A New Approach to GPS Integrity Monitoring Using Prior Probability Models and Optimal Threshold Search

Samuel P. Pullen
Boris S. Pervan
Bradford W. Parkinson

Department of Aeronautics and Astronautics
Stanford University

Abstract – Current methods for GPS receiver autonomous integrity monitoring are limited by the assumptions they make. Using published studies of navigation system reliability, this paper develops a *prior probability model* for spacecraft and receiver anomalies based on non-ideal failure models and the uncertainty present in their failure distribution parameters. With this model, the thresholds for a residuals test statistic are found to optimize an *arbitrary* objective function based on relative costs for *false alarm* and *missed detection* errors. The outputs of Monte Carlo simulations allow thresholds to be computed for each geometry case using trial-and-error optimization. The simulation outputs suggest that the assumptions of traditional RAIM may be partially invalid. These results are useful for both snapshot RAIM tests as well as multi-step integrity algorithms which use *Bayesian updating* to generate posterior failure probabilities. Multi-step algorithms may be a valuable addition to the future GIC integrity structure.

1.0 INTRODUCTION

The basic concept behind GPS receiver autonomous integrity monitoring (RAIM) is the use of additional information to verify position solutions. Since more than the minimum number of four satellites (needed to solve for 3-dimensional position and the clock bias) will be visible in almost all cases [1], redundant satellite pseudo-range (PR) information is available, and the position fix may be computed from a best fit to the overdetermined data. The consistency of the redundant measurements provides a clue as to whether a GPS satellite or some other unit is operating out of specification and whether this error makes the position solution unusable.

This general approach to RAIM has been developed into various tests of the system geometry and the pseudo-range error residuals [2,3,4] which are for the most part functionally equivalent [5]. However, these "traditional" methods only use the pseudo-range data for a single receiver sample, ignoring previous samples that are only minutes old. Even more importantly, these methods assume no prior knowledge of satellite, receiver, and ground station failure modes. Instead, failure mode and effect models are represented in a very simplified way. These features suggest that the approximations on which RAIM is based should be carefully checked by a more general model that attempts to simulate uncertainties in the GPS system.

This paper presents the results of such a study. These results are based on a simulation of the GPS constellation geometry along with distributions which model user uncertainty regarding the probability of failures in the satellites and ground equipment. The outputs of the simulation are observed probabilities of errors given knowledge of the current geometry. Thus, given a cost function which models the negative utility of RAIM decision errors for a given application, we can not only check the theoretical error predictions assumed by current RAIM methods; we can also directly choose an optimal threshold for each geometry case.

The results of this study suggest a further look at multi-step RAIM algorithms which use prior as well as current measurements. Bayesian updating (starting with the prior probability model) is the most general way yet considered of tackling this problem. Also, the addition of a GPS Integrity Channel (GIC) should result in a coordinated RAIM system whose parameters can be optimized from the top down by simulation-based search methods such as simulated annealing. A basic framework for these concepts is presented here along with suggestions for future work.

2.0 TRADITIONAL RAIM METHODS

As mentioned above, the traditional RAIM methods are based on variations of the parity-vector based threshold tests described in [2,3,4]. Most of these algorithms have been shown to be functionally identical in [5]. Because they are expressed as practical, usable algorithms, the CFAR and CPOD algorithms given in [4] will be taken as representative of this methodology.

Given the basic measurement equation:

$$z = Hx + n \quad (1)$$

where z is the vector of pseudorange (PR) measurements for M satellites in view, x is the 4 x 1 navigation state vector in three position coordinates and the clock bias, n is the vector of gaussian PR noise with mean 0 and variance σ_n^2, and H is the M x 4 observation matrix consisting of M line-of-sight row vectors to visible satellites augmented by a column vector of 1's for the clock bias state. Assuming 10 second averaging of PR measurements sampled at 1 Hz, the noise standard deviation σ_n for the case of uncorrected C/A code is computed in [4] to be about 32.4 meters due to selective availability, satellite clock and ephemeris errors, propagation uncertainties, and receiver noise and multipath. Note that (1) can express the errors in **z** and **x** since the equation and assumptions are linear.

A least-squares estimate of the true navigation state vector is given by:

$$\hat{x} = H^* z = \left(H^T H\right)^{-1} H^T z. \quad (2)$$

A least-squares residual statistic of the form:

$$D = z^T \left[I_M - HH^*\right] z \quad (3)$$

is computed anew at every RAIM decision step. Normally, D is computed and compared to a predetermined threshold T. If D > T, an integrity alarm is issued. Otherwise, normal (no-fault) operation is assumed. In addition, cases of "bad" GPS satellite geometry are designated as non-available RAIM cases, i.e. no failure determination can be issued.

