Reduced Subset Analysis for Multi-Constellation ARAIM

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

ABSTRACT

Advanced Receiver Autonomous Integrity Monitoring (ARAIM) seeks to exploit the large number of ranging signals provided by the multiplicity of Global Navigation Satellite System (GNSS) core constellations. Including more satellites increases the number of potential fault modes and the number of subset solutions that need to be evaluated by the receiver. In this paper we describe methods for reducing the total number of subsets to be evaluated. This approach neglects many of the less likely scenarios and uses subsets that exclude entire constellations to evaluate other multiple satellite fault scenarios. We further specify which subsets to evaluate given certain ranges for the probability satellite and constellation fault modes. By specifying a reduced number of overall evaluations, and the conditions that separate using one set of subsets versus another, we can dramatically simplify the overall user airborne algorithm. This approach reduces the computational load for each given set of satellite and constellation fault probabilities.

INTRODUCTION

ARAIM performance greatly improves by having more satellites in view to each user [1]. There is considerable availability and continuity to be gained by tracking more than two constellations. The Global Positioning System (GPS) has proven to be extremely reliable over at least the last ten years. However, many of the new core constellations do not yet have an established long-term track record of performance. It is quite likely that initially the new constellations will be less reliable than GPS. Accordingly, the ARAIM algorithms may be configured to assume a high likelihood of faults for these new constellations.

ARAIM ensures integrity by comparing the position solution estimate made with all satellites in view to estimates from subsets that have removed some of the satellites. Integrity is ensured provided one of the evaluated subsets has no faulted satellites. Subsets are formed by removing all sufficiently likely satellite combinations that could contain faults. Having a large number of satellites with a relatively high fault rate can lead to a very large number of satellites subsets to evaluate. Potentially, thousands of combinations need to be evaluated.

ARAIM does not only consider cases that postulate independent satellites faults. It also considers the possibility that a single fault affects more than one satellite within the constellation. Therefore, it may be essential that ARAIM evaluate subsets that remove whole constellations. This paper demonstrates that such subsets can be used to test many possible fault cases at once, greatly reducing the total number of subsets that need to be evaluated.

After determining the smaller groups of subsets to be evaluated for different scenarios, we then analyze the resulting algorithm availability given the subset selection. We use our Matlab Availability Analysis Simulation Toolset (MAAST) [2] to investigate ARAIM availability. These results are compared to an algorithm that implements all of the sufficiently likely subsets. We demonstrate that reducing the computational cost does not lead to significantly reduced availability. We further describe the best strategies to use in selecting subsets given the number of constellations and satellites, in view along with the fault probabilities. This approach provides a significant simplification of the overall ARAIM user algorithm.

MULTIPLE HYPOTHESIS SOLUTION SEPARATION USER ALGORITHM

An example ARAIM user algorithm has been described in literature called the Multiple Hypothesis Solution Separation (MHSS) algorithm [3]. We will use MHSS to investigate the performance impacts caused by reducing the number of subsets to evaluate. The method elaborated here can be equally applied to other user algorithm approaches, as the key element is that tests can be constructed that evaluate many fault cases at once. The specific algorithm that we use has been developed in cooperation with several other research organizations. Matlab scripts implementing this algorithm can be downloaded from the Stanford University GPL lab website [4].

The key inputs to the MHSS algorithm are:

- User geometry matrix,
- Satellite ranging accuracy values (σ_{URE}),
- Satellite ranging integrity bounds (σ_{URA}),
- Satellite bias bounding values (b_{nom}),
- Satellite pseudorange residuals,
- Satellite probabilities of fault ($P_{sat,i}$), and
- Constellation probabilities of fault ($P_{const,j}$)

In turn the algorithm produces many outputs:

- Position estimate
- Horizontal Protection Limit (HPL)
- Vertical Protection Limit (VPL)
- Effective Monitor Threshold (EMT)
- Expected vertical accuracy (σ_{acc})
- Consistency metric (χ^2)

Internally, the algorithm determines which subsets to evaluate and computes position estimates and matching covariance matrices for each subset. The covariance matrices are used to form thresholds on the possible differences between the subset position estimates and the all-in-view position estimate. These thresholds contribute the protection levels and EMT. They are also used to ensure that the subsets solutions are consistent with the all-in-view. If any subset solution differs from the all-in-view by more than the threshold, the protection levels are set to infinity. It is possible to perform exclusion and attempt to identify and remove the faulty satellite(s), however exclusion is beyond the scope of this paper.

The outputs may be compared against required limits to ensure that the input conditions lead to a good position solution estimate that meets the needs of the intended operation. The ARAIM studies to date have focused on meeting the requirements for LPV-200 procedures [1]. These requirements are:

- $VPL \leq 35\ m$
- $HPL \leq 40\ m$
- $\sigma_{acc} \leq 1.87\ m$
- $EMT \leq 15\ m$

Provided all of the above inequalities are met, the LPV-200 operation is considered available for the corresponding conditions.

