Integrity for Non-Aviation Users

Moving Away from Specific Risk

Non-aviation users of satellite- and ground-based augmentation systems do not require the conservative level of integrity built into these systems for aviation users. Removing it can produce substantial benefits in terms of smaller error bounds and improved availability.

Sam Pullen, Todd Walter, and Per Enge, *Stanford University*

Both space-based and ground-based augmentation systems (SBAS and GBAS, respectively) are designed to enhance standalone GNSS navigation to meet the requirements of civil aviation. SBAS and GBAS corrections and integrity information are also available to the non-aviation user population, such as automobiles, buses, and trains on land as well as ships near shore. This much larger user base can benefit as much from the integrity components of SBAS and GBAS as from the increased accuracy obtained from applying SBAS and GBAS pseudorange corrections. However, there are significant differences between the aviation interpretation of navigation integrity and the interpretation that would be natural to most users.

SBAS and GBAS provide integrity in a multi-step procedure that is laid out in the RTCA Minimum Operational Performance Standards (MOPS) for the FAA versions of both systems: DO-229D for the Wide Area Augmentation System (WAAS) and DO-253C for the Local Area Augmentation System (LAAS). These systems indicate which ranging measurements should be excluded as unsafe to use and provide bounding error standard deviations, or sigmas, for the remaining usable measurements. Each aircraft uses this information to compute vertical and horizontal protection levels that define position-domain error bounds at desired probabilities. This process is straightforward, logical, and is not limited to aviation users. However, the requirements and assumptions underlying it make it very conservative.

SBAS and GBAS are designed to meet integrity requirements defined in terms of what is known as specific risk. Briefly, this means that all safety requirements must be met for the worst combination of knowable or potentially foreseeable circumstances under which an operation may be conducted. Some variable factors important to safety, such as the user's satellite geometry, are known by defini-

▲ **FIGURE 1** Illustration of 95 percent accuracy bounds and $1–10^{-7}$ protection levels.

tion. Others, such as receiver thermal noise, are random and unpredictable. But several factors that are critical to GNSS performance, such as multipath and ionospheric errors, are neither completely random nor deterministic. Specific risk typically treats all error sources that are not completely random in a worst-case manner. SBAS and GBAS are designed to mitigate specific risk to support civil aviation, and the resulting conservatism makes SBAS and GBAS less attractive to non-aviation users who expect tighter protection levels relative to nominal system accuracy.
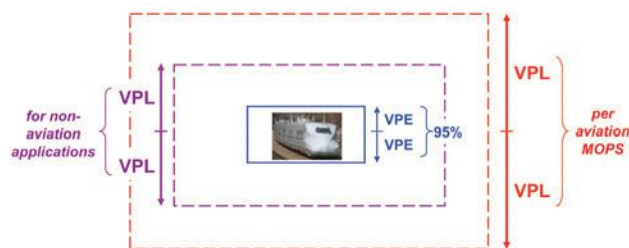
Fortunately, non-aviation users need not apply all MOPS procedures required of aviation users if their own safety requirements differ. Most users define integrity in average or ensemble terms, meaning that everything not known in practice is treated as random and is probabilistically mixed (or convolved) together. The protection levels valid for these users would be much lower than for aviation users, even though the stated bounding probability is the same. This contrast is illustrated in **FIGURE 1**, which shows example bounds on 2-D vertical errors at a probability of 0.95 (the 95th percentile, or 95 percent) for accuracy and a probability of $1–10^{-7}$ for integrity. The term VPE stands for vertical position error, while VPL stands for vertical protection level. Analogous terms (HPE and HPL) and a similar picture exist in two dimensions for horizontal errors.

Only one 95 percent error bound is shown in Figure 1 because this probability can be observed, estimated, and modeled with theory and reasonable amounts of data (hundreds or thousands of independent samples). This is not at all the case at the very small probability of 10^{-7} that applies to aviation precision approach: it is roughly equivalent to one event in 47.5 years per 150-second precision-approach interval. Both theory and data fall far short of being able to predict such rare-event errors. Extrapolating from available data to $1–10^{-7}$ using Gaussian distributions is perilous because the Gaussian distribution almost never applies at such small probabilities. Mixed-Gaussian models, other so-called fat-tailed distributions, and inflation of Gaussian parameters help address this, but the uncertainty regarding the true error distribution results in significantly different error bounds depending on the assumptions that are made. The same is true regarding the effects of faults and anomalies that are more probable than 10^{-7} but are still rare and poorly understood.

