

Shaping Aviation Integrity

April 1, 2008 By: [Boris Pervan](#), [Juan Blanch](#), [Livio Gratton](#), [Per Enge](#), [Todd Walter](#) GPS World

Two RAIMs for Safety

Modernization of GPS with new civil signals, and new GNSS constellations that will offer a much larger number of satellite navigation signals, bring new capabilities for navigating airplanes.

The Federal Aviation Administration (FAA) initiated the GPS Evolutionary Architectural Study (GEAS) in late 2006 to plan future navigation architectures, with the goal of creating an architecture capable of providing a service to bring airplanes within 200 feet of the ground anywhere on the globe. The architecture of choice will have implications for near-term planning for the Wide Area Augmentation System (WAAS) and the Local Area Augmentation System (LAAS) and their potential incorporation of the new signals.

In today's augmentation systems, integrity monitoring takes place on the ground. The WAAS and LAAS programs use reference receivers to measure the signal to correct small errors and alert users when faulted conditions may be present. The GPS Operational Control Segment (OCS) also monitors the satellites, identifies faulty satellites, and removes them from service. However, the OCS can take hours to respond to satellite faults, where WAAS and LAAS send alerts within seconds. Further, WAAS and LAAS protect against a larger class of faults and provide firm integrity assurances.

GPS modernization will provide a new civil signal at 1176.45 MHz called the L5 signal. By combining measurements at this frequency with ones from the original L1 frequency at 1575.42 MHz, a user can eliminate the largest current source of uncertainty, in which the ionosphere creates a variable amount of delay between the satellite and the airplane. Measuring the signals at both frequencies allows the removal of this error source. Therefore, future users will avoid this significant error source. As a result, they will enjoy higher availability and be able to operate in regions that are currently unavailable due to extreme ionospheric conditions.

Integrity determination can be made in one of three locations: on the GPS satellite using redundant components and sensors; on the ground using reference monitors; or in the aircraft using redundant signals or sensors. Currently, the satellites perform little integrity monitoring. However, the GPS III program wants to expand that capability, so that future satellites may detect the vast majority of errors and prevent their transmission.

The GEAS seeks to determine the best way to assure integrity of these modernized signals for aviation users. Currently integrity is provided for the L1 signal either by exploiting redundant signals on the aircraft using a technique called Receiver Autonomous Integrity Monitoring (RAIM) or through ground monitoring by WAAS. RAIM is used for Lateral Navigation (LNAV) of aircraft at altitude. WAAS

provides both lateral and Vertical Navigation (VNAV) and can be used to bring aircraft to within 200 feet of the ground.

Future architectures may shift more of the integrity monitoring responsibility to the satellite or the aircraft. This article will investigate the relative advantages of certain architectural concepts over others. In particular, we will focus on the issues of time-to-alert (TTA) and required constellation strength.

Architectures

The GEAS has focused on two architecture classes: one where integrity assurance is entirely external to the aircraft, and one where the aircraft exploits redundant signals to meet the TTA requirement. In the first class, integrity messages are broadcast to the airplane within the TTA. For this analysis, all such architectures that achieve this are labeled GPS integrity channels (GICs). The key feature of a GIC is that the signals arrive at the aircraft containing integrity information that meets the TTA. The aircraft does not perform a separate evaluation requiring redundant signals.

The other general architecture class still has integrity information arriving at the aircraft. However, this information arrives outside of the TTA requirement. The aircraft has to make its own integrity determination using this delayed information combined with its current measurements. Here we investigate two forms of RAIM to make this timely integrity determination on the aircraft.

There are many possible architectures to assure integrity external to the aircraft. **FIGURE 1** provides the notional concept of ground-based monitoring and satellite based messaging. WAAS and LAAS are two examples of ground-based architectures. WAAS uses a large geographic network and concentrates the information into a master station to evaluate all of the measurements and determine the necessary corrections and integrity parameters. Conceptually WAAS could be expanded to cover more of the globe. However, it would be difficult to do this and meet the TTA. The existing North American network already has nearly reached the limit for getting information to the user as required. Information from monitoring receivers placed even farther away, with longer communication times, would be challenging to incorporate within the TTA.



