69th International Astronautical Congress (IAC), Bremen, Germany, 1-5 October 2018.

IAC-18-A6.7.8

The Need for Comparative SSA

T.S. Kelso[a]*, Daniel L. Oltrogge[b]

[a] *Center for Space Standards and Innovation (CSSI), United States*, ts.kelso@centerforspace.com
[b] *Center for Space Standards and Innovation (CSSI), United States*, oltrogge@centerforspace.com
* Corresponding Author

Abstract

In Space Situational Awareness (SSA), an oft-advocated assertion by some is that "you never want to wear two watches, because you'll never know what time it is." The equivalent of this analogy in the SSA arena is to state that "there will only be one authoritative Resident Space Object (RSO) catalog" or "there is only one official SSA system."

However, drawing upon decades of space operations, we have found quite the opposite to be true: If you only have a single watch, then you don't know what time it is because you have no insights that your system may be systematically wrong or be imprecise.

In this presentation, we cite examples where comparative assessments between multiple SSA and CA systems have led to significant discoveries of systemic biases, process deficiencies, undersampling, ineffective fault tolerance and recovery mechanisms and characterizations of system errors by type.

The very existence of duplicative, parallel systems permits such comparative assessments to be done. And while our focus has been on comparing SSA results from separate SSA systems, we also advocate for such comparative assessments to be accomplished within the context of single SSA systems as well: An ideal SSA system continually and tirelessly seeks to assess its performance and accuracy. Comparative SSA, orbit and CA assessments using a variety of observation data sets, analysis methods, force models, data fusion approaches and positionally well-known "truth" depictions is fundamental to ensuring SSA system accuracy and precision, thereby minimizing Type 1 and 2 SSA errors, such that decision-quality SSA products can be generated.
Keywords: (Space Situational Awareness, Safety of Flight, Conjunction Assessment)

1. Introduction

Segal's Law states: "A man with a watch knows what time it is. A man with two watches is never sure." In reality, we all check our watches against other sources—whether it is to manually check them against the clock on the wall or automatically via the Network Time Protocol over the Internet against high-precision timekeeping devices.

The US Naval Observatory (USNO) Master Clock is one of those high-precision devices, which compares dozens of independently operating cesium atomic clocks and a dozen hydrogen maser clocks every 100 seconds to ensure reliable and stable time and time interval. The USNO submits its clock measurements to the Bureau International des Poids et Mesures (BIPM) to produce TAI or International Atomic Time. BIPM uses GPS to compare clocks around the world.

Comparing data against *independent* sources of similar data is a time-tested way of assessing and minimizing uncertainty and identifying problematic data sets. A robust network of collaborative data sharing in the SSA realm can similarly help to reduce uncertainty and risk and identify and eliminate problems that might lead to adverse events affecting safety of flight.

2. Background

Many satellite operators today want to develop and maintain their own independent satellite tracking systems to support orbit determination for mission assurance and safety of flight (e.g., conjunction assessment or collision avoidance). These efforts appear to be motivated by a perceived lack of quality in the current data needed for these tasks and a general unwillingness to share the best data available by some organizations.

Many of these operators seem to believe that if they just have good knowledge of their satellites' orbits and maneuvers—together with a 'comprehensive' catalog from some central authority (such as the one currently provided by the United States government)—they will be able to meet their mission requirements, without the need to share data directly with other satellite operators or tracking networks. While some of this reluctance to share may be motivated by national security or commercial competitiveness considerations, the resulting behavior can work counter to the overall goal of safety of flight and preserving the space environment.

The reality is that there are a number of implicit assumptions being made in this approach that do not hold up in real-world operations. As in any operational system, all assumptions should be specifically stated and attempts should be made to validate each assumption. If an assumption cannot be validated, an assessment should be made for how to best identify when an assumption may be violated. One key way to test assumptions of orbital data quality is through operational sharing of this data and regular, automated comparative assessments against independent data sets.