In the end, different means of assessing these uncertainties and various degrees of user risk aversion result in different $1–10^{-7}$ protection levels, as shown in Figure 1. It is this difference that we wish to quantify and exploit in this article.

Average Versus Specific Risk

The concept of average or ensemble risk is intuitive to those with a background in probability and is one of the key prin-

ciples of probabilistic risk assessment (PRA). Thus, it helps to examine it first.

Average risk is the probability of unsafe conditions based upon the convolved (averaged) estimated probabilities of all unknown events. More specifically, probability distributions are derived (based on the best available knowledge) for all unknown parameters relevant to user safety, and these are combined (by probabilistic convolution) to create an overall distribution that represents safety risk as a function of the known parameters. This straightforward, natural interpretation of probability and uncertainty has a major advantage in that it cleanly separates the probabilistic calculation of safety risk from users' aversion to risk. By keeping risk probability and risk aversion (or severity) separate, a final risk consequence measure can be derived that supports apples-to-apples comparisons of alternatives. One useful result of this is known as the value of information (VOI). By comparing the risk outcomes of two scenarios in which the latter case has additional information (for example, from an additional sensor or integrity monitor), the risk-reduction benefit of the added information can be traded off against the cost and complexity that it introduces to the system. Similar comparisons can be made for any definition of risk, but the definition and use of VOI in an average-risk framework makes the most sense in both theory and practice.

Turning to specific risk, no single definition exists within the aviation safety community, to our knowledge. This is partially because of the uniqueness and complexity of the concept and partially because multiple inconsistent interpretations appear to exist. Therefore, we provide our own definition: Specific risk is the probability of unsafe conditions subject to the assumption that all credible unknown events that could be known occur with a probability of one (on a risk-by-risk basis).

To understand how specific risk differs from average risk, it helps to start with a fault-tree representation of risk in which loss of integrity (LOI) can result from any of the nodes of the tree. **FIGURE 2** shows a simplified example of a fault tree for CAT I GBAS. It shows the allocation of the CAT I total integrity risk requirement of 2×10^{-7} per approach to the various possible causes of integrity loss. In specific-risk

analysis, each type of failure shown in the tree, if deemed to be a credible failure (meaning, in practice, that its assumed prior probability is larger than compared to its allocation in the fault tree), is assessed that the failure is guaranteed to occur in a worst-case fashion. This means that the variables that describe this particular failure scenario take the values that maximize the hazard to users. In an average-risk analysis, these variables would take many values according to their own probability distributions, and these distributions would be convolved together to provide an overall representation of risk under that scenario. Instead, one scenario drives the specific risk assessment for a particular user class, and it is the worst one possible from that user's standpoint. (Another user class would be evaluated under a different set of parameters corresponding to the separate worst case for that user.) The improbability of the worst-case combination of parameters is not considered as long as the probability of the failure scenario as a whole is deemed high enough to be of concern.

Since GNSS augmentation systems contain multiple levels of health monitoring, the worst-case scenario is usually the one that maximizes the probability of an undetected hazardous error for a particular user class. Hazardous error is typically defined as any error that exceeds a predefined safety zone known as an alert limit (AL) or any error that exceeds the computed protection level (PL), which allows integrity to be defined separately from the intended application. Both definitions are conservative in that all errors exceeding AL or PL are treated as equally hazardous. In other words, an error just above AL is treated as just as dangerous as an error of $10 \times$ AL. They are also misleading when used in specific-risk analyses because the resulting worst-case conditions are those that give errors just above AL or PL, as these are the generally hardest for monitoring algorithms to detect.

