

# On the Availability of Fault Detection and Exclusion in GNSS Receiver Autonomous Integrity Monitoring

Jinling Wang and Pieter B. Ober

(*The University of New South Wales*)

(Email: Jinling.Wang@unsw.edu.au)

Global Navigation Satellite System (GNSS) Receiver Autonomous Integrity Monitoring (RAIM) is essential for safety-of-life and liability critical applications. This paper discusses two fundamentally different ways to assess the integrity risk of an operation with RAIM, based on a different amount of information available: the expected (or average) performance that is computed using the GNSS models only and the real-time (or actual) performance, which also uses information on the internal status of a GNSS receiver. It is shown both theoretically and by simulation that the real-time integrity risk significantly exceeds the expected risk after the detection and exclusion of a failing satellite. Therefore, while most published RAIM algorithms base their performance assessment on the expected performance only, this is only correct when the requirements allow the risk evaluation to be averaged over multiple operations. However, when the GNSS integrity requirement is to be applied on a ‘per operation’ basis, real-time integrity measures are more appropriate.

## KEY WORDS

1. Integrity.
2. RAIM.
3. GNSS.

1. INTRODUCTION. As more and more human activities are relying on the use of satellite navigation technology, the integrity of navigation solutions has become a major issue, especially for life-safety-critical applications (e.g., Ochieng *et al.*, 2003) and liability-critical applications (Beech *et al.*, 2005). Therefore, a reliable integrity monitoring procedure must be used to eliminate hazardous and misleading navigation information caused by failure(s) within the navigation system and provide a timely warning message to the user if the navigation information is not good enough for certain applications at the specific time. To address the satellite navigation integrity risks, two strategies have been proposed (e.g., Kinal *et al.*, 1992): GNSS Integrity Channel (GIC); Receiver Autonomous Integrity Monitoring (RAIM).

The GIC strategy is based on an independent network of ground monitoring stations and a means of conveying the results of the integrity monitoring to the users, e.g., geostationary satellites for regional and global coverage. However, apart from the high costs and time delay of the GIC operations, the GIC cannot check the full integrity at the user level because its monitoring stations cannot detect the local

Table 1. Signal-In-Space Integrity Requirements from ICAO (2006).

Operation	Integrity risk	Alert Limits	
		horizontal/vertical	TTA
En-route	$10^{-7}$ /hour	4 NM/ -	300 s
En-route, Terminal	$10^{-7}$ /hour	2 NM/ -	15 s
NPA	$10^{-7}$ /hour	556 m/ -	10 s
APV-I	$2 \times 10^{-7}$ /approach	556 m/50 m	10 s
APV-II	$2 \times 10^{-7}$ /approach	40 m/20 m	6 s
CAT-I	$2 \times 10^{-7}$ /approach	40 m/15 m to 10 m	6 s

ranging biases/failures at the user receiver, such as extreme user multipath, and local signal interferences etc.

The RAIM strategy is based on a consistency check of satellite measurements used in a navigation solution. If a faulty measurement/satellite (failure) is detected, a procedure may be activated to identify and exclude the failure from the navigation solution, which will therefore remain fault-free and reliable for use in defined applications. Thus, a RAIM procedure is self-contained and can be used as the ultimate integrity monitor (Stansell, 2000). However, the integrity risk associated with the use of RAIM should satisfy the defined requirements.

When specifying integrity requirements, one should determine which position errors are still acceptable, and which errors are not. Therefore, to ensure that the position error is acceptable, an Alert Limit (AL) is defined to represent the largest position error which results in a safe operation. The position error may not exceed this alert limit without annunciation with a probability higher than that stated by the integrity requirements.

For short term operations, the maximum allowable integrity risk is usually specified on a 'per operation' basis. Long lasting operations are customarily thought to be conceptually consisting of multiple segments, and the risk is typically specified per segment rather than for the whole operation. As an example, aviation's Signal-In-Space (SIS) integrity requirements for GNSS-based navigation systems are summarised in Table 1. It can be seen that the integrity requirements are either defined on a 'per approach' basis (for the approach phases of flight) or 'per hour' (for the en-route, terminal and NPA phases). The integrity risk figures are defined to support an appropriate Target Level of Safety (TLS) during each of the operations, and are related to the Alert Limit, which is determined on the basis of obstacle clearance or aircraft separation requirements for the operation at hand.