The validity of the protection levels and the EMT depend on at least one of the subsets containing only unfaulted satellites. Provided such a subset exists, then the nominal covariance for that position estimate should correctly describe a region containing the true position. This subset solution is compared to the all-in-view solution, which may contain faulted satellites (and therefore an incorrect covariance estimate). If they disagree by more than the internal threshold, the all-in-view position estimate is declared invalid and not used. If they agree to within the threshold, the sum of the threshold and the subset covariance error bound is sufficient and the all-in-view position error will be bounded by the protection levels. The EMT verifies that the internal thresholds as being sufficiently tight.

It is possible that the accuracy estimate is not correct under faulted conditions. The accuracy estimate assumes that no fault is present in the all-in-view solution. Forming a fault-free subset does not affect this metric. However, faults should be sufficiently unlikely so as not to affect the 95% accuracy requirement and the driving requirement on σ_{acc} is specifically defined for fault-free conditions. Nevertheless further evaluations of these metrics should be made to ensure that the error distribution for ARAIM meets the true operational goals.

ERRORS, FAULTS, AND FAULT PROBABILITIES

ARAIM allows for the possibility that some of the ranging measurements that are made to the satellites may be incorrectly described by the parameters σ_{URE}, σ_{URA}, and b_{nom}. If the satellite is in a state where these are not accurate descriptions of the expected pseudorange errors, the satellite is considered faulted. This paper is primarily concerned with the possibilities of having one or more faulted satellites and the likelihood of being in such a state.

An unfaulted satellite can also be described as being in a nominal state. The user will experience errors on their pseudorange measurement to the satellite, but the probability distribution of these errors can be well described with a few simple parameters. The pseudorange may be biased due to look-angle dependent code-phase biases caused by the satellite and user antennas [5]. The pseudoranges may also be biased by small differences in the signal structure known as signal deformations [6]. Other bias sources may also be possible. The sum of these biases is expected to be no greater than b_{nom} in magnitude.

The pseudoranges are also expected to be affected by random errors that change over time and location. The sum of these errors can nominally be described by a zero-mean Gaussian with a standard deviation of σ_{URE}. These errors may vary over time and location. Under the very worst conditions, the error distribution can be overbounded [7] by a zero-mean Gaussian with a larger standard deviation of σ_{URA}. These three parameters may be used to describe the nominal signal-in-space errors from the satellites. There are also propagation errors from the troposphere [8] and local errors due to multipath and receiver noise at the aircraft [9] that are described by standard models. For the purposes of this paper, these latter sources only produce nominal errors and do not have faulted states. When describing fault-free maximum-likelihood distributions, the pseudorange error may be defined by a zero mean Gaussian using the URE. When describing fault-free integrity bounds, the error is defined by a Gaussian with bias b_{nom} and standard deviation using the URA.

When a fault is present, we assume that the pseudorange can take on any possible value. It is often modeled as a bias that would create a positioning error right at the threshold value, but it could be either smaller or larger. In practice, it can be difficult to distinguish a small fault from nominal conditions. However, for purposes of this paper, we can assume that the majority of the time that the satellite is in an unfaulted state as described by the nominal parameters, and a small fraction of the time it is in a faulted state described by the same parameters, but also with an arbitrary bias.

There are many different possible causes of faults on the pseudorange measurements. We will group them into three fault classes; Narrow, Wide, and Ultra-Wide. Narrow faults are ones that affect a single satellite only. These may be caused by a faulty component on the satellite that in no way affects any of the other satellites. These faults are considered to be independent from one satellite to another. That is, an observed narrow fault on one satellite does not indicate that faults on other satellites are now more likely. Narrow, or satellite, faults are described by $P_{sat,i}$. This is the probability of satellite being in a faulted state at any given time. It can be determined from the product of the probability of fault onset (P_{sat_onset}) and the mean time to flag and remove the fault (MTTR):

$$P_{sat} = P_{sat\_onset} \times MTTR \tag{1}$$

For example, GPS specifies an onset rate of close to 10^{-5}/hour and has an MTTR below one hour [10]. Therefore, the probability of finding the satellite in a faulted state at any given instant is 10^{-5}.

Wide faults are ones that can affect more than one satellite within a constellation at a given time. These are also called constellation faults. They can be caused by control segment, by a design flaw in the satellites, etc. Their distinguishing characteristic is that a single event can cascade and affect multiple satellites within the same constellation. Such faults should be made extremely rare. Simultaneous faults have not been observed on the GPS consteallation since it has been declared fully operational in 1995. The probability of constellation fault is described by $P_{const,j}$.