The use of specific risk in aviation is an evolution of deterministic guidelines for tolerable risk that date back to an earlier era when flying was more dangerous. It remains dominant in aviation safety assessment because it is partly responsible for the development of safer and more reliable air transportation. However, it has several important weaknesses compared to average risk. The first is that the degree of risk aversion preferred for aviation is buried within the hazard probabilities generated by specific risk — it cannot be separated out. This means that specific-risk results do not translate well to other classes of users, as very few users would happen to have the same risk preferences that have evolved within aviation over several decades. In addition, specific risk makes a distinction between unknown events that could be known and those that are both rare and completely unknowable. A very risk-averse value of information is much different than the risk-neutral one built into PRA, as it severely penalizes systems that do not include all po-
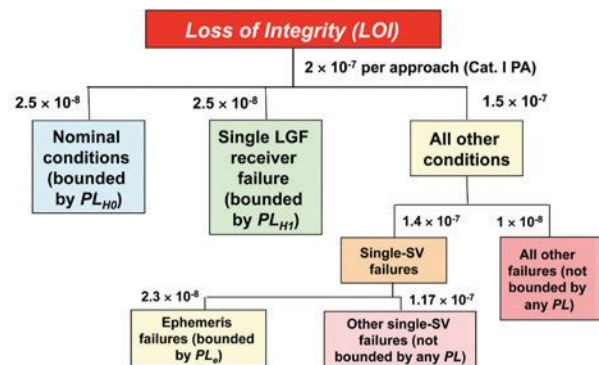
▲ **FIGURE 2** Fault tree for CAT I GBAS integrity.

tentially-informative sensors. Since each sensor added to a system provides less benefit than the last, almost all cost-effective systems choose to include less than the maximum possible number of sensors.

The conservatism implicit in specific-risk assessment severely penalizes users. Although PRA would show that the combination of factors (shown in an example induced by extreme ionospheric spatial decorrelation; see the original Institute of Navigation International Technical Meeting presentation upon which this article is based at *www.gpsworld.com/specific_risk*) needed to produce a 40-meter error in a CAT I GBAS system is exceedingly improbable (almost certainly below 10^{-10} per approach), specific risk forces a significant part of the GBAS risk-mitigation effort to be targeted at this scenario. In this case, since monitoring is not guaranteed to detect the anomaly in time, the only recourse is geometry screening, a cumbersome technique in which the ground system continually evaluates the worst-case error and, if it exceeds a 28-meter tolerable limit at the CAT I decision height, determines which broadcast parameters to inflate such that all satellite geometries causing worst-case errors exceeding 28 meters are made unavailable (because the inflated VPL is larger than the 10-meter CAT I VAL). The result of this procedure is much lower user availability than would be achieved without inflation. SBAS pays a similar penalty, as we will see later. The broadcast grid ionospheric vertical error values that bound worst-case ionospheric errors (and thus the resulting protection levels) are much higher than they would be if the unusual combination of factors needed to create the worst-case error scenario were not the dominant concern.

To the extent that loss of availability represents a safety issue at the airspace level, the worst-case focus that results from specific risk is not optimal even from a safety standpoint. But this is not the only concern. Specific risk requires a great deal of development and testing to identify and mitigate a handful of very peculiar, non-representative conditions. When schedule and resources are limited, other

▲ **FIGURE 3** WAAS PAN reference station network.

potential threats that are easier to foresee but seem extremely improbable are often neglected. One example is the treatment of multiple hardware failures. If individual failures are assumed to be statistically independent, the probability of multiple simultaneous failures is very small. However, while statistical independence is a common assumption in probability classes because it makes calculations easier, it rarely applies in the real world. Because satellites and ground receivers are similar, if not identical, the presence of a failure in one unit may suggest a common cause or at least a common vulnerability, meaning that the probability of additional failures is much higher than independence would suggest. Thus, assuming independence by default could lead to neglecting entire categories of risk that are more threatening than the worst-case events that dominate specific risk.

Maximum WAAS Errors, Protection

To investigate the conservatism built into SBAS and GBAS specific risk assessment, maximum WAAS horizontal and vertical position errors over time (as measured by the Performance Analysis Network (PAN) maintained by the William J. Hughes FAA Technical Center) have been examined and compared to the protection levels computed when the maximum errors occurred. This study begins with PAN

Report #8 (covering January to March 2004 — shortly after WAAS commissioning in mid-2003) and extends through PAN Report #34 (covering July to September 2010). Each PAN report covers three months of observed WAAS performance.