Although not always explicitly stated, integrity requirements are typically to be understood as a maximum allowed risk for each operation (or segment), rather than being interpreted as an average risk. This paradigm is, however, not always followed rigorously. As this paper shows, the existing RAIM Failure Detection and Exclusion (FDE) algorithms determine their availability on the basis of an average integrity risk, rather than using the risk that is encountered during a particular operation. As a result, these algorithms could underestimate the integrity risk for an operation, and claim availability of the system while in fact they should not. While this might be acceptable for some applications, it might not be for others. Therefore, user communities should be aware of the conceptual difference between the 'average integrity risk' and the actual or 'real-time integrity risk'.

Only when this difference is understood, can these communities ensure that the integrity risks they specify match the integrity risks that users estimate to determine the availability of the system for their particular application. This paper therefore sets out to clarify this difference, which has not been discussed in literature. In particular, great importance is attached to a clearer definition of the integrity risk by showing why and how current ‘average integrity’ based RAIM-FDE algorithms can underestimate the real-time integrity risk after a successful exclusion, both theoretically and in practice through a simple simulation.

**2. GNSS PERFORMANCE.** The positioning performance of a GNSS-based navigation system depends on the following system characteristics:

- the nominal ranging performance of each satellite;
- the satellite geometry as observed by the user;
- the occurrence rate and nature of (rare) failure conditions.

A unique aspect of GNSS when compared to other navigation systems is the time-varying performance caused by changes in the core satellite geometry, even for a stationary user. GNSS performance also varies across the service volume as a result of the varying satellite-user geometries that are concurrently experienced at different locations.

The nature and occurrence of satellite failures (here roughly defined as ‘significant deflections from the nominal performance’) is an important driver of the integrity that the system can provide. Knowledge of the exact fashion in which failures manifest themselves is usually limited, as failures are rare by nature and can be generated by a variety of different mechanisms. The common effect of interest here is that failures are expected to generate considerably larger-than-usual ranging errors, which has generally been modelled by the addition of a bias to the nominal zero-mean pseudorange noise. The size of the bias usually remains unspecified to reflect the lack of knowledge on the failures, but assumptions on this size may still be necessary to enable evaluation of the system’s performance in the presence of failures.

*2.1. Real-time/measurement-based and predicted/model-based performance.* Due to the variations in GNSS performance over time, use of GNSS for critical operations typically requires users to assess the integrity risk they are exposed to in (more or less) real time while performing a certain operation. When the integrity risk exceeds the requirements for the operation, the system is to be considered ‘unavailable’ for the operation, and the user is to be alerted that he should revert to some other navigation means or abort the operation. For a number of operations, performance prediction is required in addition to the real-time assessment; see for example FAA (2007) and EASA (2003). The prediction calculates the expected availability of the system by predicting the satellite geometry that will be observed, in combination with other intrinsic performance parameters of GNSS, such as the probability of satellite failure and the pseudo-range noise levels.

The results of a real-time performance assessment can differ from the predicted performance due to unforeseen changes in the conditions, such as the availability of fewer satellites due to masking effects. While some of these effects might be predictable in principle, forecasting them is often too hard to be practical. More fundamentally, there are circumstances that are unpredictable by their very nature, such as

the occurrence of system failures. Performance prediction is necessarily based on models and ‘a priori’ assumptions about the likelihood that a failure occurs. It therefore provides performance estimates that represent an ‘expected value’ or ‘ensemble average’ view of the integrity risk, based on the expectation that only a small fraction of the operations will encounter a satellite failure: the associated integrity risk is the one that is experienced on average over a large number of operations that are performed under the given conditions.