2.1 Space Data Center

The Space Data Center (SDC) was developed under the premise that the operational sharing of the most current orbital data—including planned maneuvers—from all satellite operators, can be used to improve overall SSA when used for conjunction assessment and collision avoidance. While standard data products are encouraged as input, the SDC works with each operator to confirm proper interpretation of data format, units, time systems, and coordinate frames. Comparative analyses are used throughout both this initial interpretation of each operator's data and daily operations to identify any anomalous behavior in the independent data sets from either Space Data Association (SDA) satellite operators—in the form of ephemerides that incorporate future maneuvers—or the Joint Space Operations Center (JSpOC)—both for their two-line element (TLE) sets and special perturbations (SP) ephemerides. Automated machine-to-machine interfaces help to further reduce the chance of injecting errors into the process.

The SDC works closely with 30 satellite operators from around the world to ingest, assess, and analyze orbital data from those operators for over 700 satellites, together with the JSpOC TLE and SP data for roughly 17,000 objects. As part of those operations, the SDC performs comparative assessments throughout, in an attempt to identify inconsistencies which might be an indication of underlying errors in one or more of these data sets. Looking at this data for ~270 GEO and ~440 LEO/MEO satellites provides insights into patterns of behavior that might be impossible to detect for smaller satellite constellations and can help to more quickly focus on whether a problem is limited to a single satellite, an entire constellation, or is the result of systemic problems in the JSpOC data. Feeding those results back to the affected parties allows for quick remediation and subsequent changes in processes and/or procedures to avoid repeat occurrences.

This paper will discuss several case studies where these comparative assessments led to uncovering serious flaws or misunderstandings in how the data was processed somewhere along the way. In some cases, we will not name an affected party, since that is not key to understanding the underlying need for comparative assessment. In no case should the fact that an error or misunderstanding was detected be taken as placing blame, since all systems will fail at some point. Our sole motivation here is to quickly identify and eliminate errors via comparative SSA to improve overall SSA and avoid jeopardizing safety of flight.

3. Case Studies

3.1 Case Study #1

As each new SDA satellite operator is added to the SDC, a process is undertaken to ensure a proper understanding of the data being provided. The data is provided as ephemeris with a timestamp for each Cartesian position and velocity. SDA operators are encouraged to follow best practices for time steps between ephemeris points and precision of their data to avoid problems with interpolation that could adversely affect computations.

While the SDC recommends the use of standard ephemeris interchange formats, such as the Consultative Committee for Space Data Systems (CCSDS) Recommended Standard CCSDS 502.0-B-2, *Orbit Data Messages* [1], this may not always be possible due to limitations resulting from legacy hardware or software. In such cases, the SDC will accept the legacy data product and work with the satellite operator to ensure a complete understanding of the time system, coordinate frame, and units being used.

In one particular case, we discovered that the operator supported a number of disparate constellations, where each group used different flight dynamics software and produced ephemeris in a different format. As it turns out, this situation is not that uncommon, with the merger or acquisition of satellite operators having become somewhat commonplace. We were asked if we wanted to independently convert the existing ephemerides from each group or receive the data in the JSpOC format, since all of the groups already generated that format for use by JSpOC in screenings for close approaches. The 'right' answer seemed to be to accept the data in the JSpOC format, realizing that there could still be some groups that might show issues.

For the most part, the comparative analyses were unremarkable, with only small discrepancies discovered. But in one case for two satellites, we discovered a discrepancy of over 10,000 km between the satellite operator's ephemeris for these satellites and the associated TLE data for the same time period. Obviously, something was seriously wrong with one or both data sets.

Since it seemed unlikely that either the satellite operator or JSpOC could be off by that much in tracking these satellites (there were no maneuvers in the time span being evaluated), the results suggested there had to be a problem with the interpretation or generation of the ephemeris data. While nothing obvious seemed amiss in a close visual examination of the data, visualizing the data in 3D using Systems Tool Kit (STK) quickly made it apparent that something wasn't quite right. Instead of the orbits closing on themselves to form an ellipse, they were open—like one might expect from a ground track. A quick check in STK to change the orbit of the TLE data to display it in the International Terrestrial Reference Frame (ITRF) coordinate system—which is Earth-fixed—now showed good agreement with the operator ephemeris.