FIGURE 3 shows the 38 WAAS reference stations (WRSs) used by the PAN to collect position error and protection level information (some of these stations were not active in 2004 and thus were not used in earlier PAN reports). While measurements from these stations are used to generate WAAS corrections and error bounds, they are also used by the PAN as static pseudo-users that compute WAAS-corrected positions and protection levels according to the aircraft user algorithms specified in the WAAS MOPS. The resulting positions are compared to the known, pre-surveyed positions of each station to derive estimates of vertical and horizontal position errors (VPE and HPE) once per second.

Figure 3 groups these stations into three sets of stations based on their presumed quality of WAAS coverage. These sets are unofficial and were created for the purposes of this study. The seven stations in the inner set are expected to have good WAAS coverage at all times because they are surrounded by other stations. The 13 stations in the outer set are expected to only have
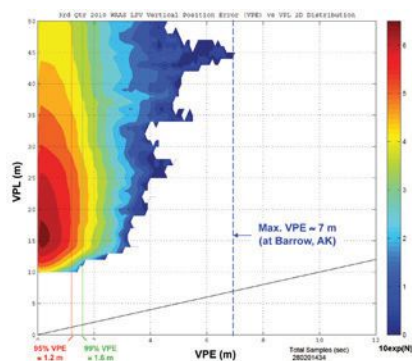
▲ **FIGURE 4** WAAS vertical protection level versus vertical position error (June–September 2010).

acceptable coverage because some of them are at the edges of CONUS. The remote stations provide coverage to the inner and outer regions as well as the best possible coverage of their own regions. Because the remote stations extend beyond the primary coverage region of WAAS in CONUS, errors at these stations are not considered here.

FIGURE 4 is a 2-D plot of position error versus protection level in the vertical axis (that is, VPE versus VPL) for all epochs and stations during the three months covered by the recent WAAS PAN Report #34 (July 1–September 30, 2010). These results are typical of the entire period since WAAS commissioning in 2003, particularly the last several years. The vertical lines on the plot indicate the 95th-percentile, 99th percentile, and maximum VPEs in this period (1.2, 1.8, and 7 meters, respectively). The maximum VPE occurred at Barrow, AK, which is one of the most remote stations in the WAAS network (see Figure 3). In comparison, the lowest VPLs (intended to be $1-10^{-7}$ bounds on VPE) are in the range of 10–15 meters, and values as high as 40 meters are not uncommon. The most demanding approach operation that WAAS supports, LPV-200, allows approaches to a 200-foot minimum decision height and requires that VPL be below a vertical alert limit (VAL) of 35 meters. HPL must also be below a horizontal alert limit (HAL) of 45 meters. When this is not the case, the

approach operation is not available; thus these higher VPLs extract a significant cost.

FIGURE 5 and **FIGURE 6** (for vertical and horizontal errors, respectively) span the entire period of WAAS PAN Reports used in this study. VPL represents the VPL at the station and time of the maximum VPE; it is not the largest VPL recorded at a particular station. The horizontal errors shown in Figure 6 are defined analogously. Note that the station that observes the largest horizontal error in a given PAN report may differ from the one that observes the largest vertical error.

Figures 5 and 6 demonstrate that, while both 95 percent and maximum errors are quite low and are within the expected range of each other, the protection levels associated with the maximum errors greatly exceed them. This pattern is clearer in Figure 5 for vertical errors because maximum VPL tends to be more consistent across PAN reports, but it is true for horizontal errors as well.

FIGURES 7 AND 8 clarify this relationship by plotting the ratio of VPL to VPE and HPL to HPE for the station and time of the maximum error. The mean of this ratio is very high and is about the same in both cases: 5.38 for vertical and 5.21 for horizontal. Figure 7 shows a steady upward trend in the ratio that is mostly due to WRS improvements that resulted in maximum VPE being reduced over time. This trend is clearly visible in Figure 5 and appears to exceed the weaker trend of lowering VPL due to WAAS algorithm enhancements. The same trend is visible in the horizontal Figures 6 and 8 but is weaker due to the greater variability of maximum HPL over time.