Generally, traditional RAIM algorithms choose T either by Monte Carlo sampling [2] or by the chi-square probability distribution [3,4] to give desired false alarm (FA) and missed detection (MD) probabilities:

$$P_{FA} = P\left[D > T \mid \delta x < RPE\right] \quad (4)$$

$$P_{MD} = P\left[D < T \mid \delta x > RPE\right] \quad (5)$$

where δx is the true (unknown) position error and RPE is the allowed position error limit. Official RTCA specifications of maximum probabilities are given in [4]. Non-availability thresholds are chosen based on geometry dilution of precision parameters (such as HDOP and PDOP) which relate the noise standard deviation σ_n to the position error standard deviation. These limits may correspond to a specification on P_{MD} when the threshold T is chosen to meet a false alarm requirement [4]. More complicated availability criteria are described in [5], such as those which utilize delta-H_{max} as described below.

Because the residual statistic D is a quadratic function of z, which is itself a linear function of the normal random vector n in (1), the false alarm probability is given by the chi-square distribution:

$$P_{FA} = Q\left(T^2/\sigma_n^2 \mid M-4\right). \quad (6)$$

In traditional RAIM, only bias failures on a single satellite are considered; thus a constant bias vector b (with only one non-zero entry) would be added to the right-hand side of (1). The missed detection probability is then bounded by the non-central chi-square distribution:

$$P_{MD} \le P\left(T^2/\sigma_n^2 \mid M-4,\ RPE/\sigma_n\ \delta H_{\max}\right). \quad (7)$$

in which δH_{max} is a geometry parameter defined in [4] which measures the least-detectable satellite failure.

In theory, this RAIM approach allows one to choose T to meet specifications on either P_{FA} or P_{MD} while checking to ensure that the requirement on the other is not exceeded [4]. However, there are many assumptions embedded in the above equations, such as neglecting multiple-fault cases. Note that P_{FA} in (6) implicitly assumes that any bias fault will lead to a position error exceeding the RPE. More to the point, the true utility of the RAIM algorithm should be a function of more than just P_{FA} and worst-case P_{MD} as defined here.

3.0 PRIOR PROBABILITY MODEL

The assumptions made by traditional RAIM not only contain simplifications; they do not consider any prior probability information that might influence P_{FA} and P_{MD}. They essentially consider that there is one possible spacecraft failure mode of unknown but small probability. While it is true that only limited data on navigation satellite and receiver integrity has been made public, it is sufficient to at least attempt to express a prior probability model (PPM) and examine its usefulness in computing optimal RAIM thresholds. This PPM would express our uncertainty regarding whether or not system failures are present before any measurements are conducted.

In an attempt to predict the likelihood of GPS position fix availability for various classes of users, two studies have been conducted from which our PPM can be drawn. The more useful study is [6], which is based on failure models drawn from previous space-based and inertial navigation systems. A spacecraft failure/renewal model was assumed, and ten separate parameter sets were selected as representative of the uncertainty of the GPS constellation. Using the model in [6], we divide failures into “hard” and “soft” cases, “hard” meaning complete spacecraft shutdowns and “soft” meaning spacecraft operating with GPS signal

errors. The other study [7], conducted by IBM for the Department of Defense, does not break down failure modes by satellite and in any case seems to assume that the DoD specifications for GPS system availability are automatically met.

Rather than using ten discrete failure parameter sets as in [6], our PPM uses continuous Gamma(a,b) and Normal(μ,σ) distributions to model the uncertainty in the failure parameters based on the data in [6]. For spacecraft failures, the key parameters in a simple renewal model are the mean-time-between-failures (MTBF) and the mean-time-to-restore (MTTR). These distributions and the parameter uncertainty models are shown in Table 1 below. In addition, models for GPS receiver soft failures are given, although it is assumed that most receiver faults will be hard failures which make GPS position fixes impossible. From these models of our failure parameter uncertainty, we can generate probabilities of system failures for each trial and then apply them to the GPS constellation.

Parameter	Dist.	mean (μ)	s.dev. (σ)
SV hard MTBF	Gamma	100 mo.	12 mo.
SV hard MTTR	Gamma	1.5 mo.	0.6 mo.
SV soft MTBF	Gamma	16 mo.	6 mo.
SV soft MTTR	Normal	28 hr.	8 hr.
SV soft PR bias	Gamma	300 m	240 m
SV soft fail. flag	Uniform	400 m	230.9 m
RCR failure prob.	Normal	0.0015	0.0025
RCR channel bias	Gamma	30 m	24 m
RCR noise multiple	Normal	8.0	2.5

Note: Normal distribution outcomes < 0 are taken to be 0.

