Using Outage History to Exclude High-Risk Satellites from GBAS Corrections

Sam Pullen and Per Enge
Stanford University

ABSTRACT

GNSS augmentation systems that provide integrity guarantees to users typically assume that all GNSS satellites have the same failure probability. The assumed failure probability is conservative such that variations among satellites in a given GNSS constellation are not expected to violate this assumption. A study of unscheduled GPS satellite outages from 1999 to present shows that, as expected, older satellites are much more likely to fail than younger ones. In addition, satellites that have recently experienced unscheduled outages are more likely to suffer additional unscheduled outages. Combining these two factors suggests that it is possible for a subset of GPS satellites to violate the overall satellite failure probability assumption, although this has not yet been demonstrated. Potential rules for GPS satellite exclusion based upon satellite age and recent outages are investigated, and suggestions for including satellite geometry are explored.

1.0 INTRODUCTION

GNSS applications with demanding requirements for real-time integrity verification must make a series of assumptions regarding the performance of the satellite constellation(s) that they are using. One key assumption is the probability of unexpected satellite outages or failures. Integrity monitors that operate directly on standalone (uncorrected) GNSS measurements, such as Receiver Autonomous Integrity Monitoring (RAIM), as well as systems that provide differential corrections such as Space Based and Ground Based Augmentation Systems (SBAS and GBAS, respectively), rely on this assumption to determine the false-alert and missed-detection probabilities that their integrity monitor algorithms must achieve.

Regarding satellite failures, two different probabilities are important. One is the probability of any unexpected satellite outage, which makes the affected satellite unusable and thus affects continuity. The other is the probability of events that pose a potential integrity risk to SBAS and GBAS users. These probabilities can be represented as rates (e.g., probability of outage per

satellite per hour) or as state probabilities, meaning the long-term average probability that a given satellite is in an "outage" or "failed" state.

For the purposes of verifying that integrity and continuity requirements are met, the systems mentioned above assume that all satellites have the same probability of outage or integrity failure. This assumption is made for simplicity, and the probabilities assumed are typically very conservative; thus little or no risk arises due to potential violations of this assumption. However, it is known from the history and planning of GPS satellite operations (and satellite operations in general) that older satellites are much more likely to fail than younger ones. This was captured earlier in the history of the GPS constellation by former GPS Joint Program Office director, Col. Gaylord Green (USAF, Ret.), who noted that "*GPS satellites are operated to failure.*" [1] By this, he meant that GPS satellites were not retired when they first began experiencing problems or approaching the end of their expected useful life but instead when they failed in a manner that was not recoverable or was recoverable but no longer maintainable. This means that older satellites, and those which have recently experienced outages, will generally keep being used despite their higher propensity for further failures.

This paper examines the degree to which unexpected GPS satellite outages and failures vary with satellite age and prior outage history. It uses the archive of GPS Non-Availability notices to NAVSTAR Users (NANUs) to compile a history of unexpected satellite outages from January 1999 to August 2011 [2,3]. These outages are examined to identify the likelihood of outages as a function of satellite age and recent outage history. These results show that, as expected, older satellites are much more likely to experience outages than younger ones, as are satellites with a history of recent outages (in the last year or two).

While these results do not immediately suggest that the satellite-fault-probability assumptions made by GBAS and other systems are violated for specific satellites, they at least raise the possibility. To address this risk, two potential satellite-exclusion heuristics are examined in

terms of their impact on GPS user satellite geometry quality. The need to select a subset of satellites from the full set that is in view is not new: it was a common need of 4 and 6-channel GPS receivers in the early 1990's, when GPS rapidly expanded to a full constellation of 24 satellites by 1993. Since more than 6 satellites were commonly in view, some means was needed to down-select the 4 or 6 satellites that were most worth tracking. Various methods were applied to do this, including differentiation by satellite quality [4] and selection of the subset that gives the best satellite geometry in terms of dilution of precision (DOP) [5]. Computationally efficient methods for computing DOP and thus finding optimal satellite subsets are still important [6].

Past and present experience can be combined to infer two general cases where satellite sub-selection is desirable. In the *channel-limited scenario*, not all satellites can be used due to hardware limitations. This includes both receiver limitations and data bandwidth constraints for augmentation systems like GBAS, which can only send differential corrections and integrity information for a limited number of satellites. In the *channel-unlimited scenario*, hardware constraints do not apply, but satellites are down-selected to remove those with poor performance or other negative traits. The former scenario is primarily a concern to future systems that make use of GPS in combination with other GNSS constellations such as GLONASS, Galileo, Compass, and/or QZSS. This paper will mostly consider the latter scenario with the intent of excluding satellites whose fault likelihoods might violate the integrity assumptions made by GBAS or other high-integrity services.

Section 2.0 of this paper describes the expected GPS satellite fault probabilities and the more conservative numbers assumed by GBAS. Section 3.0 provides the results of the unexpected satellite outage study, showing the degree to which satellite age and number of recent outages affect the likelihood of future outages. Section 4.0 uses this information to propose example satellite-exclusion heuristics and examine their impacts on GPS-only satellite geometries as measured by Vertical DOP at two U.S. user locations. Section 5.0 explains how the multiple-hypothesis protection level approach utilized in Advanced RAIM (ARAIM) can be used to make real-time trade-offs between the integrity risk posed by a weak satellite and the satellite geometry benefit that it provides. Section 6.0 summarizes this paper and briefly examines the impact that the future use of multiple satellite constellations may have on satellite exclusion.

2.0 Expected GPS Satellite Fault Probabilities

The primary open source of GPS satellite outage information is the latest (2008) *GPS Standard Positioning Service (SPS) Performance Standard* document [7],

which gives both historical information on satellite outage rates and states the minimum performance requirements that the U.S. Government, the GPS Wing, and the 2nd Space Operations Squadron (2 SOPS) hold themselves to in managing and maintaining the system.

Section 3.5.1 of [7] expresses the SPS Signal-in-Space (SIS) Integrity Standard as a probability of 10^{-5} per hour or less that the SIS User Range Error (URE) exceeds ±4.42 times the upper bound on User Range Accuracy (URA) corresponding to the URA integer broadcast by the satellite in question without a warning that prevents use of the affected measurement (this event is called a "major service failure"). Over a 32-satellite constellation, which is the maximum number of satellites that can be supported in the broadcast almanac, the above probability implies an average of 3 "major service failures" per year [7]. This is confirmed by doing the reverse calculation as follows:

$$\frac{3\,\text{events/year}}{8766\,\text{hours/year}} \cdot \frac{1}{32\,\text{satellites}} = \boxed{1.07 \times 10^{-5}\,\text{events/SV/hour}} \quad (1)$$

When applying this number to a civil system such as GBAS, it is important to note that "major service failures" for SPS do not include all events that represent potential integrity threats to GBAS. GBAS separates potential integrity threats due to satellite faults into five classes [8]:

- Clock failures (excessive acceleration)

- Low signal power

- Code-carrier divergence (separate from that caused by the ionosphere)

- C/A-code signal deformation

- Navigation data failures (e.g., large ephemeris errors)

Because the "major service failure" definition does not cover these, and because the numbers cited in [7] cannot be taken as guaranteed, GBAS assumes a satellite integrity fault probability of 10^{-4} per satellite per hour instead of 10^{-5}. Furthermore, this probability is assumed for *each* of the above five failure classes instead of all five combined [8]. This is very conservative for fault classes that are known to have been observed once at most (e.g., signal deformation and code-carrier divergence). On the other hand, the same probability is assumed for all usable GPS satellites regardless of their age or current status.

