Prior Probability Model Development to Support System Safety Verification in the Presence of Anomalies

Sam Pullen, Jason Rife, and Per Enge
Dept. of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305-4035

***Abstract* - Assessments of "prior probabilities" of faults that must be mitigated are essential to integrity and safety verification for safety-critical systems. In the case of Global Navigation Satellite System (GNSS) augmentations, many integrity threats come from faults within GNSS or anomalies in the atmosphere that are outside the control of the augmentation system designers. For a variety of reasons, including the rarity of the fault modes of concern, insufficient data exists to derive the fault probabilities directly from data without impractically-large confidence intervals. This paper addresses these concerns by illustrating a general approach derived from several recent examples in which data and judgment can be combined to produce usable and certifiable prior probabilities. Examples are shown for both Global Positioning System (GPS) satellite faults and ionosphere (atmospheric) anomalies.**

I. INTRODUCTION

One of the challenges in verifying that GNSS and its augmentations support the required levels of user safety is addressing system vulnerability to anomalies that are beyond the control of the system designer. In addition to internal equipment or software failures, GNSS augmentations such as SBAS (Wide Area Augmentation System, or WAAS, in the U.S.) and GBAS (Local Area Augmentation System, or LAAS, in the U.S.) are also potentially vulnerable to GNSS satellite failures and anomalies in the ionosphere and troposphere that GNSS signals must travel through to reach most users. In some cases, the worst-case impact of the set of "credible" anomalies is bad enough that, if it were assumed to be "nominal" (i.e., present all the time), the system would have little value to users. Threats in this category can only be accepted and mitigated if they are demonstrated to be rare.

To quantify the degree of rarity of a given anomaly, a prior probability model (or PPM) must be developed and validated for that anomaly. Equipment internal to the system being developed often supports "handbook" failure-rate analysis that, while far from perfect, is accurate to within an order of magnitude and is usually conservative. However, for external events, and particularly for atmospheric anomalies, even this level of accuracy is difficult for several reasons:

- Because the event of concern is rare, few examples of it are likely to be available in past observations;
- Data from past events usually comes from multiple different sources and eras, and harmonizing them into a single data-quality standard is often impossible;
- In some cases, where little or no data exists, probability distributions must be generated based on expert opinion;
- As money-management firms disclose in their advertisements: "Past performance is no guarantee of future results."

This paper illustrates the degree to which these obstacles can be overcome by generalizing an approach to PPM development from two recent examples: ionosphere spatial anomalies (for LAAS) and satellite signal deformation (for Local Area Monitoring, but also applicable to LAAS and WAAS). In both cases, prior probabilities must be generated based on observations of only fewer than five anomalous events. Statistical confidence intervals are generated in each case, and conservative upper confidence bounds are chosen instead of "best-estimate" point probabilities (i.e., 1 / number of events) in order to include margin against incorrect assumptions and future changes in fault behavior as well as limited data. This margin is relatively larger in the case of signal deformation for two reasons:

1. GPS satellites have not been continuously monitored for signal deformation; thus there is a possibility that more than one event has occurred;
2. The chance of the future being different from the past is higher for new generations of GNSS satellites than it is for changes in atmospheric behavior.

In both cases, however, the amount of margin is somewhat arbitrary and thus should be chosen by a consensus of experts after review and discussion. However, care should be taken in forming a consensus to mitigate the "least common denominator" problem: achieving consensus based on the amount of margin preferred by the most conservative person in the group. The paper concludes with a discussion of how the differences between "average risk" and "specific risk" philosophies affect prior probability assessment and the degree of anomaly risk mitigation required of LAAS and WAAS

II. ROLE OF PRIOR PROBABILITIES IN AUGMENTED GNSS INTEGRITY RISK ASSESSMENT

Augmented GNSS landing systems are designed to meet integrity requirements standardized by ICAO internationally and by the FAA and RTCA in the U.S. For LAAS systems designed primarily to support precision approaches under Category I (200' decision height) conditions, the probability of loss of integrity (occurrence of unsafe conditions without warning from the system) must be no greater than 2×10^{-7} per (150-second) approach [1]. One-quarter of this total allocation (5×10^{-8} per approach) is allocated to nominal conditions (H0 hypothesis) and single-receiver-failure conditions (H1 hypothesis), while the remaining 75% is allocated to all other "failure" conditions, including satellite faults and atmospheric anomalies [1,2]. This allocation must be further subdivided among all foreseen satellite, atmospheric, and multiple-receiver fault modes.

Given a sub-allocation to fault mode *i*, LAAS must mitigate the threat posed by fault mode *i* such that the following constraint is met [3]:

$$P_{alloc,i} \geq P_{PL,i}\, P_{MD,i}\, P_{prior,i}\,; \tag{1}$$

where $P_{alloc,i}$ is the sub-allocated per-approach integrity risk probability for this event, $P_{prior,i}$ is the per-approach "prior probability" of this event (i.e., the probability that fault mode *i* occurs during any separate 150-second approach interval), $P_{MD,i}$ is the probability that LAAS monitors fail to detect the anomaly and exclude the affected measurements within the time-to-alert (6 seconds for CAT I LAAS [1]) after the anomaly becomes potentially hazardous, and $P_{PL,i}$ is the probability that an undetected fault combines with nominal errors to push the resulting user position error outside of the protection level computed by the user. Note that $P_{MD,i}$ is conditional on fault *i* occurring, and $P_{PL,i}$ is conditional on fault *i* occurring and a LAAS missed detection. Clearly, the lower the value of $P_{prior,i}$, the less demanding the requirements are on LAAS to quickly detect and exclude this fault condition.

III. PREVIOUS WORK ON GPS ANOMALY PROBABILITIES

While the Global Positioning System was not originally designed to support safety-of-life applications, development of integrity-focused "add-ons" to GPS has made it important to characterize, to the degree possible, the probabilities of GPS satellite and control-segment faults. Sections 3.3 and A-4.2 of [4] provide a definition of GPS "service failures" for the Standard Positioning Service (SPS) as a range error on a satellite flagged as "healthy" exceeding 30 meters when the broadcast URA parameter times 4.42 (the multiplier for a two-sided 10^{-5} probability for the standard Normal distribution) is below 30 meters. It then goes on to say that no more than three such failures are expected per year across the entire GPS constellation and that each such failure state should last no more than 6 hours. Under the assumption of a 24-satellite standard constellation, this results in an estimated per-satellite, per-approach (150 s) probability of a service failure of no greater than:

$$\begin{aligned} \frac{3 \text{ events/yr}}{8766 \text{ hours/yr}} \frac{1}{24 \text{ satellites}} &\cong 1.43 \times 10^{-5} \text{ events/hr} \\ 1.43 \times 10^{-5} \frac{150 \text{ sec/approach}}{3600 \text{ sec/hour}} &\cong 5.94 \times 10^{-7} \text{ events/app} \end{aligned} \tag{2}$$

The primary concern with this result is that the SPS definition of service failures does not include all (or even most) failures of concern to WAAS and LAAS because only failures leading to range errors greater than 30 meters are included. The majority of satellite failures do not reach such large error levels but are included as potential LAAS threats because LAAS integrity is affected by worst-case differential errors as small as 0.5 meters to 1 meter [7]. Even if they are not potentially hazardous, almost all satellite faults will be detected by LAAS monitoring; thus leading to potential loss of continuity.

