analysts. To meet the challenges of the passive model, the 18 SPCS realigned six analysts and contracted five additional analysts to accommodate the sudden need for this new skill set. While training for some of these analysts began in 2019, the unit effectively staffed these positions during 2020 under a new section titled the Mega-Constellation Analysis (MECA) cell. These analysts were trained on the unique techniques required for the initial separation, catalog maintenance, and administrative functions associated with maintaining SDA. Further, processes were built in the earliest iterations of passive initial separation to integrate ephemeris data from the satellite operators. This rapid innovation was managed without dedicated training systems or programs. The only effective training methods for this timeframe were to first train members on the catalog maintenance functions to acquaint themselves with the toolsets and philosophies of managing mega-constellations and orbits maintained by electronic propulsion. Once familiar with these functions, analysts would work with senior analysts on actual launch separations to learn the techniques first-hand in real time. This approach further strained the responsiveness of initial separation as each senior analyst spent most launches also teaching and observing junior analysts. By the end of 2020, the number of analysts trained in initial separation had increased from three to eight, posturing the squadron well for further increases in mega-constellation activity. The success of the 18 SPCS in managing the mega-constellation launches of 2020 is directly tied to this success in self-training and mentorship as well as leadership support in terms of personnel reallocation.

While payload separation in the initial hours after insertion was insufficient for clear delineation by the SSN or 18 SPCS analysts, satellite operators were able to communicate with their payloads and acquire on-board telemetry within the first hours of deployment. Agreements were established between the 18 SPCS and both SpaceX and OneWeb to provide this ephemeris data on a recurring basis. Based first on the predicted insertion data and then via on-board telemetry, satellite operators provided predicted ephemerides to the 18 SPCS either in 36-hour predictions delivered every 8 hours or 7-day predictions delivered every 24 hours. The ephemeris points were in turn used by 18 SPCS to create two-line element sets (TLEs) for tracking and observation association, alleviating many of the issues noted in the launch processing section above. In the initial separation phase, these TLEs aligned observations with payload owners' object identification, negating the need of 18 SPCS analyst efforts to dedicate resources to identifying the names associated with each object. In the active deployment model, these TLEs allowed for SSN sensors to directly correlate observations to the correct object within the first day of launch. While the passive model yielded less autonomous results in observation association, the impacts were still sufficient to reduce strain on tracking sites and quickly identify which tracked objects might represent non-responsive payloads. Upon receipt of ephemerides at 18 SPCS, automated processes were created to perform a general perturbations differential correction[3] on a 4-hour fit of data points centered at the current time to produce TLEs. This process occurred every two hours to produce constantly-updated element sets accounting for the electric propulsion variability which can be included in the ephemeris predictions but is not accounted for in the HAC or SGP4 prediction models available to the 18 SPCS and SSN sensors.

The satellite owner/operator (O/O) ephemeris was integrated in major ways to the 18 SPCS Astrodynamics Support Workstation (ASW) software[4] during 2020 to take advantage of the shared data. This occurs primarily in the ASW Multiple Piece Separation (MPS) software which was upgraded to allow analysts to visualize the operator-provided ephemerides alongside the SSN observation residual data. Further, the ephemeris data when provided with covariance can now be used to supplement SSN observations in updating orbital states, though this feature remains unutilized pending analyst availability for configuration. By viewing the ephemeris data alongside SSN residuals, 18 SPCS analysts are better able to overcome the SSN partial tracks in the early phases of separation and identify which payloads were not observed. The analysts are also able to determine which satellites may have begun orbit-raising which would then cause additional cross-tagging as expected and observed orbits diverged. In total, this integration of O/O data greatly reduced the analysts' workload to achieve initial object separation by at least 50% with more opportunities for improvement in the future. Aside from assisting the tracking and correlation processes, this functionality also allowed 18 SPCS to notify O/Os when their predicted satellite positions failed to align with SSN observations, generally indicating a missed maneuver. It also helped 18 SPCS provide TLEs for observations not aligned to ephemeris data for antenna pointing to attempt acquisition of non-responsive payloads by quickly separating those observations which did not align to any user-reported ephemerides. The operational integration of

[3] In line with the extrapolated general perturbations (eGP) model, this process uses the USSF SGP4 standard model.
[4] The ASW contains the special perturbations algorithms and supporting software used to maintain the HAC via observation association, observation assignment, and differential corrections.

