SSA Sensor Calibration Best Practices

Thomas M. Johnson
Analytical Graphics, Inc., 220 Valley Creek Blvd, Exton, PA 19468, USA

ABSTRACT

Best practices for calibrating orbit determination sensors in general and space situational awareness (SSA) sensors in particular are presented. These practices were developed over the last ten years within AGI and most recently applied to over 70 sensors in AGI's Commercial Space Operations Center (ComSpOC) and the US Air Force Space Command Space Surveillance Network (SSN) to evaluate and configure new sensors and perform on-going system calibration. They are generally applicable to any SSA sensor and leverage some unique capabilities of an SSA estimation approach using an optimal sequential filter and smoother. Real world results are presented and analyzed.

1. INTRODUCTION

The increased interest in Space Situational Awareness (SSA) in the last decade has led to a significant increase in the number of sensors capable of tracking objects in space and thus the need for a rigorous, repeatable process for characterizing sensor performance. This paper outlines the sensor models used in AGI's Orbit Determination Tool Kit (ODTK) and the methodology we use for configuring them. Background information on ODTK itself is available in [1]. The focus is on "metric" observations such as azimuth, elevation, range, Doppler, right ascension, declination, etc. used in orbit determination. The calibration of "non-metric" observations such as visual magnitude or radar cross-section is not addressed as these values are not traditionally used in the orbit determination process. The processes described have been successfully applied to radars, optical telescopes, RF telescopes, passive RF receivers, laser ranging systems, and many other sensors, both ground, ground mobile, and space based. At AGI we have applied these processes to SSA sensors from the Space Surveillance Network (SSN) and the AGI Commercial Space Operations Center (ComSpOC) via our sensor partners at ExoAnalytics Inc., SRI Inc., Las Cumbres Observatory Global Telescope Network (LCOGT), Raytheon Inc., and others.

Properly configured sensor models are a necessary requirement for producing an accurate ephemeris with a realistic covariance. The goal of the calibration process is to identify the proper values for the sensor parameters to support a quality orbit determination result.

It's worth noting that within ODTK we use the terms measurement or observation interchangeably to mean a single value such as range from a sensor. It's often the case that a sensor may take multiple observations simultaneously (e.g. azimuth, elevation, range, and Doppler), thus producing what we call an observation set. But this leads to confusion with people who are more familiar with the Air Force Space Command (AFSPC) Space Surveillance Network (SSN) where an "observation" is the set of simultaneous measurements from a sensor and thus equivalent to an ODTK observation set. We use the ODTK terminology within this paper.

A sensor can support multiple measurement types. A radar typically has a measurement model for azimuth, elevation, and range (and often Doppler). We want to characterize the primary parameters used in the measurement model. We'll ignore the many other inputs to the overall sensor model such as local environmental conditions (temperature, pressure, and humidity), total electron count, tidal effects, tectonic plate motion, etc.

2. SENSOR MODELS

A simple method for characterizing a measurement is to assume it has a Gaussian distribution about a constant non-zero value. The primary parameters are the white noise sigma and constant bias. While this is sufficient for many purposes, we have found the need for time dynamic bias models when working with some sensors (and in particular many of those used for SSA purposes). The ODTK stochastic models are summarized in Table 1. The proper selection, application, and configuration of these models for SSA estimation problems is discussed in [2].

Table 1. Stochastic Models

Model	Parameter	Definition
Gauss Markov	Constant bias	The constant (or nominal) value associated with the state parameter.
	Bias sigma	Root-variance of the initial error in the "Constant".
	Bias half-life	This controls how fast the state error estimate will decay in the absence of measurements; the shorter the half-life the faster the decay.
	White noise sigma	Square root of the variance representing the random uncertainty in the measurements.
Vasicek	Constant Bias	The constant (or long-term mean) value associated with the state parameter.
	Long term bias sigma	Root-variance of the initial error in the "Constant".
	Short term bias sigma	Root-variance of the short-term error where the variance is about the long term mean
	Short term half-life	The "half-life" of the short-term state error estimate. This controls how fast the state error estimate will decay towards the long term mean in the absence of measurements; the shorter the half-life the faster the decay.
	White noise sigma	Square root of the variance representing the random uncertainty in the measurements.
Random Walk	Constant bias	The constant (or nominal) value associated with the state parameter.
	Diffusion coefficient	Determines the amount of process noise to be added to the state covariance in going from time t_1 to t_2. The amount of process noise added will be $\mathbf{a}^2\Delta t$, where $\mathbf{a}$ = the diffusion coefficient and $\Delta t = \lvert t_2 - t_1 \rvert$.
	White noise sigma	Square root of the variance representing the random uncertainty in the measurements.