After some discussion with the operators of the affected satellites, it turned out they had been generating their JSpOC ephemeris data using the correct format but in the ITRF rather than the J2000 coordinate system required by that format. And it appeared that they had been sending their data this way to JSpOC for over six months, without anyone noticing this huge discrepancy. Had a comparative analysis been performed at the time the data was first received by JSpOC, the situation should have been quickly identified and resolved. Fortunately, the situation was detected and corrected before it resulted in a safety-of-flight situation.

3.2 Case Study #2

Ensuring a proper interpretation of all data up front can eliminate many problems downstream, but all systems change over time. That is why it is critical to regularly ingest fresh data—via standard machine-to-machine interfaces—and to perform comparative analyses whenever possible. Better to find and eliminate an issue caused by a change during routine operations than to have it become a complicating—or worse yet—unknown factor in the middle of a serious event.

Things change for a variety of reasons and are usually unintentional. Changes may occur due to software updates—system-wide or even regular security updates—perhaps affecting a key library. This may change the data precision or format (e.g., changing fixed decimal to scientific notation). Whether the receiving system is robust to the change can be uncertain, but is often detectable through comparative analysis.

But in some cases, the change is not a change in the system, but a manifestation of a previously unseen mode of behavior that causes a problem. One such case was detected in the SDC's value-added comparative analyses of the Conjunction Data Messages (CDMs) created by JSpOC for SDA satellite operators. For most of the satellite operators in the SDC (25 of 30), the operator authorizes the SDC to have access to their CDMs. That allows the SDC to show how a particular conjunction event looks using SDA ephemerides in addition to the JSpOC SP data.

Instead of simply showing the SP vs. SP results, these analyses also show the SDC (ephemeris) vs. SP results and—when available—the SDC vs. SDC results. As part of this process, we show a variety of comparative analysis results, such as the difference between the SP results and the SDC results at the time of closest approach (TCA). Sometimes, this result can be quite large, particularly when TCA occurs after a planned maneuver. Since SP only propagates an orbit based on past observations, it is expected that this type of situation would occur occasionally, but the check is in place to ensure the operator is aware of any discrepancy.

Some time ago, JSpOC started following the SDC's lead in using operator ephemerides for some of their screenings, to compensate for this limitation. To support this task, each operator is expected to upload their data—in the JSpOC format—for use in conjunction screening. Since the SDC already includes validated tools to convert SDC ephemerides into standard formats like the JSpOC format, an SDA operator always has the option to upload their data to the SDC and then download it in the JSpOC format to send to JSpOC. But to make this process even easier, the SDC offers to perform this service automatically.

At a time designated by JSpOC, the SDC takes each of the latest SDC ephemerides and converts them to the JSpOC format. It then uploads them to Space Track. The beauty of this approach is that the operator only has to upload the data once and the SDC takes care of the rest. It also allows the SDC to name the files in a way that makes it possible to see whether the JSpOC used the same ephemeris in their screening as is currently available in the SDC. Some operators upload ephemerides frequently enough that a new ephemeris may already be available by the time JSpOC downloads the last ephemeris from Space Track, performs their conjunction assessment analyses, and uploads the results.

In many cases, though, the ephemerides are the same, so the comparison between the JSpOC and SDC ephemerides at TCA should be near zero (within a couple of meters, allowing for slight differences in Earth orientation parameters used for coordinate transformations). However, in one case for one of our geostationary satellites, we noticed that while the ephemerides showed as being the same, the comparison between the SDC and JSpOC versions of that ephemeris at TCA showed a difference of over 100 km. Something was clearly amiss.

It seemed that there would have to be an error in the SDC's generation of the JSpOC format, but a careful examination of the original and converted ephemerides showed agreement to below the precision of the data. Closer examination of the input ephemeris showed what initially appeared to be a maneuver event around TCA. But that event was not causing a problem when manually comparing the operator and JSpOC format of the same data.

Detailed analysis showed that at what had appeared to be a maneuver event, the operator's ephemeris showed the satellite virtually stopped in its orbit for one time step (10 minutes), before catching up to where it should have been on the next. For a near-impulsive maneuver, one might expect a duplicate time point with the same position but a different velocity or something similar. So, this appeared to be a problem with how the operator's legacy software handled this situation, which hadn't been detected before. And it was only detected in this case because it happened to create a conjunction.