To evaluate the significance of the large PL-to-max-PE ratios in the WAAS PAN database, we need to approximate the number of independent samples from which the maximum errors were derived. As noted before, WAAS protection levels represent error bounds at the $1-10^{-7}$ probability
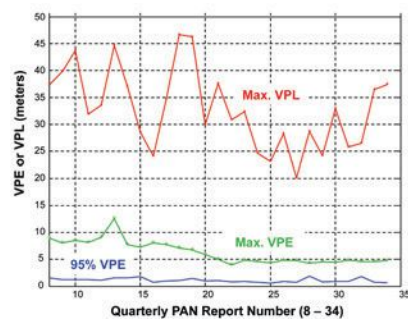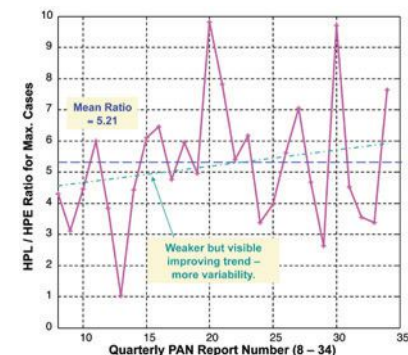

▲ **FIGURE 5** WAAS vertical errors and protection levels from 2004–2010.
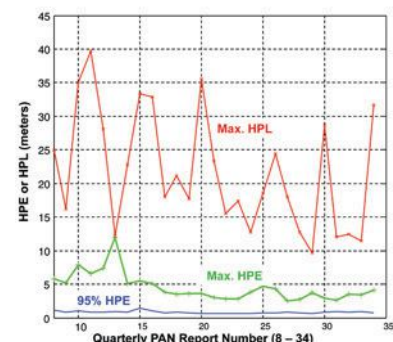
▲ **FIGURE 6** WAAS horizontal errors and protection levels from 2004–2010.
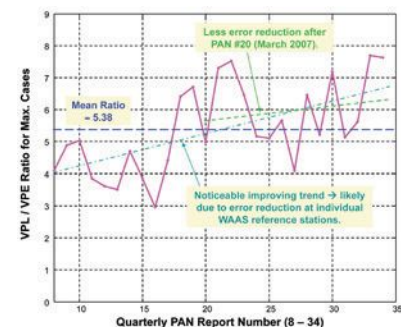
▲ **FIGURE 7** Ratio of VPL to VPE from 2004–2010.

▲ **FIGURE 8** Ratio of HPL to HPE from 2004–2010.

level based on specific risk. With one measurement being collected at each operational station every second, a total of about 4.25 billion samples were collected in the PAN reports from January 2004 to September 2010. Note that measurements from remote stations are included in this count, but they are also represented in the conclusions because their PL-to-max-PE ratios are very similar to the ones shown in Figures 7 and 8. Translating this number into the number of statistically independent samples depends on the interval between independent measurements. Because both nominal and rare-event errors affect this interval, it is hard to estimate. Our best guess is a range between roughly 30 and 150 seconds, suggesting that the PAN database contains between 2.8×10^7 and 1.4×10^8 independent samples. Both of these numbers suggest that WAAS protection levels are very conservative from the perspective of average risk.

Adjusting for Average-Risk Users
Using the above results, a preliminary estimate of the reduced WAAS protection levels that would apply to average-risk users can be made.

FIGURE 9 shows a comparison between the actual 95 percent WAAS VPL and HPL and the adjusted VPL and HPL potentially achievable with WAAS (for the same $1–10^{-7}$ bounding probability) for average-risk users. The actual WAAS VPLs are taken from the more recent WAAS PAN Reports starting from #24 (covering January to March 2008) as the period from 2008 to 2010 includes most of the WAAS algorithm improvements introduced since commissioning in 2003. The actual 95 percent VPLs and HPLs represent the largest reported 95th-percentile values among the stations within CONUS for each quarterly period. The lower adjusted VPLs and HPLs are derived by dividing each VPL by a factor of 4.0 and each HPL by a factor of 2.5. These two reduction factors are derived from Figures 7 and 8, respectively, as conservative estimates of the ratio between protection levels and maximum position errors. Note that the factor of 2.5 for horizontal errors does not include the 12-meter error in Cleveland from PAN Report #13, as this is thought to be spurious (that is, not representative of actual WAAS behavior).

