Statistical Validation Of SBAS Integrity

P.B. Ober, *Integricom, The Netherlands*, P.B.Ober@integricom.nl
L. de Haan and D. Li, *Erasmus University Rotterdam, The Netherlands*
A. van den Berg and R. Farnworth, *Eurocontrol Experimental Centre, France*

ABSTRACT

This paper summarises a study performed by Integricom and Erasmus University for Eurocontrol in the framework of the SBAS operational validation programme. It defines a structured approach towards the assessment of the integrity that is provided by the EGNOS system from measurement data. It proposes to validate the system in the range domain using the error-overbounding concept and develops a statistical test to assess the likelihood that overbounding is achieved without using unverifiable assumptions on the error distributions that are involved. Special attention is paid to a meaningful extrapolation of the tails of the observed error distributions using extreme value theory (EVT), in order to address the very small probabilities in which the integrity risk requirements are expressed using a limited amount of data.

1. INTRODUCTION

This paper summarises a study performed by Integricom and Erasmus University for Eurocontrol in the framework of the SBAS operational validation programme. To validate system performance of the EGNOS system, it will be required to demonstrate the quality of the positioning solution that the system provides at the level of very small probabilities (in the order of 10^{-7} per hour of flight for aviation applications). While an important part of the demonstration of sufficient performance will consist of simulations and system analysis, especially to assess the performance under specific rare environmental conditions, this paper rather focuses on the assessment of the level of integrity from measurement data.

In the first part of the paper, the operational requirements for the system are reviewed and the corresponding requirements for system validation are derived. In the second part of the paper, a statistical test is developed to derive the probability that the requirements are indeed met by the system. There is only room for explaining the main line of arguments laid down in [Ober04], which means that mathematical details are not provided and many subtleties of the statistical approach are not considered. Two important aspects that are not considered need at least mentioning here. First of all, in any statistical performance assessment, it is important to pay ample attention to the level of time correlation between the measurement samples as this is one of the main drivers of the uncertainty in the assessed system performance, whatever statistical method is applied. However, as the many different error sources in the system all have their own correlation properties, an in depth analysis of the influence of these dependencies is complicated and considered beyond the scope of this paper. Furthermore, SBAS system errors are known to be nonstationary: their error properties are highly varying in time [Walter03]. One should therefore establish precise sets of 'equivalent user conditions' under which the performance is well described and analyse performance for each of these sets separately. Defining such 'system modes' is a far from trivial task. In addition, rare system modes such as for example ionospheric storm conditions cannot be expected to be covered in most sets of measured data. While it is important to deal with the existence of different system modes, this is a complicated subject matter that will not be addressed in this paper; it is silently assumed that all data is collected under one particular system mode here.

2. POSITION DOMAIN REQUIREMENTS

2.1 Operational integrity requirements

According to [SARPS], a system should obey the following integrity requirement to support an operation with an alert limit *AL*: "*the probability that the position error exceeds the alert limit AL while no alert is generated within the alarm time of 6 seconds should remain below* $P_{Req} = 2x10^{-7}$ *per approach (which typically takes 150 seconds). An alert is generated as soon as the protection level PL exceeds the alert limit.*"

As the system should simultaneously support all potential operations that are allowed with the system, and these operations might have different alert limits associated with them, the conditioning on a particular alert limit is usually dropped which gives the following stronger and simpler requirement (which implies the previous one): *the probability that the position error exceeds the protection limit PL should remain below* $2x10^{-7}$ *per approach (which typically takes 150 seconds).*

To link the probability per operation to the (so-called 'marginal') probability density function of the position error at a given sample time, a translation is required to convert the operational requirements to a probability 'per sample' and vice versa. Assume that the user takes N samples during the operation. When all samples are fully independent, the probability P_{MI} that misleading information occurs during the operation (that is, the position error exceeds PL) and the probability of misleading information for each sample $P_{MI,Sample}$ are related by:

$$P_{MI} \approx NP_{MI,Sample} \tag{1}$$

When the samples are dependent, as they will be in practice, this equation becomes inaccurate and the real probability of MI will drop below the value obtained by applying this equation. Although it is important to quantify this correlation effect in practice this problem is not further discussed here: in the remainder of this paper it is implicitly assumed that all requirements are 'per sample'. The subscript 'sample' can therefore be dropped and the required integrity level for each sample is indicated generically by P_{Req}. In mathematical notation this gives the following operational requirement for the system:

$$P_{MI} = P\left(|e| > PL\right) \leq P_{Req} \tag{2}$$

in which e is the position error.