Section 3.6.1 of [7] specifies a somewhat-different satellite outage probability related to continuity of service. This is stated as a probability of 0.9998 or greater over any hour of "… Not Losing the SIS SPS Availability from a Slot Due to Unscheduled Interruption." In other words, given that a satellite (occupying an orbit slot) is healthy

and broadcasting usable signals at the start of a given hour-long period, the probability of losing usable signals from that slot (or satellite) is 0.0002 or less. This translates into a lower bound on the Mean Time Between Outages (MTBO) of 5000 hours, which is lower than the numbers assumed by GBAS based on the previously-assessed operational history of GPS satellites (e.g., 9740 hours in [9]). In general, the CAT I GBAS continuity sub-allocation to satellite failures is conservative because it allows for many more "critical satellites" than normally occur in available satellite geometries (a "critical satellite" is one whose sudden loss would lead to loss of continuity for the current operation – see [10]).

The probability given in Section 3.6.1 of [7] is directly related to the results in this paper because unexpected outages tabulated from NANUs archives cannot be separated into "outages affecting only continuity" and "outages that potentially threaten integrity" without more details. In almost all cases, the set of events that creates integrity risk is a subset of the events that cause outages. In the case of GPS, the Operational Control Segment (OCS) detects and excludes (flags as "unhealthy" or triggers a switch to an unusable PRN or to "non-standard C/A code") practically all faults of significance within a few hours [7], although its response is not fast enough to meet the time-to-alert for most applications. The fraction of unexpected outages that pose a potential integrity threat is hard to estimate. It is likely to be much less than 1.0, although a significant fraction of unexpected outages appear to occur from satellite clock anomalies, which can threaten civil user integrity if not quickly detected.

3.0 Unscheduled Outages of GPS Satellites Since 1999

As noted above, the data source for this study of unscheduled (or unexpected) GPS satellite outages is the archive of NANUs that were issued at or near the time that the outages occurred. Most outages are "scheduled" for orbit and clock maintenance and are indicated by a forecast NANU (to alert users of the upcoming outage) followed by a report of the outage after it has been completed. Unscheduled outages, of course, are not preceded by any forecast. Two NANUs typically exist for them as well – one alerting that an unscheduled outage has begun and a report on its duration after it is over and the satellite has been returned to service or has been decommissioned. Although the NANU archives cannot be regarded as complete or definitive outage records, they are sufficient to explore the degree to which satellite outages are equally spread across the constellation.

Table 1 lists the 31 GPS satellites that were commissioned and healthy as of (approximately) 20 September 2011. The rows are sorted by age, and the first two rows are highlighted to indicate that two satellites (Block IIA SVNs 23 and 24) are now more than 20 years

Table 1: Healthy GPS Satellites as of 20 Sept. 2011

SVN	PRN	Block	Launch Year	Launch J-Day	Orbit Slot (as of 9/20/11)	Age as of 9/20/11 (years)
23	32	IIA	1990	330	E5	20.81
24	24	IIA	1991	185	D5	20.21
26	26	IIA	1992	188	F5	19.20
39	09	IIA	1993	177	A1	18.23
35	30	IIA	1993	242	B5	18.05
34	04	IIA	1993	299	D4	17.90
36	06	IIA	1994	069	C6	17.53
33	03	IIA	1996	088	C2	15.48
40	10	IIA	1996	198	E6	15.18
43	13	IIR	1997	204	F3	14.16
38	08	IIA	1997	310	A3	13.87
46	11	IIR	1999	280	D2	11.95
51	20	IIR	2000	131	E1	11.36
44	28	IIR	2000	197	B3	11.18
41	14	IIR	2000	314	F1	10.86
54	18	IIR	2001	030	E4	10.64
56	16	IIR	2003	029	B1-A	8.64
45	21	IIR	2003	090	D3	8.47
47	22	IIR	2003	355	E2	7.75
59	19	IIR	2004	080	C3	7.50
60	23	IIR	2004	175	F4	7.24
61	02	IIR	2004	311	D1	6.87
53	17	IIR-M	2005	269	C4	5.98
52	31	IIR-M	2006	268	A2	4.98
58	12	IIR-M	2006	321	B4	4.84
55	15	IIR-M	2007	290	F2-A	3.92
57	29	IIR-M	2007	354	C1	3.75
48	07	IIR-M	2008	075	A6	3.51
50	05	IIR-M	2009	076	E3	2.51
62	25	IIF	2010	148	B2	1.31
63	01	IIF	2011	197	D2-A	0.18

old. This is quite remarkable considering that the Block IIA satellites had a design life of only 7.5 years, and it illustrates the truth of Col. Green's claim in the Introduction. The two highlighted satellites (among others now retired) were not only "operated to failure," but they were brought back into active service after previously being decommissioned (and retired) due to end-of-life failures of other satellites.

Figure 1 plots the ages of the satellites in Table 1 from newest to oldest (in blue) and shows a linear fit to these
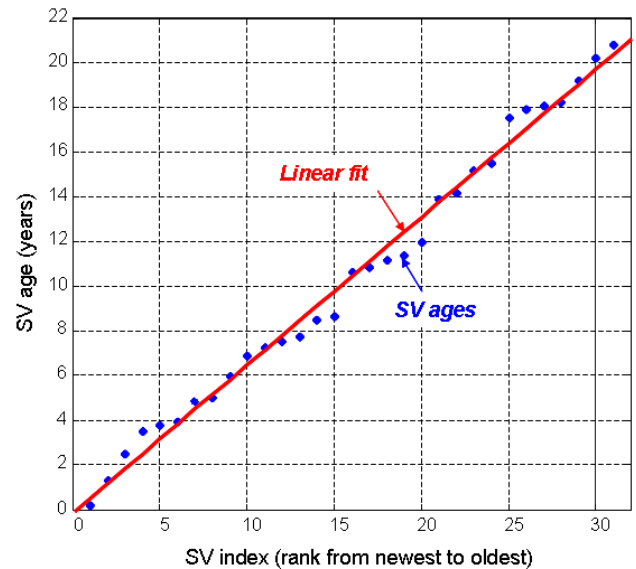