FIGURE 1 GIC architectures. Here integrity is determined external to the aircraft and supplied within the TTA. Integrity may be determined either on the ground or on the satellites or through a combination of both. The integrity information may be broadcast through either geostationary satellites as in WAAS or via the GPS satellites themselves.

As an alternative, many separate SBASs around the world could be expanded to obtain global coverage. In this case, worldwide coverage is achieved by many service providers collectively rather than just one. Another option would be to put enough integrity monitoring into the satellites that they themselves determine integrity and shut themselves off when sufficient integrity cannot be assured. Regardless of the specific implementation, the important feature of any of these GIC architectures, for this study, is that the aircraft is not required to make its own integrity determination. Thus, the GPS satellite constellation need only provide enough satellites and sufficient geometry to afford basic positioning. In contrast, the RAIM architectures require greater redundancy in the constellation. Not only must there be adequate numbers and geometry to support positioning, but also there must be enough to redundantly support it. That is, positioning must be supported for all satellite subsets formed by removing one satellite. This requires a greater number of satellites to be well distributed about the aircraft.

We next investigate these architectural concepts and their dependence on constellation strength. To do so we need to quantify their Vertical Protection Level (VPL) as a function of measurement confidence and satellite geometry. In order to be used for LPV-200, the VPL must be below 35 meters. Each architectural concept has a different VPL

formulation as a function of satellite geometry. The next sections will describe these in more detail.

<!--pagebreak-->

GIC VPL

We are not concerned with the exact method of providing integrity information to the aircraft at this stage. More important is the dependency on satellite geometry of these architectures. Future satellite constellations may change. More satellites with a better distribution may afford more options to provide integrity. Alternatively, a comparatively weak constellation limits the architectural options available. At this point, we would like to understand this dependency.

All of the architectures considered here rely on dual-frequency ranging measurements. The L1 and L5 signals are combined in a way to eliminate the first-order ionospheric delay. Unfortunately, this combination increases the impact of measurement noise and multipath. The measurement noise term for the j th satellite can be described as normally distributed with zero mean and variance,

$$\sigma_{j,DF_air}^2 = \left(\frac{f_1^2}{f_1^2 - f_5^2} \right)^2 \sigma_{L1,j,air}^2 + \left(\frac{f_5^2}{f_1^2 - f_5^2} \right)^2 \sigma_{L5,j,air}^2 \quad (1)$$

where f_1 and f_5 are the L1 and L5 frequencies, respectively, and $\sigma_{L1,j,air}^2$ and $\sigma_{L5,j,air}^2$ are the multipath and noise error variances affecting the

individual measurements. This dual-frequency term replaces the σ_{air} and σ_{UIRE} terms of Appendix J of the SBAS Minimum Operational Performance Standards (MOPS). The specific model for $\sigma_{L1,j,air}^2$ may also be found in this appendix. Although the performance of L5 for noise and multipath is expected to be better than that of L1, we will assume the same airborne model for this frequency. σ_{j,DF_air}^2 is a deterministic function of the elevation of the satellite.

A term will be broadcast to the user to overbound the errors in the satellite's clock and ephemeris. For GIC and RRAIM, this bound must protect to a fraction of the overall integrity budget as in SBAS. For ARAIM, however, the aircraft has some capability to detect absolute errors on its own, so the integrity requirements on the broadcast bound may be less stringent.