Ultra-Wide faults are those that could affect multiple satellites across more than one constellation. Also called cross-constellation faults, these are faults where a single event can create multiple faults across multiple constellations. Examples include shared designs or information. One particularly concerning example is Earth Orientation Prediction Parameters (EOPPs). These are used to convert from an inertial reference frame, used to estimate satellite orbital parameters, to an Earth Centered Earth Fixed (ECEF) reference frame. The EOPPs are internationally coordinated. If the wrong values were used by all constellations, every satellite would consistently pinpoint the user in the wrong location. This paper does not further consider ultra-wide faults. There are many processes in place to make them very unlikely to affect all constellations in a similar manner at the same time. We will assume for this paper that they can be considered as sufficiently unlikely.

Having classified the faults we now wish to determine reasonable ranges for $P_{sat,i}$ and $P_{const,j}$. These values can be determined by analysis and corroborated by data. The Constellation Service Providers (CSPs) will conduct the analyses for their own constellation. They may then elect to specify the values in a performance standard as GPS has done [11]. The GPS performance standard specifies no more than 3 major service faults per year. It does not state whether these can be concurrent or not. It further states that the maximum time to notify the user of a major service failure is six hours (although in practice the MTTR appears to be less than one hour).

The definition of a major service failure is an instantaneous error that is larger than 4.42xσ_{URA}. This does not match the fault definition given previously, however, it does appear to be a reasonable proxy for estimating how often a satellite is in a faulted state. Given that there are typically 31 healthy satellites in orbit, there are approximately 27 thousand satellite hours per year. Three faults per year corresponds to an onset rate of ~1.1x10^{-5}/hour. If the MTTR is one hour then P_{sat} for

P	Constellation Faults	Satellite Faults (24 Satellites)	Satellite Faults (32 Satellites)
10^{-3}	9 per year	210 per year	281 per year
10^{-4}	1 per year	21 per year	28 per year
10^{-5}	1 per ten years	2 per year	3 per year
10^{-6}	1 per 100 years	2 per ten years	3 per ten years

Table 1. *State probabilities given the number of observed faults per time period, assuming a 1 hour MTTR.*

GPS should be no more than 1.1×10^{-5}. However, if we use the six hour upper limit, then P_{sat} could be as large as 6.6×10^{-5}. Observations of GPS performance indicate that it averages fewer than two major faults per year and that they are typically resolved within an hour. Thus, 1×10^{-5} appears to be a sufficiently conservative value for P_{sat}.

The same specification can be used to set an upper bound on P_{const}. No more than three faulted satellites per year also implies no more than one constellation fault per year. This implies an onset rate no greater than $\sim 1.1 \times 10^{-4}$/hour. Assuming MTTR values between 1 and 6 hours yields upper bounds on P_{const} between 1.1×10^{-4} and 6.6×10^{-4}. In actuality no constellation faults have been observed on GPS since it has reached it full operational capability in 1995. Therefore it is likely that much smaller values of P_{const} could be used for GPS.

No other CSP has published a performance specification. This is not surprising as these constellations are under development. Only GLONASS has a full set of satellites on orbit. Preliminary investigations into its performance indicates that satellite faults appear to about ten times more likely and that constellation faults have occurred [12]. The observations predict rates of $P_{sat} = 1 \times 10^{-4}$ and $P_{const} = 1 \times 10^{-4}$.

Table 1 shows the approximate average number of events per given time period that correspond to a specified fault probability. These values assume an MTTR = 1 hour. A longer MTTR would correspond to fewer events (or larger probabilities if the number of events did not change). New constellations, with very little service history, may need to assume large probabilities of fault until enough operational experience is obtained to gain confidence in their analysis models. The values of P_{sat} and P_{const} used by ARAIM should overbound the true values. Therefore, they should be larger than expected and larger than indicated by the number of observed events. Initial values may be quite large until sufficient confidence is gained in the operational performance. However, very uncertain satellites and constellations probably should not be used for safety-of-life operations. A constellation with many constellation faults and hundreds of satellite faults per year is likely not suited for ARAIM. Therefore, we assume 10^{-3} is a good upper bound for the largest usable values of P_{sat} and P_{const}. 10^{-6} appears to be a practical lower bound at least as far empirical observation can determine. The CSP fault analysis may assert lower values, but it no longer becomes practical to wait long enough to observe enough fault events to validate the number. Further, constellations are not stationary over the long term. Satellites and operational software are replaced, operators and operational procedures change. After many years the older data becomes less relevant to the current system. Any number below 10^{-6} would need to be established by analysis and corroborated (but not firmly established) by having no such observed faults.

SUBSET SOLUTIONS

As described earlier, the protection levels are assured by forming at least one subset that contains no faulted satellites. Therefore, subsets are formed that remove individual satellites and individual constellations. If P_{const} is greater or equal to 10^{-7}, the constellation threat cannot be ignored and the ARAIM algorithm requires the user to track two full constellations, as it will form subsets that remove each one in turn. We shall refer to the subsets created by removing a single satellite or a single constellation, as first order subsets. They are addressing single fault events only. However, as the number of satellites tracked by the user grows, and depending on the magnitude of P_{sat} and P_{const}, it is possible that multiple independent faults become sufficiently likely to overlap in time.