Table 1: Prior Probability Model Parameters

Modeling the effects of system failures is much trickier because of the many possible failure modes that can exist. There is little information to go by; thus the Gamma distributions in Table 1 which model the likelihood of bias magnitude have a large variance. This represents considerable prior uncertainty. No detailed attempt has been made to break down the causes of failures because there are too many possibilities. Therefore, the models in Table 1 simply try to model our uncertainty regarding the possible bias outputs. Note that these failures are above and beyond the uncertainties modeled as noise in σ_n^2 [4].

4.0 SIMULATION OF GPS SYSTEM UNDER PPM

The GPS performance consequences of the prior probability model outlined above cannot be calculated analytically; Monte Carlo simulation is the best way to perform the analysis without having to make severe limiting assumptions on the PPM. For this study, a simulation of the GPS primary 21+3 satellite constellation [1], the PPM, and the results of RAIM-assisted position fixes was written in C. It is a follow-on to the simulation developed in [8].

Each simulation run executes N_{max} position fix trials. As in [8], from one trial to the next, a clock is incremented by a random time interval that is uniformly distributed between 0 and 30 minutes. The GPS constellation position is then updated accordingly. Beyond this step, however, no continuity between trials is assumed. For each trial, new values of the PPM failure parameters are randomly generated from the distributions given in Table 1. Given the MTBF and MTTR for both hard and soft satellite failures for a given trial, the probability of being in a failure state at that instant is given by renewal theory [6]:

$$P_F = \frac{MTTR}{MTBF + MTTR} . \qquad (8)$$

Each of the 24 satellites is then checked for failure against the resulting P_F's. Those that suffer hard failures are simply removed from possible visibility for that trial (only). Those which suffer soft failures receive a randomly generated bias which is assumed to be static for the purposes of a single trial (dynamic effects such as ramp errors are not considered here). This random bias is then compared to a uniform random number to determine if the bias has been "flagged" by ground control. The higher the bias, the more likely it is to be flagged. Flagged soft failure effects are the same as those for hard failures; i.e. these satellites are not used if they are visible.

Once the status of all satellites is determined, the GPS geometry is resolved to determine the number of satellites visible to a receiver positioned at San Francisco (this position does not change) with a 7.5^o mask angle. Using the PPM, it is possible that one or more of the apparently usable satellites has an unflagged bias error. Next the probability of receiver error is generated, and the receiver error state is simulated for this trial. If a (soft) failure exists, it is assumed to be single-channel based with 50% probability, inducing an apparent bias on the relevant spacecraft. Otherwise, it is assumed to be system-wide, thereby increasing the receiver noise by a randomly-generated multiplicative factor (see Table 1).

Next, white noise with variance σ_n^2 is added to the output pseudorange of each satellite using the variances given in [4] and assuming a 10 second averaging time. GPS position is then computed and compared to the known user position. In addition to the position fix error magnitude δx, the least-squares residual magnitude $D = \delta r$ is computed as the RAIM decision test statistic from (3).

These results are stored in histogram arrays that contain the number of occurrences of the discretized index values in question. One bin equals 10 meters in δr and δx and 0.1 in HDOP. Each trial is recorded according to the number of

usable (non-flagged) satellites in view (called n_view) and the HDOP of the visible spacecraft geometry. In addition to recording δr vs. δx within the prevailing n_view and HDOP, marginal distributions of δr vs. δx along with satellite failure state, PR error, receiver error, and others are stored for later analysis.

5.0 RISK-BASED RAIM COST MODEL

As mentioned in Section 2.0, traditional RAIM methods use the chi-square distribution to set thresholds to meet specifications on P_{FA} and P_{MD} for various types of missions. In [4], thresholds are set for the non-precision approach case (RPE = 550 m) to give $P_{FA} = 1.4 \times 10^{-5}$ and $P_{MD} = 10^{-4}$ based on equations (6,7). Heretofore, these RTCA specifications have been the only guidelines for RAIM threshold calculation.

The question of where these specifications came from has seldom been asked, however. Clearly, they should represent the "cost" to the user of a false alarm, missed detection, or non-availability result. These costs may not be uniform for all users. Furthermore, as pointed out in [9], absolute specifications cannot be demonstrated to 100% confidence in any case. Since uncertainty is unavoidable, it may be better to use a cost-based method that more easily models uncertainty and risk aversion.