For these reasons, a conservative prior failure probability of 10^{-4} per satellite per hour (about 7 times the value in (2)) is currently used for satellite faults in LAAS. Furthermore, an additional element of conservatism is added: *each* satellite failure mode is assigned a 10^{-4} probability instead of that probability being divided up among the five satellite failure modes defined by LAAS: clock failure (excess range acceleration), ephemeris failure, signal deformation, low signal power, and code-carrier divergence [6]. This was motivated by a lack of official information as to the proportion of satellite faults in these classes as well as the residual uncertainty regarding the applicability of SPS guidance for faults threatening to WAAS and LAAS. However, it appears to be exceedingly conservative in retrospect.

In an attempt to fill this gap, the Interagency GPS Executive Board and the GPS Joint Program Office (JPO) instituted the Integrity Failure Modes, Effects, and Analysis (IFMEA) project [5]. The goal of this effort is to make use of existing JPO GPS failure mode documents (such as the FMEA documents for the components of the GPS Block IIA and IIR spacecraft) and anomaly observations to provide more information regarding what can go wrong, what the likely effects on ranging signals are, and what the relative event probabilities are. While the results of the IFMEA study are not public yet (to the authors' knowledge), this information should allow us to materially reduce the conservatism in the current GPS satellite prior probabilities.

IV. REVISED APPROACH FOR GPS SATELLITE ANOMALIES

A revised prior probability analysis for GPS satellite anomalies was conducted as part of a recent integrity analysis for a LAAS-variant concept known as Local Area Monitoring (or LAM). In the LAM concept, WAAS corrections are received at a simplified LAAS ground station, transmitted to users via the LAAS VHF data broadcast, and monitored in the position domain relative to the known location of the ground

station [8,9]. Since LAM does not have signal deformation monitoring (SDM, also known as "Signal Quality Monitoring" or SQM – see [10]) capability and must instead rely partly on a coming WAAS SDM upgrade in 2008, the LAM integrity budget would benefit from a lower prior probability than the 10^{-4} per satellite per hour figure cited above.

The prior probability (P_F in this example) that a signal-deformation fault occurs in any 150-second span can be derived empirically to be smaller than 2×10^{-6} (for all visible satellites). This number is based on an empirical analysis of satellite SDM failures in the time that GPS has been active. The derivation of this 2×10^{-6} per 150-second number is based on a conservative extension of the straightforward P_F estimate derived by normalizing the number of observed faults by the total number of observations. For the SDM threat, only a single fault (on SVN-19) has been observed in more than a decade [10]. Over this time, the GPS constellation has continuously maintained at least 24 satellites in orbit. Thus an estimate of the single-satellite fault probability, P_{SS}, is $P_{SS} = 1/M$ where M is the number of 150-second intervals for each satellite observed over a decade:

$$M = 10 \text{ years} \cdot 365 \frac{\text{days}}{\text{year}} \cdot 24 \frac{\text{hour}}{\text{day}} \cdot 24 \frac{\text{observation intervals}}{\text{hour-satellite}} \cdot 24 \text{ satellite}$$
$$= 5.04576 \cdot 10^7 \text{ observation intervals} \qquad [3]$$

The estimated single-satellite fault probability is thus approximately 2×10^{-8}. Transforming this estimate into a prior probability for SDM requires that (1) the single-satellite probability be inflated to account for uncertainty in the above estimation and (2) the single-satellite probability be scaled to account for the presence of multiple satellites in the sky above the LAM ground station.

Evaluating the prior probability at the 95% confidence level provides margin to account for empirical uncertainty and, to some degree, uncertainty regarding future changes to GPS satellites. Since each 150-second window in the decade-long observation period is independent, P_{ss} can be modeled using the binomial distribution (see [11]). Figure 1 illustrates confidence levels for the binary distribution with one observed fault in M observations. The value of P_{ss} at the 95% confidence level is that for which 95% of possible outcomes would result in more than one observed fault. This 95% confidence bound on P_{ss} has a value just under 1×10^{-7}, which five times larger than the best estimate of P_{ss}.

For LAM use, this 95% confidence bound must be further scaled to account for multiple satellites in the sky above the LAM site. This total fault probability for all satellites is called the fault prior, P_F. Based on field data, a reasonable upper bound for the number of visible satellites at one time is 12 satellites. To remain consistent with the LAAS standard [6], however, this analysis conservatively assumes as many as 18 satellites may be visible above the LAM.

$$P_F = 18 P_{ss,95\%} \qquad (4)$$

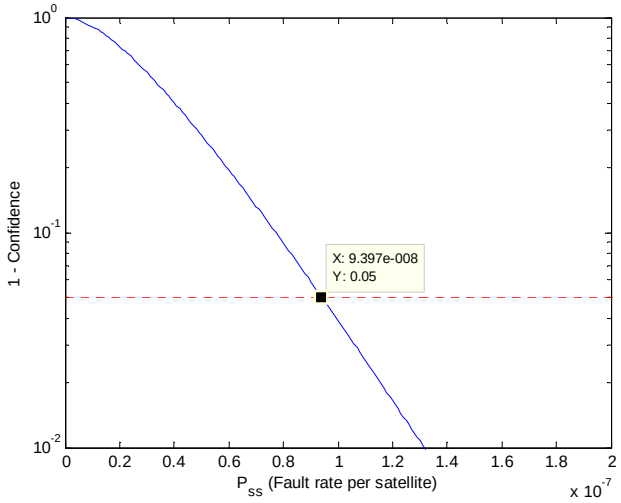

Figure 1: 95% Confidence Bound on P_{ss} Derived Using Binary Distribution CDF

Note that the single-satellite 95% confidence bound on P_F, rounded up to 10^{-7} per SV per 150-second approach, is equivalent to (multiplying by 3600/150 = 24) 2.4×10^{-6} per SV per hour, which is a factor of almost 42 lower than the 10^{-4} value currently assumed by LAAS (!).