If the combined probability of multiple faults exceeds 10^{-7}, then the user must also evaluate second order subsets. These include two simultaneous individual satellite out combinations, one-satellite and one-constellation out combinations, and two-constellation out combinations. For example, if P_{const} exceeds 3.16×10^{-4} for more than one constellation, then the user must be tracking at least three full constellations, as they will have to exclude two of them to form a safe subset.

If the user is tracking N satellites, there are $N(N-1)/2$ unique two satellite out combinations. With four full constellations, the number of tracked satellites can become very large. If the user tracks twelve satellites per constellation they would have 48 satellites in view. This number leads to 1,128 two-satellite out subsets to evaluate. This is a very large number and could require a significant increase in computational power compared to conventional RAIM with only one constellation and only first order subsets to evaluate.

In addition, if there are M constellations, there are NM simultaneous one-satellite and one-constellation fault combinations. Using the prior example of $M = 4$ and $N = 48$, leads to 192 such combinations. Finally, there are $M(M-1)/2$ two-constellation combinations. This leads to six different combinations for our example. Thus, there can be a total of $N(N-1)/2 + NM + M(M-1)/2$ unique second order fault modes, or 1,326 possible two simultaneous fault combinations for our example.

If the second order (and higher) fault modes are sufficiently unlikely, there is no need to evaluate them. This is currently the case for conventional RAIM where the constellation fault probability is set to zero, the satellite fault probability is 10^{-5}, and there are rarely more than 14 satellites in view. We now present formulas to determine the total probability of different orders of fault combinations. We start with the 0^{th} order or fault-free mode. The fault-free mode uses the all-in-view set of satellites, where no faults are considered to be present. In this case, none of the fault modes are present. That is, the N satellite faults are not present and the M constellation faults are also not present. The probability of no fault being present is $(1 - P_{mode})$ where P_{mode} corresponds to P_{sat} or P_{const} depending on which fault mode is being considered. The product of the individual probabilities gives the probability of all faults being absent:

$$P_{fault-free} = P_{0th\text{-}order} = \prod_{i=1}^{N+M}\left(1 - P_{\mathrm{mode},i}\right) \approx 1 \qquad (2)$$

The prior probability of this mode can be approximated as 1 for the user algorithm as the fault mode probabilities are quite small.

The total first order probability is the sum of the probabilities that each fault mode, and only that fault mode is present

$$P_{1st\text{-}order} = \sum_{i=1}^{N+M}\left[P_{\mathrm{mode},i} \times \prod_{j=1,j\neq i}^{N+M}\left(1 - P_{\mathrm{mode},j}\right)\right] \approx \sum_{i=1}^{N+M} P_{\mathrm{mode},i} \qquad (3)$$

The latter approximation is actually the probability that all first order *and higher order terms* are present. Rather than saying only one fault mode is present, the latter approximation says that the fault mode is present and the other satellites can be in either the faulted or unfaulted states.

The total second order probability is the double sum of all pairwise faults. From here on out we will neglect the precise specification that the remaining satellites be in fault-free states.

$$\begin{aligned} P_{2nd\text{-}order} &= \sum_{i=1}^{N+M-1}\sum_{j=i+1}^{N+M} P_{\mathrm{mode},i}P_{\mathrm{mode},j} \\ &= \frac{1}{2}\left\{\left(\sum_{i=1}^{N+M} P_{\mathrm{mode},i}\right)^2 - \sum_{i=1}^{N+M} P^2_{\mathrm{mode},i}\right\} \end{aligned} \qquad (4)$$

This probability corresponds to the total probability of all second and higher order fault modes. If this sum is sufficiently smaller than the total integrity budget, these modes can be neglected and this probability is subtracted from the overall budget.

If the above probability is not sufficiently small, it becomes necessary to evaluate some or all of the second order subsets. In this event, it is necessary to determine the probability of the next higher order subsets (third order) that hopefully do not need to be evaluated, but that do need to be included in the total integrity budget.

The total third order probability is the triple sum of all three fault mode combinations, which can be expressed as a function of three single sums:

$$P_{3rd\text{-}order} = \frac{1}{6}\left\{\left(\sum_{i=1}^{N+M} P_{\mathrm{mode},i}\right)^3 - 3\left(\sum_{i=1}^{N+M} P^2_{\mathrm{mode},i}\right)\left(\sum_{i=1}^{N+M} P_{\mathrm{mode},i}\right) + 2\left(\sum_{i=1}^{N+M} P^3_{\mathrm{mode},i}\right)\right\} \qquad (5)$$

In this paper we will assume that this probability is always below 10^{-7}. It can then be subtracted from the total allocation and the remainder is partitioned among the first and second order modes that are evaluated. In no event are any three fault combination subsets evaluated.