That still didn't fully explain the discrepancy, however. The Lagrange interpolation experienced expected "ringing" around the anomaly, but a similar interpolation using the same data might be expected to have produced similar results. We might have been justified to simply assume this discrepancy was solely the result of the problem in the operator's software.

Unfortunately, the CDMs always require a state at TCA, even if that state is somehow derived from an ephemeris that does not provide a data point at that TCA. How that generation is performed is unspecified, but could be assumed to be the output of something like Lagrange interpolation. Of course, this is an unvalidated assumption and is why CDMs should provide the source data rather than using an unspecified conversion of that data.

A close review of the CDM produced by JSpOC showed a surprising result, however. In the optional fields for apogee and perigee (altitude), it showed an apogee of 35,791 km—a reasonable value—but a perigee of -2,132 km—below the surface of the Earth. JSpOC later acknowledged that when they *fit* the given velocity (not interpolated the ephemeris), it produced a near-hyperbolic orbit. If it was determined to be hyperbolic, the perigee would be at -6,378 km.

In this case, the comparative analysis led to the discovery of a problem in the operator's ephemeris generation at a maneuver event and the realization that assumptions about how the CDM data is generated can produce wildly misleading results. The SDC advised all operators of what to look for that might indicate a similar problem: the CDM analysis report shows both SDC and JSpOC use the same ephemeris but show significantly different results at TCA (kilometers vs. meters). And we now check for negative perigee values in the CDMs. These actions should ensure that the next time we see a previously unobserved anomaly—most likely from another operator—we will be able to quickly resolve the issue.

3.3 Case Study #3

Sometimes, the things that break aren't the older legacy processes, but instead are the result of new additions to a satellite operator's constellation. While the operator may have a well-established and validated set of tools for creating standard ephemerides, the addition of a new satellite—perhaps manufactured by a different company—brings with it new software to allow that satellite to be more effectively utilized or take advantage of new capabilities. And while that shouldn't require rewriting or recreating established processes, it often does. That is why the SDC independently validates the addition of any new satellite.

In one recent case, one of our SDA operators began adding a new class of satellites. For the first satellite of that class, the ephemeris provided used a slightly modified version of the JSpOC format, which used the J2000 coordinate frame along with the JSpOC epoch of the format YYDDDHHMMSS.SSS. Since the SDC already supported the JSpOC format for a number of other operators, it was straightforward to modify that converter to the slight differences in the format and the initial validation showed good results.

But when we started receiving CDMs for this satellite, our comparative analysis as part of that process was showing a difference at TCA of ~12,000 km when JSpOC screened this satellite with the operator's ephemeris, which was being provided directly by the operator. Since the SDC wasn't providing the ephemeris to JSpOC, it was initially difficult to determine whether the difference might be due to using different ephemerides. But the size of the discrepancy suggested a more fundamental problem, perhaps in the SDC's interpretation of what should be a standard data set.

To make things even more confusing, our weekly comparisons of operator ephemerides to the JSpOC TLE and SP data was showing a difference of ~1 km. It was unclear how this could happen unless there was a problem with the SDC CDM analysis process. But detailed review of a number of ephemerides from the operator and assurances that they were sending the same data to JSpOC, together with in-depth review of the CDM analysis process, was showing no errors.

At this point, we took all the ephemerides we had from both the SDC and the JSpOC SP ephemerides and generated piecewise ephemerides spanning the previous two weeks. The piecewise ephemeris process overlays each new day's data on top of the previous ephemeris, truncating the old ephemeris where each new ephemeris begins. When we compared the two sets, we discovered that the differences would switch from near zero to ~12,000 km as a step function at the beginning of *some* days. Examining the two data sets showed that the operator ephemerides were causing the problem, not the SP ephemerides. But the switching wasn't happening each day. With uploads spanning July 9-23, we found five of the eleven uploads showed a problem.

On a hunch—based on a careful examination of the faulty data and 3D visualization of that data—we adjusted the day of the year in the faulty ephemerides by one full day and found that adjustment eliminated the problem. That is, the adjustment reduced the difference to near zero.

