

A Framework for Analyzing Architectures that Support ARAIM

Todd Walter, Juan Blanch, and Per Enge,
Stanford University

ABSTRACT

Advanced Receiver Autonomous Integrity Monitoring (ARAIM) has been studied in recent years as method of providing vertical guidance for aircraft [1]. The aircraft compares the various ranging measurements that it makes to different satellites to ensure that they are consistent with each other. However, for the aircraft to meet its integrity requirement, the satellites must perform within a certain set of expectations. While the satellite providers might be able to guarantee this performance, in order to meet aviation integrity requirements and to retain sovereign control, it is more likely that individual aviation authorities will independently perform monitoring of satellite performance.

This paper examines important components for different architectural choice such as monitoring network size and density, the method for getting information to the aircraft including latency of information, the handling of consistent faults, and the methodology for demonstrating integrity. We examine two extreme architectures to better explore the trades between these and other important performance parameters. One architecture is based on Space Based Augmentation System (SBAS) assets and relies very heavily on ground monitoring. Another architecture minimizes the ground requirements and is more similar to today's Receiver Autonomous Integrity Monitoring (RAIM). By examining these architectures, we will better understand the impact of possible choices and determine which are most important to lock down early and which can be left open to flexibility. Ultimately, this allows us to characterize other architectures that exist between the two extremes.

INTRODUCTION

Given an assumed set of fault modes, responsibility for mitigating each can be assigned to the aircraft, the ground, the space segment, or some combination of all. A particular ARAIM architecture will make such an assignment and then assess whether each segment achieves its goal. An architecture that puts little trust in the satellites, or the ground, may assume a very high

probability of signal-in-space failure and require the aircraft to have exceedingly good geometry to detect all possible faults. Another architecture that places more trust in either the ground or the space performance can operate with comparatively weaker geometries, but at potentially greater expense in the system operation.

Different entities may make different architectural choices. Those that have invested in SBAS may choose to make re-use of the monitoring networks and delivery channels. Those that have not invested in monitoring networks may opt to put more burden on the aircraft. This paper seeks to compare such architectural choices and highlight their commonalities and differences. Ideally we will identify a structure that allows such differences to co-exist without incurring significant complexity in the avionics, thus retaining options for the service providers.

The European Union (EU) and the United States (US) have an agreement establishing cooperation between GPS and Europe's Galileo system. As part of this cooperative agreement a subgroup was formed to investigate the benefits of Advanced Receiver Autonomous Integrity Monitoring (ARAIM) [2]. Over the last two years, this EU-US ARAIM subgroup has identified key issues affecting the potential use of ARAIM. This paper highlights several of the key characteristics that need to be evaluated by any architecture intended to support ARAIM. These characteristics also demonstrate how different architectures may compare against each other based on their approaches to addressing the key issues. All ARAIM architectures contain three distinct elements: the space component, the ground component, and the airborne component. The space component consists of the core Global Navigation Satellite System (GNSS) constellations and accompanying performance commitments. The ground component consists of the reference-monitoring network, a coordinating facility that collects the raw data, processes it, and sends the results to the aircraft. The airborne component collects its own raw data and processes it with the ground information to determine the aircraft position and confidence bounds. Different architectures may make different choices about how to spread responsibility across the components. This

paper seeks to identify these key choices, creating a framework for direct comparison of similarities and differences.

There are several common elements across all considered architectures that are not specifically studied as part of this paper. Included in these are the threats that must be mitigated by the system. A separate effort has been made to identify a high level list of threats that any architecture needs to evaluate [2]. An obvious exception is in the case where a particular architecture introduces a unique vulnerability that isn't present for other architectures. Other parameters that are considered external to the architecture are: constellation strength, satellite bias and confidence parameters, and fault probabilities. While these have an important influence over the performance and viability of each architecture, they are considered input parameters rather than architectural properties.

Because they will have a significant impact on evaluating the relative merits of the architectures, it is worth describing each in more detail. Constellation strength refers to the total number and distribution of useful satellites available to the user. It is often measured in number of constellations, numbers of satellites per constellation, and geometrical diversity. Although more constellations and more satellites are generally considered favorable, the satellite locations relative to each other are important also. It is not automatic that more satellites lead to better availability, although such should be the case if the satellites are well distributed.