Table 2 shows the number of subsets that need to be evaluated as a function of P_{const}, P_{sat}, M and N. The above formulas were used to determine whether or not second order terms needed to be evaluated. In some cases only some of the second order subsets are considered as others can be neglected. In most cases, the neglected subsets were those that removed the largest number of satellites and that would result in the worst geometries. Individual constellation out subsets were considered to already test the two satellite out combinations within that subset as

		4 Constellations		3 Constellations		2 Constellations	
P_{const}	P_{sat}	48 SVs	32 SVs	36 SVs	24 SVs	24 SVs	16 SVs
10^{-3}	10^{-4}	1330	634	744	354	-	-
	10^{-5}	1330	138	114	78	-	-
	10^{-6}	202	138	42	30	-	-
10^{-4}	10^{-4}	1324	628	741	351	326	154
	10^{-5}	1324	628	111	303	26	18
	10^{-6}	52	36	39	27	26	18
10^{-5}	10^{-4}	1324	628	669	303	302	138
	10^{-5}	1180	628	39	27	26	18
	10^{-6}	52	36	39	27	26	18

Table 2. *Number of subsets that need to be addressed for given values of* P_{const}, P_{sat}, M and N.

well as all of the one-constellation and one-satellite out combinations that affect the same constellation. As can be seen, as the input parameters increase, more subsets need to be evaluated. With particularly large inputs, thousands of subsets may need to be evaluated. In the next section we explore ways to reduce the total numbers of subsets to be evaluated

REDUCING THE NUMBER OF SUBSETS

Because a constellation out subset removes many satellites, it can be used to evaluate multiple fault cases at once. As indicated in the previous section, a one-constellation out subset already tests any possible fault combination within that constellation. That is, it evaluates any two-satellite out combination within that constellation and it also evaluates a constellation fault and a one-satellite fault affecting the same constellation. The subset that excludes a constellation is not affected by any combination of errors within that constellation. This means that the first order satellite faults are also included in that evaluation. The fault probabilities of the subsets that are no longer explicitly evaluated need to be included in the prior probability of this constellation out subset. The second order terms are already included, but if the first order satellite fault subsets are to be covered by the constellation out subset, the prior probability of this subset should also include the sum of the satellite faults within that constellation:

$$P^*_{const,i} = P_{const,i} + \sum_{j=1}^{N_i} P_{sat,j} \tag{6}$$

where N_i is the number of satellite used by the receiver within that constellation.

However, the vast majority of second order fault modes are for two independent faults between satellites from different constellations. These fault modes can also be tested by considering the one-constellation and one-satellite out subsets. A subset that removes one-constellation and one-satellite tests that satellite in combination with any of the satellites within that constellation. That is, it tests N_i two-satellite out subsets at once. There are only $N(M\text{-}1)$ of these subsets vs. $N(N\text{-}1)/2$. This can be a substantial reduction depending on the values for N and M.

An even larger reduction can take place by recognizing that the two-constellation out subsets also test all of the two-satellite out and all of the one-constellation and one-satellite out fault modes. Thus, all of the second order modes can be tested by the $M(M\text{-}1)/2$ two-constellation out subsets. Given that M is approximately an order of magnitude smaller than N, this results in about a two order magnitude reduction in the number of subsets to evaluate.

By applying these methods, the number of subsets that actually need to be evaluated can be dramatically reduced. However, there is a risk that by increasing the assigned *a priori* probability of subsets that remove many satellites will lead to larger protection level and EMT values. The protection levels and EMT depend on the threshold for the expected separation between the subset and the all-in-view solution. These thresholds are a function of the all-in-view and subset position estimation covariances and on their continuity allocation. More threshold tests result in a smaller continuity partition for each test. Thus, having fewer subsets to evaluate reduces each threshold.

The protection levels also depend on the required probability of missed detection. The missed detection probability is the integrity allocation for the subset divided by the *a priori* probability of occurrence. As the prior probability is increased, the missed detection probability must be decreased. A smaller missed detection probability increases the protection levels. For example, equation (6) assigns a greater *a priori* probability to the constellation subset, which tends to increase the protection levels. However, having fewer subsets to evaluate decreases the threshold, which tends to decrease the protection levels. It remains to be shown which process dominates.