Real-time assessments have access to extra information in the form of pseudo-range measurements. These measurements contain information on the presence of failures. While there are different ways to exploit this information to improve on the calculations of the integrity risk (e.g., Blanch *et al.*, 2007; Ober, 2003), this paper only considers the (potential) use of the internal state of the fault-detection and exclusion circuitry in the receiver as an extra source of information, as this is considered the most instructive way to explain the concepts behind the use of real-time information.

**3. RAIM ALGORITHMS AND THEIR PERFORMANCES.** The first RAIM Failure Detection (FD) algorithms were developed for use with GPS in the second half of the 1980s. In the 1990s, Failure Detection and Exclusion (FDE) algorithms were introduced. Both variants are concisely discussed here in the light of their respective expected (average) and real-time performances.

**3.1. RAIM-FD Performance and Availability.** When the RAIM-FD is used, both the detection of a failure, as well as the encountering of a period of RAIM-FD unavailability, would cause the system’s use to be discontinued. Receivers that are using the system for navigation therefore operate essentially in one single mode, characterised by the following conditions:

- the assessed integrity risk is sufficiently low and thus the RAIM-FD is available;
- no failure has been detected by the FD algorithm.

With the limited computing power available at the time, simple algorithms were required to evaluate the integrity risk to determine whether the system is available to support the operation. The algorithms developed at the time were based on the assumption that a satellite failure could be modelled conservatively by a bias of a predefined size (such as the Minimal Detectable Bias (MDB)), the smallest bias that can be detected with some predefined probability (or the largest bias that will remain undetected with that probability). The MDB only depends on the number of satellites in view but not on their relative geometry, and can thus easily be pre-computed and stored in a table (see for example Brown 1992). The RAIM-FD performance is evaluated by verifying whether a bias of the size of the MDB is likely to cause unacceptably large errors in the position domain. If this is the case, the integrity risk at the Alert Limit is deemed too high and the system is unavailable. The only time varying ‘variable’ in this process is the satellite geometry, as real-time receiver information remains unused. As a result, integrity prediction and real-time evaluation can be based on exactly the same ‘predictive’ algorithm that uses either the predicted or the actual real-time satellite geometry.

Now let us have a look at the way integrity is evaluated in the receiver. Assume that the operation or one of its segments takes a time  $T_S$  (in hours) to complete. When the

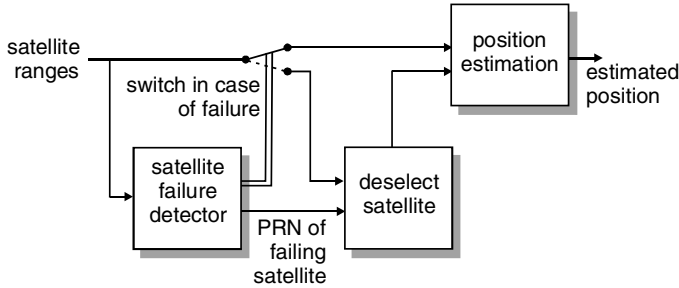


Figure 1. The operation of RAIM FDE. When a failure is detected, the satellite that is most likely to have failed is removed from the position solution. Navigation can only continue successfully when the remaining system performance is still adequate.

probability that a satellite failure occurs over this period is denoted by  $P_F$ , and the average number of satellites in view equals  $N$ , the probability of finding a failing satellite in view is approximately  $N \cdot P_F$ . (For simplicity it is assumed throughout the paper that multiple failures are too rare to be taken into account.) When the probability of not detecting the presence of a failure is  $P_{MD}$  (the probability of missed detection), the integrity risk  $IR_{FD}$  that is experienced by the receiver is given by:

$$IR_{FD} = T_S \cdot N \cdot P_F \cdot P_{MD} \tag{1}$$

The receiver can evaluate this risk and compare it against the requirements for a given operation. Let us assume that the acceptable risk for the operation  $IR_{REQ}$  equals  $10^{-7}$  per hour. When typical parameters for GPS are substituted ( $N = 10$ ,  $P_F = 10^{-5}/hour$ ), and an operation segment of 1 hour is considered ( $T_S = 1 hr$ ), it is readily seen that the integrity risk is acceptable whenever  $P_{MD} \leq 0.001$ . The receiver then checks whether the largest biases that would not meet this requirement cause no excessive position errors. If they do not, the system is considered available for the operation. We will continue to use these typical figures throughout the paper to allow for easy comparison of all outcomes.