The primary goal of the calibration process is to determine the proper values for the stochastic model for each measurement. It's also important to identify which corrections (if any) need to be applied to the measurements. The most common corrections are: light time delay, annual and diurnal aberration, troposphere refraction, ionosphere refraction, and time bias. Right ascension and declination measurements also require configuring the proper coordinate frame (J2000, ICRF, TEME of Date, etc.).

3. CALIBRATION STRATEGY

The calibration strategy is outlined below. Ideally all steps would be exercised, but in practice there are times when they cannot. We identify these situations and suggest possible alternatives.

3.1 Sensor Information

Collect the basic information about the sensor itself. The easiest way is to ask the vendor. When this fails or is unavailable then you will need to resort to on-line searching or trial and error during the rest of the calibration process. Older sensors (such as those in the SSN) often require the latter.

3.1.1 Location

The sensor position is the most critical piece of information. Get the sensor latitude, longitude, and altitude or equivalent Cartesian coordinates and the datum the coordinates are in. The datum is typically WGS-84 (given the proliferation of GPS based surveying systems) and is a reasonable assumption when precise information is not available. If you suspect that the survey position is old (10 years or more) or generated by a non-US organization then it could be a different datum. Transform coordinates between datums as necessary in order to be consistent with your OD tool (ODTK requires WGS-84).

A simple location validation check can be done by entering the location into an on-line mapping site such as Google Maps, Google Earth, or Bing Maps. A sample is shown in Fig. 1. We recommend using multiple sites because of the wide diversity of image quality. The end goal is to verify that the image shows the sensor at the location you specified. Our experience has been that both sites have pretty accurate image registration. Be careful when the image shows multiple sensors (often the case with teleport facilities or observatories). The coordinates you have may not be for the sensor you wanted.

Fig. 1. Google Earth Location for Las Cumbres Tenerife Sensor

The location can represent a mean tide location or an instantaneous location at a particular time. Sensor locations vary with solid Earth tides, ocean tides, tectonic plate motion, etc. The measurement accuracy governs how important this is. A mean tide position is sufficient (or can be assumed) for typical radar range (> 1 m accuracy) or optical angle measurements (> 0.25 arcsec) Laser ranging, GNSS measurements, or time and frequency difference of arrival (TDOA, FDOA) measurements require a higher fidelity model to account for the Earth tidal effects and tectonic plate motion.

The coordinates may represent the base pedestal location, not the antenna phase center of the dish or telescope. The latter is preferred. The antenna phase center on larger dishes may not be at the mechanical pivot point of the dish and thus introduce a geometry dependent range bias.

3.1.2 Coordinate Frame

Processing right ascension and declination measurements requires knowing the inertial coordinate frame they are referenced to. Modern optical sensors typically use J2000 or ICRF. Note that SSN optical sensor data (when stored in the B3 tracking data format) use a True Equator Mean Equinox (TEME) of Date frame.

3.1.3 Corrections

Some sensors correct the raw measurements to account for various environmental effects or modeling decisions and only provide the corrected measurements. Therefore it’s important to know what effects have already been accounted for so that we don’t apply the correction again. Typical corrections include:

a. Troposphere refraction
b. Ionosphere refraction
c. Light time delay (and the related time tag reference)
d. Annual and diurnal aberration

3.1.4 Miscellaneous

Confirm other sensor configuration information:

a. Doppler count interval
b. Units – meters vs. kilometers, are the angles in radians, degrees, arc seconds, etc.
c. Time tag reference – Does the time tag represent the time the signal was transmitted to the satellite, the time it arrived at the satellite, the time it left the satellite, or the time it was received at the sensor? Typically it's the latter, but it's important to know for sure.
d. Antenna autotrack – Is the antenna actively tracking the satellite based on the received signal power or is the antenna pointing along a pre-programmed set of pointing angles. In the latter case the angles are not useful for estimation.