		4 Constellations		3 Constellations		2 Constellations	
P_{const}	P_{sat}	48 SVs	32 SVs	36 SVs	24 SVs	24 SVs	16 SVs
10^{-3}	10^{-4}	10	10	6	6	-	-
	10^{-5}	10	10	6	6	-	-
	10^{-6}	10	10	6	6	-	-
10^{-4}	10^{-4}	10	10	75	51	26	18
	10^{-5}	10	10	75	51	2	2
	10^{-6}	4	4	3	3	2	2
10^{-5}	10^{-4}	10	10	75	51	26	18
	10^{-5}	10	10	75	3	2	2
	10^{-6}	4	4	3	3	2	2

Table 3. Recommended *number of subsets to evaluate, for given values of* P_{const}, P_{sat}, *M* and *N, given the reduction techniques described in the paper.*

TWO CONSTELLATION SCENARIOS

In the case that the receiver is using satellites from only two constellations, it is not possible to protect against simultaneous faults on both constellations. The product of their constellation fault probabilities must be below 10^{-7}. Neither is it possible to use the two-constellation out subset to evaluate the two-satellite fault cases. These cases must be evaluated directly or as part of a one-constellation out and one-satellite out subset. However, these latter subsets are also better to be avoided, as they require evaluating subset consisting of a single constellation with its most valuable satellite removed. Such a subset is more likely to have poor geometry and lead to lower availability. Fortunately, for two constellations, *N* is less likely to be a very large number. For most values, $P_{2nd\text{-}order}$ is likely to be below 10^{-7} and therefore only the first order modes need to be evaluated. In this case, the number of subsets is already a reasonable number and no further improvement is necessary. However, as P_{sat} and P_{const} become smaller, it becomes possible to use the constellation out mode to also test for the one satellite out modes. This is a good approach when the combined probability (6) is below ~$3\text{x}10^{-4}$. In either case, $P_{2nd\text{-}order}$ is subtracted from 10^{-7} and the remainder is distributed across evaluated modes.

Table 3 shows the recommended number of subsets to be evaluated for specific values of P_{const}, P_{sat}, *M* and *N*. The two constellation cases are in the rightmost columns. Two constellations cannot safely function with P_{const} = 10^{-3}. In the lower rows, the larger values of P_{sat} = 10^{-4} lead to unacceptably large values of P^*_{const}. These rows require that all first order modes be evaluated individually. When P_{sat} is reduced, it is possible to only evaluate the two one-constellation out subsets with the assigned probability of P^*_{const}.

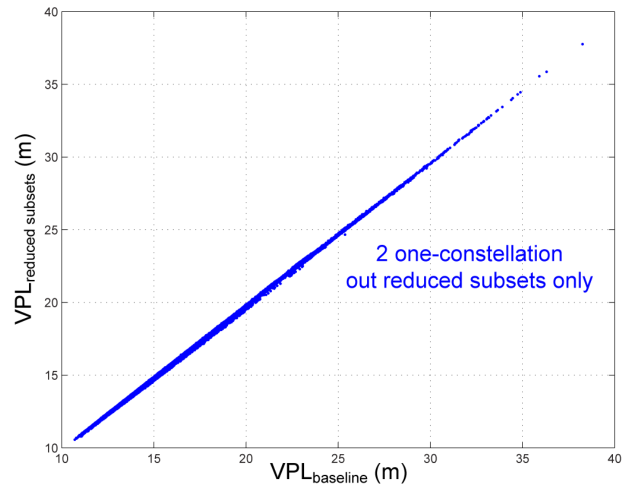

Figure 1. *Relative performance of 2 one constellation subsets vs. evaluating all first order subsets.*

Figure 1 shows a comparison of Vertical Protection Levels (VPLs) for the two methods. This analysis simulates two 30 satellite constellations using MAAST [2]. MAAST uses satellite almanacs to determine satellite positions at a predetermined set of time steps. A grid of users, every 10 degrees by latitude and longitude, look to see which satellites are visible and form corresponding geometry and measurement covariance matrices. Using our MHSS code [3], we can then determine values the outputs specified in the earlier description of the algorithm. Fixed values are used for:

- σ_{URE} = 0.5 m,
- σ_{URA} = 1 m,
- b_{nom} = 1 m,
- $P_{sat,i}$ = 10^{-5}, and
- $P_{const,j}$ = 10^{-4}.

The two approaches used were the baseline case, which evaluated all first order modes (horizontal axis) and the reduced approach which only evaluated two one-constellation out modes. Although the results are essentially identical, the reduced subset approach produces slightly lower VPLs. The effect of reducing the thresholds dominates in this example. The EMTs are similarly distributed. Both approaches yield 100% coverage of LPV-200 service in the simulation.