This example shows that even a very conservative approach to prior probability estimation based on a single observed event over a long period produces a result significantly below the one currently used for augmentation system integrity verification. Even this value is likely to be much greater than the real value. It is useful to recap the elements of conservatism in the above analysis:

- The observation interval included in (3) is only 10 years, whereas the initial SVN-19 fault observation was made in August 1993 (almost 13 years ago), and no SDM events have been reported since. If we were to count from the announcement of GPS Final Operational Capability (FOC) in 1995, giving a 10-year interval, no SDM events have occurred.

- A 24-satellite constellation is assumed, instead of a more typical average of 27 satellites over the last 10 years.

- A total of 18 possible satellites in view is assumed in (4), which is much greater than the typical value of 7 to 10 satellites visible at a single location.

- The 95th percentile probability from the binomial distribution was used to derive the final estimate. Wide confidence intervals like this are typically used when many observed events occur. When only one event occurs, the 95th percentile is very conservative – it inflates the best-estimate probability by a factor of 5. If a 75th-percentile

probability were used instead, Figure 1 gives a resulting probability of only half of the 1×10^{-7} value derived here.

However, as mentioned in Section I, a greater degree of conservatism is warranted for this threat than for the ionosphere anomaly threat to be discussed later because of the possibility that additional signal-deformation events have occurred but were not noticed before they were corrected. The performance of GPS satellites has been observed by many organizations from time to time, but signal deformation is difficult to detect because it may not appear abnormal on a single receiver. The SVN-19 fault was only detected because of the *differential* errors it created between reference and user receivers of different designs (see [10]). Therefore, a reasonable person could conclude that the conservatism built into the 10^{-7} per satellite per approach result is appropriate, whereas another reasonable person might see it as over-conservative. When many such people are organized into a group of experts charged with producing a single "consensus" integrity justification, it is easier for the group to agree on the more conservative result because, in that case, no one goes away thinking that safety was potentially compromised.

One other aspect of prior-probability conservatism deserves mention here. As shown in Figure 2, the signal deformation failure class is a set of many possible failure behaviors as represented by the signal deformation "threat model", which approximates C/A-code deformations using three parameters describing a 2nd-order-step response to a C/A-code bit transition. Two of these parameters (f_D and σ) represent a 2nd-order analog step response, while the third (Δ) represents a digital time-delay in the response of the deformed signal relative to the "correct" bit transition time (see [10,11] for details). The threat model expresses which values of these three parameters could conceivably result from a signal deformation fault (by inference, faults with parameters outside those allowed by the threat model must be extremely improbable so that their combination to overall integrity risk is negligible even if not detected). The numerical bounds on these parameters were chosen with the SVN-19 example in mind (a 2nd-order-step fit to what is known about that event lies near the center of the analog threat space) based on knowledge of the limits of the signal-generation hardware on GPS satellites and a consensus of expert judgment about where the edges of the possible might lie.

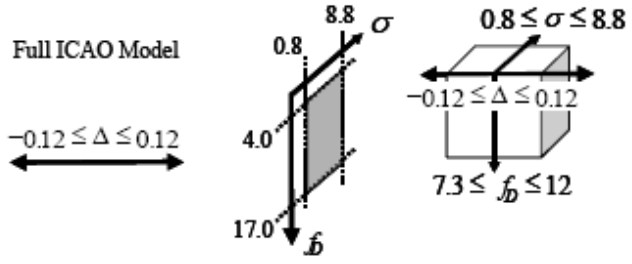

Figure 2: Signal Deformation Threat Model – Parameter Limits

Based on extensive study of how this threat model interacts with typical GPS receivers, the worst-case combination of parameters within the signal deformation threat model (and the resulting worst-case differential range error) can be identified for a given pair of reference and user receivers and a given multi-correlator signal-deformation monitor in the LGF [11]. More generally, the set of points within the threat space that are potentially hazardous to a particular reference – user receiver pair is known. However, no credit is taken in the integrity risk calculation for the possibility that a particular signal-deformation fault is not in this set and is thus not threatening to a particular user. Given a fault within the threat model, the worst-case fault for a given LGF and user pair is assumed with probability 1. Since the worst-case fault parameters change with the user receiver design (the reference receiver design being fixed, this is the most conservative possible assumption. This assumption is addressed further in Section V, but its presence combined with the conservatism in the prior probability assessment appears to result in an extremely conservative model of the signal-deformation threat to WAAS and LAAS users.

V. PROBABILITY MODEL FOR IONOSPHERE ANOMALIES AFFECTING LAAS

The most potentially-threatening anomaly to LAAS is not GPS satellite failure. Instead, it is an anomaly within the ionosphere that can create large differences (i.e., gradients) in GPS range measurements over short baselines. These events were first discovered in 2002 based on post-processed WAAS "Supertruth" data from 6-7 April 2000 [12,13]. The physics of such events is not well understood, but it appears that they are a potential component of extreme ionosphere storms, such as those induced by the enormous solar Coronal Mass Ejection (CME) in October 2003, which led to severe ionosphere gradients observable from CONUS on 29-31 October 2003 and 20 November 2003.

Figure 3 shows a Matlab-generated visualization of the large, westward-moving sharp ionosphere gradient "wave-front" on 20 November 2003 as viewed from NGS CORS reference stations in the Ohio/Michigan region. Figure 4 shows a 2-D view of estimated ionosphere delay vs. time for a subset of 7 CORS stations that saw similar ionosphere delay changes. This event generated the largest spatial gradient observed in a set of ionosphere delay data for all significant ionosphere storms since late 1999 (i.e., covering the last peak of the 11-year solar-magnetic cycle). The largest verifiable (using data from multiple sources) gradient seen in this data was between 310 and 350 mm/km – far greater than the typical one-sigma gradient (under quiet ionosphere conditions) of 1 mm/km [15]. Even with LGF monitoring, differential range errors as large as 3 – 5 meters could have occurred for the worst-located and worst-times LAAS user aircraft if a LAAS-equipped airport were hit by such a storm [13,14].