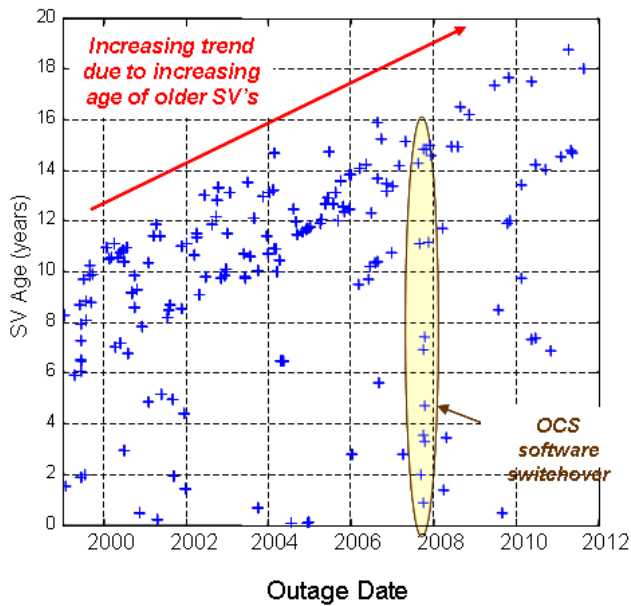**Figure 1: Distribution of GPS Satellite Ages on 20 Sept. 2011**

Figure 2: Unscheduled Outages by Year and SV Age

Figure 3: Histogram of Outages by SV Age

points (in red). This plot shows that the current age distribution among the GPS satellites is close to uniform, meaning that older and newer satellites are present with approximately equal frequencies (this appears in the plot as roughly linear from newest to oldest). A roughly uniform distribution with age is expected from a "mature" constellation, and it suggests that outage frequencies should appear uniform as well if the underlying outage probabilities are roughly the same for all satellites. It should be noted, however, that the distribution of GPS satellites ages was not necessarily uniform over the entire period in which outages were tabulated (January 1999 – August 2011).

Figure 2 plots each of the 178 unscheduled outages in the found in the NANU archives by year of occurrence and age of the affected satellite at the time of the outage. This graph makes it clear that most outages occur on older satellites (for simplicity of wording, "outages" will often be used when referring specifically to "unscheduled outages"). It also shows that the definition of "older satellites" has changed since 1999 as the oldest satellites in the constellation continued to be usable well past their previously-predicted design lives. Except for a few "infant mortality" outages among new satellites and a brief period of outages across all satellite ages in late 2007, which are known to have been caused by the OCS software changeover that occurred at that time (see [11]), outages among newer satellites (less than the original Block IIA design life of 7.5 years old [7]) are relatively infrequent and are scattered randomly across the plot.

Figure 3 shows a histogram of outages by satellite age from 1 to 20 years (note that a given bin "*x*" represents all outages of satellite ages from *x* – 1 to *x* years). This
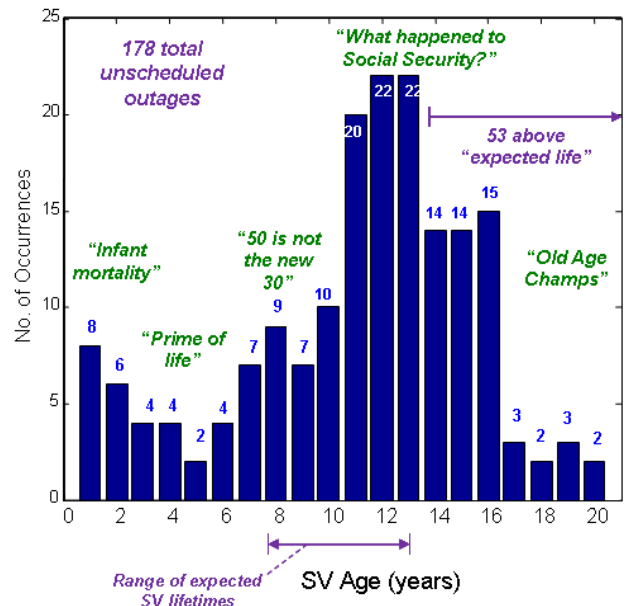
histogram not only reinforces the relationship between satellite age and outage frequency but shows the variation of outage frequency with specific phases of satellite life. As expected, satellites newer than 2 years of age experience more outages than satellites in the "prime of life" due to unexpected design or manufacturing problems that were not discovered during testing. Outages increase in frequency as the expected satellite lifetime (approximately 7.5 – 12.5 years, depending on satellite Block type) approaches and increases further as the expected lifetime is reached and passed.

The relative frequency of outages is highest for satellites that live beyond the 12.5-year expected lifetime of the most modern GPS satellites. Figure 3 does not make this clear because, as indicated in Figure 2, satellites aged 13 years or more are a relatively new component of the GPS constellation. Figure 4 emphasizes this point by re-plotting the outage histogram in Figure 3 while combining all satellites beyond 13 years of age into a single bin. This bin has 53 outages in it compared to the total of 64 outages for satellites from 10 to 13 years old, but the total prevalence of satellites over 13 years old during the data collection period is significantly lower than the prevalence of satellites from 10 to 13 years old. Figure 4 also clarifies that satellites newer than 10 years of age have significantly fewer outages than older satellites and can be treated as "nominal" relative to assumed satellite failure probabilities, even though significant outage variations within this class are visible in Figure 3.

Figure 5 changes the focus to individual satellites and shows the number of unscheduled outages experienced on each satellite observed during the data collection period. The *x*-axis is sorted by SVN instead of launch order, but
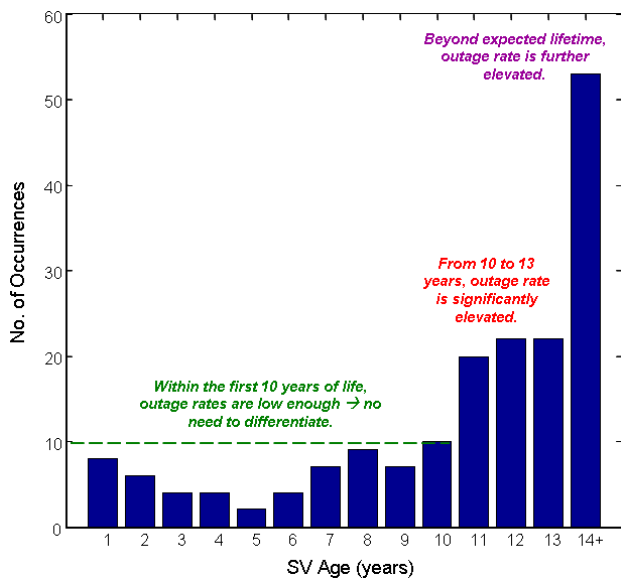
Figure 4: Histogram of Outages by SV Age
(Ages > 13 years combined into a single bin)

Figure 5: Histogram of Outages by SV Number

these two sequences are highly correlated. What is most striking is the variability of outage frequency among the satellites. A few SVNs show zero outages because they expired before 1999 (e.g. SVNs 20 and 28), have not had any unscheduled outages yet (e.g., SVNs 48 and 52), or never entered service (e.g., SVN 42, which suffered a launch failure). Otherwise, some satellites of the same type, such as Block IIA SVNs 25 and 31, experienced many more outages (12 and 13, respectively) than did other Block IIA satellites. Long life is not the sole explanation, as SVN 31 was decommissioned "for good" after 12.57 years, while SVN 25 lasted 17.82 years. While some variation of outage frequency would be expected even if all satellites had equal failure probabilities, the degree of variation shown in Figure 5 strongly suggests that "all satellites are not made equal."