O/O data was a major step forward in 18 SPCS operational capability, which benefits the SDA mission and satellite operators alike. Prior applications of this data were restricted primarily to conjunction assessment processes, but now are integrated routinely into the SDA architecture. With these modifications and acquired skillsets, the opportunities for integrating O/O data into the catalog maintenance functions of the 18 SPCS are just beginning to be realized. This data remains supplemental for catalog maintenance, though, as it represents un-validated[5] solutions on the predicted future positions of active payloads. Thruster performance and automated maneuver management remain prohibitive to prediction accuracy, but multiple phenomenology concepts of future position can now be assessed for opportunities for full catalog integration of this operator data.

5. CATALOG MAINTENANCE

Following the initial separation phase, mega-constellation satellites must be consistently updated to remain current in the HAC, a process referred to as catalog maintenance. While the initial separation phase uses many of the same techniques as the catalog maintenance phase, the catalog maintenance phase focuses less on single-launch issues such as cross-tagging and more on large-scale management of the mega-constellation satellites. As satellites enter catalog maintenance, their separation is generally large and cross-tagging is generally minimal. These generalizations are not rules, however, and some 2020 launches maintained close separation and cross-tagging well into this maintenance phase. Criteria for entering this phase of orbital maintenance is defined by analyst workload where initial separation is managed by analysts focusing solely on a single launch event, the maintenance phase of all satellites is managed by a small team of analysts focused on changes to states rather than state definition. The lifecycle of a satellite in maintenance consists of an initial orbit, orbit raising, final orbit, and orbit lowering. Each phase consists of a varying amount of maneuvers utilizing a non-constant electronic propulsion thruster. Current software was built around the concept that maneuvers would be modeled as either instantaneous or continuous. The longer burn timelines of electric propulsion violated both of these models, driving 18 SPCS to develop a new approach. Within the ASW software, the continuous thrust parameter is a solved parameter corrected in each update through an iterative approach. The software allows segmentation of the solved thrust value over time to allow for varying values during each segments of the observation span rather than a single continuous value across the entire orbit determination interval. The original intent of this was to allow for non-uniform performance or, more likely, inaccuracies of the drag model which resulted in perceived fluctuations in the thrust value. For traditional LEO payloads, this segmentation is applied to the B term to account for the variations in drag over time, but applying the segmentation technique to more than one variable has created performance and accuracy issues. As a result, orbit determination procedures were tested for segmenting each of the drag parameter and the thrust parameter before setting into this configuration.

For the mega-constellation satellites, the segmentation method was applied to the thrust parameter to enable the thruster cycles and more accurately model the environment of the satellites. In coordination with the satellite operators, it was determined that the thrust could be represented in 30-minute windows where the thrusters were discretely on or off. This window drove a segment length of 15 minutes to allow for two to three segments of thrust application with a variance set to the average output of the satellites' thrusters. While the thruster output was based on information from the satellite operators, it also aligned with independent assessments of the ASW-determined thrust values proving a use case for defining thrust values in the future. Traditional approaches of drag segmentation or continuous thrust anticipate a relatively steady value with moderate levels of change, but the mega-constellation system assumes on/off states where the value of the thrust variable in one segment may differ greatly from its neighboring segments. In the first half of 2020, segmentation was applied in this way but met with middling results. The prediction model failed to consider past performance, so predictions were always of a non-maneuvering state and resulted in poor prediction accuracy. Further, the segmentation model was built to modify a variable over time and not represent binary states, so this segmentation method resulted in many ill-represented thrust segments where non-thrust perturbations were modeled as reduced-power thrust segments. Due to these two issues, SSN sensors were unable to effectively reacquire satellites against TLEs propagated more than 36 hours as the predicted and actual states quickly diverged. In August of 2020, the ASW software was upgraded to apply a feature termed Adaptive Thrust Uncertainty (ATU). This concept is addressed separately and will not be examined in depth in this paper[6]. The impact of ATU, though, was clearly noticed in terms of HAC and TLE performance for mega-