Each satellite has an expected error distribution that can be characterized by four values: a nominal bias, an accuracy bound, an integrity bound, and a probability of fault [1] [2] [3]. The nominal bias is an upper bound on nominal, uncorrectable errors present on the satellite's ranging signal. The bias arises primarily from satellite antenna group delay variations and small deformations in the signal structure. The nominal one-sigma error about this bias bound is known as the User Range Error (URE). It is typically valid for 95% or more of the observed errors and is useful for indicating satellite accuracy. The User Range Accuracy (URA) is a one-sigma number that typically bounds 99.99% or more of the errors and is used to indicate confidence in the integrity of the satellite. The probability of satellite fault (P_{sat}) describes the probability that a fault may exist on the satellite (independently from one satellite to another). A final parameter is one that describes the probability that a fault mode may affect more than one satellite within a constellation (P_{const}). To be conservative it is often assumed that all satellites in the constellation may be faulted. These parameters may have

a significant effect on the performance of the evaluated architectures.

KEY ARCHITECTURAL PROPERTIES

The key architectural properties identified in this paper are:

- Bounding methodology
- Communication and computation latency
- Broadcast methodology
- Integrity Service Message (ISM) contents
- Handling of constellation faults
- Reference network

These properties are strongly interconnected to one another. Making a particular choice in one may strongly encourage particular choices in others. These properties have been selected because different solutions have been discussed for each during our ARAIM investigations. In some cases, there may be multiple valid approaches, but further analysis is required. These properties are useful in distinguishing the different architectures and are described in greater detail below

Bounding Methodology

The first property examines how data is used to support the airborne algorithm. The bounding methodology analyzes the trade among the integrity responsibility of the space, ground, and the airborne components. As more trust is placed in the space segment, the less the ground segment is needed to determine integrity. Conversely, if less trust can be placed in the space segment, then the ground segment is needed to meet the target level of reliability assumed by the airborne segment. As there may be four independently operated core GNSS constellations, each at different levels of maturity, and with different performance commitments, there may be differing levels of ground monitoring required within the same system.

Table 1 illustrates the range of potential ground integrity bounding requirements. At one extreme, the space segment could be trusted to fulfill the ARAIM requirements on its own, without any ground monitoring.

None	Off-Line Determination	Real-Time Determination
------	---------------------------	----------------------------

Table 1. *Potential Bounding Methodologies*

At the other extreme, the space segment is only trusted for a relatively short interval and it requires real-time ground monitoring to ensure that the assumed satellite performance characteristics continue to be met. In between the two extremes, the space segment is trusted to operate as expected for relatively long intervals. However, some performance changes may be expected to occur slowly as the GNSS ages and evolves over time. The ARAIM provider has the ability to monitor performance, but does not need to react in real-time to potential changes. All ARAIM architectures require some level of trust in the space segment, at least in the short-term.

Communication and Computation Latency

This property concerns the data collection frequency, the rate at which data is returned for processing, the amount of time it takes to make a decision on current performance, and the amount of time it takes to get this decision to the aircraft. These factors affect the overall delay of the Integrity Service Message (ISM). We define a new term here, called the Time-to-ISM-Alert (TIA). This is the time it takes for the ground network to identify an issue in the space segment and alert the aircraft to that issue. This is distinct from the normal Time-To-Alert (TTA) which is the amount of time it takes for the system as a whole to identify and remove a problem or to alert the pilot that the system can no longer safely meet its function. The TTA for the targeted operations of ARAIM is 6 seconds. This will be met through actions in the airborne component. However, the ground segment does support the air component in performing its function. The airborne component may need to assume that certain error sources can persist for only so long, that the assumed parameters are valid, or that certain slowly-growing errors are detected on the ground before their magnitude becomes a concern. The time it takes the ground to alert the aircraft of such problems is the TIA.

Table 2 illustrates the trade space for the TIA ranging from essentially no ISM update (years) to a six second value. This latter value essentially places the entire integrity burden on the ground and the communication channel making the ISM fully responsible for meeting the TTA.