Figure 6 shows the outages for each satellite in more detail by plotting each one as a function of satellite age. Each blue box represents the start of an outage for the satellite shown on the *x*-axis, while the *y*-axis indicates the age of that satellite when the outage began. Almost all blue boxes are followed by a red "plus" sign indicating when the outage ended (the exceptions are satellites that were decommissioned instead of returned to service). Because the *y*-axis has units of years, and most unscheduled outages only last a few hours or days, the plus signs tend to lie practically on top of the blue boxes in all but a few cases where outages lasted for several months. Because many outages occurred in close proximity to one another, they "overlap" on Figure 6 and are not visible as distinct events. Also, the limits of the 1999-2011 data collection period are visible, as "young" outages for earlier GPS satellites are not visible, nor are "old" outages for satellites launched in the last few years.
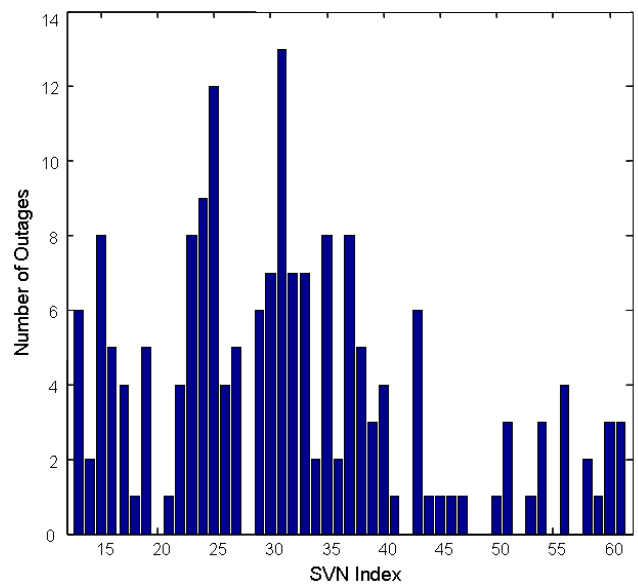
Figure 6 reiterates the fact that older satellites are more likely to have unscheduled outages, and it shows the history of outages and repairs on each satellite to help clarify what separates "well-behaved" and "troublesome" satellites. "Troublesome" satellites not only experience many outages, but these outages become more frequent as satellites age and pass their expected lifetimes. In particular, the appearance of one or two outages within 1 – 2 years on older satellites is a strong indication that more outages are coming.

Figure 7 highlights this tendency by "zooming in" on the unscheduled outage history of SVN 25, which as noted above experienced 12 unscheduled outages over a lifetime of almost 18 years. The same data shown in Figure 6 is repeated in Figure 7 for SVN 25, but additional details such are added in text form (since outage durations are now provided by text, the red "plus" signs have been removed).

As shown in Figure 7, SVN 25 experienced two outages in 1999 and 2000 as it passed the original Block IIA design life of 7.5 years, and it experienced one more outage in 2004 after passing 12 years of age. Six months later, with the satellite now more than 13 years old, it experiences a series of 5 outages from February 2005 to May 2006. Two of these outages, on 23 and 25 December 2005, occurred so close to each other that they appear as one overlapping box in Figures 6 and 7. Note that the earlier of these two outages was resolved in about 8.5 hours, but the underlying problem remained and led to the second outage shortly thereafter. That second outage lasted about a month, but further outages followed, suggesting that the GPS OCS was fighting a persistent problem on this satellite. Four additional outages occurred on SVN 25 after the period of 5 outages ended
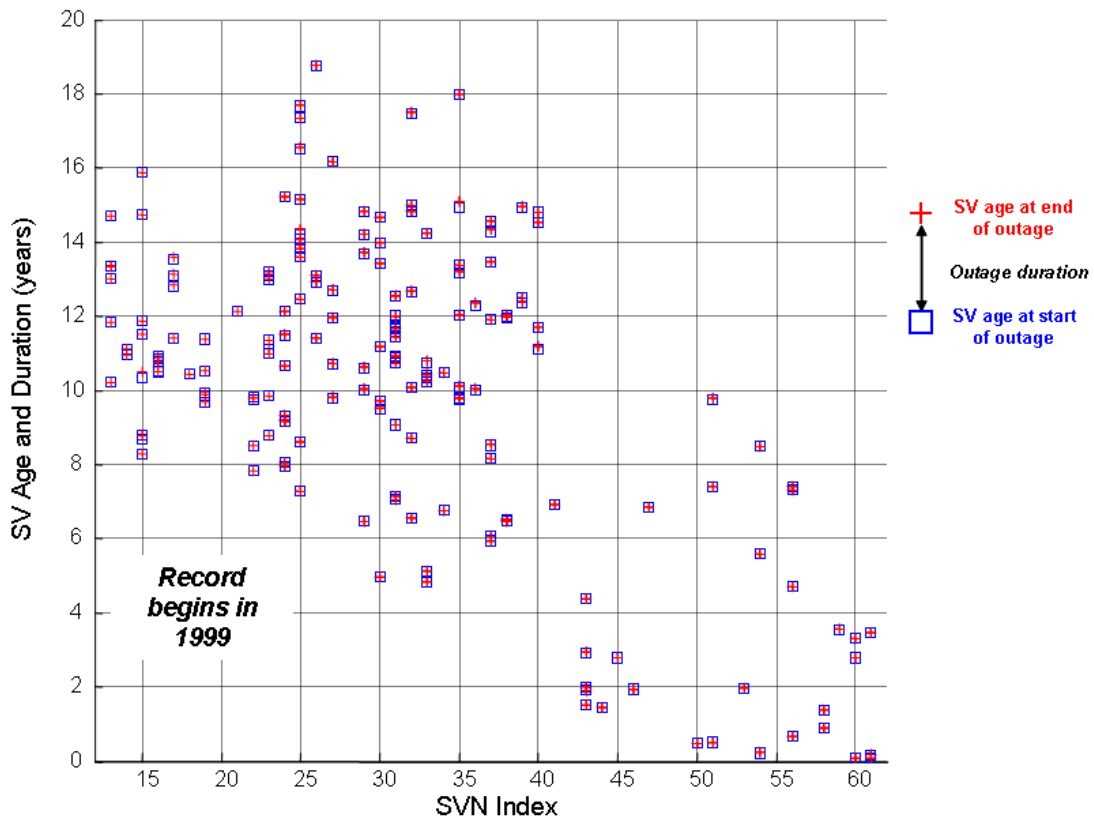
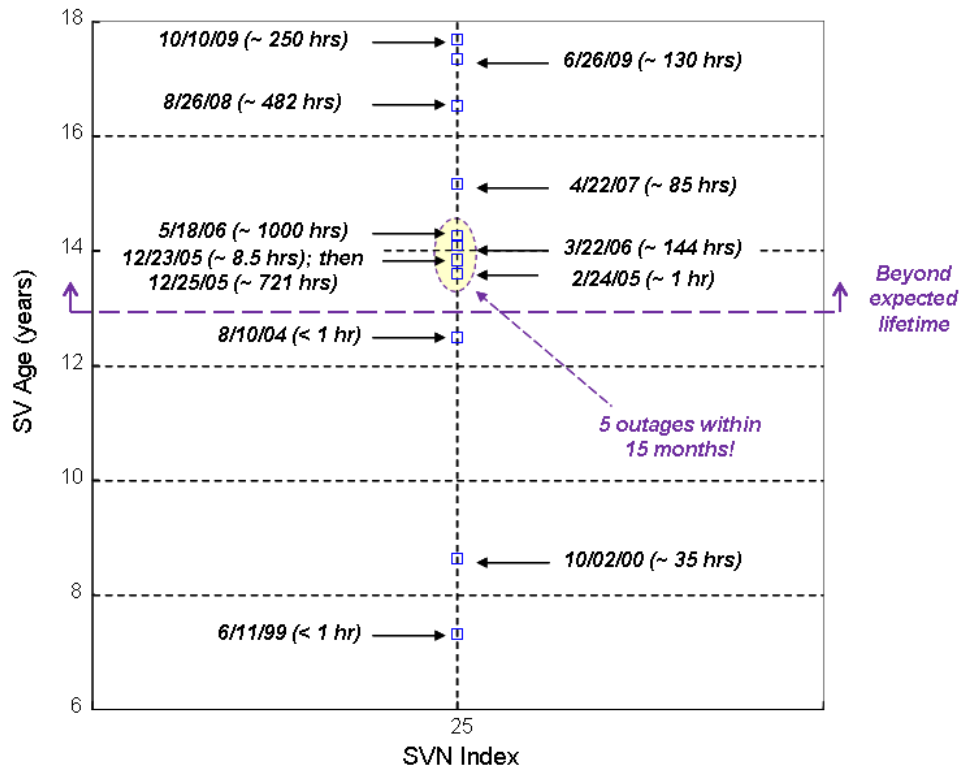**Figure 6: Age and Duration of Outages on Each Satellite**