THREE CONSTELLATION SCENARIOS

When the number of constellations used is increased to three, it becomes possible to evaluate two-constellation

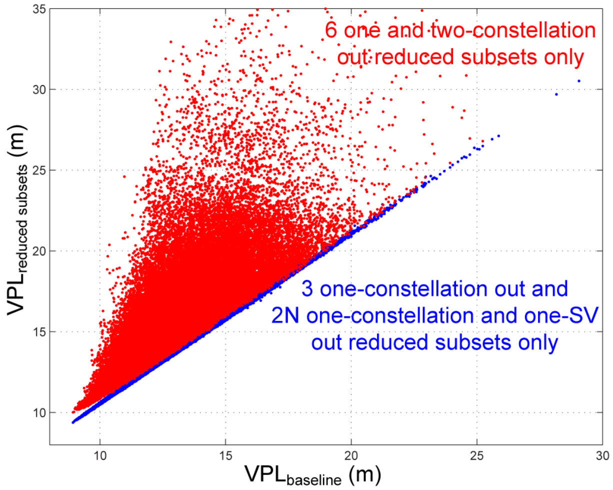

Figure 2. *Relative performance of 6 one- and two-constellation out subsets & 3 one-constellation out and 2N one-constellation and one-satellite out subsets vs. evaluating all first order (and second order as needed) subsets.*

out subsets. However, such subsets, which now only contain a single constellation may not always have good geometry, particularly if one of the constellations is weak. The middle columns of Table 3 have the recommended number of subsets to compute for specific three constellation scenarios. When P_{const} = 10^{-3}, it is mandatory to evaluate the two constellation out cases. In this case, we recommend only evaluating the three one-constellation out and three two-constellation out subsets. However, when P_{const} = 10^{-4}, we recommend evaluating the three one-constellation out and 2N one-constellation and one-satellite out subsets. In these rows, P_{const} is replaced by P^{*}_{const}. Smaller values of P_{const} and/or P_{sat} can lead to sufficiently small values of $P_{2nd\text{-}order}$. When this value falls below 10^{-7}, it becomes possible to evaluate only the three one-constellation out subsets.

Figure 2 shows a comparison of VPLs for three methods. In this plot, MAAST used three 24-satellite constellations, one of which was relatively weak. $P_{sat,i}$ was set to 10^{-4}, all other parameters match those used in the two constellation scenario. The baseline case evaluates all first order modes and all second order modes as needed. The baseline VPL is plotted on the x-axis. The red points correspond to using the one- and two-constellation out subsets only. While this works well the majority of the time, there are clearly many cases where a weak subset is used and the VPL is significantly increased. This approach does not lead to 100% coverage. The blue dots correspond to evaluating the three one-constellation out and 2N one-constellation and one-satellite out subsets. Here the VPL is nearly identical to the baseline case but using and order of magnitude fewer subsets. This case and the baseline case both provided 100% coverage.

FOUR CONSTELLATION SCENARIOS

The situation improves when the receiver can use four constellations. Now the two-constellation out subsets still have two full constellations remaining. These subsets have strong geometries even for relatively weak constellations. The left columns in Table 3 contain the recommended number of subsets to evaluate. In most of the cases the four one-constellation out and six two-constellation out subsets are recommended. When none of the second order terms need to be evaluated, only the four one-constellation out subsets need to be evaluated. In all cases P_{const} is replaced by P^{*}_{const}. For four constellations, no more than 10 subsets need to be evaluated. This number does not depend on the number of satellites in view (within reason). It is also possible to only evaluate the six two-constellation out cases as they cover the single constellation out cases. However, the probability of the two out cases are dramatically increased in order to properly protect the one out case. We considered the case where the one-constellation out probability was evenly distributed among the two-constellation out subsets that it participated in. That is, the prior probability of the subset that removes constellations i and j is given by:

$$\frac{P^{*}_{const,i} + P^{*}_{const,j}}{3} + P^{*}_{const,i} \times P^{*}_{const,j} \qquad (6)$$

Notice that this increases the probability of this subset by many orders of magnitude.

Figure 3 shows a comparison of VPLs for three methods. In this figure, MAAST used four 24-satellite constellations. $P_{const,j}$ was set to 10^{-3} and $P_{sat,i}$ was set to 10^{-4} all other parameters match those used in the two and three constellation scenarios. The red points correspond to using six two-constellation out subsets only. There is an obvious increase in the VPL, but only by a few meters. All of the VPLs remain below 35 m, but some do come close to that value. The increase in the assigned two-constellation fault probability creates this increase. Because of it, we recommend evaluating the one-constellation out subsets separately. The blue dots correspond to evaluating the ten one- and two-constellation out subsets. Here the VPL is nearly identical to the baseline case but using two orders of magnitude fewer subsets. All three cases provided 100% coverage.