A basis for both the RTCA specifications and a cost-based model can be found in the RNP Tunnel Concept for precision aircraft approaches and landings [10]. It describes a means of setting position error specifications along an approach "tunnel", and it estimates the consequences of straying outside the tunnel based on the historical record of commercial aviation accidents. For a precision approach, integrity decisions must be made close to the ground. A false alarm (causing an aborted landing) has a small but significant probability of leading to a fatal accident, while the more serious missed detection has a much higher chance of leading to a fatal accident. Since non-availability decisions are made before the final approach begins (based on satellite geometry for example), there is minimal risk involved, but there is still a small inconvenience cost that must be weighted against the fatal accident risks.

The non-precision approach case studied here does not face the same exacting requirements, so the relative costs are harder to specify. The RPE in this case is much looser, however, so violating it still poses considerable risk if not detected. Any detected violation (known as a "bad" position result) requires an alternative approach method, so there is an inconvenience cost plus some minimal risk. A false alarm should have the same cost as a "bad" position result except that the true availability of the GPS system is wasted. Thus the false alarm cost is equal to the "bad" cost plus a small non-availability (NA) inconvenience cost. Although many different cost cases have been tested, the one used for the results shown here is given in Table 2 below. Note the variable cost of the MD event means that in addition to the base cost of any MD, an additional cost is paid for each bin (10 m) beyond the RPE limit that the missed position error is located.

RAIM Result	Base Cost	Variable Cost
good position	0	0
detected bad pos.	1	0
missed detect.	200	10
false alarm	1.01	0
non-available	0.01	0

Table 2: RAIM Cost Parameters

The NA cost for GPS is perhaps the most important to specify. Here it is assumed that this cost is only 1/100th of the "bad" position cost. However, in some cases a NA result has virtually the same effect as a detected "bad" position: the approach must be aborted (and conducted by other means) at about the same position in the approach. In this case, the NA cost is as much as half of the FA cost; thus relatively few geometries should be screened out.

6.0 COST-BASED THRESHOLD OPTIMIZATION

Since each combination of n_view and HDOP for the non-precision approach represents an independent decision case, an optimal residual threshold can be chosen separately for each one. Recall that the simulation described above stores δr as a function of δx for each case. Once all trials are completed, threshold optimization is conducted simply by computing the cost of all possible discrete thresholds $T = \delta r_{limit}$ using the following equation:

$$J(T) = P_{FA}\,C_{FA} + P_{MD}(\delta_e)\,C_{MD}(\delta_e) + P_{bad}\,C_{bad} \quad (9)$$

where δ_e is the amount by which RPE is exceeded. The threshold that gives the lowest J(T) is the optimal choice.

At this point, the expected cost of the optimal threshold is compared to the GPS non-availability cost. If the optimal threshold gives a lower cost, then GPS is available for that case and the optimal RAIM threshold is set. Otherwise, that case of n_view and HDOP is designated as non-available for integrity, and GPS is assumed to be unusable. Note that the NA cost is thus a critical number, as it measures the user's risk aversion regarding the reliability of RAIM decisions (see Section 5.0).

Finally, the optimal thresholds (or zero if NA) are output in a 2-D lookup table of n_view vs. HDOP. Overall results for all of the cases, weighted by the likelihood of each case, are also computed to measure the general utility of this RAIM methodology as well as the marginal probabilities of FA, MD, and NA outcomes.

7.0 OPTIMAL THRESHOLD RESULTS

All the results shown in this section use the parameters given in Tables 1 and 2 and a sample size of N_{max} = 10 million trials. However, simulations have been run with many variations of the PPM and cost parameters; thus we can draw some conclusions regarding solution sensitivity to parameter changes while we study the nominal results.

The simulation outputs contain important marginal output probabilities that deserve study so as to illustrate the effects of the prior probability model. Table 3 shows the observed probability of spacecraft and receiver failures:

Observed Event	Probability
no SV failure	0.98237
SV hard failure	0.01492
SV soft failure-unflagged	0.00173
SV soft failure-flagged	0.00098
RCR failure (separate)	0.00191

Table 3: Observed Failure Probabilities

These results contain few surprises. Most spacecraft failures are "hard"; they only degrade GPS geometry. The issue of flagged soft failures could use further study. Although receiver (soft) failures are not too uncommon, their effects are usually not very significant (see Table 1).