Research on how to mitigate this rare but potentially dangerous threat to LAAS is ongoing (see [14]). As always, a

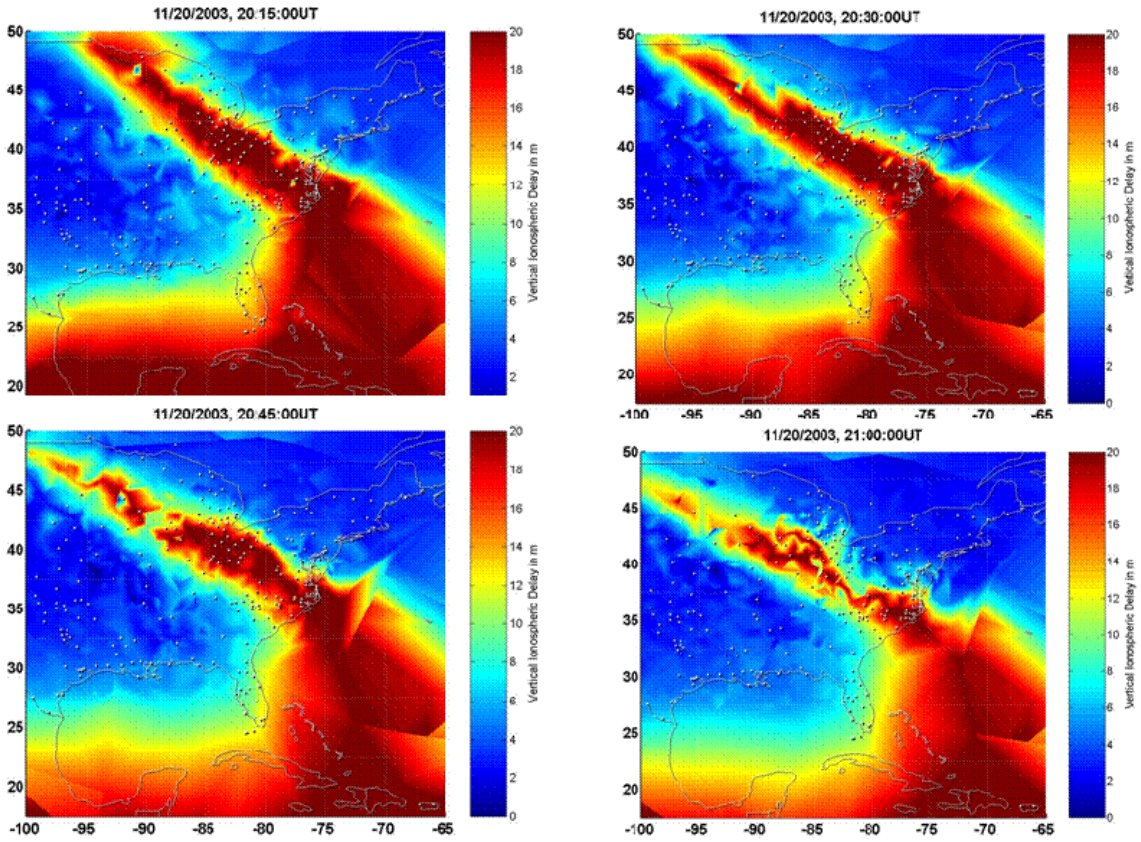

Figure 3: Ionosphere Delay Evolution over CONUS during Peak of 11/20/03 Ionosphere Storm [14]

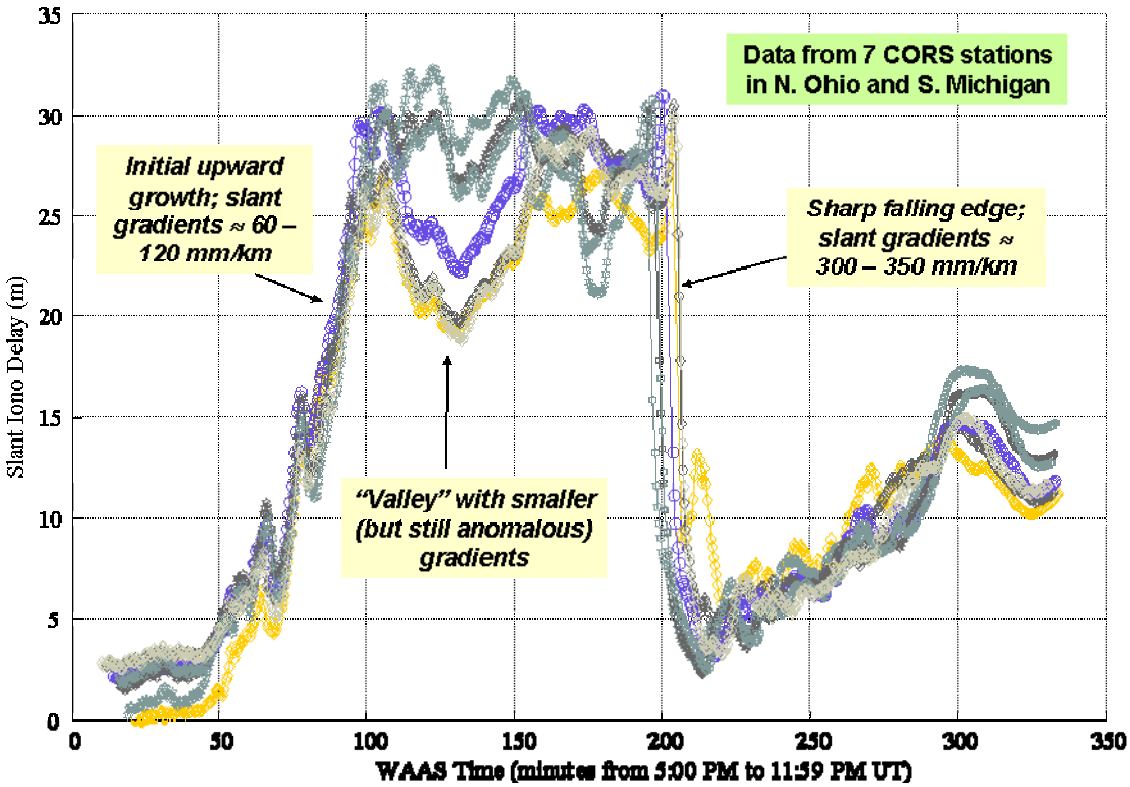

Figure 4: Ionosphere Delay Evolution during 11/20/03 Ionosphere Storm in Northern Ohio and Southern Michigan [14]