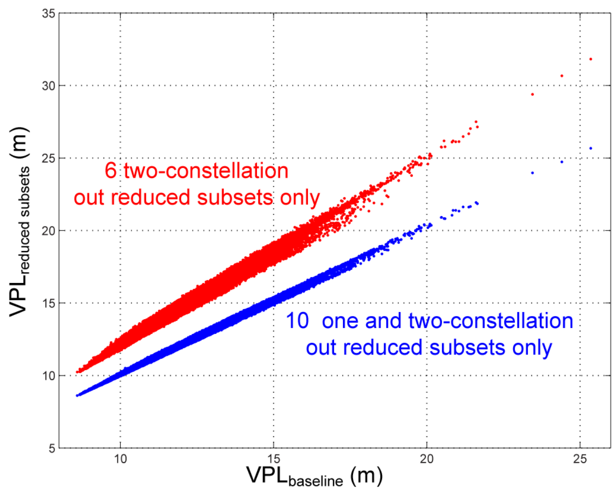

Figure 3. *Relative performance of 6 two-constellation out subsets & 10 one- and two-constellation out subsets vs. evaluating all first order (and second order as needed) subsets.*

CONCLUSIONS

This paper has proposed a method to reduce the number of subsets that an ARAIM receiver needs to evaluate in order to protect user integrity. This method utilizes subsets that remove constellations as a block to evaluate many fault modes at once. It was shown that the reduction achieved could be more than two orders of magnitude for four (or more) constellations, with many satellites in view. For fewer constellations the reduction may be smaller but still very significant. This reduction directly translates into computational effort required to calculate the protection levels and EMTs. We observed a significant reduction in the run time of MAAST when implementing this approach.

This approach can be used even if the specified constellation fault probability is zero. In that event, "constellations" could be formed by grouping together satellites and aggregating probability as in (6). However, we expect that, in practice, the constellation fault probability will likely be above 10^{-7} and therefore these subsets will need to be evaluated. It is important that when a single subset is used to evaluate many different fault modes, that all of their prior probabilities be properly taken into account.

We have seen that for two constellations, availability is driven by performance of the weakest constellation. This is also true for three constellations with high values of P_{const}. When P_{const} is smaller or when four constellations are used, the user essentially always has good satellite geometry and should have high availability. Further, for four constellations, the number of subsets to evaluate is smaller than the numbers often used in today's single constellation RAIM. Thus, using more constellations and more satellites leads to an overall reduction in the computational effort to calculate the protection levels and EMTs.

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 authors' and this paper does not represent a government position on the future development of ARAIM.

REFERENCES

[1] Blanch, J., Walter, T., Enge, P., Wallner, S., Fernandez, F. A., Dellago, R., Ioannides, R., Hernandez, I. F., Belabbas, B., Spletter, A., and Rippl, M., "Critical Elements for a Multi-Constellation Advanced RAIM", NAVIGATION, Journal of The Institute of Navigation, Vol. 60, No. 1, Spring 2013, pp. 53-69.

[2] Jan, S.S., Chan, W., Walter, T., and Enge, P., "MATLAB Simulation Toolset for SBAS Service Volume Analysis," in proceedings of ION GPS-2001, Salt Lake City, UT, September 2001.

[3] Blanch, J., Walter, T., Enge, P., Lee, Y., Pervan, B., Rippl, M., Spletter, A., "Advanced RAIM user Algorithm Description: Integrity Support Message Processing, Fault Detection, Exclusion, and Protection Level Calculation," Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012), Nashville, TN, September 2012.

[4] http://waas.stanford.edu/staff/maast/maast.html

[5] Shallberg, K. and Grabowski, J., "Considerations for Characterizing Antenna Induced Range Errors," Proceedings of the 15th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2002), Portland, OR, September 2002, pp. 809-815.

[6] Wong, G., Chen, Y.H., Phelts, R. E., Walter, T., and Enge, P., "Measuring Code-Phase Differences due to

Inter-Satellite Hardware Differences," Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012), Nashville, TN, September 2012.

[7] DeCleene, B., "Defining Pseudorange Integrity - Overbounding," Proceedings of the 13th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2000), Salt Lake City, UT, September 2000, pp. 1916-1924.

[8] Collins, J. P., Langley, R. B., "The residual tropospheric propagation delay: How bad can it get?," Proceedings of the 11th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 1998), Nashville, TN, September 1998, pp. 729-738.

[9] Murphy, T., Harris, M., Booth, J., Geren, P., Pankaskie, T., Clark, B., Burns, J., Urda, T., "Results from the Program for the Investigation of Airborne Multipath Errors," Proceedings of the 2005 National Technical Meeting of The Institute of Navigation, San Diego, CA, January 2005, pp. 153-169.

[10] Heng, L., Gao, G., X., Walter, T., and Enge, P., "GPS Signal-in-Space Integrity Performance Evolution in the Last Decade," IEEE Transactions on Aerospace and Electronic Systems Vol. 48, No. 4, October 2012.

[11] GLOBAL POSITIONING SYSTEM STANDARD POSITIONING SERVICE PERFORMANCE STANDARD, 2008 available at http://www.navcen.uscg.gov/ gps/geninfo/2008SPSPerformanceStandardFINAL.pdf.

[12] Heng, L., Gao, G. X., Walter, T., and Enge, P., "GLONASS Signal-in-Space Anomalies Since 2009," Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012), Nashville, TN, September 2012.