3.2 Get Measurements

As counter-intuitive as this sounds, obtaining real measurements can be the source of the longest delay. The measurements that are often readily available are usually not suitable for calibration purposes. While this doesn't prevent you from doing orbit determination runs, the credibility of the resulting OD solution is at risk, particularly when trying to assess whether you have a realistic covariance. A dedicated calibration campaign should be attempted for each sensor. Ideally the sensor should be tasked to collect multiple tracks on calibration class satellites and preferably under a variety of geometrical conditions (low elevation, high elevation, time of day) and for multiple satellites. An ideal calibration satellite has the following characteristics:

1. Constant area – e.g. spherical in shape. This avoids introducing issues around a varying area that might affect the drag or solar radiation pressure model. Since most OD tools default to a spherical model area anyways it's convenient to be consistent with this to eliminate a source of error.
2. Not maneuvering. This doesn't have to mean that it never maneuvers, just not during the calibration campaign. Be careful because "maneuvering" includes anything that results in a thrust (or equivalent force). This includes momentum dumps, spin or attitude control, active nutation damping, solar sails, etc. Some of this may be unavoidable, in which case use the satellites that minimally do this.
3. High enough to avoid drag. Generally 1200 km or higher perigee is desired. Not all sensors reach this high (or this low). The primary concern with drag is the atmospheric density model. Each model has its own inaccuracies that introduce uncertainty into the calibration process. Avoid low altitude satellites when possible.
4. Known mass. This goes hand in hand with the desire for constant area. The overall goal is to be able to state with some confidence that we know the area to mass ratio A/m used as part of the drag or solar radiation pressure calculations.

3.3 Reference Ephemeris

Obtain a high accuracy ephemeris for one or more satellites covering the time period when you have measurements. The ephemeris will be used as a reference or "truth" model and ideally will have position, velocity, and a realistic covariance (in both position and velocity). This is often the most difficult step in the calibration process as suitable reference ephemerides are hard to come by. Table 2 outlines the most common satellites we use and ephemeris sources.

Table 2. Reference Ephemeris

Satellites	Source	Notes
GPS, GLONASS	International GNSS Service www.igs.org	High accuracy ephemerides in SP3 format is readily available for GNSS satellites and for the GPS and GLONASS satellites in particular.
GPS	National Geospatial Agency earth-info.nga.mil/GandG/sathtml/ephemeris.html	High accuracy ephemerides in SP3 format is readily available for GPS (referenced to the antenna phase center or center of mass).
STELLA, STARLETTE, LARES, LARETS, AJISAI, LAGEOS 1 and 2, ETALON 1 and 2	International Laser Ranging Service (ILRS) ilrs.gsfc.nasa.gov	Predicted ephemeris is available in CPF format. Normal point and full rate observations are available if you want to produce your own reference ephemeris (this is the method we use at AGI).
WAAS satellites GALAXY-15, ANIK F1-R	Federal Aviation Administration (FAA) www.nstb.tc.faa.gov/DisplayNSTBDataDownload.htm	Broadcast ephemeris messages can be downloaded.

GNSS satellites are generally preferred whenever possible because of the easy access to high quality ephemeris. The SP3 files don't contain covariance, but you can assume with confidence that the position error is less than 0.5 m (even for the ultra-rapid solutions).

We prefer to produce our own reference ephemerides for the satellite laser ranging (SLR) satellites using the ILRS observations. This allows us to calculate a full position and velocity covariance. There are other SLR satellites that could be used such as GRACE A and B, CRYOSAT 2. These occasionally maneuver though or are low altitude, thus complicating the calibration process.

Three GEO satellites are currently used as part of the FAA's Wide Area Augmentation System (WAAS). The WAAS ground processing system estimates the ECEF position of each satellite and uploads it to the satellite for subsequent broadcasting on an L-band downlink as a Type 9 GEO Navigation Message. These messages are available on the FAA web site but require you to write a parser to extract the ephemeris. The format is described in [4]. Note the NTSB Type 5 message are wrappers around the Type 9 WAAS message (you must parse the outer format data, then the inner message). The Type 9 message format is described in [5].

In practice we have found that the GALAXY-15 and ANIK F1-R positions are reliable. The third WAAS satellite INMARSAT-4 F3 is not used for calibration because we have found that the WAAS ephemeris is NOT reliable – we suspect this is because the satellite is performing daily north-south station keeping. This is evidenced by a declination residual error of as much as 15 arcseconds. Definitive WAAS OD accuracy numbers are difficult to come by, although [4] cites an accuracy of 30-40 m. The use of ANIK F1-R can be difficult when processing optical measurements because it is co-located in orbit with ANIK F1 and ANIK G1 leading to possible mis-correlation of the measurements being validated.

The European Geostationary Navigation Overlay Service (EGNOS) is similar to the US WAAS system and also supports a broadcast GEO Navigation Message. However, we have not found a source for obtaining these messages. But this shows promise for providing GEO calibration satellites over Europe.

3.4 Residuals vs. Reference Analysis

Generate the measurement residuals by running the measurement models using the reference ephemeris. No orbit or other parameter estimation occurs in this step. This analysis is so frequently performed that we simply call it an "RvR run." This task can be accomplished in ODTK using the Residuals vs. Reference HTML utility. The goal is to validate the sensor information and modeling decisions from the earlier steps and get a general sense of the sensor performance. Plot each measurement independently and examine the residuals and draw the initial conclusions. Ideally the residuals should have a nice normal distribution centered on zero.