The user will also calculate the overbound for unmodeled tropospheric effects. The tropospheric model and uncertainty we use here are identical to those specified in Appendix A of the SBAS MOPS. The error is defined to be normally distributed with variance specified by $\sigma_{j,tropo}^2$. This variance is also a function of the elevation of the satellite. The three error components are independent, so the variance of the j th line of sight for our smoothed pseudorange measurements will be described as

$$\sigma_{j,p}^2 = \sigma_{j,clk_eph}^2 + \sigma_{j,DF_air}^2 + \sigma_{j,tropo}^2 \quad (2)$$

As part of the safety certification for WAAS and LAAS, it was determined that the protection-level equations were not ideally suited for the actual error distributions encountered. The PL

equations were based on the convolution of zero-mean Gaussians. There was no explicit provision for non-zero-mean or non-Gaussian errors. Several methods were developed to mathematically account for these shortcomings. All methods inflate confidence values to protect the worst-case user as opposed to the typical user. This imperfect matching has led to an inflation of the protection-level values that may be as much as 20 percent. As we move forward, we want to explicitly include terms to account for non-zero-means and non-Gaussian behavior. This could include broadcasting bias terms or excess mass terms. For this study, we investigate the inclusion of bias terms in the protection-level calculation that can be used to account for non-Gaussian behavior through a technique known as paired bounding.

This term bounds errors that may appear random, but that affect users in the same way repeatedly. Examples of such biases are antenna biases or nominal signal deformations. These error sources affect a particular geometry identically each time it is encountered. Thus, a maximum bias term, $b_{j, max}$, is broadcast to bound the effect of these error sources.

$$VPL_{GIC} = K(P_{HMB}) \sqrt{\sum_{i=1}^n (S_{u,i}^p \sigma_{i,p})^2 + \sum_{i=1}^n |S_{u,i}^p| b_{i,max}} \quad (3)$$

This VPL equation will be used to investigate the availability of any architecture that determines integrity external to the aircraft and provides integrity alerts within the required TTA.

<!--pagebreak-->

RRAIM VPL

Because RAIM is sensitive to both the confidence values and the geometries, a promising alternative investigates the use of RAIM with carrier phase. The uncertainty of the carrier phase is dramatically lower than for carrier-smoothed code, therefore Relative RAIM (RRAIM) based on carrier will be available for many more geometries than Absolute RAIM (ARAIM). Unfortunately, the carrier is not an absolute measure. We need a starting position and confidence in order to use it. This can be provided by a GIC architecture.

RRAIM uses any of the GIC architectures as a starting point. It then forms a position solution and protection levels. However, the latency of the GIC is such that these values are only valid for some time in the past. This is acceptable because we can use the carrier measurements to update the information to the current time. With the carrier-phase measurements, we can form a precise trajectory from the prior time, when the GIC measurements are valid, to the current time, as shown in **FIGURE 2**. If we have redundant carrier measurements, then we can cross check them using a RAIM technique to assure that the trajectory calculation is valid. Because there is some uncertainty in the satellite clock values and the troposphere over the update time interval, the protection levels will need to be increased. The longer the time interval of projection, the greater the increase. For short times, 30 seconds and less, the increase can be small and RRAIM availability will be comparable to GIC availability. For long intervals, five minutes and longer, RRAIM availability approaches ARAIM availability and the advantage of the carrier measurements diminishes.




FIGURE 2 Relative RAIM concept. A pseudorange-based position with integrity assurance is calculated for a past time. Carrier measurements are used to determine the aircraft trajectory between the past time and the current epoch, updating the position estimate. RAIM is performed on the carrier

trajectory to assure its integrity. New protection levels are calculated based upon the original values and the accumulated uncertainty over time. Faults can affect either the initial pseudorange-based position or the carrier-based trajectory. Therefore, the integrity allocation is divided between these two conditions. If the fault occurred at or before the time of the initial solution and the external GIC failed to raise an alert, the initial VPL would be in error. The likelihood of this occurrence must be below a fraction of the total integrity allocation. If the fault occurs after the initial position fix, then the carrier trajectory may be in error. Here the carrier-phase RAIM algorithm must raise an alert. The likelihood of it failing to do so must be below the remaining fraction of the total integrity budget.