3.2. *RAIM-FDE Performance and Availability.* The RAIM-FDE combines a failure detection function with a failure exclusion function, which has the capability to exclude failing satellites. The exclusion algorithm determines which of the satellites has been failing and excludes this satellite from the position solution, thereby allowing further navigation using the GNSS. The basic operation of the exclusion algorithm is depicted in Figure 1.

In the 1990s, FDE algorithms were extensively investigated by the RTCA’s GPS Integrity Working Group SC-159. A baseline algorithm was developed in 1993 by Van Graas and Farrell (1993) and Van Graas (1996). Van Graas’ approach has been adopted in many studies on the availability of RAIM-FDE, see for example Van Dyke and Lee (2002). This baseline algorithm also seems to form the basis of most of the previously performed studies into the use of RAIM for combined Galileo and GPS systems by for example Lee (1999), Van Dyke (2001) and Hewitson *et al.* (2004), Hewitson and Wang (2006; 2007) as well as the report from Eurocae (2003) and the studies therein described. In case of a detected failure with  $N$  satellites in view, this algorithm searches for a subset of  $N - 1$  satellites in which the failure detection

algorithm no longer detects a failure. When multiple subsets without detection are present, two ‘exclusion’ options are possible:

- navigation is continued with the first subset that is found (successful exclusion);
- navigation is continued with the best subset that is found, that is, the subset with the smallest failure detection test statistic (successful exclusion).

It is assumed here that the second of these options is chosen, as it will provide the lowest probability of wrong exclusion.

In the baseline algorithm, RAIM failure detection and exclusion (FDE) is declared to be available *whenever the failure detection function is still available after exclusion*. The reasoning behind this criterion, which remains unexplained in most papers, seems to be as follows: for a one hour operation, there is an a priori probability of  $T_S \cdot N \cdot P_F$  ( $10^{-4}$ ) that a failure occurs, with a (low) probability  $P_{MD}$  ( $\leq 0.001$ ) of remaining undetected. Consequently, there is an a priori probability of  $T_S \cdot (1 - P_{MD}) \cdot N \cdot P_F \approx T_S \cdot N \cdot P_F$  ( $10^{-4}$ ) that a user navigates with a subset of  $N - 1$  satellites, which has at most a probability of  $IR_{FDE} = T_S \cdot (N \cdot P_F) \cdot P_{MD}$  ( $10^{-7}$ ) of still containing the bad satellite. As long as this risk is associated with a position error that does not exceed the alert limit, the system is said to be available. This requires that the MDBs be related to a probability of missed detection of  $P_{MD} = 0.001$ , and that their influence on the position error are computed for the satellite geometries of the subsets, rather than the full set. As the probability of missed detection will differ for each of the subsets, it is conservatively assumed that the satellite with the worst detectability over all subsets is the one that failed. Lee and Van Dyke (2002) show that, although this assumption is obviously conservative, it doesn't degrade RAIM-FDE availability significantly.

The baseline approach described above is based on a predictive performance perspective: no real-time information is used to assess the chances of a failure during the operation. The integrity risk that is considered is the risk that is therefore associated with the average integrity risk that the user will experience. The line of reasoning can be summarised as follows: only a small fraction of 1 out of 10,000 operations is expected to encounter a failure, which remains undetected in the subset used for further navigation in 1 out of 1000 cases. Therefore, on average, an integrity risk of  $10^{-7}$  can be associated with that condition. However, during an actual operation the receiver *knows* when it has encountered and detected a failure. Once a failure has been detected by the receiver, the risk of continuing that particular operation with one of the subsets might become much higher than  $10^{-7}$ . In fact, as is shown below, based on the criterion used, the integrity risk that is experienced when navigating on a subset that still contains the failing satellite is only guaranteed to be  $10^{-3}$ , although it will typically be lower than that due to the fact that navigation typically continues with the subset that is least likely to contain that failing satellite. However, when the baseline RAIM-FDE availability criterion is used to assess availability during an operation, it can still happen that the actually experienced integrity risk is orders of magnitude higher than the integrity requirement, while the receiver still claims the system to be available.