Gross errors are often indicative of an incorrect sensor location, incorrect measurement units, or using the wrong satellite ephemeris. Smaller residual errors are often caused by errors in the coordinate frame, ionosphere or troposphere model corrections, light time delay corrections, or aberration corrections. A systematic study of each these models is often necessary, turning each one on and off and examining the impact on the residuals. Examining the residuals as a function of elevation angle, range, range rate, and time of day can often be enlightening as well. Trends in the overall pattern of behavior can suggest initial values for the measurement white noise sigma, measurement bias, and bias stochastic model parameters (sigma and half-life). Apply these to the sensor model as appropriate.

Figures 2-34 illustrate a typical analysis (vertical axis units are in arcseconds). Fig. 2 shows the RA/DEC residuals from an RvR run for an optical sensor tracking several GPS satellites with no aberration or time bias applied. In Fig. 3 we have corrected for annual and diurnal aberration, and in Fig. 4 we have also applied a time bias correction.

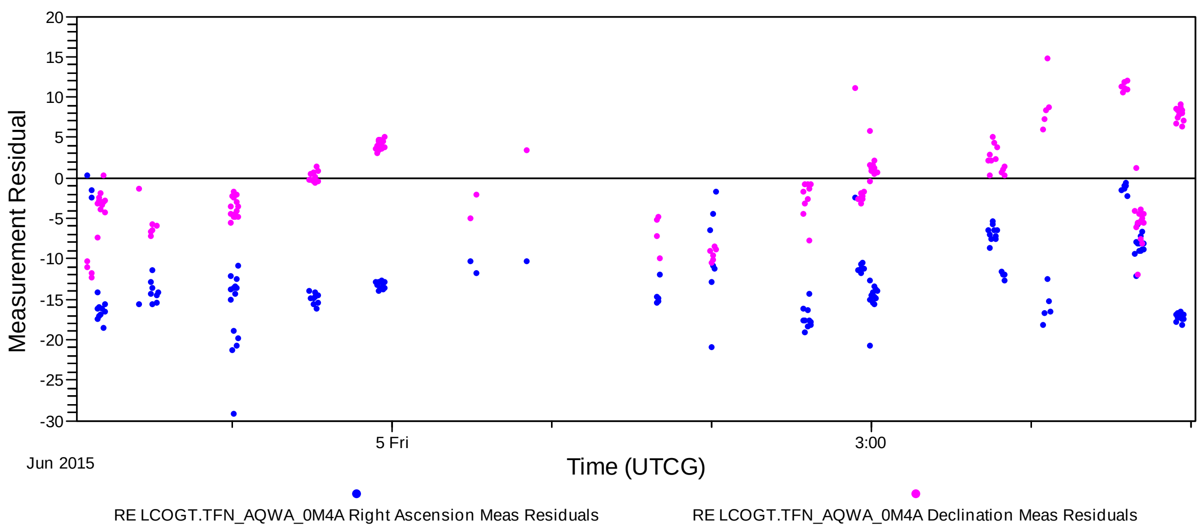

Fig. 2. RA/DEC Residuals Without Corrections

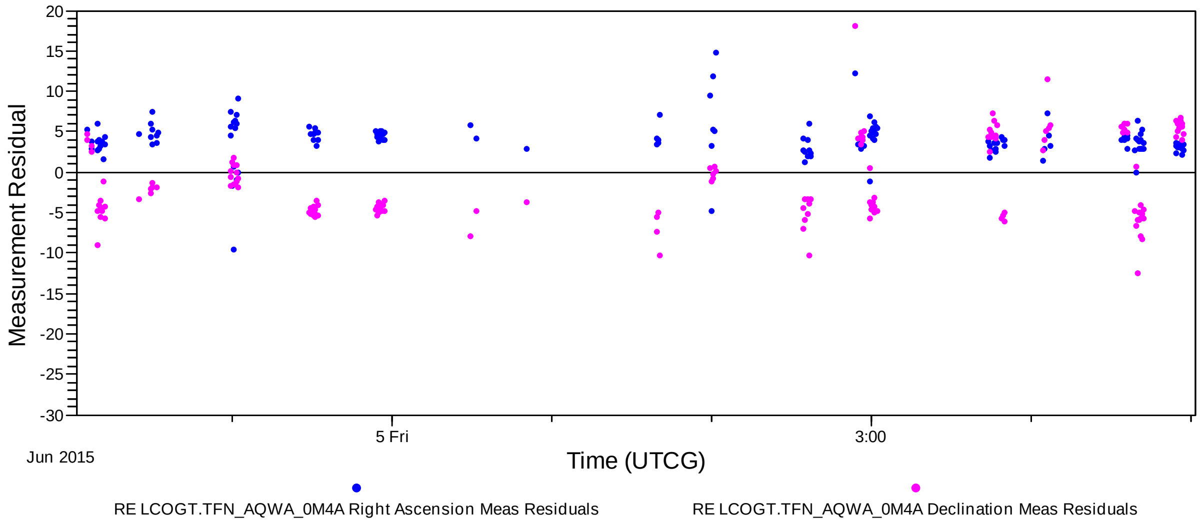

Fig. 3. RA/DEC Residuals With Aberration Corrections

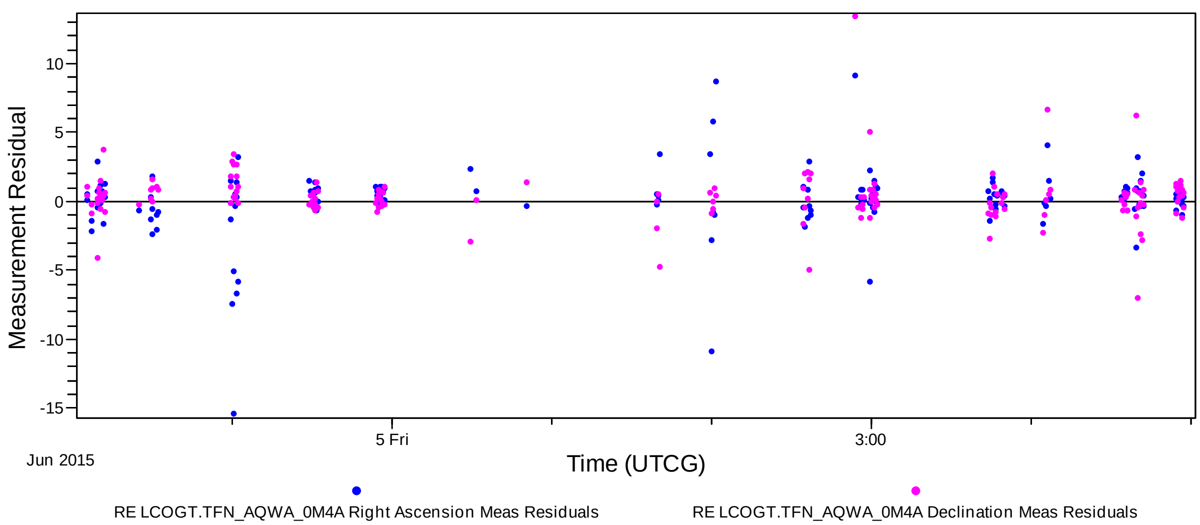

Fig. 4. RA/DEC Residuals With Aberration and Time Bias Corrections

3.5 Sensor Model Estimation