Years	Months	Days	Hours	Minutes	6 Seconds
-------	--------	------	-------	---------	-----------

Table 2. *Time-to-ISM-Alert (TIA) Trade Space*

None	At Dis-patch	At Arrival	Intermittently En Route	Continuously
------	--------------	------------	-------------------------	--------------

Table 3. *Broadcast Methodology Trade Space*

Broadcast Methodology

The next key property outlined in this document is closely related to the previous two. This property is the method for broadcasting the ISM to the aircraft. The method chosen will have an impact on TIA. Also affected will be the coverage area. Further there may be sovereignty issues as the aircraft may originate in one country and land in another. Several possible methods for broadcast have been discussed: cockpit communication data channels, local area VHF broadcasts, geostationary satellite downloads, etc. It is possible that more than one broadcast channel may be chosen. In this case, it is important to define message packets for the ISM that may be easily accommodated on different channels. Table 3 describes some potential broadcast strategies for the ISM.

Integrity Service Message (ISM) contents

Another important aspect is to be broadcast from the ground to the aircraft. This will be closely tied to the level of ground responsibility in providing integrity. The greater the responsibility, the more information likely will need to be broadcast (and the more often the data may need to be updated). Table 4 shows the range of options for message content. At the simplest level nothing needs to be broadcast because all of the data coming from the satellites themselves have sufficient trust. One step up from that, the ground monitoring either confirms or rejects that the satellite data is currently valid. The next two options potentially enhance performance by sending more information about current level of performance supported by the ground monitoring (a slightly degraded satellite can be so indicated rather than forcing a binary yes/no decision). Finally, certain threat modes may be

None	1 bit Health per Satellite	Health per Satellite, Other Parameters per Constellation	Parameters per satellite, & P_{const} per Constellation	All Parameters Plus Ephemeris per Satellite
------	----------------------------	--	---	---

Table 4. *ISM Content Trade Space*

None	Slow, Independent, and/or < 3D	Fast, Independent, and 3D	Common Across All Constellations
------	--------------------------------------	---------------------------------	--

Table 5. *Constellation Fault Trade Space*

eliminated altogether (e.g. EOPP) if the full ephemeris for each satellite originates from a trusted source (at the cost of greater required bandwidth and ground processing).

Handling of Constellation Faults

An important aspect of the integrity methodology is the handling of constellation wide faults. Although this can be seen as a subset of the overall integrity approach, it is of sufficient importance to warrant its own discussion. Constellation wide faults are faults that may affect more than one satellite within a core constellation. Such faults are distinct from satellite faults, which affect each satellite independently. In a constellation fault, a single cause will lead to significant errors on more than one satellite. Several potential mechanisms for constellation faults affecting only a single satellite have been identified [2] [3]. These are in the process of being evaluated for their likelihood and effect, but if accepted will require some special actions on the ground, in the air, or at both.

Of even greater concern is the potential for cross-constellation faults, that is, a single fault leading to common errors across all constellations. To date, only a single cross-constellation fault has been identified: errors in the Earth Orientation Predication Parameters (EOPP). This fault is particularly damaging because there is no means to detect it in the air if all satellites are identically affected. The simultaneous EOPP fault must either be ruled out as a viable threat or there must be some form of other monitoring to eliminate it.

Certain fault properties have been identified as more easily detected by one component vs. another. Depending on which faults are accepted as valid and on their behavior, certain architectural choices may become preferable. Some key features that have been identified are the rate of growth of the error (sufficiently slowly growing errors may effectively be mitigated by the ground), whether they affect just one constellation or can affect all, and if they can have any error signature or if their impact is limited (e.g. to a rotation about the Earth's axis rotation). Table 5 illustrates this range of possibilities.

None	Single	Sparse Regional	Dense Regional	Sparse Global	Dense Global
------	--------	--------------------	-------------------	------------------	-----------------

Table 6. *Reference Network Trade Space*

Reference Network

This property describes the overall approach to collecting the raw data from the core GNSS constellations. The reference network properties includes aspects such as: the number of stations, the geographical spread of the network, and the level of redundancy and reliability at each station. Other important considerations include the maintenance of the network, i.e., if it is a dedicated network for ARAIM, or if these are shared receivers that primarily serve another function. ARAIM architectures could span a wide range of possible densities. As shown in Table 6, the range could go from having no dedicated real-time ground monitoring all the way to having very dense global coverage.