Figure 7: Unscheduled Outage History of SVN 25

in 2006, showing that the problems affecting SVN 25 were never completely resolved before it was decommissioned. Despite this, SVN 25 was continuously reactivated and kept in service until the level of maintenance and monitoring difficultly that it posed to OCS was no longer worthwhile.

The history of outages by individual satellite shown in Figures 6 and 7 suggests that recent outage history is at least as relevant as age in predicting future satellite outages. Since older satellites experience a significant majority of outages, the two effects are highly correlated, but considering prior outage history helps to separate older satellites with a particularly high risk of future outages from those whose risk is not high enough to require additional mitigation. The benefits of this will become apparent when example satellite-exclusion rules are examined in the following section.

Finally, what comparisons can be drawn between these results and the expected satellite failure and outage probabilities described in Section 2.0? If the 178 outages observed in this study are divided by the number of hours covered by the data collection period (about 111,035 hours from January 1999 through August 2011, inclusive), an outage rate of 1.16×10^{-3} per hour is obtained. Dividing this result by 24 satellites (a conservative estimate of the number of satellites present in the constellation over this period) gives an outage rate of 6.51×10^{-5} per hour per satellite, which is well below the SPS Performance Standard of 2.0×10^{-4} per slot per hour for unscheduled outages. Furthermore, even if all of these outages represent potential integrity threats for GBAS, the observed outage rate is below the 10^{-4} per hour per satellite fault probability assumed by GBAS for each fault mode.

While these numbers suggest that the *average* outage and failure rates over the GPS constellation are sufficiently low, the focus of this study is the possibility that particular satellites might exceed what is assumed by GBAS or other systems. SVN 25, highlighted in Figure 7, is a good example. It experienced 12 outages over the period from January 1999 to its decommissioning in mid-December 2009 (a period of about 96,060 hours), giving an average outage rate over its lifespan of about 1.25×10^{-4} per hour. This is just under twice the overall outage rate for all satellites – it is elevated but not a major concern. However, if we focus on the period from the 4th outage (on 10 August 2004, when SVN 25 was about 12.5 years old) until decommissioning, 9 outages occurred within a period of about 46,920 hours, giving an outage rate of 1.92×10^{-4} per hour. A more pessimistic number is obtained by focusing on the 5 outages that occurred between 24 February 2005 and 18 May 2006 (a period of about 10,750 hours). This gives an outage rate of 4.65×10^{-4} per hour, which exceeds the SPS outage rate and would exceed the SPS failure rate if the fraction of unscheduled outages that pose an integrity threat exceeds $10^{-5} / 4.65 \times 10^{-4}$, or about 2.15%. As noted above, the actual fraction of outages that pose integrity threats to GBAS or other civil systems is unknown but is almost certainly higher than this.

4.0 Example Satellite-Exclusion Heuristics

The possibility that individual GNSS satellites might exceed the failure probabilities assumed by high-integrity civil systems motivates the consideration of satellite-exclusion heuristics, or rules that would prevent high-risk satellites from being used based on characteristics that are observable to these systems. Two example heuristics are proposed and evaluated in this section:

Heuristic 1: Exclude all GPS satellites greater than 13 years of age.

Heuristic 2: Exclude all GPS satellites greater than 10 years of age that have experienced one or more unscheduled outages within the last two (2) years.

Heuristic 1 is suggested as a conservative example only and is not intended to be practical when using only GPS for satellite navigation (i.e., without also using other GNSS satellites from GLONASS, Galileo, etc.). Recall from Figure 1 that the satellites in the current GPS constellation shown in Table 1 (as of 20 September 2011) are relatively uniformly distributed in age between 0 and 20 years. Therefore, excluding all satellites older than 13 years would eliminate a significant fraction of the constellation and would be impractical for most users. In particular, applying this heuristic to the GPS satellites in Table 1 would eliminate the following 11 satellites (in descending order of age): SVNs 23, 24, 26, 39, 35, 34, 36, 33, 40, 43, and 38. Six of these 11 satellites are in "spare" orbit slots and thus contribute less to user geometry, but the other five remain in "primary" slots and are heavily relied upon.