TABLE 1
Frequency of Severe Ionosphere Days

	Number of Days in Database	Fraction of Days in Database (2038)	Fraction of Days from NOAA Storm Scale (over 11-year = 4017 day cycle)
Storm Days with Max Kp ≥ 5 ("Minor")	96	0.04711	0.22405
Storm Days with Max Kp ≥ 6 ("Moderate")	81	0.03974	0.08962
Storm Days with Max Kp ≥ 7 ("Major")	65	0.03189	0.03236
Storm Days with Max Kp ≥ 8 ("Severe")	23	0.01129	0.01494
Storm Days with Max Kp ≥ 9 ("Extreme")	9	0.00442	0.00100
Storm Days known to be threatening in CONUS (6 April 2000, 30-31 October 2003, 20 November 2003)	4	0.00196	N/A

key aspect of this analysis is the prior probability that can be assigned to this threat from the point of view of a LAAS-equipped aircraft in CONUS. The starting point for this analysis is the fraction of days within which gradients large enough to threaten LAAS might occur. Table 1 shows this based on the aforementioned database of all significant anomalous ionosphere days in CONUS since October 1999 [16]. Only four such days exist in this database, compared to 9 days where the globally-averaged ionosphere "Kp index" reached the maximum possible value of "9". Table 2 compares the observed fraction of Kp = 9 days from Table 1 (9 / 2038 = 0.0044) to two other estimates derived independently: the one indicated on the NOAA Space Weather Storm Scale [17] and the one derived in the integrity verification material for WAAS [18].

TABLE 2
Comparison of Probabilities of Anomalous Ionosphere Days

	P_{irreg} Model (1932-2000)	NOAA Storm Scale (one solar cycle)	Observed Since October 1999
Kp = 8 ("severe")	0.0026	0.01494	0.01129
Kp = 9 ("extreme")	0.0004	0.0010	0.0044

The comparison in Table 2 between the Table 1 results and two other estimates of the frequency of Kp = "9" days suggests that the value in Table 1 is conservative, which is not too surprising since the data used to generate Table 1 comes from the "more severe" half of the 11-year solar cycle (ionosphere storms are more likely in the years following the last solar peak than in the years approaching the next solar peak). The comparison in Table 2 supports the idea that the results in Table 1 are sufficiently conservative. Therefore, the frequency of potentially threatening days to LAAS in Table 1 (4 out of 2038) can be used as a conservative estimate of the probability of having a sufficiently threatening ionosphere storm on any given day. Instead of using this ratio (4 / 2038 = 0.00196) directly, the 60% upper confidence limit from the binomial distribution used for signal deformation in Section IV is used instead. A direct means of computing this probability is given by the probability $1 - L(x)_\alpha$ at the 40^{th}-percentile level (i.e., $\alpha = 0.4$), where $L(x)_\alpha$ is given by [20]:

$$L(x) = \frac{x}{x + \left[(n - x + 1) F_\alpha (2n - 2x + 2, 2x)\right]} \quad (5)$$

where n is the number of samples (2038 in this case), s is the number of observed faults (4 in this case), and $x = n - s$ (2034 in this case). "F_α" refers to the statistical F-distribution at the probability given by α. This equation gives the same results as a direct application of the binomial probability equation and was used in the past when evaluations of binomial probabilities were much more time-consuming than table lookups of the F-distribution.

For the numbers in this case, the resulting 40^{th}-percentile lower-bound probability (which is the same as the 60^{th}-percentile upper-bound probability) of a threatening storm on a given day is 0.00257, or 31% higher than the mean estimate of 0.00196. A 60^{th}-percentile upper-bound is used here instead of the more-conservative 95^{th}-percentile bound for signal deformation in Section IV because of the greater quantity and quality of data for ionosphere storms (as opposed to only one observation of serious signal deformation) as well as the fact that the mean result from Table 1 is already conservative when compared to Table 2 (see [16]).

Table 3 uses the conservative probability of 0.00257 for a storm on a given day as the starting point for determining the probability of a threatening storm affecting the LAAS-supported precision approach of an individual aircraft at a single airport. Because severe gradient wave fronts created by ionosphere storms almost always move quickly (at 90 m/s or more – faster than an approaching aircraft), a single airport will only be affected briefly (see [14,21]). For all such storms in the database represented by Table 1, a given location was affected by a potentially-threatening storm only once per threatening storm day. As shown in Figure 4, the 20 November 2003 storm in the Ohio / Michigan region created two large gradients at observing CORS sites: an earlier one with ionosphere delay rapidly rising and a sharper later one with ionosphere delay falling dramatically in a few minutes. Both the rising and falling events were clearly anomalous, but only the later "falling" one created a potential hazard to LAAS (see [14,23]). Despite this, the model in Table 3 assumes two potentially hazardous gradients per threatening storm day on two different GPS satellites (after two satellites are impacted by a threatening anomaly, LGF executive monitoring will notice detected anomalies on two separate satellites and will stop broadcasting differential corrections until the anomaly disappears [22]).

The good news is that the worst impact of a given gradient passing by will only be experienced by one approaching

aircraft (at most) because the threatening impact is limited to one 150-second approach interval. Before this worst-case interval, the differential range error has not grown enough to be threatening, and after it, LGF detection and exclusion has almost certainly occurred (see [23]). The result is that only four separate 150-second periods over one 86,400-second day could be threatening, which gives a probability of hazardous impact at a given airport (given a threatening storm day) of $4 \times 150 / 86400 = 0.006944$. Multiplying this number by the probability of a threatening day reduces the ionosphere threat prior probability for a given aircraft approach to about 1.785×10^{-5}. An additional factor-of-5 probability reduction is taken credit for near the bottom of Table 3 because the threatening time interval for the worst-case approach is not the entire 150-second approach but instead the worst 30-second time slot within that approach [23]).

TABLE 3
Ionosphere Gradient Threat Probability for a Single Aircraft Approach

Probability of Threatening Storm Day (60th pct)	**0.00257**
Prob. over 1 day that specific CONUS airport affected	**1.7847E-05**
(for a given airport, only 2 * 2 = 4 approach periods per day could be threatened): Pr ~ 150 * 4 / 8640(0.006944	
Probability of Worst-Case Approach Direction (1)	**1.7847E-05**
(1/6 = 60/360 for a given approach, but assume many approaches, at least one of which will have worst-case direction)	
Probability of Worst-Case Timing for a given aircraft (0.2)	**3.5694E-06**
(1/5 = 30 / 150 second approach)	
Probability of No Early LGF (i.e. Precursor) Detection (0.2)	**7.1389E-07**
(conservative precursor credit based on > 80% data rejection during iono. anomalies)	