EXAMPLE ARCHITECTURES

To illustrate the extreme ends of potential architectures we will analyze the characteristics of two different architectures already approved for GNSS guidance of aircraft. These are: Receiver Autonomous Integrity Monitoring (RAIM) [4], which is approved for lateral navigation and Spaced Based Augmentation Systems (SBAS), of which the Wide Area Augmentation System (WAAS) [5] is an example. SBAS is approved for both lateral and vertical guidance.

RAIM Based Architecture

RAIM is approved for the en route, terminal, and non-precision approach phases of flight. Generally the aircraft is at higher altitude and the consequences of inaccurate navigation are considered major, but not hazardous. The aircraft is kept well separated from obstacles and other aircraft. RAIM places complete trust in the GPS performance commitments, i.e. that there will not be simultaneous very large on multiple GPS satellites. If large satellite faults do occur, they are assumed to affect only a single satellite and are identified by the Master Control Segment (MCS) of GPS, typically within an hour and at most six hours.



Figure 1. An overview of the RAIM architecture. The satellites are monitored by the GPS operators. No separate ground monitoring is needed.

RAIM has the aircraft perform a consistency check and compute a protection radius based on the largest undetectable error affecting any one satellite. There is no independent ground monitoring of the satellite performance and no need for reference networks. Therefore, there is no integrity message to the aircraft (save the navigation data from the GPS satellites themselves) and no separate TIA (the MCS has a maximum TIA of 6 hours). RAIM is also global as it can be made to work wherever there is sufficient satellite coverage (see Figure 1).

Although there is no ISM, the FAA does monitor GPS performance using reference receivers from other networks. It evaluates the performance of the satellites vs. the performance commitments and RAIM assumptions. If GPS is found to be deficient, the approval of RAIM could be revoked. However any such action may take from days to months depending on the urgency. Fortunately, GPS has always historically been observed to significantly exceed its performance commitments.

However, not all Air Navigation Service Providers (ANSPs) have approved RAIM for air use. This is because liability of the service rests with the ANSP and not the MCS. Not all ANSPs are willing place the required level of trust in the operation of GPS or do not want to become dependent on a service over which they have no control.

SBAS Based Architecture

The SBAS architecture consists of a dense regional network of monitoring reference stations, as illustrated in Figure 2. These stations send their information back to one or more master stations that evaluate all of the data and forms corrections and confidence bounds. These are sent to the user via geostationary satellites using a signal

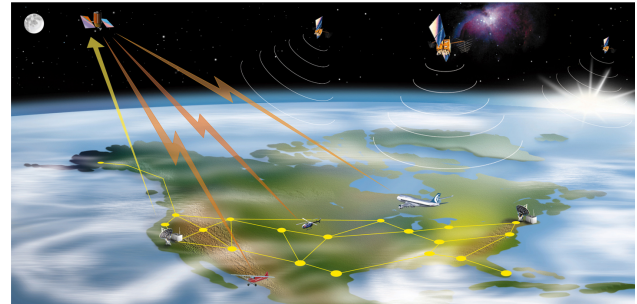


Figure 2. An overview of the SBAS architecture. The satellites are monitored by a regional monitoring network. Corrections and integrity are continually sent to the users including any necessary alerts within 6 seconds.

very similar to the GPS L1 civil signal. The SBAS is able to alert the aircraft within six seconds if any of the satellite signals or previous information is deemed unsafe.

The TIA of the ground monitoring is six seconds and therefore meets the TTA for the system as well. The corrections and integrity messages effectively form the ISM and although the safety proof for SBAS relies on GPS system performance, there is much less trust required in the operation of the satellites as the ground can detect far more GPS fault conditions than the airborne algorithm in RAIM.

COMPARISON OF ARCHITECTURES

The RAIM and SBAS architectures occupy opposite extremes of the parameters identified in the previous sections. In almost all cases RAIM corresponds to the far left box in Tables 1-6 while SBAS corresponds to the far right box. Due to the lower cost of operation and global nature of RAIM there is a strong desire to push ARAIM architectures as far to the left as possible. However, the increased hazard level of vertical guidance and increased liability to be assumed by the ANSP tends to push architectures to the right. The final ARAIM architecture (or architectures) will likely be in between these two extremes.