Put another way: it has been shown above that RAIM-FD receivers can successfully base their real-time performance assessment on the model-based integrity risk. This is due to the fact that the small fractions of operations that encounter a failure (and thus an increased integrity risk) actually stop using the system. As a result, all

receivers that are using the system for navigation operate in one single mode, as indicated above. On the other hand, RAIM-FDE receivers can use a navigation system under two fundamentally different sets of operational conditions:

- The RAIM-FD(E) function was evaluated to be available for the operation;
- no failure has been detected by the FD algorithm ('no detection' case);

or

- The RAIM-FD(E) function was evaluated to be available for the operation
- A failure has been detected by the FD algorithm and a satellite has been successfully excluded ('satellite excluded' case)

As demonstrated, operations under these two conditions can yield highly different integrity risks, which the receiver should take into account when it is to evaluate its real-time integrity risk for an operation, rather than the average risk it would encounter when repeating the operation many times.

In mathematical notation, the integrity risk  $IR_{FDE|NoDet}$  in the 'no detection' case approximately equals the probability that a failure remains undetected, just as in the RAIM-FD only case:

$$IR_{FDE|NoDet} = T_S \cdot N \cdot P_F \cdot P_{MD} \tag{2}$$

which for typical values of the probabilities equals  $10^{-7}$ . After the detection of a satellite failure, the integrity risk  $IR_{FDE|SatEx}$  in the 'satellite excluded' case equals:

$$IR_{FDE|SatEx} = P_{WEX} \tag{3}$$

in which  $P_{WEX}$  is the probability that the wrong satellite has been excluded and the failing satellite is still in the subset that is to be used for further navigation. While the probability of wrong exclusion is hard to compute exactly, see for example Ober and Harriman (2006), simple conservative approximations exist. Within the baseline algorithm that is discussed, the probability of wrong exclusion is controlled by the probability of missed detection in the subset without the suspected satellite. As one knows that after wrong exclusion the failing satellite is still in the subset that is used for further navigation, while the probability that it remains undetected in that subset is at most  $P_{MD}$ , one obtains the following upper bound:

$$P_{WEX} \leq P_{MD} \tag{4}$$

With the traditional GPS parameter values substituted it is thus only guaranteed that  $IR_{FDE|SatEx} \leq 10^{-3}$ . Therefore, using the missed detection probability of subsets as a basis for RAIM-FDE integrity risk and availability assessment potentially significantly underestimates the actual risk. As this upper bound is typically far from tight, typical performance will be significantly better than that for most satellite geometries, but the exact performance that is achieved is not explicitly controlled to remain within the requirements.

3.3. *A Numerical Example.* Above, it has been demonstrated that the availability criterion can underestimate the integrity risk that a user experiences once a failure has been detected and successfully excluded. Whether such underestimation actually occurs will depend on the interaction of the various optimistic and conservative assumptions underpinning the availability criterion. To further investigate their effects

Table 2. Number of geometries with observed missed detection and no exclusion fractions.

	$0-10^{-7}$	$10^{-7}-10^{-6}$	$10^{-6}-10^{-5}$	$10^{-5}-10^{-4}$	Worst value observed
Missed detection	228	2	178	98	$7.3 \times 10^{-5}$
No exclusion	394	1	109	0	$2.5 \times 10^{-6}$

in a limited but realistic setting, a Monte-Carlo simulation to assess the performance of the baseline RAIM-FDE availability criterion has been performed.

In the simulations, a total of 504 satellite geometries are investigated by computing world-wide visibility of the optimized 24 satellite GPS constellation from RTCA/DO-229D (2006) on a grid of 504 locations at one particular moment in time. The performance of RAIM-FDE for each of these geometries is investigated in the presence of a failure on each of the satellites (one at a time). The bias in the range measurements to the failing satellites is chosen such that a probability of missed detection of at least  $10^{-3}$  is achieved for all of the subsets that contain the failing satellite, assuring that RAIM-FDE is ‘available’ when using the baseline criterion. The detection threshold used by the failure-detection algorithm assures that the probability of false detection equals  $0.333 \times 10^{-6}$ , in accordance with the value in RTCA/DO-229D.