It turned out that the operator had some kind of boundary-condition problem when converting their ephemeris timestamps to day of the year that was occasionally off by one full day. The fact that the problem only occurred intermittently made the problem even more difficult to discover. What further complicated the effort to detect the problem was that the two comparative analysis processes happened to be looking at the data at different times. The CDM comparative analysis process (which ran several times a day for that operator) that caught the problem just happened to be using the faulty data while the weekly review was looking at periods where the data just happened to be good. Without the use of regular comparative analyses as part of SDC operations, this problem might not have been detected and could have led to serious safety-of-flight issues.

3.4 Case Study #4

In early October 2013, we started noticing unusual behaviors in our comparison of the SDC ephemerides to the JSpOC TLEs for our 200+ geostationary satellites. Typically, we expect to see the TLEs track quite well except for periods following a maneuver, when orbit solutions will slowly diverge until the maneuver effect is incorporated in the JSpOC TLE orbit determination solution. This correction might take several days, but at that point, the TLE typically jumps back to the SDC solution.

In this case, we were seeing the opposite occur. The TLE would be tracking along with the SDC ephemeris quite well and then suddenly jump away and start drifting. This behavior was most noticeable in the 7-day longitude histories we reviewed, as seen in Fig. 1.

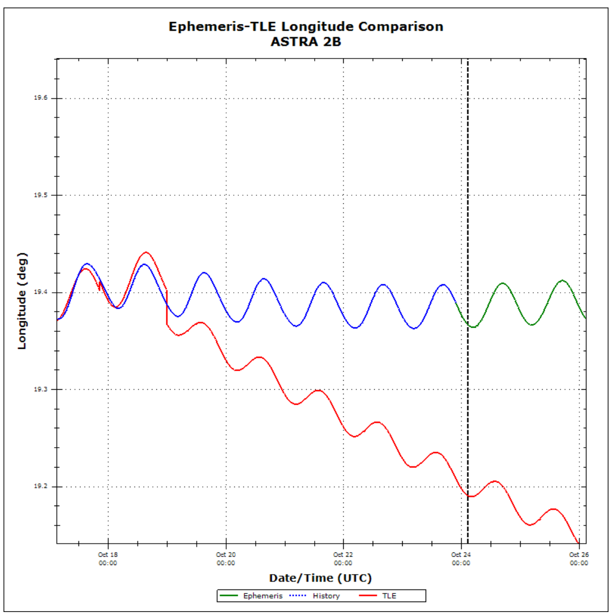

Fig. 1. Sample Longitude History

What was particularly unusual was that while we might have seen the occasional occurrence of this type of behavior randomly over all of our geostationary operators in our previous weekly reviews, we were now seeing a dozen or more each week and only for SES. The other distinguishing characteristic was that all of the TLEs had an epoch with a fractional day of exactly zero (e.g., 281.00000000). Because of the way that TLEs are generated, this would not be expected to occur naturally.

Upon reporting the problem to JSpOC, they initially responded, "When the TLE has an epoch of 0000Z, this is usually an element set we have ingested from the operator directly." But we were unaware of SES providing TLE data to JSpOC. JSpOC later told us, "SES provides ephemeris data to Millstone who then sends it in to the JSpOC in the form of a TLE." Since the JSpOC analysts considered it to be "truth data," it was added directly to the catalog "as long as it is reasonably close to the previous set." Our results seemed to show otherwise. Our SES contacts subsequently confirmed that they did send data to Millstone, but were unaware that it was being used to produce TLEs that then went into the public catalog.

It wasn't until late November that we were finally able to track down and obtain a sample of the data being provided by SES to Millstone, which turned out to be in the form of classical Keplerian elements at one-day intervals. The epoch of each element set was at 0000 UTC. SES confirmed that the data were osculating elements in the True of Date (TOD) reference frame. We were also able to extract the corresponding TLEs for Day 313 (2013 Nov 9) from the public catalog for comparison.

It turned out that the TLEs being generated by Millstone were a simple reformatting of the data being provided by SES. The mean motion (n) was calculated using Kepler's Third Law:

$$n^2 \cdot a^3 = \mu$$

where a was the semi-major axis provided by SES and:

$$\mu = 398600.5 \text{ km}^3/\text{s}^2$$

The remaining orbital elements (except for eccentricity) were rounded down from six decimal places to three—even though the TLE format has room for four decimal places.