Heuristic 2 is still very conservative for users of only GPS satellites but is less so than Heuristic 1 because it requires both advanced age (10 or more years instead of 13) and one or more unscheduled outages in the last two years. Table 1 shows that the 10-year age cutoff adds 5 more satellites (SVNs 46, 51, 44, 41, and 54) to the above list of those eligible to be excluded. However, the additional requirement of an outage within the last 2 years allows the majority of these satellites to be used despite their age. Only 5 satellites are excluded by Heuristic 2: SVNs 26, 35, 38, 40, and 51. Three of these five satellites are in "spare" orbit slots, and all but one (SVN 51) is included in the larger set of satellites excluded by Heuristic 1.
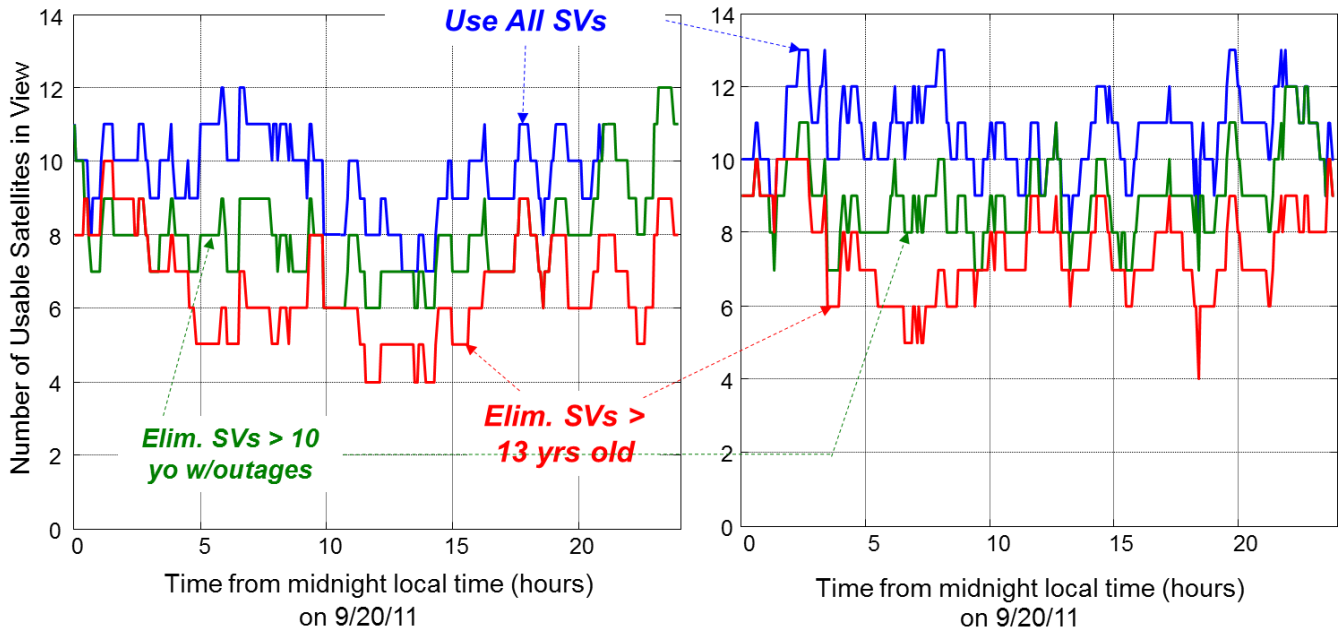
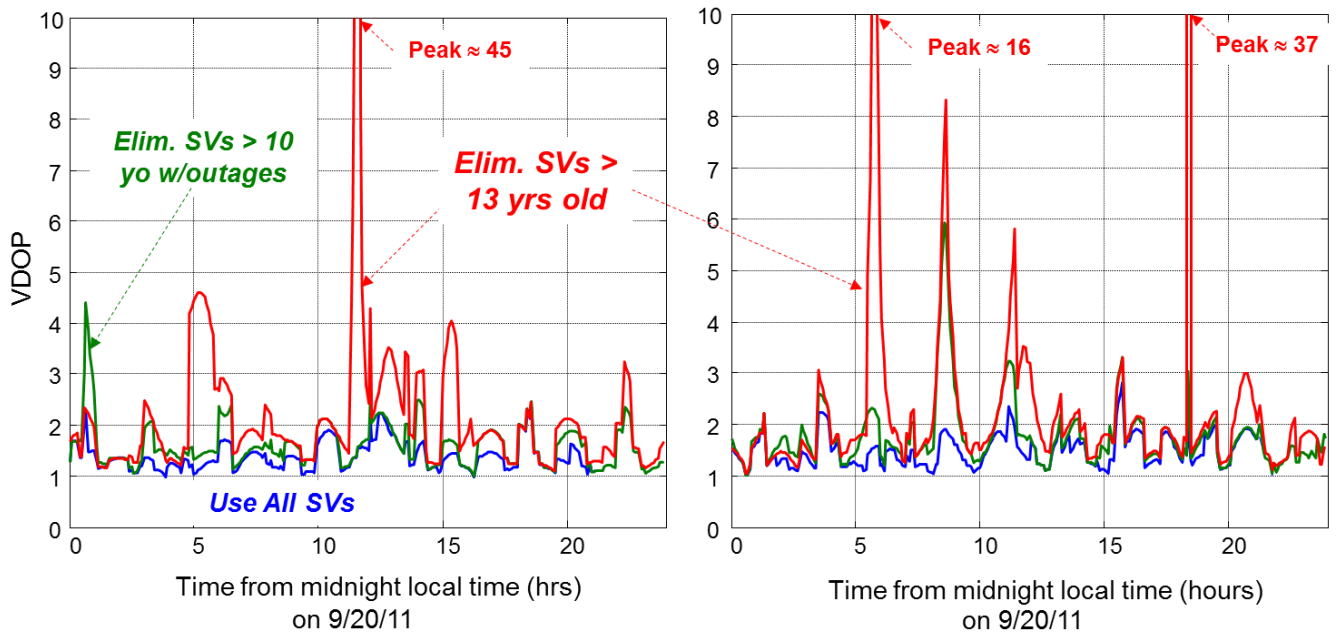**Figure 8: Number of Usable Satellites at Palo Alto and Fairbanks**

Figure 9: Vertical DOP (VDOP) at Palo Alto and Fairbanks

These two heuristics were evaluated using satellite geometry simulations at two locations in the U.S.: a mid-latitude site (Palo Alto, California, at 37.4° North latitude) and a high-latitude site (Fairbanks, Alaska, at 64.8° North latitude). One day of GPS satellite geometries was simulated at 5-minute intervals over 20 September 2011 (local time), and a satellite visibility elevation "mask angle" of 5° was used for both locations. Trimble's freely-available "Setup Planning" software for Windows (Version 2.9) was used to perform this analysis [12].

The two plots in Figure 8 show the number of visible and usable GPS satellites on 20 September 2011 at both Palo Alto (left-hand plot) and Fairbanks (right-hand plot) under three scenarios. The first scenario, shown in blue, includes all visible satellites without any exclusion rules. The second scenario, shown in green, implements the satellite-exclusion rules of Heuristic 2, while the third scenario, shown in red, implements the satellite-exclusion rules of Heuristic 1. The same pattern is followed in Figure 9, which plots Vertical Dilution of Precision

(VDOP) at both locations. VDOP translates the quality of the usable satellite geometry into a measure of user error in the vertical dimension. Multiplying the current VDOP by the one-sigma accuracy of each ranging measurement provides an approximate one-sigma estimate of vertical position accuracy.