The main interest of the simulations lies in estimating  $P_{WEX}$  (or  $IR_{FDE|SatEx}$ ) by determining the fraction of the simulated data points in which a wrong satellite is excluded. To this end, the internal state of the RAIM-FDE algorithm for each particular satellite-bias-geometry combination is recorded for a total of  $10^5$  samples, for which the nominal ranging errors are drawn from a zero mean normal distribution with a standard deviation as prescribed by the WAAS model from RTCA/DO-229D. Statistics for the failure of a random satellite and for the failure of the worst-case satellite (the satellite that showed the highest probability of remaining in the solution after exclusion) are collected separately. The possible outcomes that are recorded for each sample include:

- missed detection;
- correct detection, followed by the exclusion of the correct satellite;
- correct detection, followed by the exclusion of a wrong satellite;
- correct detection, followed by an inconclusive situation in which no exclusion could be made.

It is assumed that exclusion is only performed correctly when the subset without the failing satellite (a) is not detecting a failure and (b) has the smallest failure-detection test statistic over all subsets. As a result, a wrong exclusion occurs when one of the subsets that still includes the failing satellite has the smallest failure-detection test statistic; furthermore, this subset should not detect a failure. Using the baseline exclusion criteria, the only occasion in which no exclusion can be made is the one in which all of the subsets still detect the presence of a failure.

The results are summarised in Tables 2 and 3 as well as Figures 2 and 3. They show that the integrity risk that is experienced after a successful exclusion  $IR_{FDE|SatEx}$  is typically well below the probability of missed detection  $P_{MD}$  of  $10^{-3}$ , and the upper bound in Equation (4) is very conservative. However, for all but a few geometries, the integrity risk that is observed in the simulation is well above the target risk of  $10^{-7}$



Table 3. Number of geometries with observed integrity risks after an exclusion.

	0	$10^{-7}$ – $10^{-6}$	$10^{-6}$ – $10^{-5}$	$10^{-5}$ – $10^{-4}$	Worst value observed
Integrity risk after a random satellite failure and exclusion	189	2	299	14	$2.5 \times 10^{-5}$
Integrity risk after worst satellite failure and exclusion	189	0	0	315	$6 \times 10^{-5}$

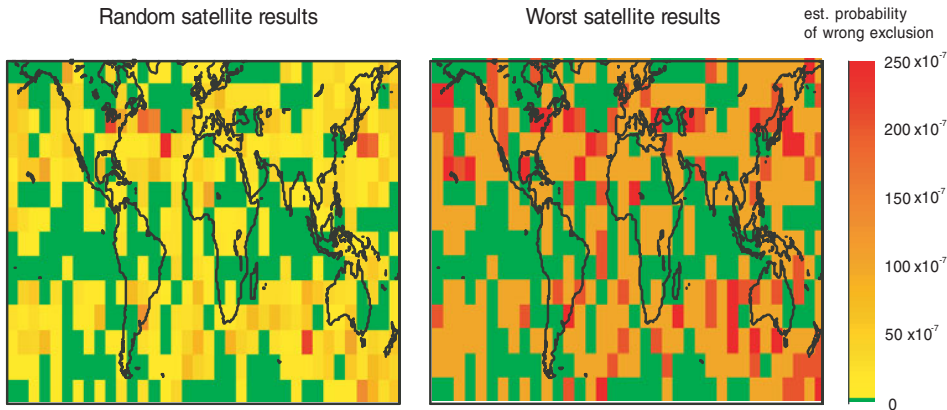


Figure 2. The estimated integrity risk after the exclusion of a satellite for 504 locations on a grid on the earth (latitudes:  $-65, -55, -45, \dots, 55, 65$ , latitudes  $-180, -170, -160, \dots, 170$ ) for a single point in time (GPS Week 703, 344 064 seconds).