While projections based on these
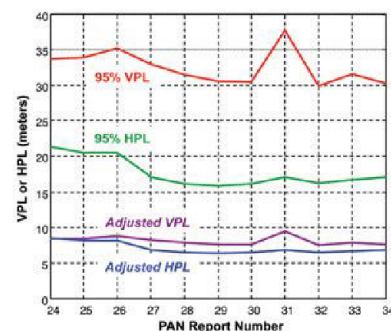
▲ **FIGURE 9** Projected WAAS protection level reductions for average-risk users.

reduction factors are imprecise, they demonstrate the much lower error bounds that non-aviation users with an average-risk safety perspective could achieve. Most non-aviation users operate on land or sea and will be primarily concerned with horizontal error bounds. Figure 9 suggests that the typical 95th percentile WAAS HPLs of 15–20 meters (for the worst location in CONUS) can be reduced to 6–8 meters and still provide a confident $1–10^{-7}$ error bound.

It is important to emphasize that these preliminary projections for average-risk users are just that. In order to formally establish new integrity requirements and protection levels for existing systems, the hazardously
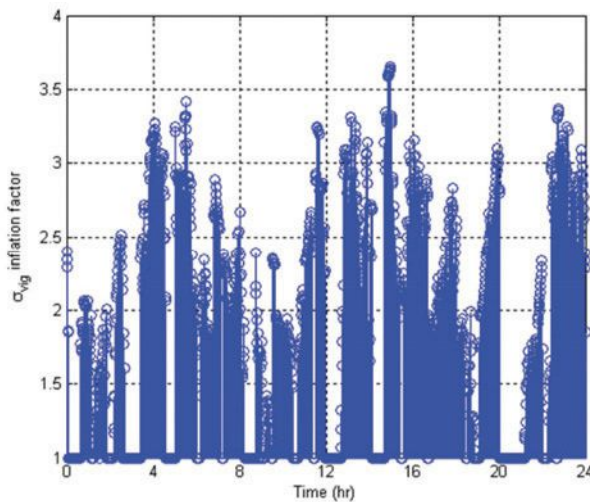
▲ **FIGURE 10** Typical σ_{vig} inflation factors for CAT I LAAS.
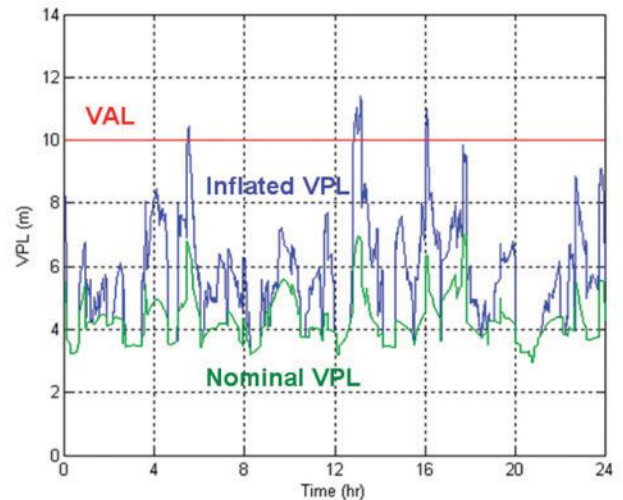
▲ **FIGURE 11** Impact of σ_{vig} inflation on LAAS VPL.

misleading information (HMI) analyses previously done for these systems need to be redone using the principles of PRA and average risk. While the original development of the WAAS and LAAS HMI analyses was lengthy and resource-intensive, almost all of the detailed work is already complete. As long as the original analyses are available, it is a much smaller task to take these results and create PRAs out of them by extracting the original specific-risk assumptions and applying average-risk principles instead.

LAAS Users. Since the first GBAS ground station design (the Honeywell SLS-4000 LAAS Ground Facility) was certified for CAT I use in 2009 and has not yet been approved for operations at a specific airport, much less data is available to do a preliminary analysis for GBAS similar to the one done for WAAS above. However, the degree of sigma inflation in the parameters broadcast by CAT I LAAS is approximately known, meaning that it can be more-precisely removed from the current LAAS protection levels to estimate what they would be for average-risk users.