The two approaches to bounding methodology are very different. RAIM relies entirely on the performance commitments of GPS and on its performance history. Real-time data is not actively used to ensure its integrity. Instead data is collected off-line and evaluated by hand to verify that it conforms to expectations. SBAS, on the other hand, collects real-time data and assumes that the constellation is performing no better than what it can confirmed that instant. Performance commitments and service history inform the threat models used to formulate confidence. However, SBAS is not bound by these

factors, it anticipates worse performance and threats than have actually been observed.

For ARAIM it would be ideal to perform the integrity analysis off-line as is done for RAIM, but that would require the ANSPs to have full confidence in each core constellation that they approve. Currently only GPS has substantive performance commitments and an established history of meeting them. Such confidence is unlikely to be placed in the other constellations until several years after achieving their full operating capability and publishing meaningful commitments. Even GPS has not published commitments for its L1/L5 service nor established any service history for it yet. Further, different political alliances of the ANSPs may cause inherent distrust of some or all core constellations. The fact that the ANSPs have to assume liability for any services they approve using ARAIM, may strongly push them in the direction of having real-time monitoring under their own sovereign control.

The type of monitoring ties closely with the TIA. Since RAIM has no independent real-time monitoring it effectively has a very long TIA. RAIM approval could be removed if problems are identified with GPS performance, but there is not a specific mechanism to enact this removal with a strict upper bound on the amount of time taken. Likely, problem reports would come from the pilots themselves or due to off-line observations and then a manual decision would have to be made to issue a notice to airmen (NOTAM) de-authorizing the use of RAIM. Fortunately, this has never happened nor is it ever expected to occur due to the inherent trust in the operation of GPS.

As mentioned previously, RAIM is only authorized for lateral guidance with large protection regions around the aircraft. In order to provide vertical guidance and bring the aircraft close to obstacles (including the ground and other aircraft), a very different method is used. SBAS and the Ground Based Augmentation System (GBAS) [6] use ground monitoring that have upper bounds on the TIA measured in seconds. For these systems, the satellite performance is observed in real-time and an automatic and near instantaneous decision is made on the level of performance of the GPS satellites. The levels of confidence that WAAS is willing to assure are larger than those guaranteed by the satellites and the number and types of failures go beyond what is described in the performance commitment. If ARAIM is to be used for vertical guidance, ANSPs are likely either require independent monitoring or stronger performance assurances from the core constellation providers.

As mentioned above, RAIM has no dedicated broadcast methodology. If a problem were to be found, it would likely be announced through a combination of NOTAMs that are checked prior to dispatch and communication from air traffic controllers. The SBAS method is through a continuous broadcast from a geostationary satellite. Thus all aircraft in the service region are continually updated as to the integrity status of the satellites within seconds. The SBAS contents include specified levels of confidence per each satellite used to form the position solution. It also includes corrections to reduce moderate errors and achieve tighter confidences. This is required because vertical errors must be kept below tens of meters at the very most.

RAIM only uses one constellation and assumes that no more than one satellite will have a corrupted ranging signal at any given time. Therefore, it considers the possibility of a consistent constellation wide fault to be sufficiently unlikely as to be negligible. Of course, to be a threat to lateral operations, the fault would have to lead to a position error hundreds of meters in error. SBAS on the other hand, has the ability to detect and flag constellation wide faults as it has independent ground truth data and can observe consistent faults.

The RAIM network is potentially global as data from all available reference sites may be used in the analysis of historical performance. These sites are not specific to RAIM and indeed exist to serve other purposes. SBASs generally have dense, dedicated, regional reference networks that offer both security and sovereign control over the operation of the system. The type of reference network is likely determined by the bounding methodology and the size by the intended coverage region.

POTENTIAL ARCHITECTURE

The two extreme architectures described above each have significant advantages and disadvantages. It may be impossible to create a single architecture that meets all desired goals. The RAIM architecture is superior in its simplicity and global coverage. The SBAS architecture is superior in its ability to handle integrity, liability, and sovereignty issues. The best architecture may include elements of each approach. One method to accomplish this goal would consist of an international global mode to support lateral navigation and separate local modes to support vertical navigation.