Two additional mitigating conditions are included in Table 3. One is the probability of the worst-case approach direction. No credit is taken for this (the assumed probability of the worst direction is 1) because a LAAS site at a busy airport such as Chicago/O'Hare will simultaneously support approaches to multiple runway ends and multiple approach directions. If only one direction were being supported, a factor-of-6 probability credit would be justified because the worst-case approach direction is within ± 30 degrees of the direction the aircraft is approaching from (this maximizes the growth of differential error before the LGF is affected and is able to detect and exclude the affected satellite(s)). The other mitigating condition is the likely occurrence of "precursor" events, such as ionosphere scintillation impacts on measurement amplitude and phase, prior to the arrival of a severe ionosphere gradient [24]. LAAS Test Prototype (LTP) data collected by the FAA Technical Center supports our expectation that severe gradients are unlikely to "sneak up" upon a LAAS-equipped airport without the appearance of other symptoms of anomalous behavior [21]. However, only a handful of LTP observations during ionosphere storms exist, and precursors are hard to identify in CORS station data (with updates only every 30 seconds), so we have chosen to take only a factor-of-5 credit for precursors (our best guess would be at least a factor-of-10 risk reduction).

Taken together, the mitigating factors in Table 3 result in a probability of hazardous ionosphere gradient of 7.14×10^{-7} per approach. Despite the conservatism already introduced into the conditional probabilities in Tables 1 and 3, further conservatism is likely to be introduced by "rounding up" to 10^{-6} per approach (thus adding about 40% margin) before this number is used in LAAS integrity risk assessment. Despite this, concern has been raised regarding taking credit for the probability of a specific approach being affected during an ionosphere storm day. Without this credit, the resulting prior probability (without the "round-up" to 10^{-6}) would rise to $7.14 \times 10^{-7} / 0.006944 = 0.0001028$ per approach, which is ridiculously conservative if taken at face value. The rationale for such a change is discussed in Section VI on "specific risk", but the resulting probability is too large and too misleading to be of any use in assessing integrity risk.

We should also note that, as with signal deformation, no credit is taken for the fact that only a small segment of the total ionosphere-spatial-anomaly threat space would actually be hazardous to users. While an argument could be made for this in the case of signal deformation, where only one not-entirely-understood event has taken place, Figure 5 shows several dozen validated ionosphere anomalies (for satellites above 12°) elevation from the dataset encompassed by Table 1 in terms of two of the threat model parameters: (linear) gradient slope (in mm/km) and wave front propagation speed (in m/s) [25]. Since the observed points span the majority of the threat space that will eventually be selected, strong justification exists for treating each point in the eventual threat space as equally likely to occur given that an ionosphere anomaly has occurred. If this were done, less than 10% of the resulting threat space would actually be hazardous to a worst-case LAAS user; thus at least another factor-of-10 reduction in the hazardous-anomaly prior probability would result. Even

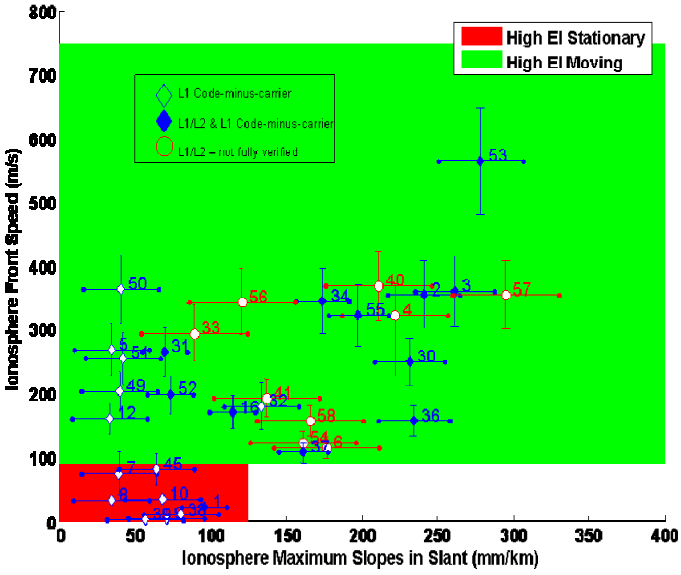

Figure 5: Validated Severe Ionosphere Spatial Anomalies from 1999 – 2005 CORS Data [25]

this model is likely to be conservative because the study of ionosphere anomalies in [13,14,23] focused on the largest gradients rather than smaller but still anomalous gradients in the range of 30 – 100 mm/km. If a valid random sample of anomalous days on the CORS dataset covered by Table 1 were taken, smaller gradients would likely appear much more often than larger (and more-threatening) gradients, which would indicate that a 2-D uniform distribution within the threat model is conservative than reality.

In any case, probability credit for "averaging" over the parameters within the ionosphere anomaly threat model has not been seriously considered. Instead, as with satellite faults, the worst point in the threat model is still assumed to have a probability of one given that a potentially-hazardous ionosphere anomaly has occurred. The resulting conservatism leads to over-conservative risk estimates and over-design of the monitors needed to mitigate these risks. This aspect of prior-probability modeling should be carefully reconsidered in the future.

VI. "SPECIFIC" VS. "AVERAGE" RISK CONSIDERATIONS

The term "specific risk" has been introduced into civil aviation certification and safety assessment by the FAA over the past several years. This term is meant to distinguish itself from "average" or "ensemble" risk, in which probabilities of anomalies that affect the overall population of aircraft flights are averaged together to obtain "mean" safety risk in such a way that some flights are allowed exceed the maximum-risk requirement because other flights are below it and thus help "bring the average down" to an acceptable level. "Specific risk", in contrast, attempts to insure that *every* flight allowed to proceed meets the maximum-risk requirement. Some flights will still have lower risk than others, but no flight would be allowed to exceed the maximum acceptable safety risk just because others are well below it (see [19] for additional details).

One consequence of this approach is that "latent" risks, or off-nominal conditions that occur more often than the allowed safety risk and may persist because they are below detectable levels, must be treated as "nominal" (i.e., treated as always present with probability 1) if such conditions could be detected with additional investment in risk mitigation. As a hypothetical example, if a given aircraft component needed for safe flight were vulnerable to malfunction or shutdown when exposed to external (ambient) temperatures exceeding 40^{o} Celsius, it would not be acceptable to take credit for (i.e., "average over") the relatively low probability of temperatures exceeding 40^{o} over all times and locations. Instead, the relevant system and/or an external FAA service would need to measure the temperature and alert aircraft when the temperature approaches 40^{o} so that no aircraft is exposed to unacceptable risk. Measuring external temperature is standard in almost all existing aircraft and is straightforward, so there would be no tolerance for allowing this risk to be averaged over. Also note that there would be no tolerance for allowing significantly increased risk in very warm locations (such as Phoenix and Las Vegas) just because most locations in CONUS almost never exceed 40^{o} Celsius.