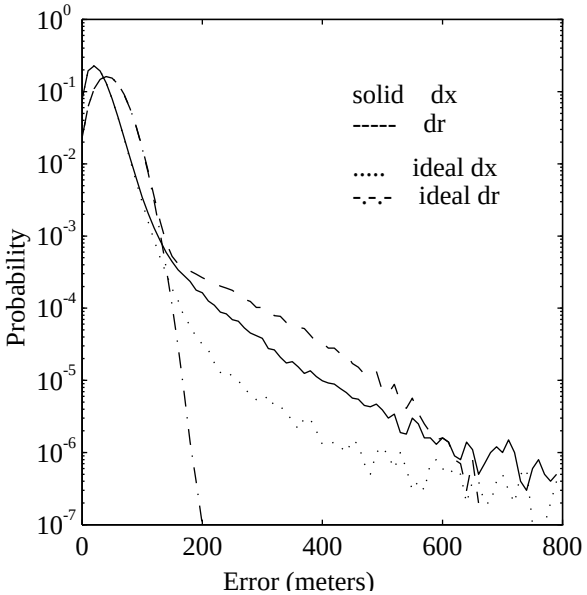

Figure 1: Marginal Densities of δr and δx

Figure 1 shows the marginal probabilities of δr and δx over all geometry cases. The "theoretical" probability curves generated from an "ideal" simulation (with no soft failures) are included for comparison. Agreement is good for the low error, high probability cases, but above 150 meters, the ideal plots dive well below what is observed using the PPM. Oscillations in the lower right-hand corner are due to limited sample size for the lowest probabilities.

Figure 2 shows the relative likelihoods of the possible n_view and HDOP geometries. This 3-D line plot illustrates the high likelihood of having 6 or more usable satellites in view (despite flagged failures) with low HDOP's. Many cases with only 5 satellites usable (the minimum for RAIM) and high HDOP's have sample sizes that are too small for reliable threshold optimization; thus they are declared to be non-available.

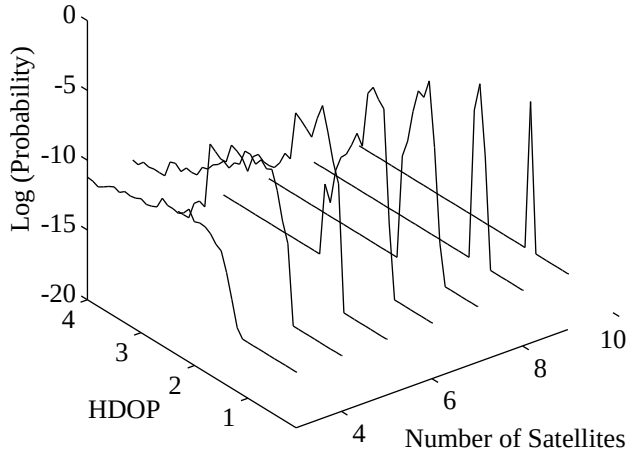

Figure 2: GPS Geometry Joint Probabilities

Figure 3 is a 3-D plot of the optimal thresholds as functions of GPS geometry. All "zero" regions of this plot are non-available either due to an optimal decision of NA for that case, insufficient sample size (see Fig. 2), or non-redundancy (for n_view < 5). These thresholds tend to increase with n_view and fall with HDOP as expected, although there are many exceptions and embedded NA results (T = 0) which most likely are due to limited sample size. For actual use, many more simulation trials would be conducted to gain further statistical significance.

Figure 4 follows the same 3-D format but shows the optimal costs per case. Here NA results have a fixed cost of 0.01 from Table 2. The valley features in the plot show the relative improvement gained by using RAIM with the optimal thresholds from Figure 3. Cost reductions of 80% are typical for the highest-likelihood geometry cases.

Table 4 summarizes the overall results for the PPM and cost model used for these tests. Note that results from setting RPE = 275 meters (half of the non-precision specification) are included for comparison purposes. The results for the 550 m case seem quite good except for the