2.2 Position-level integrity validation requirements

Statistical validation should assess whether the probability of misleading information - as to be estimated from the data - obeys the requirements from the SARPS. For a specified type of operation the integrity requirement is given to be P_{Req} per sample. Due to the random nature of the measurements, data analysis cannot prove with certainty that the requirements are met; it can, however, provide some level of confidence P_{Conf} that this is the case. When the minimum acceptable level of confidence is set to some value $P_{Conf,Min}$, this implies that the system validation requirement is of the form:

$$P(P_{MI} \leq P_{Req}) \geq P_{Conf,Min} \tag{3}$$

In a statistical validation campaign one should therefore somehow compute the level of confidence $P(P_{MI} \leq P_{Req})$ that is provided by the data.

3. POSITION-DOMAIN OVERBOUNDING

In the position domain, the operational integrity requirement tells that the probability that the actual position error exceeds the protection level as computed by the user should remain below some value P_{Req}. One could turn this argument around and state that the protection level should be computed such, that (2) is indeed obeyed: the protection level computation relies on a model distribution that should *overbound* the actual error distribution. Overbounding is a concept that describes the relation between the actual error distribution and this model distribution and is further discussed in the second part of this section, as it will be required to address the protection level computation first.

3.1 Protection level computation

The SBAS system provides the user with both corrected pseudoranges and with standard deviations σ_i that describe the quality of these corrected pseudoranges in terms of a *model distribution* - that has been chosen to equal a zero-mean normal distribution. The model's standard deviation that is broadcast should be such, that the use of the model distribution always leads to conservative performance estimates.

range domain overbounding

⇓

position domain overbounding

⇓

protection level exceeds the actual position error with a sufficiently small probability

Figure 1. Requirement summary: in this scheme, the range domain overbounding requirement is the most conservative, but also the most generally valid one: it applies to any kind of operation, regardless of the specific integrity requirements, and to any satellite geometry; the requirement that the protection level exceed the actual position error is the least conservative but is only valid for a given type of operation with given integrity requirements; furthermore, it might be obeyed for some geometries but not for others.

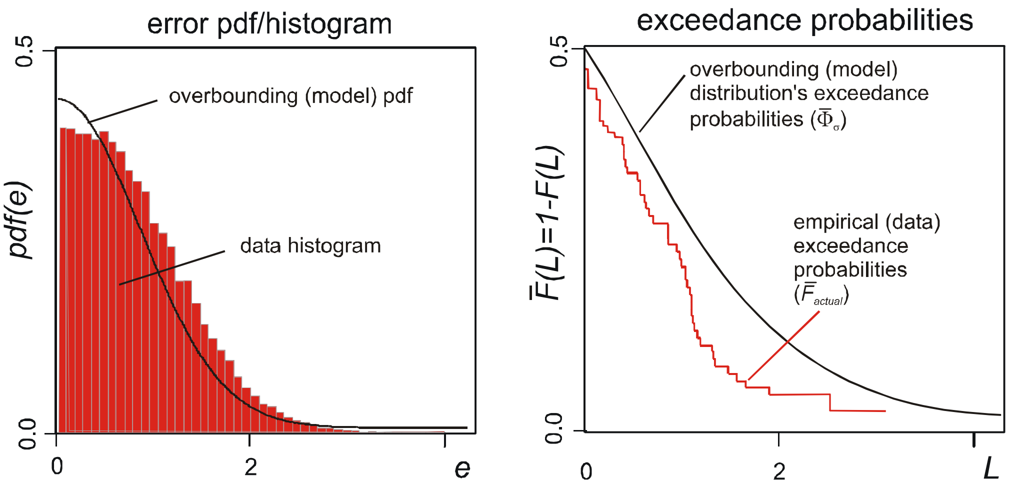

Figure 2. The concept of cdf overbounding: while the pdf of the model distribution cannot exceed the pdf of the data (as approximated by a histogram on the left-hand side), the overbounding distribution has a larger 'exceedance probability' (the probability that the error exceeds some value L) for all positive values of L. In the plot of the exceedance probability, the overbounding model should therefore always lie above that of the distribution of the actual errors (as approximated by the empirical exceedance probabilities).