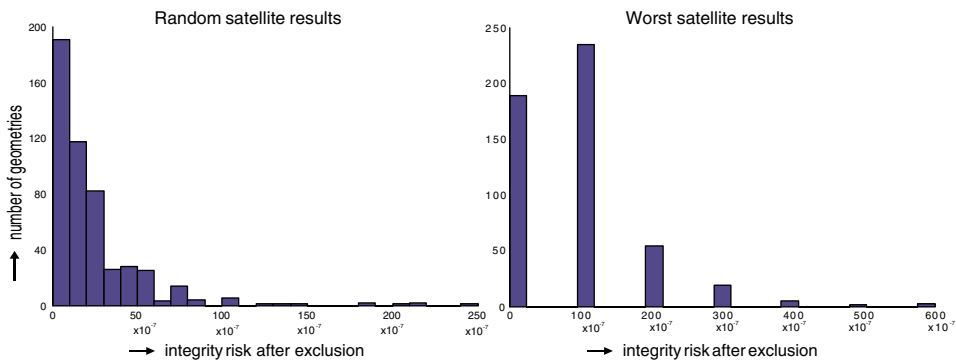


Figure 3. The number of geometries for which a certain integrity risk was found after the exclusion of a satellite—for failures in a random satellite and failures in the worst-case satellite respectively.

that was specified to drive the baseline algorithm to determine RAIM-FDE availability. For the worst location, the risk becomes as high as  $250 \times 10^{-7}$  for those users that experience a failure on a random satellite, increasing to a risk of  $600 \times 10^{-7}$  for those who have to cope with a failure of the worst-case satellite.

4. **CONCLUSIONS.** The paper demonstrates that there are two fundamentally different ways to assess the integrity risk of an operation, based on a different amount of available information. In particular, the availability of real-time receiver information alters the assessment after the detection and exclusion of a failing satellite with respect to the case in which only a priori models of the system can be used. After an exclusion, the real-time integrity risk and system availability should be evaluated conditionally on the fact that there has been a detection to avoid underestimation of this risk, unless it is acceptable that risks are being averaged over multiple operations.

To avoid any misunderstanding, the requirements for any application should clearly state whether the specified integrity risk is to be understood to be applied to each and every operation separately, or whether it is understood that it only needs to be met on average. The following addition to ICAO Annex 10 that was proposed by DeCleene (2005) shows that this importance has already been recognised by the aviation community:

*“The approach integrity requirements apply in any one landing and require a fail-safe design. If the specific risk on a given approach is known to exceed this requirement, the operation should not be conducted”.*

In the same proposal, DeCleene argues that *“The continuity requirement should be applied as applying to the average risk of loss of service”*. It is felt that other user communities should consider similar clarifications where applicable.

Once the requirements are clear, it should be guaranteed that system performance parameters are calculated in a manner that is compatible with these requirements. As this paper shows, the baseline RAIM-FDE algorithm that has been used in many studies to assess the availability of GNSS is unfit to guarantee that the integrity risk for each and every operation remains within the requirements: once a user encounters a satellite failure which is successfully excluded by the receiver, this calls for the use of a different availability criterion.