Having accounted and corrected for gross errors, the next step is to estimate sensor parameters and fine tune the sensor model configuration. Still using the reference ephemeris (to keep the orbit locked down), run the estimator (the filter and smoother in ODTK) and solve for the measurement bias. If you suspect that there is a time tag bias then we recommend solving for the time tag bias first (and do not solve for measurement biases at the same time). Examine the measurement residuals and bias estimate. Useful analysis tools for doing so include measurement histograms and quantile-quantile (QQ) plots [6]. Iterate, fine tuning the white noise sigma, bias, bias sigma, and half-life settings.

- Examine the bias estimate(s).
- Verify we can pass the bias consistency test for each measurement
- Examine the QQ plots. If the residuals are non-Gaussian then you may have to artificially inflate the white noise sigma to get the consistency test to pass

- Include multiple satellites if possible so that you don't get a solution that favors a time gap between passes but will have problems when doing a simultaneous run with back to back or overlapping passes on the same tracker.

The bias consistency test is a particular application of the McReynolds filter-smoother consistency test as applied to the measurement bias. The test evaluates a one-dimensional test statistic using results from the filter (F) and the smoother (S)

$$\Psi = \frac{x_F - x_S}{\sqrt{\sigma_F^2 - \sigma_S^2}}$$

where x is any state parameter (the measurement bias in this case) and σ is its uncertainty from the covariance. We test if $|\Psi| \leq 3$, which McReynolds showed was a necessary condition for having a realistic covariance.

Figs. 5 and 6 show the smoother estimate of the RA/DEC measurement bias and the measurement bias consistency test. There appears to be a 0.25 arcsecond bias in both RA and DEC. The consistency test passes with the statistic being within ±3.