Figures 8 and 9 clearly show the disadvantages of implementing the proposed satellite-exclusion heuristics when only using GPS satellites for navigation. Heuristic 1, in particular, is clearly impractical. Although 7 or more satellites are always visible at both locations, and 9 or more satellites are visible at most times (10 or more at Fairbanks), the eliminations required by Heuristic 1 bring the number of usable satellites down to 4 – 6 at times. The resulting penalty is most obvious in Figure 9, where the red curve representing Heuristic 1 has multiple undesirable "spikes" where VDOP grows to very high values, and a significant fraction of the day has VDOP above 3.0 for this case. In comparison, VDOP for Heuristic 2 (the green curve) stays much closer to the ideal VDOP for the no-SV-elimination case (the blue curve), although significant deviations are visible and will lead to lower availability if Heuristic 2 is implemented.

Figure 9 suggests that, while Heuristic 1 is too conservative, Heuristic 2 is a good starting point for refinement of the exclusion rules for GPS satellites. An improved heuristic likely needs to exclude fewer satellites to be optimal for users of GPS satellites only. On the other hand, if GPS were combined with another full or nearly-full satellite constellation of similar ranging accuracy (potentially Galileo and/or GLONASS in the next 10 years or so), Heuristic 2 might be usable as is, as many more satellites would be available to make up for the excluded ones. Use of GPS with other satellite constellations remains speculative, and the outage characteristics of new GNSS satellites remain to be observed, but it is reasonable to expect that their behavior over time will follow patters similar to those of GPS satellites.

Note that, while Fairbanks typically has more satellites in view (on average) than Palo Alto in Figure 8, VDOP at Fairbanks is worse than at Palo Alto in Figure 9. This is true for all three cases shown but is most visible for Heuristic 1, and it is due to the weakness of satellite geometries viewed from latitudes above the inclination angle of the GPS satellite constellation (55°). In general, locations with weaker satellite geometries under nominal conditions (e.g., due to local obstructions causing higher elevation mask angles) will make the loss of performance due to satellite exclusions that much worse, since "less margin" for these exclusions is present to begin with.

Another consideration regarding satellite-exclusion rules is how to implement them in civil systems in real time.

The study used to develop these heuristics was based on access to archived NANUs via the Internet [2,3]. Internet access is one straightforward means of obtaining NANUs, but some means external to GPS receiver hardware is needed because NANU information is not broadcast in the GPS navigation data messages. Augmentation systems such as GBAS and SBAS, could obtain this information externally, but are not necessarily designed to do so. For example, the Honeywell SLS-4000 LAAS Ground Facility, which achieved CAT I System Design Approval in 2009 (but not yet airport site approval), mirrors the much older Instrument Landing System in allowing no external contact (beyond GPS and VDB signals) except for manual interaction with FAA maintenance personnel.

Without access to NANUs or other external information about satellite constellation health, even tracking satellite age is difficult because satellites do not broadcast their SV numbers – only their PR numbers, which can move from satellite to satellite (usually when a new satellite replaces an older one that has been decommissioned). Therefore, some external input is needed to update a maintenance file when a new GPS satellite is approved for use. An opportunity for this exists for the SLS-4000 because manual input is required to approve the C/A-code signal quality of each new satellite.

However, as this paper establishes, tracking recent satellite outages is key to limiting the number of satellites that are excluded. Without regular NANU inputs, individual GBAS sites could observe satellites being flagged unhealthy when in view of that site, but they would miss the times when this occurs out of view, and more importantly, they could not easily distinguish between an unscheduled outage and a scheduled outage for maintenance unless that satellite was maneuvered. Ideally, augmentation systems (and ARAIM, which will rely on external integrity messages to some degree [13]) will expand their access to NANU information as they mature, but for now, its absence significantly limits our ability to implement satellite-exclusion rules.

5.0 Real-Time Exclusion Method that Incorporates Satellite Geometry

The satellite-exclusion rules proposed in Section 5.0 are simple "open-loop" rules that would be applied regardless of the usefulness of each satellite to a user's position solution. This makes sense if the intent is to remove satellites whose potential harm to user integrity is large and exceeds the assumptions made by user integrity algorithms to an unknowable extent. While this could be the case, the results in this paper suggest that any such violations of existing assumptions will be a matter of degree rather than an immediate threat. If this is the case, and the degree of violation can be estimated and/or bounded, the methods currently used to verify integrity in

civil-aviation navigation systems can be adapted to determine in real time (or near-real-time) when the value of a given satellite to the user's positioning geometry outweighs its failure risk. This approach would blend the risk-sensitive heuristics proposed above with the common means of channel-limited satellite selection described in Section 1.0, which selects the satellite subset that minimizes the position DOP of interest. Minimizing DOP may be sufficient when all satellites have similar failure risks, but a generalized approach would handle the results shown in Section 3.0, where this is not the case.

The concept that makes this possible is known as the "multiple-hypothesis" (MH) approach to estimating integrity risk, which was first developed in [15] and has been used more recently as the basis for computing real-time user protection levels in ARAIM [16]. ARAIM computes protection levels by evaluating the probabilistic impacts of nominal conditions and each faulted or anomalous threat mode whose probability is significant relative to the overall loss-of-integrity probability that the system aims to protect (e.g., 10^{-7} per operation for the most demanding flight modes to be supported by ARAIM). Threats due to satellite faults are weighted by the assumed prior probability of that fault. Protection levels are output at each epoch that bound the maximum position errors at the required loss-of-integrity probability (or "PHMI") based upon the current measurements and a prior allocation of that probability to each fault and to the fault-free case [16]. While this approach is used directly by ARAIM, the theory behind it underlies the calculation of protection levels for all civil-aviation applications and is thus widely applicable [17].

An addition to the ARAIM multiple-hypothesis method could handle potential satellite exclusions by treating them as potential "failures" whose risk could be *accepted* by including the satellites in question or *rejected* by excluding them. First, a list of "satellites of concern" with higher-than-normal failure probabilities would be built using a refinement of Heuristic 2, for example. The nominal protection-level calculation would exclude all of these "troubled" satellites, meaning that it would enforce the exclusion heuristic as an "in or out" rule. Next, each satellite excluded by the heuristic is re-included in a separate protection level calculation in which failures of the re-included satellite are evaluated with a higher probability based on the analysis in Section 3.0. If the resulting protection levels are meaningfully lower than for the nominal case (with all "troubled" satellites excluded), then that satellite should be included in the position solution because its improvement to user geometry outweighs the additional failure risk that it brings with it. To the degree that computational resources permit, this would be repeated for all satellites excluded by the heuristic being used.

The advantage of using heuristics as "in or out rules" is that there is no need to assess how much higher the probability should be for "satellites of concern" relative to the assumed probability for all satellites. However, the multiple-hypothesis approach requires a numerical assessment to support the quantitative trade-off that it performs. The results in Section 3.0 suggest that assuming an integrity failure rate for "satellites of concern" that is two orders of magnitude greater than the number given by the SPS Performance Standard (i.e., about 10^{-3} per "troubled" SV per hour) would be sufficiently conservative for almost all of these satellites (note that the ARAIM multiple-hypothesis method does not require that the same failure probability be used for each satellite).