FIGURE 10 shows the degree of inflation applied to the broadcast $\sigma_{vertical\_iono\_gradient}$ (or σ_{vig}) parameter in order to protect against the worst-case ionospheric anomaly described previously. This result is for the SPS-standard 24-satellite constellation over a 24-hour period at the LAAS installation at Newark Airport, New Jersey (the method used by the Honeywell SLS-4000 is somewhat different). While not all epochs require inflation, a majority cause the nominal σ_{vig} value to be increased by a factor of 2 or more, which significantly decreases CAT I availability and currently makes it impossible to take advantage of the Differentially Corrected Positioning Service (DCPS) for non-CAT-I operations.

Because of the extreme rarity of the worst-case event that dictates this inflation, it would likely not be needed for average-risk users. **FIGURE 11** shows how much the σ_{vig} inflation in Figure 10 increases the LAAS VPL at Newark for the

standard 24-satellite constellation. The VPL reduction from removing the inflation is not as dramatic as the potential reductions shown for WAAS in Figure 9, but they are significant relative to the 10-meter VAL for LAAS CAT I approaches. Furthermore, the pre-inflated nominal value of σ_{vig} for LAAS is 6.4 millimeters/kilometer, which is much higher than the actual one-sigma nominal gradient value of 1–2 mm/km because, under specific risk, the very worst nominal data must be bounded (also, worst-case tropospheric gradients must also be bounded by σ_{vig}). Other broadcast parameters that affect VPL, such as σ_{pr_gnd} and the ephemeris P-value that bounds worst-case ephemeris failures, would also be reduced significantly by switching to average risk. Overall, it is likely that LAAS protection levels based on average risk would be reduced from the current specific-risk PLs by about the same range of factors (2–5) observed from WAAS data.

User Performance Improvements

This discussion assumes that most non-aviation users who are not encumbered by the history of aviation standards development will prefer to quantify risk using PRA and the average-risk approach. As noted earlier, average risk better matches most users' intuitive understanding of uncertainty and has the enormous advantage of separating risk quantification from risk aversion. Regardless of how risk-averse or conservative a given operator is, his or her model of risk aversion can be applied most efficiently to a risk-neutral calculation of risk that fairly represents all aspects of uncertainty. Inserting risk aversion into the calculation of risk, as done in the specific-risk approach, is both inefficient and non-optimal from a safety perspective because extensive focus on a few extreme worst-case events drives attention away from other events.

The HPL reductions for average-risk users illustrated

here would be significant for many classes of ground and marine transportation users. They would allow operations with tighter physical safety margins to be supported. Users who gain no particular benefit from tighter protection levels would still obtain much higher availability of integrity, as a 25-meter HPL could be supported by much poorer satellite geometries than would otherwise be the case. In other words, users that can tolerate 25-meter horizontal error bounds would be able to operate safely a much higher percentage of the time, because the degree of GNSS constellation deterioration needed to exceed this limit would occur much less often. These benefits do not only apply at the $1–10^{-7}$ probability level, as they would scale to the higher probabilities ($1–10^{-4}$ to $1–10^{-6}$) that many non-aviation applications would be most concerned with.

While very few non-aviation users of GNSS today have real-time safety requirements similar to those of civil aviation, the number of such users will likely increase as the coverage of augmented GNSS (and the availability of integrity from standalone receiver-autonomous integrity monitoring, or RAIM) expands. The evolution of standalone civil GPS usage provides a precedent: as basic GPS accuracy improved from tens of meters to several meters, and the cost of user equipment dropped, more and more uses were dis-

14 m (±7 m) using average risk
65 m
36 m (±18 m) using WAAS MOPS

▲ **FIGURE 12** Example of reduced protection levels for harbor/marina access.

covered. A similar, although smaller-scale, trend is likely to occur as the advantages of augmented GNSS become more available and better understood. The primary beneficiaries are likely to be intelligent road-transport systems, train services, and marine transportation in restricted waters.

One application where tight real-time integrity bounds would be useful is in harbor and marina entry and exit; see **FIGURE 12**, taken from a Google map of a marina in San Diego, California. Based on the earlier analysis, two typical $1-10^{-7}$ horizontal protection levels are shown: 18 meters using the unchanged WAAS MOPS approach, and 7 meters based upon modifying the broadcast bounding parameters to represent average risk (these HPLs are bounds on error in either direction, positive or negative; thus the 2-D error bounding circle has a diameter of twice the HPL).