[5] User ephemeris data has not been analyzed by the USSF in regards of quality of user data nor the impact of integrating that data into the HAC

[6] Casali, S. 24 Jul 2020. "Orbit Determination Updates for Starlink Satellites". US Space Force/MOSSAIC

constellation satellite predictions. As a generalization, the ATU assesses thrust state changes to determine whether to apply all or none of the thrust in a given segment, thus creating a binary setting rather than the graduated approach. The impact on TLE prediction accuracy is dramatic and has reduced errors at the 36-hour mark by more than an order of magnitude. The research enabling ATU, though, was focused on the steady-state orbits of a mega-constellation satellite such as the initial and final orbits. Prediction accuracies were improved to a lesser extent during orbit-raising and lowering periods. This difference is likely due to the decision to not propagate predicted maneuvers in the space catalog combined with the frequency and duration of maneuvers during these phases of orbital change.

Aside from segmentation, other maintenance parameters were considered and adjusted to allow for automatic differential correction updates of mega-constellation objects. These changes were made to reduce the amount of manual interventions in the ever-growing space catalog while maintaining a sufficient level of accuracy for all SDA functions. The primary change in these settings were considerations for the weighted root mean square (WRMS) calculated for a given satellite update. This measurement is a ratio of observational residuals to the standard deviation for that sensor's ability to collect data and represents a general quality of fit for a state update. The larger WRMS allowance and ATU application yield a larger and less representative covariance field for the mega-constellation satellites, but nonetheless offer a large improvement on previous modelling attempts.

6. TASKING

The initial tasking of a satellite is based on an initial launch message sent by 18 SPCS to the sensor sites which contains the tasking instructions as well as pointing information for sensor sites. Due to the low prediction quality and small inter-payload separation of passive deployments, sites bypass the association of all deployed objects and instead force the assigned satellite to a satellite number for the initial phase of launch tracking. This approach prevents the sensors from providing additional tracking on uncorrelated objects while consolidating data for 18 SPCS analysts in the initial separation phases. Though this approach simplifies the tasking and sensor operations, it adds to the manual requirements of observation assignment in the initial separation phase. Once operator ephemeris data is first received for these passive deployments, tracking data for individual satellites are distributed to the SSN sensors and observations are correlated to the nearest match. Sensors also provide increased tracking for the newly launched objects for the first week after launch to allow time to integrate the satellites into the Special Perturbations Tasker (SPTasker) software responsible for assigning sensor tracking assignments. The initial process was not used for active deployments in 2020 due to the large amount of time between orbital insertion and payload deployments. Instead, the active deployment launches deferred to traditional launch and tasking management techniques.