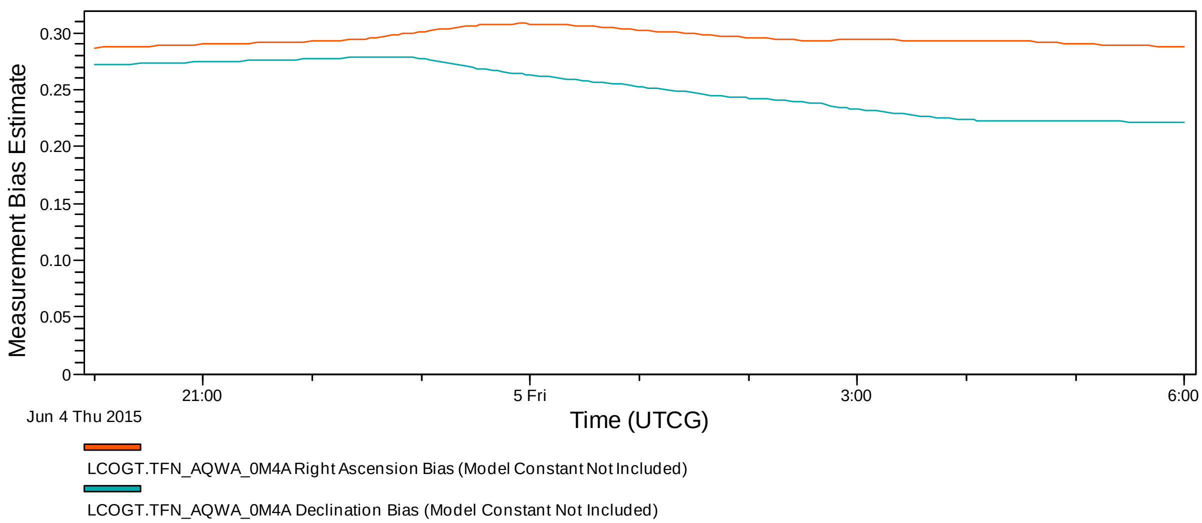

Fig. 5. RA/DEC Bias Estimates

The histograms and QQ plots must be generated separately for each sensor and measurement for proper interpretation. For example, avoid combining all the right ascension residuals from multiple sensors into one QQ plot as this will make it very difficult to separate out which sensors are properly calibrated. QQ plots are the preferred analysis tool because they are easy to interpret, even with few data points. They provide insight into the proper white noise sigma, measurement bias, and how Gaussian the residuals are. Fig. 7 shows the RA and DEC QQ plots. The residual sigmas are nicely contained within the acceptance bounds meaning they have a reasonably Gaussian behavior and have a slope of 1 indicating the white noise sigma is configured correctly.