The RRAIM VPL is then the maximum of two separate calculations. One looks at a fault-free coasting (FFC) condition where the concern is bounding the initial pseudorange based solution. This VPL calculation is essentially identical to the GIC equation except that the overall integrity allocation is smaller. The second VPL calculation looks at a fault during coasting (FDC). Here the initial position is good, but a fault occurs during the trajectory update interval. The FDC evaluation uses a standard RAIM technique to determine the maximum vertical impact of an undetected fault. This is then added to some nominal uncertainty around the code-phase solution (see **FIGURE 3**). The full details of the VPL calculation are beyond the scope of this article. This algorithm is still being evaluated and further improvements are likely to be found.

<!--pagebreak-->

ARAIM VPL

Traditional RAIM compares smoothed pseudorange measurements to one another to ensure that they are all consistent with each other. If there is a faulty measurement, then it should stand out from the others. Unfortunately, not all such faults are always readily apparent. Only strong geometries lead to small protection levels.

Absolute RAIM investigated here is very similar to traditional RAIM. The primary difference will be external to the aircraft. RAIM for LNAV has only one significant threat: bad satellite clock/ephemeris information. No other error source can create horizontal position errors of hundreds of meters (for any reasonable geometry). Ionospheric errors are measured in tens of meters. Worst-case multipath or tropospheric errors are even smaller. Satellite signal deformation similarly leads to constrained errors. Therefore, today's RAIM user need only worry about very large satellite clock or ephemeris errors. Historically these errors have been rare, supporting RAIM's underlying assumption that only one satellite will be faulted at a time.

Today, RAIM is used only for LNAV. The smallest HAL allowed falls just below 200 meters. RAIM is not currently authorized for VNAV. To apply ARAIM for vertical navigation, significantly smaller errors must now be considered potentially hazardous. Although the ionospheric threat can be eliminated through the use of two frequencies, the other error sources remain. It has not been established that meter-level threats to VNAV are rare. The OCS does not currently monitor the L1 CA signal and cannot assure that it is unaffected by small errors that may threaten a 35-meter VAL.

ARAIM will require that the civil signal be monitored and be free of errors greater than a few meters. The monitoring will have to be updated to include signal deformations. It will also have to be demonstrated that errors more than a few meters are extremely rare so that ARAIM too may be primarily concerned with no more than one fault at time. This will require considerably more effort in monitoring of the GPS civil




FIGURE 3 RRAIM VPL concept. The VPL calculation against a fault during coasting (FDC) combines the projected vertical error for the largest undetected fault with the nominal vertical position uncertainty of the pseudorange-based solutions.

signals than is performed by the current OCS. The level of monitoring will need to be akin to WAAS or LAAS.

However, one significant advantage will be that the external monitoring will not have to alert the aircraft within the TTA. The primary goal of the monitoring will be to maintain the extremely low fault rate. Since the aircraft can identify and remove a single fault, the external monitoring must ensure that single faults are rare and multiple faults are exceedingly rare. Unlike current augmentation systems, the monitoring to support ARAIM has tens of minutes to identify a fault and alert the aircraft (or remove the satellite from service). Thus, these monitors will be able to evaluate minutes worth of data instead of one or two seconds. Additionally, the monitors also have several minutes to get this information to the aircraft. These looser requirements create a much simpler architecture, in both the monitoring algorithms and the messaging channel to the user.

The VPL equation for ARAIM evaluated here resembles the VPL for the GIC. However, it is also evaluated for each single satellite out subset as well as the all-in-view solution. An additional term accounts for the likely difference between the all-in-view vertical position estimate and the estimate for the subset. The final VPL is the maximum value across all such geometries.

Results

The performance of the algorithms was evaluated using a set of MATLAB scripts (including scripts from the publicly available Matlab Algorithm Availability Simulation Tools) to compute the predicted VPLs for the set of users distributed over the world during one day. **TABLE 1** shows our results.

Architecture	Satellite Constellation	Satellite Constellation	Satellite Constellation	Satellite Constellation	Satellite Constellation	Satellite Constellation
GIC	24	27	30	99.5%	99.5%	99.5%
RAAIM	24	27	30	99.5%	99.5%	99.5%
ARAIM	24	27	30	99.5%	99.5%	99.5%

TABLE 1 Fraction of the Earth that achieves 99.5 percent availability or better for the different architectures as a function of constellation. Three constellations, optimized for 24, 27, and 30 satellites, were considered in this study. To further investigate sensitivity to constellation, a satellite was removed from each that had a significant impact on availability. We can see in this table that the GIC architectures operate well with 24 or more satellite constellations. ARAIM requires many more satellites to achieve the same performance. RRAIM sits in between. For short times, it is comparable to GIC; for longer times it is closer to ARAIM. All are vulnerable to outages. It is interesting to note that a constellation optimized

for 27, but missing an important satellite, performs worse than the constellation optimized for 24 despite having more satellites. Thus, it is not simply a question of the number of ranging sources. Their distribution is also very important.

Availability is calculated as the fraction of time that the VPL is below the 35-meter alert limit. Users are placed on a five-degree by five-degree grid around the world from -70 to 70 degrees (2088 locations). Geometries are evaluated every minute for a full 24-hour period (1440 epochs). Coverage is calculated as the fraction of the users that meet a 99.5 percent availability goal. The 99.5 percent availability goal was chosen as a trade between simulation time and expected fidelity of the models. Accurately determining higher availabilities often requires modeling additional effects beyond geometry and require longer simulation runs. To account for the fact that grid spacing becomes closer

at larger latitudes, each user grid contribution to coverage is weighted by the cosine of the latitude. Table 1 gives the fraction of the globe between -70 and 70 degrees where users would enjoy 99.5 percent availability of vertical guidance. The availability calculations are based on specific satellite constellations in combination with assumed numerical models for the error bounds.

Several satellite configurations are considered, and **TABLE 1** contains coverage results for three different six-plane GPS constellations optimized for 24, 27, and 30 satellites. It includes the cases with all satellites available and cases where one of the most important satellites has been removed. The latter cases are to investigate the vulnerability of the performance of each architecture to satellite outages.

In general, the clock/ephemeris and maximum bias values will be functions of the ground networks and algorithms. For this analysis, a simpler estimate of performance is obtained by using constant values that are close to the expected values for well-observed regions. For this analysis, we will assume the following values:

These values are based on performance of the satellites best observed by WAAS today and possible contributions of nominal deformations and antenna biases.

$$\sigma_{j,clk_eph} = 0.75m, b_{j,max} = 1.125m \quad (4)$$

For each time and location, the VPL was computed. As shown in Table 1, performance for the GIC is very good for all constellations considered. The 24-

satellite constellation is near the lower limit for performance, however, as even a single satellite outage can cause large regions to suffer some outage periods. Notice also that the 27-satellite constellation also has some vulnerability although the availability outages only affect a very small subset of users. Note that the 26-satellite constellation arranged sub-optimally performs worse than the optimal 24-satellite constellation despite having two more satellites. This holds true for the other two architectures as well. It is not simply a matter of the number of healthy satellites in the constellation; their orbital location in relation to one another is also very important. A single outage can create a gap in coverage.

As expected, ARAIM is more sensitive to the constellation quality. It does not achieve high values for the current 24-satellite optimized constellation. It requires a constellation optimized for 27 or 30 to obtain good performance. RRAIM with short latency is much closer to the GIC performance. The additional fault screening causes a small loss in coverage, but overall, RRAIM performs well for all three constellations. Like ARAIM, it strongly benefits from having a stronger constellation.

<!--pagebreak-->

Conclusions

From this analysis, it is possible to gain some relief against the six-second TTA requirement through the use of RAIM. However, this relief comes at a cost. The number of satellites needed to support high-availability RAIM is greater than for a GIC architecture. While GIC provides high availability with the

current constellation, either of the RAIM solutions will require more satellites. RRAIM with short latency can come close to GIC performance and only require a few more satellites. RRAIM with long latency and ARAIM require something approaching 30 satellites to achieve very high availability.

Of course, it is not just a matter of the number of satellites. The current constellation has 32 satellites. However, they are not optimally arranged and several may be marked unhealthy at a given time. This constellation matches the one optimized for 24 satellites and the extra eight are merely redundant to certain positions. To truly take advantage of the extra satellites, we would need them more optimally placed for the larger number. As can be seen in Table 1, 26 satellites sub-optimally arranged perform worse than 24 optimally placed.

One of the most important areas of study for the GEAS in the next phase is the development of the transition strategy. As the GEAS develops the details for the final architectural construct, it is crucial to map the path from the current infrastructure to this desired endpoint. It is critical that the GEAS does not impede the increased use of satellite navigation currently underway. Ideally, the eventual implemented architecture would maximize use of the existing user base. We must investigate paths that encourage the current uptake of L1-only GPS, WAAS, and LAAS and ultimately provide even further capability when there is a full constellation of L5 satellites. The final architecture must be compatible with existing users and support them as part of a reversionary mode should L5 not be available. The GEAS will develop specific requirements associated with reversionary mode to support LPV operations as part of the definition of transition to L1 and L5 operations. The transition from L1 to L1 and L5 is the most important consideration in choosing our recommended architecture and must be evaluated carefully in the upcoming years.

In the longer term, the GEAS will investigate combining dual-frequency GNSS with other sensors. Integration with a precise clock and/or inertial sensors may provide better continuity performance, by providing resistance to interference and scintillation. Altimeters (barometric, radar, or laser) add an additional measurement and provide relief against constellation weakness, but at the cost of complexity. The GEAS will examine these trades to determine if they have merit and warrant more detailed investigations. A long-term goal of the GEAS is the provision of Cat II/III capability worldwide. It is highly desirable that the final architecture offer an upgrade path to support this goal. The requirements on Cat II/III are very stringent and will be very challenging to meet. The GEAS must investigate the implications of these requirements and understand how the recommended architectures contribute to the satisfaction of these needs.

<!--pagebreak-->

Acknowledgments

The authors gratefully acknowledge contributions from all GEAS participants. Many of the discussions at these meetings helped shape this article, originally presented at the Institute of Navigation (ION) National Technical Meeting, January 2008. The work was supported by the FAA Satellite Product Team under research grant 95-G-005.

Further Reading

For RRAIM VPLS: "Worldwide Vertical Guidance of Aircraft Based on Modernized GPS and New Integrity Augmentations" by Walter, Enge, Blanch, and Pervan, submitted to *Proceedings of IEEE*.

For ARAIM VPLS: "An Optimized Multiple Solution Separation RAIM Algorithm for Vertical Guidance," by Blanch, Ene, Walter, Enge, *Proceedings of the ION GNSS 2007*, Fort Worth, Texas, September 2007.

TODD WALTER is a senior research engineer at Stanford University. He has been active in the development of the Wide Area Augmentation System and related systems around the globe. His focus is on the provision of certified integrity for aviation applications.

JUAN BLANCH is a research associate at Stanford University, where he works on integrity algorithms for GNSS. A graduate of Ecole Polytechnique in France, he holds a Ph.D. in aeronautics and astronautics from Stanford.

PER ENGE is a professor of aeronautics and astronautics at Stanford, where he directs the Stanford GPS Research Laboratory. He received the Kepler, Thurlow, and Burka Awards from the Institute of Navigation, and his Ph.D. from the University of Illinois.

BORIS PERVAN is associate professor of mechanical and aerospace engineering at the Illinois Institute of Technology and is editor of ION's *Navigation* journal.

LIVIO GRATTAN is a Ph.D. candidate at the Illinois Institute of Technology. He has developed several algorithms for ephemeris and atmospheric threats mitigation for precision approach systems.