When the resulting error bounds are compared, the relative advantage of the smaller bound for this application is immediately apparent. In general, when HPL is significant compared to potential obstacles, its significance varies with the square of HPL rather than HPL itself, as the area being protected matters more than either linear direction. In this example, the ratio of HPLs being compared is 18/7, or 2.57, but the ratio of HPL-squared is much larger: $18^2/7^2 = 6.61$.

When real-time integrity is not needed, augmented GNSS provides an easy means to guarantee or certify vehicle locations after the fact with great precision and reliability, without the need for post-processing. Vehicle and cargo tracking based on stand-alone GPS is common today, a certification of the correctness of the tracking data to probabilities suitable for legal or commercial guarantees is lacking. For this, error bounds at $1-10^{-4}$ to $1-10^{-6}$ probabilities are likely sufficient, and would allow HPLs of below 5 meters from WAAS and below 3 meters from LAAS. In some scenarios, the difference between a 5-meter and a 15-meter guarantee would be minor, but in others, it could make a substantial difference.

As noted earlier, even for uses where the required HPL (as represented by the safe error limit, or HAL, for a particular application) is satisfied by the existing WAAS and LAAS protection levels, the use of modified average-risk protection levels increases the availability of integrity, which is most often expressed as the probability or percentage of time (over all satellite geometries and other variable system states) that the integrity requirement is met throughout an operation (in simple terms, that HPL ≤ HAL). For user locations within good WAAS or LAAS coverage, the most variable element over time is satellite geometry. Decreasing HPL by a factor of 2.5 or more substantially increases the margin between HPL and HAL and makes it far less likely that the satellite geometry will degrade to the point where HPL exceeds HAL. For example, if the un-modified WAAS HPL equals HAL at an (un-weighted) HDOP of about 1.5, the resulting satellite availability (an upper bound on overall availability) for the SPS-standard 24-satellite GPS constellation would be roughly 98.5 percent. This means that the satellites in view (in this case, all satellites above 5 degrees elevation at a location in CONUS) would provide HDOP ≤ 1.5 about 98.5 percent of the time. However, the modified average-risk HPL (using the factor-of-2.5 reduction) would roughly translate into a limiting HDOP of about 3.75. This allows the required integrity bound to be satisfied by much poorer GPS geometries and gives a satellite availability of greater than 99.9 percent. Thus, when integrity is needed, this much greater availability of integrity is a major advantage.

Summary

SBAS and GBAS broadcasts are freely available to all GNSS users, most of whom will have different definitions of acceptable risk. These users are not optimally served at present and may hesitate to take advantage of SBAS and GBAS as a result.

Using years of collected data for the FAA WAAS system and analysis of the inflation factors built into the CAT I version of the FAA LAAS system, it appears that average-risk users of WAAS and LAAS would be adequately supported by protection levels that are 2 to 5 times lower than those currently derived by aviation users. The fact that two different approaches used to examine WAAS and LAAS suggest similar levels of over-conservatism lends credence to these estimates. While further validation by full-scale probabilistic risk assessments is necessary, we conclude that non-aviation users willing to accept average risk would obtain much better performance and availability from simple modifications to the existing SBAS and GBAS protection level calculations specified for aviation users.

SAM PULLEN is a senior research engineer at Stanford University, where he is the director of the Local Area Augmentation System (LAAS) research effort. He has supported the FAA and others in developing GNSS system concepts, requirements, integrity algorithms, and performance models since obtaining his Ph.D. from Stanford in Aeronautics and Astronautics.

TODD WALTER is a senior research engineer in the Department of Aeronautics and Astronautics at Stanford University. He received his Ph.D. from Stanford and is currently working on the Wide Area Augmentation System (WAAS), defining future architectures to provide aircraft guidance, and on assuring integrity on GPS III.

PER ENGE is a professor of aeronautics and astronautics at Stanford University, where he is the Kleiner Perkins, Mayfield, Sequoia Capital Professor in the School of Engineering. He directs the GPS Research Laboratory and received his Ph.D. from the University of Illinois.