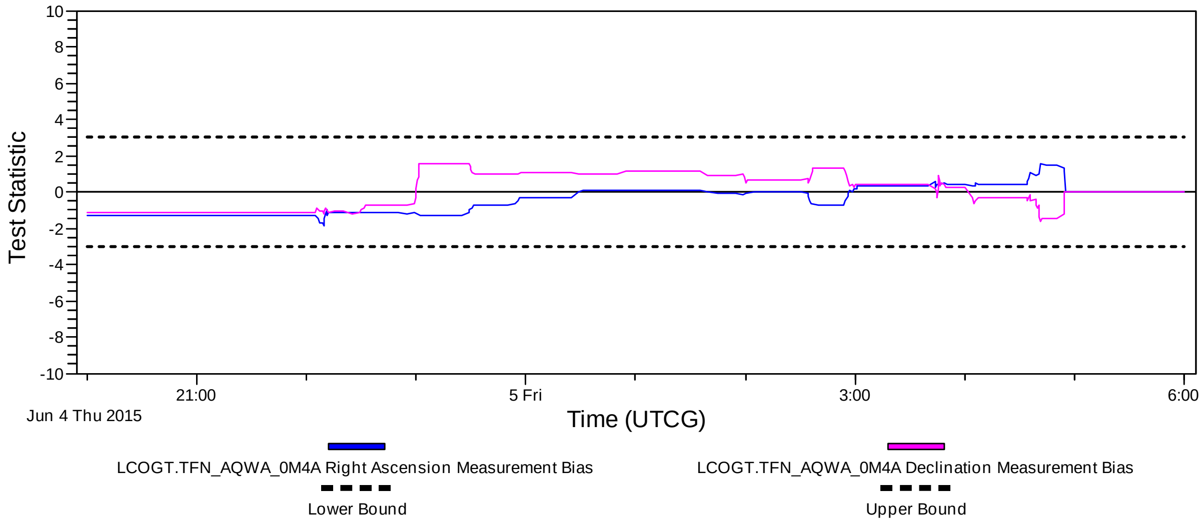

Fig. 6. RA/DEC Bias Consistency Statistics

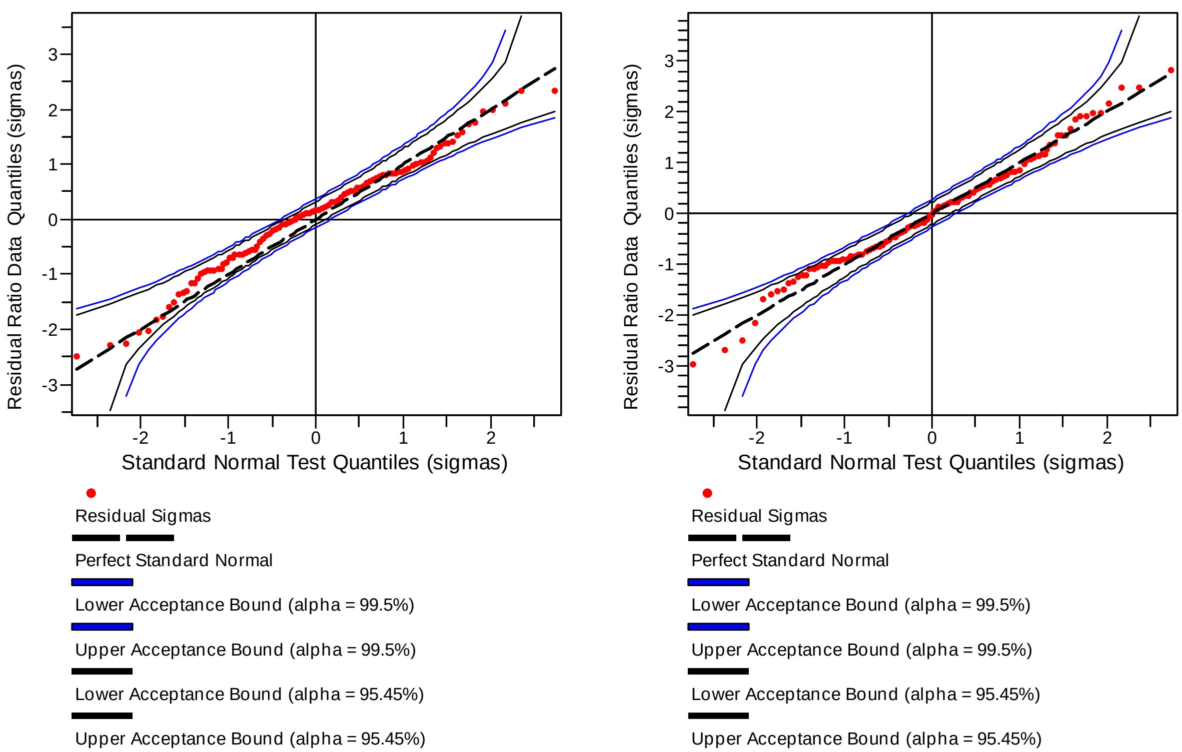

Fig. 7. RA/DEC QQ Plots

3.6 Satellite Estimation

Now that we have a reasonable sensor configuration we can begin estimating the satellite orbits (no longer using the reference ephemeris). The initial orbit state and covariance and force model configuration are very important. Using an initial state and covariance (if available) from the reference ephemeris is generally recommended.

Run the estimator, solving for the satellite orbits and measurement biases. The more you know a-priori about the satellites you are using the better (the fewer unknowns the better so we can focus on the sensors themselves). Do single satellite estimation runs first, then repeat the analysis solving for multiple satellites simultaneously. The runs do not have to be limited to just the calibration satellites, but using them first is highly recommended.

After each estimation run verify:

- Passing all filter-smoother consistency tests
  - Satellite position, velocity, ballistic coefficient, SRP coefficient, transponder bias
  - Tracker bias
- QQ plots look good
- Compare your filter and smoother ephemeris with the reference ephemeris

The ephemeris comparison should examine the overall position difference between your estimates and the reference satellite as well as how the difference compares with your covariance estimate. Evaluating the position difference and covariance is done in STK using the RIC Coordinates – Position Cov Sigma data provider. This reports N where the current relative position vector lies on an N sigma ellipsoid, where the ellipsoid is obtained from the position covariance of the primary satellite. Ideally N should be less than 3. This test is highly sensitive to differences along the narrow axis of the covariance ellipsoid (typically cross-track). The covariance of the reference ephemeris is typically much smaller than your estimated covariance and thus has a negligible effect on this metric. But if desired it can be accounted for using the Position Combined Cov Sigma data provider which will add the covariance of the reference ephemeris as well (under the assumption that its covariance is independent of your estimated covariance).