The fundamental problem, however, was that Millstone had taken osculating orbital elements in the TOD reference frame and reformatted them as TLEs, which are propagated by SGP4 as Brouwer mean elements in the TEME (True Equator Mean Equinox) reference frame [2]. The result was that these TLEs were often off by 50-100 km within a day or two of epoch.

JSpOC and Millstone were using this approach in an attempt to sort out geostationary satellites in clusters, which can often be difficult to do with the current Space Surveillance Network (SSN) sensors. However, without using proper comparative analysis in their validation, they injected bogus TLEs directly into the operational public catalog, which may have adversely affected numerous operational processes.

Ultimately, the SDC was able to assist JSpOC with their need for supplemental TLEs—TLEs derived from operator data rather than SSN observations—by setting up a process to generate supplemental TLEs for SES satellites using SES public data and performing the necessary comparative analysis to validate the results against the ephemerides provided by SES to the SDC. Those supplemental TLEs are now provided on CelesTrak along with those for Intelsat, EUMETSAT, GPS, GLONASS, and others and are updated daily [3].

3.5 Other Cases

Of course, the daily comparative analyses conducted by the SDC are used to support and improve SSA in the interest of safety of flight for all operators—whether they are SDA members or not. Oftentimes, the problems detected in our analyses may not be errors, but simply be the result of delays caused by JSpOC temporarily losing track custody following a recent maneuver. The magnitude of the maneuver and when it was conducted (local day or night) can make detection difficult, especially for geostationary satellites. Limited optical sensors and bad weather can further complicate this process. When the SDC detects large discrepancies that indicate a likely loss of track custody, we generate a supplemental TLE to assist JSpOC with reacquiring the orbit. Often this is as simple as making the association with an 'unknown' object being tracked in their analyst catalog, but the supplemental TLE can also be used by the SSN to relocate the object, if needed. The SDC has supported JSpOC with supplemental TLEs hundreds of times as part of this open collaboration.

Sometimes the problems detected point to the misassociation of observations for one satellite with the track of another—causing something referred to as a cross-tag. These situations often occur in geostationary clusters, particularly following stationkeeping maneuvers. Gaps in monitoring due to daylight, weather, and limited optical observing time all contribute to the situation. Of course, each SDA operator knows the identities of their satellites and their associated orbits, so it is a simple matter to report the problem to JSpOC so that they can take the appropriate corrective action. Otherwise, the potential exists for misleading reports of things like close approaches that could result in confusion that adversely impacts safety of flight.

4. Conclusions

All systems fail at one point or another. As we have seen in SDC operations, failures can occur for a variety of reasons: bad input data, incorrectly formatted data or unanticipated changes in the data format, faulty algorithms or the application of algorithms in situations outside of their expected use, problems converting data from one format to another, misunderstanding of coordinate systems or other metrics being used, or even unnoticed failures in data transfer that may lead to data latency. The regular reporting of comparative analysis results in our standard products allows us to detect and eliminate many of these problems before they become an issue affecting safety of flight.

Over the years, as we have gained experience examining millions of analysis products for satellites in all orbital regimes, we have been able to discern patterns of behavior pointing to families of problems and have those problems automatically identified for review by both our SDA operators and our analysts. But the potential for out-of-the-ordinary cases like those described above require review of the analyses to ensure those cases don't slip by unnoticed.

References

[1] The Consultative Committee for Space Data Systems (CCSDS), *Orbit Data Messages*, Recommended Standard CCSDS 502.0-B-2, Nov 2009, https://public.ccsds.org/pubs/502x0b2c1.pdf, (accessed 2018 Sep 16).

[2] F.R. Hoots and R.L. Roehrich, *Spacetrack Report No. 3: Models for Propagation of NORAD Element Sets*, Dec 1980, http://celestrak.com/NORAD/documentation/spacetrk.pdf (accessed 2018 Sep 16).

[3] T.S. Kelso, Supplemental Two-Line Element Sets, http://celestrak.com/NORAD/elements/supplemental/, (accessed 2018 Sep 16).