Within SPTasker, mega-constellation satellites are managed through a separate tasking category to ensure that data is collected routinely but not at high density rates. During 2020, the standard practice of tasking newly launched objects at a high observational density and revisit rate yielded inefficient requests on the limited resources of the SSN. While no direct impacts were measured, anecdotal evidence implies that the increased request rate was pushing lower-priority objects out of daily tracking cycles. Due to the self-reported data within the user ephemerides, tracking density could be reduced for mega-constellation satellites without noticeable impacts to SDA. Daily tracking is still assigned, though, to validate the orbits in near constant change and allow for independent verification of orbits. This approach is precautionary and intended to ensure rapid escalation of a purely internal tracking and orbit updating process in the event of a satellite anomaly. Further research is being done to determine the tracking density required to maintain sufficiently accurate states for an active satellite which did not provide predicted ephemeris to the 18 SPCS. It is also unknown how little tracking would suffice for objects when greater reliance on O/O ephemeris data is incorporated into orbital update procedures. These questions are difficult to answer as only a single tracking acquisition can be applied to real-world operations with only simulated models capable of showing theoretical performance for the non-used methods. These models are not currently available and would require a large amount of resources to develop with minimal operational benefit. Due to the lack of alternative tasking potential models, the balance of timely orbit updates and SSN resource utilization is likely to be reevaluated with each new constellation. Additionally, non-SSN tracking data could be used to supplement the space catalog and avoid the constraint of network saturation. These options are possible in the immediate term, but not actively employed for the sake of mega-constellations today.

7. CONJUNCTION ASSESSMENT

As part of the current on-orbit conjunction assessment (CA) process, 18 SPCS is generating over 200,000 Conjunction Data Messages (CDMs) per day which is an increase from 2019's starting average of 3,000 per day. Of these daily messages, 180,000 are based on Starlink satellites. A large majority of these messages identify close approaches between two Starlink satellites. The inter-constellation close approaches are a result of the satellite operator proactively sending predictive ephemeris with planned maneuvers to 18 SPCS for conjunction assessment screening. 18 SPCS screens these predictive ephemeris files against the existing HAC states, which do not include predicted maneuvers. The high rate of movement within the constellation creates a litany of close approach reports. Since SpaceX has exquisite knowledge of their constellation, a majority of the messages that are sent to SpaceX provide no meaningful data or further insight, but the reports are provided in the event that maneuvers are not executed. The satellite operator also benefits from a second-party identification of the state for their satellite which can be used for operational tuning of their telemetry and planning systems.

The increase of active satellites also has dramatic impacts on the data throughput and computer network requirements of the 18 SPCS. The number of ephemeris files submitted to the 18 SPCS per day from any satellite owner/operator (O/O) has increased from 500 to over 6,000 with many of them coming from SpaceX. Beyond quantity of files transmitted, there are data implications on the systems that store and utilize this ephemeris information which must constantly scale with demand. In 2020, 18 SPCS generated 31.5 million CDMs; assuming a linear growth, 18 SPCS is estimated to produce 70 million CDMs in 2021. Again looking to data management, these reports are government records and as such must be retained. All of this data must also flow to and through 18 SPCS' public-facing website, Space-Track.org. The growth in data creates a strain on all information technology systems, driving requirements to upgrade technology sooner than expected.

By the end of 2020, SpaceX was responsible for over 54% of the conjunction data that was output by 18 SPCS, with over 74% being dedicated to O/O-provided ephemeris processing as a whole[7]. During 2020, 18 SPCS CA operations increased 470% and the data output increased 1400%. Given this sudden surge in data throughput requirements, 18 SPCS rapidly redefined how to conduct CA operations. Through an effort named Reduction of Conjunction Assessment Processing (ReCAP), 18 SPCS incorporated a 4-year data analytics study by the NASA Conjunction Assessment and Risk Analysis (CARA) program to better understand the current space domain and how operations could be made more sustainable, scalable, and ultimately provide more meaningful data to global satellite operators. The hypothesis of the NASA/CARA study was that 18 SPCS screening volumes may be too large, resulting in the dissemination of conjunction predictions which represented no credible threat to the satellites. NASA/CARA found in this study that a reduction in the radial component of 18 SPCS screening volumes could provide significant relief in terms of data output while still capturing all conjunctions that were of a significant probability of collision. The screening volume adjustments reduced unnecessary data generation and dissemination, but the growth in processing remained an issue. Under project ReCAP, CA operations were retrofitted with an automated application. This change converted manually-intensive functions requiring two 18 SPCS analysts at all times into a more hands-off series of functions, which resulted in manning reductions while expanding the efficiency of remaining analysts. The manning reduction allowed 18 SPCS to relocate analysts to the MECA Team to focus on mega-constellation initial separations and maintenance.