While the intent of "specific risk" is laudable and has broad support in the navigation community, extreme conservatism can result if insufficient tolerance is given to averaging over conditions that cannot easily be observed or foreseen. As cited in Section V, averaging over the probability of an ionosphere storm affecting a particular airport given a threatening storm day is a good example. Given that LAAS sites, like the Instrument Landing System (ILS) transmitters that they will replace, do not communicate with each other, there is no means for LAAS sites that appear to have been affected by a severe gradient to warn others "downstream" of the gradient's apparent motion to take preventive measures (one such measure would be raising the broadcast sigmas to make marginal satellite geometries unusable). A communication link could be added to each LAAS site for this purpose, but only at great expense and schedule delay (reliance on broadcast SBAS Grid Ionosphere Vertical Error, or GIVE, values is a more feasible option, but only in locations with good SBAS coverage [9]). Without such a major system architecture change, it seems reasonable to average over this probability – the resulting risk is truly rare, random, and impacts all locations approximately the same. Not averaging over this risk presents an unduly pessimistic picture of the risk to LAAS that it poses, and decisions made based on this picture would almost certainly assure that the benefit-to-cost ratio for LAAS would always be well below unity.

LAAS is not the only system that would have severe cost-benefit problems with a literal, inflexible interpretation of "specific risk". Manufacturers of other systems that have already been certified, such as aircraft engines, have raised similar concerns with the FAA [26]. As a result, a study has been commissioned into the precise definition and application of "specific risk" (see [27]). The outcome of this study is of significant importance to the design and integrity verification of both present and future augmented GNSS systems.

VII. SUMMARY AND CONCLUSIONS

This paper illustrates how prior probabilities for anomalous conditions affecting augmented GPS users have been calculated. In all such calculations, a significant degree of conservatism is included to account for statistical uncertainty in the collected data, uncertainty in interpreting the collected data, uncertainty in the models created to relate anomaly occurrence to LAAS or WAAS threat impact, and differing opinions among experts as to each of these factors. This level of conservatism is needed because GNSS and augmentation are too new to provide multiple examples of each possible failure mode. Because no comprehensive rule exists that can be confidently applied to all possible anomalies, expert judgment is needed to find the appropriate level of conservatism for each anomaly.

While significant conservatism is required in assessing prior probabilities of GNSS anomalies at present, additional conservatism comes from the reluctance to take credit for the conditional probability of worst-case anomalies as a fraction of all possible anomaly conditions. In cases where sufficient data exists to demonstrate that the worst-case anomaly is not the only one that can occur, it would be wiser to take at least some credit for the fact that, for most anomaly classes, most specific examples of these anomalies pose a lesser threat than the worst possible example. Separately, the "specific risk" interpretation of when credit may be taken for uncertain events can, if applied inflexibly, also lead to extremely conservative prior probability assessments. From a safety standpoint, "more" conservatism is preferable to "less", but the practical consequence of excess conservatism is excess expenditure, delay in system commissioning, and lowered system availability resulting from over-conservatism in the monitor algorithms needed to mitigate the worst-case threat.

VIII. ACKNOWLEDGMENTS

The authors would like to thank many people inside and outside of Stanford for helping us with this work in the past or present: Ming Luo, Jiyun Lee, Godwin Zhang, Hiro Konno, Youngshin Park, Todd Walter, and Eric Phelts at Stanford as well as Boris Pervan (Illinois Inst. Of Technology), Barbara Clark (FAA), Bruce DeCleene (FAA), Arun Murthi (FAA), Rick Cassell (Rannoch), Ron Braff (MITRE – now retired), JP Fernow (MITRE), and Curt Shively (MITRE). The authors also gratefully acknowledge the FAA Satellite Navigation LAAS Program Office for supporting this research. However, the views expressed in this paper belong to the authors alone and do not necessarily represent the position of any other organization or person.

IX. REFERENCES

Note: *Hyperlinks split across multiple lines have excess spaces in them at the points of splitting – please remove these spaces when following the links.*

[1] *Minimum Aviation System Performance Standards for Local Area Augmentation System (LAAS)*. Washington, D.C., RTCA SC-159, WG-4, DO-245A, Dec. 9, 2004. http://www.rtca.org

[2] S. Pullen, T. Walter, and P. Enge, "System Overview, Recent Developments, and Future Outlook for WAAS and LAAS," *Proceedings of 2002 Japan GNSS Symposium*, Tokyo, Japan, Nov. 11-13, 2002. http://waas.stanford.edu/~wwu/papers/gps/PDF/PullenTokyo02.pdf

[3] J. Rife and R.E. Phelts, "Time-Varying MERR/Conditional Risk Analysis for LAAS," Stanford University, GPS Laboratory Group Meeting, April 14, 2006. http://waas.stanford.edu/research/laas.htm

[4] *Global Positioning System Standard Positioning Service Performance Standard*. Washington, D.C., U.S. Department of Defense, October 2001. Internet URL: http://www.navcen.uscg.gov/gps/geninfo/2001SPSPerformanceStandardFINAL.pdf

[5] K. Van Dyke, K. Kovach, J. Lavrakas, and B. Carroll. "Status Update on GPS Integrity Failure Modes and Effects Analysis," *Proceedings of ION National Technical Meeting (NTM) 2004*, San Diego, CA., Jan. 26-28, 2004, pp. 92-102. http://www.ion.org

[6] *Specification: Performance Type One Local Area Augmentation System Ground Facility*. U.S. Federal Aviation Administration, Washington, D.C., FAA-E-2937A, April 17, 2002. http://gps.faa.gov/Library/Data/LAAS/LGF2937A.PDF

[7] T. Zaugg, "A New Evaluation of Maximum Allowable Errors and Missed Detection Probabilities for LAAS Ranging Source Monitors," *Proceedings of the ION 58th Annual Meeting*, Albuquerque, NM, June 24-26, 2002, pp. 187-194. http://www.ion.org

[8] J. Rife, S. Pullen, T. Walter, P. Enge, "Vertical Protection Levels for a Local Airport Monitor for WAAS," *Proceedings of the ION 61st Annual Meeting*, Cambridge, MA., June 27-29, 2005. http://waas.stanford.edu/~wwu/papers/gps/PDF/RifeIONAM05Monitor.pdf