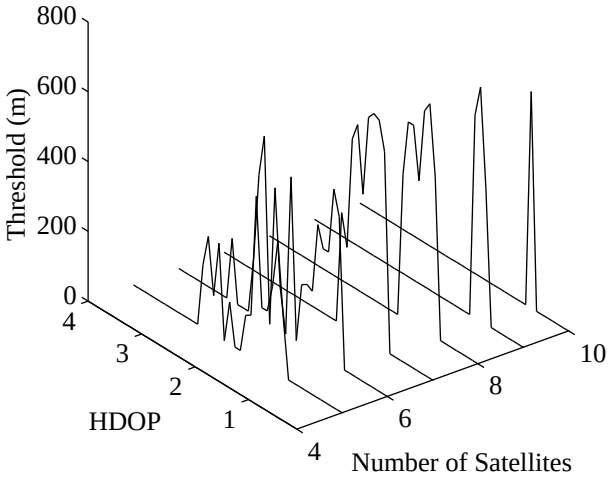

Figure 3: Optimal Thresholds for GPS Geometries

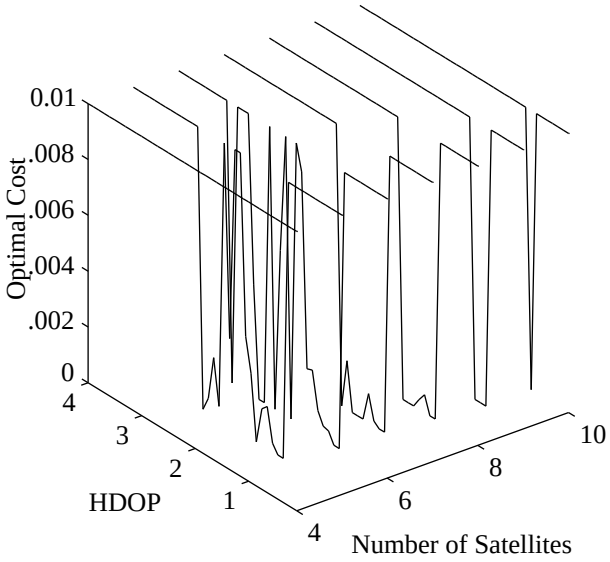

Figure 4: Optimal Costs for GPS Geometries

high observed P_{MD}. Since the probability of exceeding 550 meters error is so small (5×10^{-6}), MD results are very rare; thus the optimization routine can "afford" to miss some of those which do occur. Availability is 95%, but the optimal cost is only 17% below the NA cost due to these MD's. As NA increases (i.e. as risk aversion decreases), this ratio will generally improve.

We see significant differences for the 275 meter case precisely because only RPE was changed. Bad position fixes are more probable, so the thresholds must "tighten up" to reduce P_{MD} at the cost of increasing P_{FA} by almost

Sim. Output	RPE = 550 m	RPE = 275 m
Total Cost	0.00832	0.00708
FA prob.	0.000165	0.003921
bad pos. prob	0.000005	0.000166
MD | bad prob.	0.484	0.060
availability prob.	0.951	0.651

Table 4: Simulation Output Summary

24 times. Availability also suffers as a result. But it is interesting that the cost is slightly below the 550 meter case. This indicates inconsistency between the requirements and the cost model - a recurrent problem. For a tighter RPE, the NA cost should be increased (since greater risk is inevitable) which will improve availability but increase the resulting absolute cost above that for the 550 meter case.

But it is significant to note that the observed P_{FA}, P_{MD}, and P_{NA} do not agree with the results in [4]. While one of the three specifications can be achieved, at least one of the other two must be sacrificed. This is partially due to an inconsistent cost model, but it also indicates that under "real world" conditions, the RTCA specifications will be difficult to meet given a realistic PPM and cost model.

8.0 COMPARISON WITH TRADITIONAL RAIM

It is difficult to compare directly the performance of this PPM-based cost optimization method with that of traditional RAIM methods in [4,5] because their thresholds were set to meet arbitrary specifications. In order for a fair comparison to be made, the same cost model must be used to set thresholds for both methods. This is very difficult for the case of traditional RAIM (although it might be possible using global optimization methods).

Lacking a fair bottom-line performance comparison, another approach is to determine the accuracy of the chi-square model (6,7) used in traditional RAIM given the PPM used in our simulations. The RAIM method outlined here chooses optimal thresholds based on the simulation outputs instead of relying on (6,7). It is also possible to compute the observed P_{FA} as a function of the detection threshold for each n_view and HDOP.

Using (6), we can compute the expected P_{FA} given T and n_view (it is independent of HDOP) and then use a chi-square statistical test to determine the significance (Σ) of the chi-square assumption; that is, the likelihood that the simulation output P_{FA} was produced by the theoretical distribution (6). For comparison, a "control" simulation was run in which all soft failure models were removed. In this case, (6) should hold almost exactly, although bin discretization error and limited sample size prevent 90-100% confidence from being achieved in practice.