## REFERENCES

- Beech T.W., Martinez-Olague M.A., Cosmen-Schortmann J. (2005) Integrity: A key Enabler for Liability Critical Applications. *Proceedings of the US ION 61<sup>st</sup> Annual Meeting, 27–29 June, Cambridge, MA*, pp. 1–10.
- Blanch J., Ene A., Walter T., Enge P. (2007) An Optimized Multiple Hypothesis RAIM Algorithm for Vertical Guidance. *Proceedings of ION-GNSS 2007*, pp. 2924–2933.
- Brown, R.G. (1992) A baseline RAIM scheme and a note on the equivalence of three RAIM methods, *Proceedings of the US ION National Technical Meeting*, San Diego CA, January 27–29, 1992
- DeCleene B. (2005) Review of Continuity and Integrity Guidance Material, ICAO NSP WG1&2/WP29 Working Group Meeting, Montreal, Canada, 11–21 October 2005, Agenda Item 9
- EASA (2003) AMC 20-4 Airworthiness Approval and Operational Criteria For the Use of Navigation Systems in European Airspace Designated For Basic RNAV Operations, 2003
- Eurocae (2002) Eurocae WG62 working paper WG62-04-10, Thales Avionics, 4th Eurocae Working Group 62 meeting, 8–9 October 2002
- FAA (2007) Advisory Circular 90–100A on U.S. Terminal and En Route Area Navigation (RNAV) Operations, 2007, <http://web.nbaa.org/public/ops/ac/AC90-100A.pdf>
- Hewitson, S., H.K. Lee and J. Wang (2004) Localizability analysis for GPS/Galileo Receiver Autonomous Integrity Monitoring, *The Journal of Navigation*, **57**, 249–259.
- Hewitson S. & J. Wang (2006) GNSS Receiver Autonomous Integrity Monitoring (RAIM) performance analysis. *GPS Solutions*. 10(3): 155–170.

- Hewitson S. and J. Wang (2007) GNSS Receiver Autonomous Integrity Monitoring (RAIM) with a dynamic model, *The Journal of Navigation*, **60**, 247–263.
- ICAO (2006) International Civil Aviation Organization (ICAO) Annex 10 to the Convention on International Civil Aviation – Aeronautical Telecommunications – Volume 1, Radio Navigation Aids., 6<sup>th</sup> edition, July 2006
- Kinal G.V., Nagle J. & Lipke D.W. (1992) INMARSAT integrity channels for Global Navigation Satellite Systems. *Proceedings of the IEEE*, No. 3, pp. 5–8.
- Lee, Y.C. & K.L. Van Dyke (2002) Analysis performed in support of the ad-hoc working group of RTCA SC-159 on RAIM/FDE issues, *Proceedings of the ION's National Technical Meeting, San Diego CA*, January 28–30, 2002
- Lee, Y.C. (1999) Investigation of New Criteria for Fault Detection and Fault Detection and Exclusion Functions to Obtain Higher Availability, *Proceedings of the 55th Annual Meeting of The Institute of Navigation, Cambridge, MA*, June 1999, pp. 771–778.
- Ober, P.B. (2003) Integrity Prediction and Monitoring of Navigation Systems, *Integricom Publishers*, ISBN 90-77267-01-8, 2003
- Ober, P.B. and D. Harriman (2006) On the Use of Multiconstellation-RAIM for Aircraft Approaches, *Proceedings of ION-GNSS 2006*, pp. 2587–2596.
- Ochieng W.Y., Sauer K., Walsh D., Griffin S., Denney M. (2003) GPS Integrity and Potential Impact on Aviation Safety, *The Journal of Navigation*, **56**, 1–15.
- RTCA/DO-229D (2006), Minimum Operational Performance Standards for Global Positioning System/ Wide Area Augmentation System Airborne Equipment, DO-229D.
- Stansell T. (2000) Current Issues in GPS: GPS after Selective Availability, <http://www.volpe.dot.gov/ourwork/friends/session6.html> (visited on 26/2/2007).
- Van Dyke, K. (2001) Use of Standalone GPS for Approach with Vertical Guidance, *Proceedings of the ION's National Technical Meeting*, Long Beach, CA, January 22–24, pp. 301–309.
- Van Dyke, K. and Y.C. Lee (2002) Analysis Performed in Support of the Ad-Hoc Working Group of RTCA SC-159 RAIM/FDE Issues, *Proceedings of the ION's National Technical Meeting, San Diego, CA*, January 28–30, pp. 639–654.
- Van Graas, F. (1996) Signals Integrity, *NATO AGARD Lecture Series No. 207, System Implications and Innovative Application of Satellite Navigation*, 12 July.
- Van Graas, F. and J. Farrell (1993) Baseline Fault Detection and Exclusion Algorithm, *Proceedings of the 49th annual meeting of the Institute of Navigation*, June 21–23, 1993, pp. 413–420.