[9] *Algorithm Description Document for Local Area Monitor*. U.S. Federal Aviation Administration, Washington, D.C., Unpublished Draft, March 3, 2006. http://gps.faa.gov

[10] A. Mitelman, *Signal Quality Monitoring for GPS Augmentation Systems*. Ph.D. Dissertation, Stanford University, Dept. of Aeronautics and Astronautics, Dec. 2004. http://waas.stanford.edu/~wwu/papers/gps/PDF/Thesis/AlexanderMitelmanThesis04.pdf

[11] R.E. Phelts, A. Mitelman, S. Pullen, D. Akos, and P. Enge, "Transient Performance Analysis of a Multicorrelator Signal Quality Monitor," *Proceedings of ION GPS 2001*, Salt Lake City, UT., Sept. 11-14, 2001, pp. 1700-1710. http://waas.stanford.edu/~wwu/papers/gps/PDF/ericion01.pdf

[12] S. Datta-Barua, T. Walter, S. Pullen, M. Luo, J. Blanch, and P. Enge, "Using WAAS Ionospheric Data to Estimate LAAS Short Baseline Gradients," *Proceedings of ION 2002 National Technical Meeting*, Anaheim, CA, January 28-30, 2002, pp. 523-530. http://waas.stanford.edu/%7Ewwu/papers/gps/PDF/DattaBaruaIONNTM02.pdf

[13] M. Luo, S. Pullen, D. Akos, G. Xie, S. Datta-Barua, T. Walter, and P. Enge, "Assessment of Ionospheric Impact on LAAS Using WAAS Supertruth Data", *Proceedings of The ION 58th Annual Meeting*, Albuquerque, NM, June 24-26, 2002, pp. 175-186. http://waas.stanford.edu/~wwu/papers/gps/PDF/LuoIONAM02.pdf

[14] M. Luo, S. Pullen, S. Datta-Barua, G. Zhang, T. Walter, and P. Enge, "LAAS Study of Slow-Moving Ionosphere Anomalies and Their Potential Impacts," *Proceedings of ION GNSS 2005*, Long Beach, CA., Sept. 13-16, 2005. http://waas.stanford.edu/~wwu/papers/gps/PDF/LuoIONGNSS05.ppt

[15] J. Lee, S. Pullen, S. Datta-Barua, and P. Enge, " Assessment of Nominal Ionosphere Spatial Decorrelation for LAAS," *Proceedings of IEEE/ION PLANS 2006*, Coronado, CA., April 24-27, 2006 (forthcoming). http://waas.stanford.edu/~wwu/papers/gps/PDF/LeeIONPLANS06.pdf

[16] S. Pullen, "LAAS Ionosphere Anomaly Prior Probability Model: Version 3.0," Stanford University, GPS Laboratory, Oct. 14, 2005. http://www-leland.stanford.edu/~spullen/Iono_PProb_Model_v3-1.ppt

[17] *NOAA Space Weather Scales* website: http://www.sec.noaa.gov/NOAAscales/

[18] *Algorithm Contribution to HMI for the Wide Area Augmentation System*. The Raytheon Company, Fullerton, CA., CRDL Sequence No. A014-011, Sectiom A.29, Sept. 30, 2002 (unpublished work). http://gps.faa.gov

[19] T. Walter, P. Enge, and B. DeCleene, "Integrity Lessons from the WAAS Integrity Performance Panel (WIPP)," *Proceedings of ION 2003 National Technical Meeting*, Anaheim, CA., Jan. 22-24, 2003, pp. 183-194. http://waas.stanford.edu/~wwu/papers/gps/PDF/WalterIONNTM03.pdf

[20] H.F. Martz and R.A. Waller, *Bayesian Reliability Analysis*. New York: John Wiley and Sons, Inc., 1982. http://www.amazon.com/gp/product/0471864250/ref=ed_oe_h/103-4040504-1160621?%5Fencoding=UTF8

[21] T. Dehel, F. Lorge, J. Warburton, and D. Nelthropp, "Satellite Navigation vs. the Ionosphere: Where Are We, and Where Are We Going?," *Proceedings of ION GNSS 2004*, Long Beach, CA., Sept. 21-24, 2004. http://www.ion.org

[22] G. Xie, S. Pullen, M. Luo, P.L. Normark, D. Akos, J. Lee, P. Enge, and B. Pervan, "Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed," *Proceedings of ION 57th Annual Meeting and CIGTF 20th Guidance Test Symposium*, Albuquerque, NM., June 11-13 2001, pp. 681-693. http://waas.stanford.edu/~wwu/papers/gps/PDF/xieionam01.pdf

[23] M. Luo, S. Pullen, A. Ene, D. Qiu, T. Walter, and P. Enge, "Ionosphere Threat to LAAS: Updated Model, User Impact, and Mitigations," *Proceedings of ION GNSS 2004*, Long Beach, CA., Sept. 21-24, 2004. http://waas.stanford.edu/~wwu/papers/gps/PDF/LuoIONGNSS04.pdf

[24] S. Datta-Barua, P. Doherty, S. Delay, T. Dehel, and J. Klobuchar, "Ionospheric Scintillation Effects on Single and Dual Frequency GPS Positioning," *Proceedings of ION GPS/GNSS 2003*, Portland, OR., Sept. 9-12, 2003, pp. 336-346. http://waas.stanford.edu/%7Ewwu/papers/gps/PDF/DattaBaruaIONGPS03.pdf

[25] S. Pullen, M. Luo, J. Lee, G. Zhang, and P. Enge, "Update on Mitigation of Ionosphere Spatial Anomaly Threat to GBAS," RTCA SC-159, WG-4 Meeting, Washington, D.C., March 22-23, 2006. http://sc159.tc.faa.gov/wg4/032506/2006-03_SC-159_WG4_Working_Papers_DCA.zip (inside Zip file)

[26] "Aviation Rulemaking Advisory Committee (ARAC) Transport Airplane and Engine (TAE) Issues Area: Meeting Minutes," Tukwila, WA., The Boeing Company, Oct. 19, 2005. http://www.faa.gov/regulations_policies/rulemaking/committees/arac/minutes/media/TAE_OCT_05.pdf

[27] "Dept. of Transportation / Federal Aviation Administration: Aviation Rulemaking Advisory Committee; Transport Airplane and Engine Issue Area—New Task," *The U.S. Federal Register*, Vol. 71, No. 54, Notices for Tuesday, March 21, 2006, pp. 14284-14286. http://frwebgate3.access.gpo.gov/cgi-bin/waisgate.cgi?WAISdocID=1419106375+0+0+0&WAISaction=retrieve