To assess the integrity of his position, the user computes the standard deviations σ_{hor} and σ_{ver} of the model's horizontal and vertical position error components. To avoid duplication, the remainder of this paper will consider the vertical position error component only; similar arguments also hold for the horizontal error, and the 'ver' subscript will often be dropped to reflect the general validity of the line of argument.

Because the pseudorange errors are (almost) linearly related to the position errors, σ_{ver} can be computed as linear combination of the pseudorange sigma's:

$$\sigma_{ver} = g_{ver,1} \cdot \sigma_1 + \ldots + g_{ver,N} \cdot \sigma_N \quad (4)$$

In this equation, the *g*-factors constitute the measurement geometry. The standard deviations are multiplied by a factor K_{ver} to arrive at the vertical protection level:

$$VPL(t_i) = K_{ver} \cdot \sigma_{ver} \quad (5)$$

For the vertical position component the model distribution is a zero mean normal distribution with standard deviation σ_{ver}. The *K*-factor is defined such that (2) is met with equality for the normal model and are therefore closely related to a particular integrity requirement P_{Req}, see [Roturier01] for details. Whether (2) is also met for the actual distribution of the position error obviously depends on the relationship between the actual and the model distributions: as stated, the model should be such that it always gives conservative performance estimates. In technical terms, one expresses this by saying that the model distribution should *overbound* the actual error distribution.

3.2 Overbounding

This section briefly explains the technical details of the exact meaning of overbounding. When the ICAO SARPS speak of overbounding, this implies overbounding in the sense that could more precisely be called 'cdf-overbounding'. The cumulative distribution function (cdf) represents the energy under the distribution rather than the probability density and will for general distributions be denoted as F_{pos} here. The value $F_{pos}(L)$ equals the area under the distribution from -∞ to L, and is thus defined as:

$$F_{pos}(L) = P(e \le L) = \int_{-\infty}^{L} pdf_{pos}(e)de\,. \quad (6)$$

When the cdf of the actual position error distribution is written as $F_{pos,actual}$ and the cdf of the overbounding Gaussian distribution with standard deviation σ is denoted by the special symbol Φ_σ, one can say that the actual distribution is overbounded by this Gaussian model when the following conditions hold:

$$\begin{aligned} &\textit{left tail}\ :\ \Phi_\sigma(-L) \ge F_{pos,actual}(-L) \text{ for all } L > 0 \\ &\textit{right tail}:\ \bar{\Phi}_\sigma(L) \ge \bar{F}_{pos,actual}(L) \text{ for all } L > 0 \end{aligned} \quad (7)$$

in which the right tail's exceedance probabilities are defined using the 'bar' notation as:

$$\begin{aligned} \bar{\Phi}_\sigma(L) &= 1 - \Phi_\sigma(L) \\ \bar{F}_{pos,actual}(L) &= 1 - F_{pos,actual}(L) \end{aligned}. \quad (8)$$

From this definition it can be seen that overbounding does not rely on the existence of a special point (like an 'alert or protection level') on the distribution where 'overbounding' starts. Conceptually, overbounding exploits the fact that it is always advantageous to make errors smaller. An attempt to depict the meaning of overbounding has been made in Figure 2. It can be seen that the overbounding pdf does not exceed the actual one everywhere, although it necessarily does in the 'far out' tail area, but the exceedance probabilities of the model should always dominate those of the actual error distribution.

It can readily be shown that overbounding implies that the operational integrity requirement of equation (2) is met regardless of the particular integrity requirement:

$$\begin{aligned}&\textit{position domain overbounding} \Rightarrow \\ &\quad P\left(|e| > PL\right) \leq P_{Req} \textit{ for all } P_{Req} \geq 0\end{aligned} \tag{9}$$

As a result, the validation requirement (3) can conservatively be written in terms of the probability that the position's error distribution is overbounded by the SBAS model distribution:

$$\begin{aligned}&P\left(\textit{position domain overbounding}\right) \geq P_{Conf,Min} \\ &\quad \Rightarrow P(P_{MI} \leq P_{Req}) \geq P_{Conf,Min} \textit{ for all } P_{Req} \geq 0\end{aligned} \tag{10}$$

There are two important reasons to prefer validation of this overbounding relationship to using (3). First of all, (10) does not rely on any particular value of the integrity requirement and makes the validation effort applicable to all possible applications simultaneously. Furthermore, as shown in [DeCleene00], (9) and (10) can be translated to equivalent requirements in the range domain, while there exists no range domain equivalent of (3).