As the SpaceX constellation has grown, there has been continuous dialogue between 18 SPCS and SpaceX representatives in order to build a common understanding of operational practices. This dialogue has created load-sharing opportunities between the two organizations while also allowing each to optimize their operations. Currently each SpaceX satellite generates approximately 50 conjunction data messages which scales exponentially as the constellation grows due to intra-constellation conjunctions. By the end of 2020, discussion had begun to enable internal intra-constellation CA screenings by SpaceX with HAC data support from the 18 SPCS, ensuring that core services continue to be provided to the satellite operator while shifting the burden of constellation management to the constellation owner. OneWeb's satellite fleet is likely to draw the same level of resources as their satellites fill the intended orbital shells. The initial data and mission-sharing lessons learned from Starlink constellation management and CA mission sharing will likely apply to the OneWeb constellation growth plan.

[7] SpaceX was the first organization to deploy a mega-constellation fleet. As such, these numbers do not represent SpaceX's specific impact on SDA capabilities, but rather inform on what will happen as more satellites are launched given no change in the CA architecture and procedures.

Future companies are also likely impacted by these decisions and reactions. Amazon's projected Kuiper constellation is anticipated to introduce a daily average of 500,000 CDMs concerning their constellation alone. With the multitude of mega-constellations, further reductions of data generation, transmission, and storage are needed to continue to balance the need for safety of spaceflight with operational responsiveness and resource management.

8. RE-ENTRY PROCESSING

The 18 SPCS is charged with predicting and announcing the re-entry of space objects based on pre-defined thresholds. The procedures are generally started one week before final decay of the space object and include regular reporting requirements containing updated predictions of final entry points and times. In recent years, these thresholds were modified to allow for non-reporting of CubeSat re-entries as a result of SmallSat proliferation and the impacts of constant re-entry assessment (RA) processing and subsequent reporting. This exception is not extended to the mega-constellation satellites, some of which have already begun their final descent. With an operational altitude of 1200 kilometers, the OneWeb constellation growth has not shown an impact in 2020 on re-entry processing. While these satellites will surely reach end of life in the future, their mass re-entry is expected to emulate the challenges posed by SpaceX's Starlinks but is not expected to do so in the near future. SpaceX, on the other hand, has decided to de-orbit their initial round of satellites, which has provided insight into the impacts of increased satellite populations in LEO.

Starlink satellites are expected have an approximate five-year operational life span. Thereafter, SpaceX plans to safely deorbit them to a low-altitude disposal orbit for natural decay into the atmosphere. The controlled descent phase is intended to reduce impact on the transited orbits and enable collision avoidance if necessary, though the drag design of the constellation permits an estimated one-year decay cycle for uncontrolled decay scenarios. Since the Starlink satellites will decay by natural perturbations, 18 SPCS is required to perform RA procedures at regular intervals in the final week of a satellite's life which will significantly increase the 18 SPCS reentry workload. While the initial Starlink satellites have demonstrated the pattern for decay, the majority of re-entries from this constellation remain years in the future. Starlink re-entries will presumably occur at approximately the same cadence as their launches with the first major round of decays beginning roughly in 2025. Thus far, the cadence of launch has been approximately 60 satellites every 2 weeks which would imply a potential 1500 Starlink reentries per year if de-orbited in bulk. This would represent an average of 30 Starlink satellites reentering every week in contrast with the current RA workload of just over 100 RAs each year. This potential 15-fold increase is a significant increase in work for the 18 SPCS and will surely drive systems and procedural changes in the coming years. Further complicating this growth projection, is the possibility that SpaceX may increase their ability to launch up to 300 satellites at a time. Even if this does not come to fruition, the chances of unanticipated RAs increase with each mega-constellation launch. While no operator intends to launch non-responsive payloads, the increase in occurrence is likely to come with the launch cadence. This only adds to the complicating factors of additional satellite decays and commercial space traffic which shows no signs of slowing. Though the impacts of mega-constellations have already been felt in the other phases of satellite life cycle management, re-entry assessment has only begun to realize the full potential of the proliferation of satellites in low-Earth orbit.

9. CONCLUSIONS AND FUTURE WORK

In response to rapid industry growth in the low-Earth orbit satellite field, the 18 SPCS has made major progress in overcoming the challenges of providing SDA. For decades, many aspects of the current SDA architecture were designed around low-volume satellite deployments. Despite this, no major disruptions have occurred in the SDA enterprise. Rather, industry cooperation and internal innovation have enabled rapid adoption of new techniques to maintain continuous, comprehensive and combat-relevant SDA. This work is not complete, however, and leaves room for additional improvement. Starting from launch, O/O ephemeris can be further integrated into the tracking and piece identification systems. This data is readily integrated to the 18 SPCS systems, but further integrating it into sensor sites to gain more efficiency would require redesign of fundamental data transfer systems and formats. These functions will also need to be performed without reliance on external data, so the integration would need to allow for automatic assessment of available data to determine whether tracking is even required. To rely upon the ephemeris data, systems will also need to account for data validity to ensure a true accounting of satellites launched.

The O/O ephemeris data can also be better integrated into the conjunction assessment process to enable risk assessment for the large constellations rather than instantaneous close approaches. By integrating the maneuver-inclusive O/O ephemeris data and the perturbation-inclusive HAC predictions, alternate future positions can be contrasted in terms of likelihood. This merging of conflicting data could define a continuum of possibilities whereas today it is interpreted simply as two competing realities. These upgrades will likely need to take place outside of the 18 SPCS, though, as the unit resolves to provide state-based data and not risk assessment products. Finally, user ephemeris could be further integrated into catalog maintenance to reduce the need for SSN tracking. This final step would remove the ability to rapidly initiate a HAC-modeled state but would reduce the systems burden of tracking, identifying, and predicting these mega-constellation orbits.

Looking to improvements internal to 18 SPCS, there is a need to better refine the timespans and data densities used in converting O/O ephemerides into element sets. The current four-hour span with two-hour disseminations is likely not optimized and exacts a processing toll on the 18 SPCS systems. In response to this toll, the TLE generation is only enabled for satellites which are not sufficiently served by the HAC's ATU function. Further streamlining the process of TLE generation based on O/O data would create more opportunity to provide greater numbers of O/O ephemerides to the SSN for tracking without major network upgrades. There is a need for refined guidance on which data sources are allowed to drive the SSN tracking, but the ability to offload data processing surely contains potential gains for SDA capabilities.

Similar to the TLE generation, the internal model and maintenance parameters remain un-tuned for nearly all mega-constellation satellites. Due to the prevalence of their satellites, SpaceX has served as the model for many settings such as drag, thrust, and ATU parameters. Similar settings must be defined for each payload and further defined for times of stability or change within the satellite's orbit. Differences have been noted in the maintenance parameters of orbit-raising and on-orbit Starlink satellites, and it stands to reason that the same will be true of other constellations. Alternate orbits can also present different perturbations which must be assessed against the segmented thrust model defined in 2020. OneWeb satellites, for example, are likely subject to radiation pressure from the Sun. Because of the singular mega-constellation model currently in place, these perturbations go un-modeled and likely result in increased state and prediction errors.

Unfortunately, these future studies and innovations will require additional time, resources, and analysts to execute. This will be required, of course, at a time when mega-constellations continue to launch more satellites and increase the workload for the same analysts required to solve the current problems. There are many possible ways to alleviate the resource requirements of this work to include further industry collaboration, increased workforce, and improved processing systems and software. No one of these changes alone will provide the relief needed to both maintain the expanding operations while adapting for future problems, so it is likely a combination of all is needed to continue the success of the SDA mission.