Be careful not to "over-tune" the sensor configuration for any single satellite. A single satellite (non-GEO) has an implicit time gap between passes because of the orbit geometry. During this time gap a bias correction estimate can decay per the stochastic model and find a new correction at the next pass. When the corrections are reasonably constant this works well, but if the bias is geometry dependent (typically elevation for a radar) then it's possible to have high and low elevation passes with significantly different biases. Adjusting the bias half-life can accommodate this. When the same sensor sees multiple satellites at the same time, each with a unique elevation dependent bias then the single satellite half-life you just tuned may still be too long. This is a case where an improved bias model would be warranted. The SSN phased array radars often exhibit this type of behavior [7]. We have seen similar results when calibrating the SSN radar models as part of AGI's Joint Space Operations Center (JSpOC) Mission Systems contract and numeric validation work.

4. REAL WORLD ISSUES

Real world data analysis exposes issues that complicate an otherwise normal calibration process. A few of these are discussed below along with possible mitigation strategies.

The sensor didn't see any calibration class satellites. Generate reference ephemeris on other satellites as best you can using only well calibrated sensors (not the one under investigation). This allows you to sort out the satellite specific settings (ballistic coefficient, SRP). Perform the calibration process using this reference ephemeris.

The uncalibrated sensor is collecting data on an unknown object. Similar to the situation above. If nothing is known about the satellite or sensor then you can only do so much to produce a self-consistent solution (self-consistency should not be interpreted as correct, it's easy to have biases that result in the former but not the latter).

All you have are measurements for Geostationary satellites. GEO satellites are very poor calibration objects because they have very little orbit dynamics to help out the solution and are often maneuvering. Typical active GEO satellites (including the WAAS satellites mentioned earlier) have an inclination less than 0.1 degree. The result is that a time tag bias, right ascension bias, range bias, and transponder bias all alias directly into each other and into the Earth fixed longitude estimate of the satellite position. These can be difficult if not impossible to separate out. Objects with larger inclinations are much preferred in order to mitigate this because the inclination introduces a time varying declination offset and declination rate. Using older GEOS that no longer perform north-south stationkeeping is recommended.

Make sure the measurements are independent. This is particularly a problem with a sensor generating azimuth and elevation or X and Y angles. The sensor may simply be returning the pointing angles the antenna was given and not an actual measurement. Verify that the antenna *can* auto-track and actually *was* auto-tracking.

The tracking data format may introduce errors due to limited precision. The JSpOC B3 format is a classic example of this. Range measurements are formatted as 7 digits and an exponent digit. When the range exceeds 10,000 km, the least significant digit switches from a value of 1 m to 10 m. If the sensor white noise sigma is less than 10 m this will change the observed white noise sigma (and the distribution of the residuals will become non-Gaussian).

5. REFERENCES

1. Vallado, D.A., Hujsak R.S., Johnson T.M., Seago J.H., Woodburn J.W., "Orbit Determination using ODTK Version 6", International Conference for Astrodynamic Tools and Techniques, European Space Astronomy Centre (ESA/ESSAC), Madrid, Spain, 3 – 6 May 2010.
2. Johnson, T.M., "Orbit Prediction Accuracy Using Vasicek, Gauss-Markov, and Random Walk Stochastic Models," Paper AAS 13-280, 2013 AAS/AIAA Space Flight Mechanics Meeting, Kauai, Hawaii.
3. Tombasco, J., *Orbit Estimation of Geosynchronous Objects Via Ground-Based and Space-Based Optical Tracking*, Dissertation. University of Colorado, 2011.
4. Federal Aviation Administration, *NTSB Data Format*, November 2003. Web. 9 Sept. 2015. <ftp://ftp.nstb.tc.faa.gov/pub/NSTB_data/0-NstbDataFormat.doc>
5. Federal Aviation Administration, *Specification for the Wide Area Augmentation System (WAAS)*, Section 4.4.11.
6. Wright, J. R., "QQ-Plot for Sequential Orbit Determination," Paper AAS 11-607, 2011 AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska.
7. Barker, W. N., Casali, S.J., Walker, C. A. H, "Improved Space Surveillance Network Observation Error Modeling and Techniques for Force Model Error Mitigation," Paper AAS 99-420, 1999 AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska.