CLOUD SERVICES FOR SPACE SITUATIONAL AWARENESS

Sarah Law, Jared Dorny, Jared Stallings, Jabari Loving
Raytheon

ABSTRACT

On Feb. 14, 2017 the Indian Space Research Organization launched 104 satellites with a single rocket. This single launch illustrates the increasingly difficult problem of maintaining space situational awareness. Space based constellations are becoming much larger, with some in current design planned to be in the thousands. In addition, the number of countries and organizations that are participating in space is growing at an unprecedented rate, creating difficulties with maintaining the satellite ancillary data including intended use and ownership.

To manage this difficult problem, Raytheon has developed a scalable distributed data repository for entity state information along with cloud services to generate orbit events. These orbit events including ascending node, descending node, perigee, apogee, umbra, and penumbra are critical for mission and contact planning, as well as situational awareness. In legacy systems the orbit event data is based on position data which is frequently single-sourced, generated, and stored as the ephemeris is generated. This data is then presented to the user for all satellites in the system, which is generally only a few satellites. However, in this new system, use of cloud scalability patterns, together with dynamic features of micro services provides for a highly scalable and distributable orbit event service which can handle datasets for thousands of vehicles in near real time and allows for distribution to many users.

Beyond just ephemeris, the Entity State Information Service will store information gathered from a variety of sources, together with its relationships to other entities. Examples of source data include Space-track.org, AGI COMSPOC, and Gunter's Space Page. Not surprisingly, data from these sources may have varied accuracy and are in a variety of formats. As such this repository provides linkages to other data stores, or to the source of such data. Changes in the data, as well as the time when they occurred, are also tracked.

1. BACKGROUND

Traditional satellite constellation management has consisted of a few vehicles per constellation, with larger constellations ranging up to the low dozens. This paradigm is rapidly changing due to the increasing accessibility of space, as cost of launch goes down, and smaller satellites offer real mission value. For example, BlackSky had 10 operational satellites as of Dec. 2016 and will have a 60-satellite constellation when they are fully operational [3]. This is only the beginning and much larger constellations have been proposed. OneWeb plans to build a constellation for broadband, which has an initial operating set of 900 satellites. Airbus has already begun building some of the initial satellites [4].

Many existing satellite command and control (C2) systems are built around assumptions in the size of the constellation. In the timeframe that these C2 systems were designed, the benefits of limiting the scalability outweighed the limitations. While the number of vehicles is going up, the number of operators is often steady or even decreasing, meaning that the system must be highly automated, and situational awareness must be presented to the operator in a highly efficient manner.

The Space Situational Awareness (SSA) mission, though different from the C2 mission, was conceived under similar pretenses. When SSA first began, in the 1950s, there were very few objects to track in space. In the 1970s there were about 1800 objects being tracked by the Space Surveillance Network (SSN) and now there are over 22,000. While the numbers of items being tracked have greatly increased, new techniques are needed to fully integrate the situational awareness about these objects into operations. Figure 1 shows the number of objects tracked by the SSN over time, through 2015, including debris and rocket bodies. These numbers are expected to continue to increase substantially. Figure 2 shows a projection from SpaceWorks on the number of nano/microsatellites that will be launched through 2022 – up to 3,000 [5]. The likelihood of conjunctions will continue to rise as more and more objects are put into orbit, potentially creating even more objects to track.

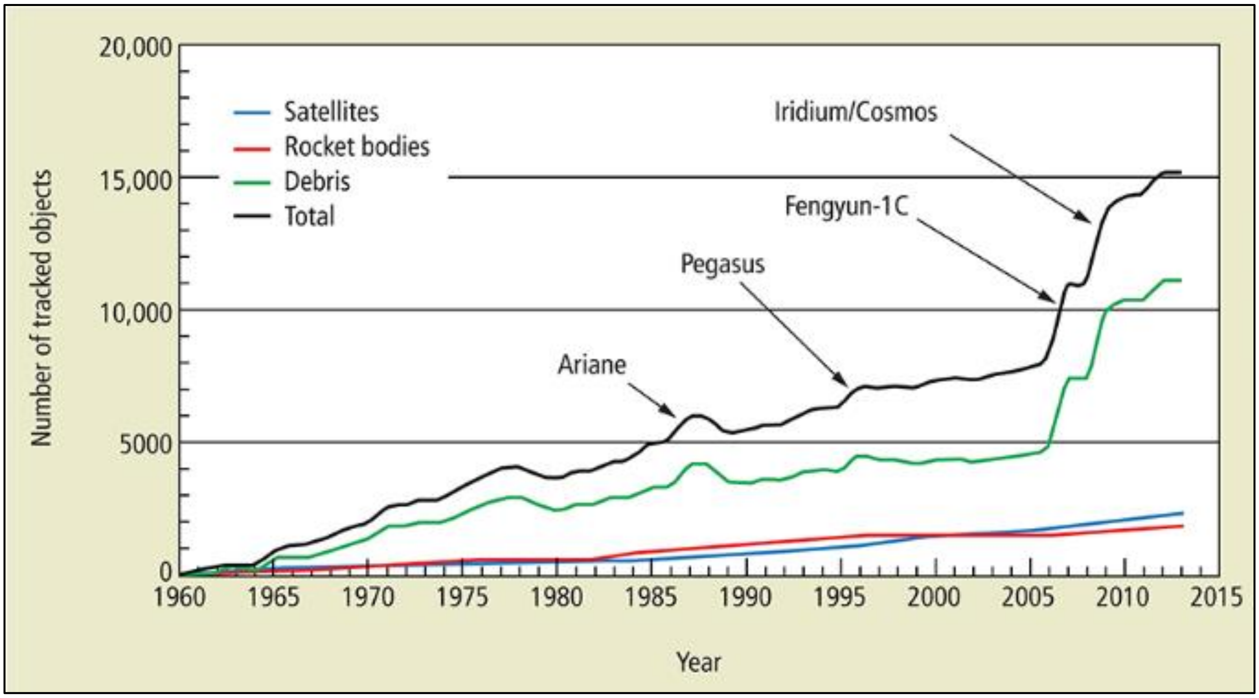

Figure 1. The number of objects being tracked by the SSN is continuing to increase, both due to a greater number of launches, as well as increasing sensor sensitivity allowing smaller objects to be tracked. In addition, there have been several conjunction events that have created large amounts of debris that must also be tracked. [6]

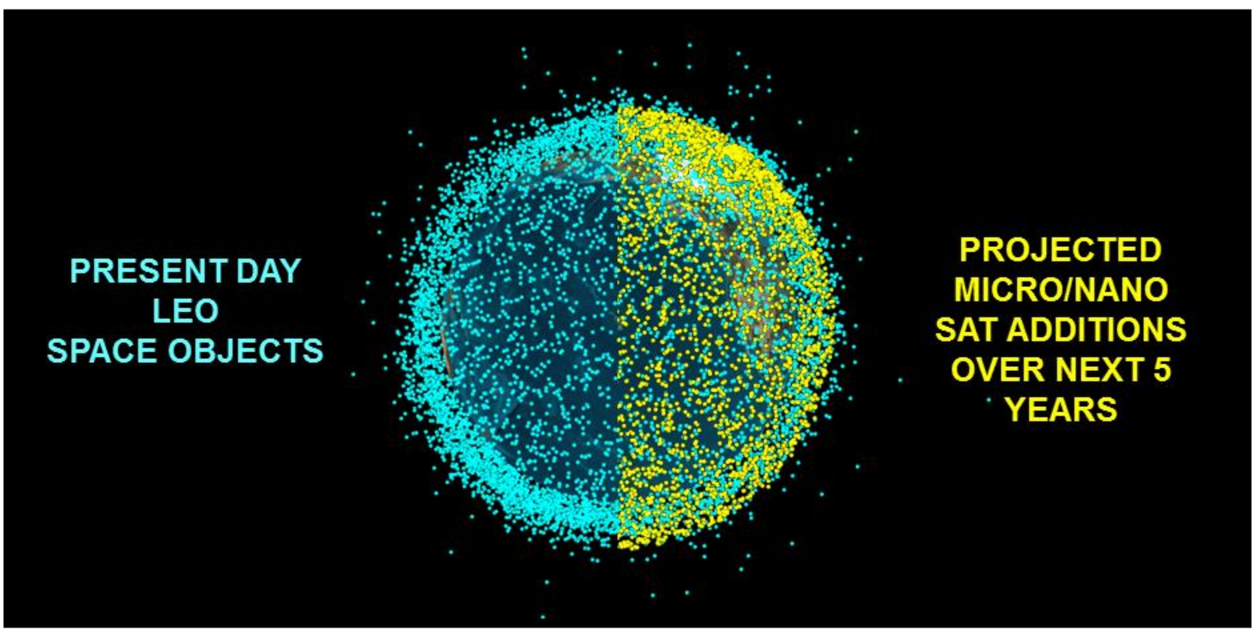

Figure 2. Nano-, micro-, and small satellites are also increasingly being launched. A single launch may now release hundreds of objects into space. [5]

This rapid increase in the number of space objects creates constellation management concerns including conjunction calculations, determining umbra and penumbra for certain mission operations, and quickly assessing mission planning and scheduling, among others.

In addition to tracking space objects, the SSN is also responsible for characterizing the aforementioned 22,000 objects. This characterization aspects adds even more complexity and "Big Data" challenges for future SSA systems. In this paper we discuss the benefits for transitioning to a cloud infrastructure to manage, maintain control, and provide situational awareness for both mission constellations and the Space Object Catalog. We will also introduce services which allow the users to securely leverage the cloud to meet their mission and, if necessary, easily distribute orbit event information to many users.

2. THE SPACE SURVEILLANCE NETWORK

Historically the Space Situational Awareness mission has been performed by specialized networks that are designed to observe and track the locations of space related objects. For example, the Space Surveillance Network (SSN) falls under the responsibility of the Joint Functional Component Command for Space (JFCC Space) which is one of the components of United States Strategic Command (USSTRATCOM). The SSN is responsible for tracking and cataloging man-made objects in space. The SSN consists of a network of 30 sensors that are tasked by the Joint Space Operations Center (JSpOC) to observe space objects [9]. The JSpOC uses the SSN to collect approximately 400,000 observations each day and is responsible for using these collects to provide services for the space community. The measurements are used to create orbital element sets which can be used to represent the position of the satellite. With this information the JSpOC compares current satellite position to predicted positions and updates predictions for future satellite locations. This information is used to update the space object catalog. The catalog can then be leveraged by Department of Defense (DoD) and civilian agencies to prevent collisions and conjunctions or to improve situational awareness about possible threats in space.

Over the last few decades the SSN has been very effective at providing information for SSA missions. However, when we analyze how the network is constructed some limitations become evident. One of the benefits of the SSN is that it is a closed network with fully vetted information sources. This closed network ensures that every node is producing high quality, trust worthy, authoritative data. A tradeoff of this approach is the rigidity of the network. The system cannot quickly evolve to incorporate new sensors or intelligence sources and, when new sources need to be added, the associated costs can be staggering. Also, the system loses out on the value of some of the less vetted information about satellites. There is a vast amount of valuable open source SSA information that is not produced by SSN nodes. If we could augment SSN information with this open source information and apply levels of trust while providing responsive ephemeris and orbital event generation we could greatly improve the process of tracking and categorizing objects in space and providing this information to a wider user base.

3. COMMAND AND CONTROL

C2 systems are focused around providing a mission and maintenance for specific constellations. As stated earlier, many of these systems are limited to being able to control smaller constellations. However, in the last few years we have seen an increasing trend towards leveraging low cost microsatellites which are part of very large constellations. Constellations in the hundreds and even thousands have been proposed [8]. Some of the challenges are similar to those seen by the space surveillance network.

Generally speaking, most C2 systems are primarily concerned only with their own assets and are unaware of what the rest of the space objects are doing. They rely on notification produced by agencies such as the JSpOC to warn them if other space object catalog items may impact their mission. When these systems produce ephemeris, it is typically only for their constellation and it is used to generate schedules or plan for in-constellation orbital events. The ephemeris generation is accomplished using local computing power which is adequate given the size of most systems. As the size of these constellations grow the demands put on using dedicated computers to generate ephemeris may lead to long wait times or large hardware procurement bills. Another aspect to consider is the idea that as space becomes more congested C2 operations may want to generate orbital information about objects that are outside of their own constellation. For example; if your C2 system is managing a constellation of 24 satellites in low earth orbit there may be some value in gaining first hand understanding of how other constellations operating in LEO could possibly impact your constellation.

4. CHALLENGES

The following challenges are presented by the growing complexity of SSA.

- **Data storage and access** – the types and kinds of data and information needed to effectively identify, track and characterize entities in space or related to space come from various sources and services: some information is static, some is calculated, some is temporal, some is well-known, and some is predicted. Keeping track of this information and providing common services to store, manage and make it available to a suite of disparate and decentralized users (services, humans, portals, analytics) supports the ability to use the data for effective SSA.

- **Exponential growth in computations** – as the number of space objects grows, the number of computationally intensive processes grows exponentially. A simple example is conjunction analysis. As a new object is added to the catalog, that object's orbit must now be compared to every other object in a similar orbit to search for potential conjunctions.

Raytheon developed an architecture that tackles these key problems, called Ptolemy. Ptolemy was designed to meet the scalability and flexibility needs of the increasingly challenging problems presented by SSA. At its foundation, Ptolemy leverages open-source capabilities in use in the commercial industry and applies them to the challenges associated with the management of SSA data.

5. DATA STORAGE AND ACCESS

In today's data-driven world, several data repository technologies have emerged to solve data store and access problems for large enterprise problems. We have investigated several of these technologies, including Apache Cassandra, MongoDB, Neo4J and MariaDB. These data repositories are designed to provide scalable data storage and retrieval across a distributed set of cloud-based servers. Each provides specific flexibility and scalability for certain types of data management use cases. We then selected the technology that best maps to the type of data management issues SSA needs. Secondly, we created a data ontology that allows us to store the disparate types and kinds of information available in a common, normalized way. Finally, we developed data adapters that allow us to ingest data from the main external data sources to populate and maintain the repository.

To simplify access to this data, we have create RESTful web service endpoints to allow access to the underlying data from any client type (i.e. scripts, web applications, micro services), or users who need more direct access can use the API provided by the underlying technology.

To validate use of these technologies, we have ingested the entire satellite catalog open to the public, which includes 40K satellite catalog entries and over 110M TLE records for the past 10 years. We have also ingested satellite entity information from other open source data, including ITU, WMO, NSSDC, UNOOSA, SatsList, and other open source data sources. We have merged this data into single "Entity" records that represent the known knowledge about each entity.

6. EXPONENTIAL GROWTH IN COMPUTATIONS

With a core suite of scalable data management technologies, we are able to tackle the problem of the growth of computationally intensive operations. The foundation of the solution lies in the adoption of software container technology. Docker is becoming the de-facto software container platform and was selected for this problem. Combined with Rancher, a container management platform, we have developed a suite of container-ized software applications that can perform specific computationally intensive operations for SSA.

Then, Raytheon developed an application, called Brigata, which responds to demands for computationally expensive types of operations to be performed and dynamically configures and spins up a software application through Rancher and Docker to meet the need.

Three examples of when the Brigata, Rancher and Docker capabilities are leveraged are provided:

- **Ephemeris Generation** – where precise orbit state prediction information is necessary, orbit states must be propagated. As orbit state updates are received, objects that require precise orbit prediction data can have an ephemeris generation job sent to Brigata. Brigata dynamically spins up Docker containers that can perform the ephemeris generation job and store if off in the underlying data repository. The ephemeris generation is now easily available for any user to query.
- **Conjunction Analysis** – as state updates about space objects are received, the set of conjunction analysis runs that are required are defined and handed off to Brigata. Brigata then manages the dynamic creation of a software container to run a conjunction analysis for each job specified. The results are stored in the underlying scalable data repository, notifications are generated, and the job completes.
- **Visibility Analysis** – Several SSA questions can be answered by performing visibility analysis. For example, the question might be asked, when is space object A visible by space object B. Or, when is ground object C visible to space object D. Performing visibility analysis on demand can slow down the process asking the question. As orbit states are updated, potential visibility questions that may be asked can be inferred, and a job request for each sent to Brigata. Brigata will spin up a job to pre-generate and store the visibilities in the

underlying data repository. Then as the processes that use visibility data execute, they can quickly query the cached store of visibilities rather than preform the computations necessary.

7. CONCLUSIONS

Given the increasing complexity of SSA and C2 missions in today's changing space environment, Raytheon has developed a set of tools that can manage the difficulties of data storage, data access, and scalability. Leveraging open sources tools allowed us to create easily scalable solutions that provide rich information sets about the increasingly large space catalog.

8. REFERENCES

1. Ward, A., *There may soon be a new US military service — for space*, Vox, https://www.vox.com/world/2017/7/5/15905018/space-corps-air-force-rogers-house-senate-armed-services, Jul 5, 2017.
2. Harbottle, A., *Earth observation to capture the mainstream market*, ViaSatellite, http://interactive.satellitetoday.com/via/may-june-2017/earth-observation-to-capture-the-mainstream-market/, May 2017.
3. *Spaceflight Industries unveils its BlackSky platform for observing the planet*, BlackSky, https://www.blacksky.com/2016/12/14/spaceflight-industries-unveils-its-blacksky-platform-for-observing-the-planet/, Dec 14, 2016.
4. Foust, J., *Airbus start assembly of first satellites for OneWeb's broadband constellation*, http://spacenews.com/airbus-starts-assembly-of-first-satellites-for-onewebs-broadband-constellation/, Jun 28, 2017.
5. *PSLV-C37 Successfully Launches 104 Satellites in a Single Flight*, Department of Space Indian Research Organisation, http://www.isro.gov.in/update/15-feb-2017/pslv-c37-successfully-launches-104-satellites-single-flight
6. Sorge, M.E. and Peterson, G.E., *How to Clean Space: Disposal and Active Debris Removal*, Aerospace, http://www.aerospace.org/crosslinkmag/fall-2015/how-to-clean-space-disposal-and-active-debris-removal/, Dec 10, 2015.
7. *The satellite market is changing… and so are we*, AAC Microtec, http://aacmicrotec.com/spacecraft-bus-and-subsystems/, Aug 6, 2017.
8. Mosher, D., *SpaceX just asked permission to launch 4,425 satellites — more than orbit Earth today*, Business Insider, http://www.businessinsider.com/spacex-internet-satellite-constellation-2016-11, Nov 16, 2016.
9. *USSTRATCOM Space Control and Space Surveillance*, U.S. Strategic Command, http://www.stratcom.mil/Media/Factsheets/Factsheet-View/Article/976414/usstratcom-space-control-and-space-surveillance/, Oct 17, 2017.