The international version could be coordinated through the International Civil Aviation Organization (ICAO) and

consist of a common ISM that may be used for lateral navigation. This ISM would be supported by off-line monitoring and would change very slowly. The ISM may take days to months to change. However, the main advantage is that it would include all core constellations and enjoy international support. It also would allow for the aviation community to specify the integrity status of satellites separately from the constellation operators. This would allow constellation operators to permit worse performing satellites to be used by GNSS communities who can benefit from the extra ranging signals without harming aviation users. Further, many countries today do not allow the use of GNSS either for political reasons or due to a lack of trust. Having an ISM that includes all core constellations and is backed by ICAO may encourage greater adoption than RAIM.

While many ANSPs may be willing to use this approach for lateral navigation, they may be less willing to extend it to vertical navigation. Therefore a second mode with its own ISM could be used for the approach phase of flight. Here the ISM could be transmitted locally (for example a VHF link at the airport or a geostationary satellite). It will be provided by the local ANSP and be matched to the desired approach. The ISM could be generated from either SBAS or GBAS assets and therefore have a relatively short TIA (perhaps minutes to hours). This method will provide the ANSPs with the ability to monitor satellite performance in real-time and therefore be more willing to accept liability for approving the procedure.

The local monitoring is less suited to en route navigation as different ANSPs may monitor and allow the use of different constellations. Consequently an aircraft at altitude may need to change which satellites it is tracking and how it uses them each time it enters a different airspace. This introduces unnecessary complexity into the airborne algorithms. Adoption is likely to be higher with a common set of parameters used for en route and switching to the local ISM only when planning for approach.

SUMMARY

Since the airborne algorithms for ARAIM are comparatively mature [3], it is time to start investigating the full architecture needed to support operation. As different architectures are proposed, it is important to understand the possible trade-offs in terms of cost and performance. It is also essential to know which approaches will be acceptable to certification authorities.

This paper proposes some key properties to be evaluated and examines two example architectures where GNSS is used for navigation in aviation today. These two architectures: RAIM and SBAS, occupy opposite ends of the spectrum on several of these properties.

By examining relative advantages of each example architecture we can better understand the trade-offs selected in each case and see which direction we would like to pursue for ARAIM. Finally, we propose including elements of each method depending on the phase of flight. While this has the disadvantage of multiple selected architectures, it does appear that each method is much more ideally suited to specific modes of flight.

ACKNOWLEDGMENTS

The authors would like to gratefully acknowledge the FAA Satellite Product Team for supporting this work under Cooperative Agreement 2008-G-007. However, the opinions expressed in this paper are the author's and this paper does not represent a government position on the future development of ARAIM.

REFERENCES

- [1] Walter, T., Blanch, J., Enge, P., Pervan, B., Gratton, L., "Future Architectures to Provide Aviation Integrity," Proceedings of the 2008 National Technical Meeting of The Institute of Navigation, San Diego, CA, January 2008, pp. 394-401.
- [2] Blanch, J., et al., "A Proposal for Multi-constellation Advanced RAIM for Vertical Guidance," Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2011), Portland, OR, September 2011, pp. 2665-2680.
- [3] Blanch, J., Walter, T., and Enge, P., "Advanced RAIM user Algorithm Description: Integrity Support Message Processing, Fault Detection, Exclusion, and Protection Level Calculation," Proceedings of ION GNSS 2012, Nashville, TN, September 2012.
- [4] Institute of Navigation, "Global Positioning System: Papers Published in NAVIGATION Volume V, RAIM: Requirements, Algorithms, and Performance," 1998.
- [5] Lawrence, D., Bunce, D., Mathur, N. G., and Sigler, C. E., "Wide Area Augmentation System (WAAS) - Program Status," Proceedings of the 20th International Technical Meeting of the Satellite Division of The

Institute of Navigation (ION GNSS 2007), Fort Worth, TX, September 2007, pp. 892-899.

[6] Swider, R., Braff, R., and Wulschleger, V., "The FAA's Local Area Augmentation System (LAAS)," Proceedings of the 1997 National Technical Meeting of The Institute of Navigation, Santa Monica, CA, January 1997, pp. 33-42.