One concern of using a probability this high for "troubled" satellites is that independent, simultaneous failures of two "troubled" satellites would have a probability of roughly 10^{-6} per satellite per hour, and independent, simultaneous failures of one troubled satellite and one nominal satellite (using 10^{-4} per satellite per hour as conservative for nominal satellites) would have a probability of roughly 10^{-7} per satellite per hour. These probabilities are significant compared to the 10^{-7} PHMI requirement for many civil operations, which means that the multiple-hypothesis risk calculation might not be able to neglect dual-satellite failure scenarios for "troubled" satellites that it might have neglected previously because their probabilities were low enough to neglect. This is not a problem for the algorithm, but it points out one reason why the additional risk of using "troubled" satellites might not be worth the benefit.

6.0 Summary and Future Steps

This paper has examined the history of unscheduled GPS satellite outages from January 1999 to August 2011 (based upon NANU archives) and isolated two factors that can be used to identify satellites with much higher than average risk of future failures. As expected, satellite age is the primary observable factor. GPS satellites routinely outlive their predicted design lives by many years, but as they do so, they become much more susceptible to failures that require them to be declared unhealthy and repaired by the GPS OCS. In addition, certain "problem" satellites have a much greater propensity for recurring unscheduled outages than others. Therefore, once one or two recent outages have been observed on an older satellite, the probability of further outages increases significantly.

Since satellite age and number of outages are observable to GPS users (at least those who have access to NANUs or who independently monitor the constellation on a semi-worldwide basis), they can be used in heuristics that pre-emptively exclude "troubled" satellites from use to avoid

the potential integrity risk that they might pose. Doing this with today's GPS satellite constellation is painful in terms of the number of satellites that would be excluded and the resulting degradation in positioning geometry, but this may change in the future as GPS is combined with other satellite constellations such as Galileo, GLONASS, Compass, or regional satellites such as SBAS GEOs or QZSS. In addition, a modification of the existing multiple-hypothesis protection level algorithm used in ARAIM would allow the positioning benefit of using troubled satellites to be traded against the increased failure risk that they add.

Much work remains to refine the satellite-exclusion heuristics proposed in this paper and to develop the multiple-hypothesis (MH) exclusion method into a practical algorithm. Computational workload is a concern for the MH method because the current proposal requires that a separate set of MH protection levels be computed for each satellite that is a candidate for exclusion. The performance benefit of the MH approach relative to using heuristics alone (meaning exclusion of all "troubled" satellites) should be evaluated for both the GPS-satellites-only case and one or more GPS+GNSS-satellites cases. Simulations of GNSS satellites in addition to GPS are complicated by the fact that most potentially usable non-GPS satellites are very new and have limited outage histories; thus the exclusion rules that should apply to them must be extrapolated from those that have been derived for GPS satellites.

ACKNOWLEDGMENTS

REFERENCES

[1] Gaylord Green, personal communication, c. 1992.

[2] USNO GPS Timing Operations website, http://tycho.usno.navy.mil/gps.html

[3] U.S Coast Guard Navigation Center (NavCen) website, http://www.navcen.uscg.gov/?pageName= gpsNanuInfo

[4] J.W. Lavrakas and C. Shank, "A Study into the Use of URA in Satellite Selection," *Proceedings of 49th ION Annual Meeting*. Cambridge, MA., June 21-23, 1993, pp.

261-270. http://ion.org/search/view_abstract.cfm?jp=p& idno=4518

[5] C.W. Park and J.P. How, "Quasi-optimal Satellite Selection Algorithm for Real-time Applications," *Proceedings of ION GPS 2001*. Salt Lake City, UT., Sept. 11-14, 2001, pp. 3018-3028. http://ion.org/search/ view_abstract.cfm?jp=p&idno=1978

[6] N. Blanco-Delgado and F. D. Nunes, "A Convex Geometry Approach to Dynamic GNSS Satellite Selection for a Multi-Constellation System," *Proceedings of ION GNSS 2009*. Savannah, GA., Sept. 22-25, 2009, pp. 1351-1360. http://ion.org/search/view_abstract.cfm ?jp=p&idno=8544

[7] *Global Positioning System Standard Positioning Service Performance Standard*. U.S. Department of Defense, Washington, DC, 4th Edition, Sept. 2008. http://www.pnt.gov/public/docs/2008/spsps2008.pdf

[8] S. Pullen, J. Rife, and P. Enge, "Prior Probability Model Development to Support System Safety Verification in the Presence of Anomalies," *Proceedings of IEEE/ION PLANS 2006*, San Diego, CA, April 24-27, 2006, pp. 1127-1136. http://waas.stanford.edu/~wwu/ papers/gps/PDF/PullenIONPLANS06.pdf

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

[10] P. Kline, T. Murphy, and K. Van Dyke, "LAAS Availability Assessment: The Effects of Augmentations and Critical Satellites on Service Availability," *Proceedings of ION GPS 1998*, Nashville, TN, Sept. 15-18, 1998, pp. 503-516. http://ion.org/search/view_ abstract.cfm?jp=p&idno=2980

[11] K. Kovach, J. Berg, and V. Lin, "Investigation of Upload Anomalies Affecting IIR Satellites in October 2007," *Proceedings of ION GNSS 2008*, Savannah, GA, Sept. 16-19, 2008, pp. 1679-1687.

[12] Trimble Planning Software website, http://www.trimble.com/planningsoftware.shtml

[13] M. Brenner, F. Liu, *et al*, "Natural Signal Deformations Observed in New Satellites and their Impact on GBAS," *Proceedings of ION GNSS 2009*. Savannah, GA, Sept. 22-25, 2009, pp. 1100-1111. http:// www.ion.org/search/view_abstract.cfm?jp=p&idno=8519

[14] J. Blanch, S. Wallner, *et al*, "A Proposal for Multi-constellation Advanced RAIM for Vertical Guidance," *Proceedings of ION GNSS 2011*. Portland, OR, Sept. 20-23, 2011 (forthcoming).

[15] B. Pervan, S. Pullen, and J. Christie, "A Multiple Hypothesis Approach to Satellite Navigation Integrity," *NAVIGATION*, Vol. 45, No. 1, Spring 1998, pp. 61-71.

http://waas.stanford.edu/~wwu/papers/gps/PDF/Pullen1998_Multiple%20Hypothesis%20Integrity%20v45n1-6.pdf

[16] J. Blanch, T. Walter, *et al*, "An Optimized Multiple Hypothesis RAIM Algorithm for Vertical Guidance," *Proceedings of ION GNSS 2007*. Fort Worth, TX, Sept. 25-28, 2007, pp. 2924-2933. http://waas.stanford.edu/~wwu/papers/gps/PDF/BlanchIONGNSS07.pdf

[17] S. Pullen, "A Comparative Overview of the Protection Level Concept for Augmented GNSS and LORAN," Stanford University, GPS Laboratory Meeting, Dec. 20, 2002. http://www.stanford.edu/~spullen/GNSS_PL_concept_2002.ppt