The random pattern of these results is important, as it suggests that the errors mentioned above are significant; thus high significances are unlikely even for the ideal case. It can also be observed from the test data itself that, as expected, agreement is better for lower thresholds (where many FA's will result) than for the thresholds likely to be selected as optimal. A summation of these results weighted by the number of position fixes in each bin is given in Table 5 below:

Simulation Model	Overall Significance S_T
PPM case	0.032
Ideal case	0.555

Table 5: Chi-Square Test Significance

The moderate significance for the control simulation suggests a strong likelihood that the outputs follow the chi-square distribution (6). In sharp contrast, the low PPM-based significance means that we can reject the hypothesis that the PPM simulation outputs are chi-square with 96.8% confidence. We must conclude that using what we believe to be a reasonable PPM, the chi-square distributions used in RAIM are likely to give incorrect results. A fair bottom-line cost comparison is needed before we can be more specific about the relative penalty incurred by traditional RAIM methods.

9.0 BAYESIAN UPDATING AND GIC POTENTIAL

While the method used to compute the optimal thresholds shown in Figure 3 can be used as part of a "snapshot" RAIM algorithm (which uses only the current sample of pseudorange data), even more benefit may be gained by applying the PPM and the cost model to a multi-step updating algorithm. A Kalman filter algorithm for this role is derived in [11]. It relies on "censoring" out the estimated bias from a single source. This idea utilizes past sensor information and requires fewer assumptions than does traditional RAIM, but failures are still assumed to come one at a time, and random effects are presumed to be perfectly white and gaussian once the bias estimate is removed. Another approach is implied by current plans for GIC, which uses WADGPS to broadcast precise PR corrections [12]. WADGPS requires observing past PR data; so it indirectly conducts multi-step updating.

We can avoid the white noise limitation built into Kalman filters by conducting Bayesian probability updates at each time step. This type of algorithm is presented in [13] using a hybrid version of Bayes' rule:

$$h(\theta_i|e_t) = \frac{f(e_t|\theta_i)\,h(\theta_i)}{\sum_i f(e_t|\theta_i)\,h(\theta_i)} \qquad (10)$$

where e_t is the residual vector, θ_i is a fault hypothesis, $h(\theta_i|e_t)$ is the posterior probability of θ_i given e_t, $f(e_t|\theta_i)$ is the likelihood of the observed residual vector given an assumed fault state, and $h(\theta_i)$ is the prior probability distribution given by:

$$h(\theta_i) = \begin{cases} 1 - m\alpha & for\ i = 0 \\ \alpha & for\ 1 \le i \le m \end{cases}. \qquad (11)$$

Here, α is the assumed prior probability of a single spacecraft soft failure and m is the number of visible satellites. This model thus has only m+1 independent failure modes (either no failure or one of m single-satellite bias failures). In [13], $f(e_t|\theta_i)$ is computed based on the least-squares fit model in Section 2.0 after applying a correction for the most likely bias estimate.

This formulation is very promising, and it can be modified to incorporate a PPM and more general fault assumptions. In [13], α is not set from prior information; it is chosen to give a desired P_{FA} just as in traditional RAIM. Introducing the PPM from Section 3.0 does complicate the problem greatly, but it can be simplified by adding just one additional fault mode to (10): a "grab bag" of receiver and multiple satellite faults. The overall prior likelihood of these miscellaneous cases can be estimated from our PPM simulations, and $f(e_t|\theta_i)$ could be obtained by storing residuals conditioned on failure causes θ_i. If the simulation does not provide enough samples of θ_{m+1}, we can arbitrarily represent our uncertainty here by:

$$f(e_t|\theta_{m+1}) = 2\,\|e_t\| / \|e_{\max}\|^2; \qquad \|e_{\max}\| > 2. \qquad (12)$$

This is simply a ramp-shaped probability density which suggests that the likelihood of a residual vector e_t given the occurrence of this "unexplained" fault class increases linearly as the magnitude of e_t increases. The higher the expected "maximum" residual e_{max} is chosen, the more uncertainty exists (because the probability band is spread over more possible results). Of course, an attempt to isolate a single satellite failure can give more information about this likelihood (see [2,3]). Finally, a "loss function" decision cost model (which does not need to set thresholds *per se*) is proposed in [13] and could be augmented to model risk aversion using the concepts in Section 5.0.

An interesting system design problem results from the introduction of GIC: how should system-wide algorithms and parameters be designed to optimally utilize all available information and transform it for use by independent receivers with simple RAIM test capability? A GIC station should be able to carry out the complex filtering or probability updating algorithms outlined here and then transform its posterior results to test thresholds which users can vary based on their distance from the station and their

own specifications and risk aversion. The optimization of the overall shared RAIM algorithm is a system design problem suitable for global search optimization algorithms such as simulated annealing and genetic algorithms. Top-down optimization of this type is one way to avoid relying upon simplifying assumptions.

10.0 SUMMARY AND RECOMMENDATIONS

In this study, a prior probability model of the GPS system was constructed and applied to RAIM algorithm threshold optimization. Monte Carlo simulations of the PPM and the GPS non-precision approach application produced output probability distributions from which optimal decision thresholds were computed to minimize a risk-based cost model. The outputs were also used to show that the chi-square assumption of traditional RAIM is questionable in real GPS operating conditions. The PPM-based thresholds represent a more general solution to the RAIM problem, and with more design-stage simulation and cost modeling effort could be turned into a high-fidelity snapshot RAIM algorithm.

In addition to modeling precision approaches using DGPS and carrier-phase ambiguity resolution, the future potential of this methodology is illustrated by multi-step RAIM algorithms based on generalized failure and/or bias probability updates in real time. While complex algorithms of this type may not be practical for all GPS users, the planned introduction of GIC and the resulting centralization of RAIM should allow these updates to be computed at GIC stations and used to optimize the failure status information to be disseminated to users.

ACKNOWLEDGMENTS

The authors would like to thank the following people for their help with this research and the software on which it is based: Dr. Clark Cohen, Dr. Per Enge, Dave Lawrence, E. Glenn Lightsey, Dr. Hiro Uematsu, and Dr. Todd Walter. The advice and interest of many other people in the Stanford GPS research group is also highly appreciated.

REFERENCES

[1] Col. G.B. Green, P.D. Massatt, and N.W. Rhodus, "The GPS 21 Primary Satellite Constellation", *Navigation*, Vol. 36, No. 1, Spring 1989, pp. 9-24.

[2] B. W. Parkinson and P. Axelrad, "Autonomous GPS Integrity Monitoring Using the Pseudorange Residual", *Navigation*, Vol. 35, No. 2, Summer 1988, pp. 255-274.

[3] M. A. Sturza, "Navigation System Integrity Monitoring Using Redundant Measurements", *Navigation*, Vol. 35, No. 4, Winter 1988-1989, pp. 483-501.

[4] M. A. Sturza and A. K. Brown, "Comparison of Fixed and Variable Threshold RAIM Algorithms", *Proceedings of ION GPS-90*. Colorado Springs, CO., Sept. 19-21, 1990, pp. 437-443.

[5] R. G. Brown, "A Baseline GPS RAIM Scheme and a Note on the Equivalence of Three RAIM Methods", *Navigation*, Vol. 39, No. 3, Fall 1992, pp. 301-316.

[6] J-M. Durand and A. Caseau, "GPS Availability, Part II: Evaluation of State Probabilities for 21 Satellite and 24 Satellite Constellations", *Navigation*, Vol. 37, No. 3, Fall 1990, pp. 285-297.

[7] A. G. Gower, "Putting a Number on GPS Integrity: The IBM GPS Integrity Study for the DoD", *Proceedings of ION GPS-91*. Albuquerque, NM. Sept. 11-13, 1991, pp. 753-760.

[8] B. S. Pervan, C. E. Cohen, and B. W. Parkinson, "Autonomous Integrity Monitoring for Precision Approach using DGPS and a Ground-Based Pseudolite", *Proceedings of ION GPS-93*. Salt Lake City, UT., Sept. 22-24, 1993.

[9] J. L. Farrell and F. von Graas, "Statistical Validation for GPS Integrity Test", *Proceedings of ION GPS-91*. Albuquerque, NM., Sept. 11-13, 1991, pp.773-79.

[10] J. M. Davis amd R. J. Kelly, "RNP Tunnel Concept for Precision Approach with GNSS Application", *Proceedings of the ION 49th Annual Meeting.* Cambridge, MA., June 21-23, 1993, pp. 1-20.

[11] P. W. McBurney and R. G. Brown, "Self-Contained GPS Integrity Monitoring Using a Censored Kalman Filter", *Proceedings of ION GPS-88.* Colorado Springs, CO., Sept. 19-23, 1988, pp. 441-450.

[12] P. Enge, A.J. van Dierendonck, and G. Kinal, "A Signal Design for the GIC Which Includes Capacity for WADGPS Data", *Proceedings of ION GPS-92*. Albuquerque, NM., Sept. 16-18, 1992, pp. 875-884.

[13] S. Bancroft and S. S. Chen, "Integrity Monitoring Using Bayes' Rule", *Proceedings of ION GPS-91*. Albuquerque, NM., Sept. 11-13, 1991, pp. 781-787.