4. RANGE DOMAIN REQUIREMENTS

Under certain conditions on the pseudorange's error distribution, the position-domain requirements can be translated into range domain requirements. Note that there is no known way to translate the requirements for a particular type of operation with a particular integrity into a range domain equivalent; current knowledge only provides for a way to translate overbounding of the position error *position domain overbounding* into a *range domain overbounding* requirement, provided that a number of extra conditions are met. In particular, to guarantee that the translations works, the (pseudo)range error's pdf needs to be [Pincus97], [Schempp02]:

- symmetrical around its mean;
- unimodal;
- overbounded by the Gaussian model distribution

Furthermore, the mean cannot be too large in relation to the standard deviation of the overbounding distribution - that is zero mean by design; for details on the exact conditions, see [Schemp02]. As noted, under these conditions the translation works for all possible integrity requirements and thus for all possible phases of flight - as it essentially ensures that the SBAS model distribution is guaranteed to be conservative in both the range and the position domain. Moreover, the conditions above ensure that overbounding translates correctly to the position domain for all possible satellite geometries. Note that the conditions on the error distribution are *sufficient* but *not necessary*. This means that overbounding in the range domain is only -as the SARPS put it – 'one way' of proving that the position domain requirements are met. While overbounding in the range domain is thus sufficient, it is not necessary for meeting the position domain requirements. However, there is currently no knowledge on how the conditions in the range domain could be relaxed – hence, from a practical viewpoint, one needs to prove overbounding when validating the system in the range domain.

When considering the distribution $F_{PR,actual}$ of the absolute value of the normalised pseudorange error (that is, the error divided by the broadcast standard deviation of the error model), Gaussian overbounding in the range domain implies that the following holds:

$$2\bar{\Phi}(L) \geq \bar{F}_{PR,actual}(L) \text{ for all } L \geq 0\,. \tag{11}$$

in which Φ equals the cdf of the overbounding zero-mean Gaussian distribution with unit variance (which is used to the normalisation of the errors); the factor two stems from the fact that the absolute values of the errors are considered instead of the errors themselves, which makes sense because the distribution is assumed to be symmetric around zero.

Validation of the system in the range domain requires showing that the overbounding relation (11) holds for the actual but unknown error distribution with a sufficiently high probability $P_{Conf,Min}$:

$$P\begin{pmatrix}\left(2\bar{\Phi}(L) \geq \bar{F}_{PR,actual}(L) \text{ for all } L \geq 0\right) \\ +\textit{extra conditions on pdf are met}\end{pmatrix} \geq P_{Conf,Min} \tag{12}$$

which is the range-domain equivalent of (10).

5. RANGE VERSUS POSITION DOMAIN

When validating the system, both the position and the range domain can be taken as a starting point. As the range domain errors themselves are a sum of multiple error sources (related to satellite ephemeris, satellite clock, ionospheric and tropospheric delay and the receiver) one could even consider analysing each of these sources separately for overbounding, although that will not be considered here.

It can be viewed from the discussion above that each validation domain has its pros and cons. The main disadvantage of the position domain is that the results that are obtained are not valid for arbitrary geometries. Geometry is one of the major drivers of the position domain performance of the system and will vary for each operation. As the system requirements should be met for each individual operation, one should ensure that the system integrity suffices regardless of the geometry. For that reason, the position domain is considered to be inappropriate as the main validation domain and it is recommended to consider the range domain as a starting point instead: even though validation using the range domain is more conservative, it delivers generally valid results and as such is apter for validation than the position domain.

The major disadvantage of range domain validation is that it is build upon the assumption that the range error distributions are symmetric, unimodal and only slightly biased. One could consider to test formally whether all these conditions are actually satisfied. However, symmetry and unimodality are not easily testable, and testing them against a high level of confidence might make the integrity assessment extremely conservative. Furthermore, these conditions might not even be required as they are only sufficient but not necessary. It is therefore proposed in [Ober04] to rather check explicitly whether overbounding is indeed translated to the position domain, for a necessarily limited -but representative- number of geometries. There are number of different options here. One could for example select geometries that have been observed during a measurement campaign (by either using the all-in-view solution or any other 'rule' to compute positions), essentially building up a position error distribution that corresponds to the particular mix of geometries as they have occurred during the campaign. One could also generate position error distributions from the range errors and their models that have been observed, by 'off-line' computing of the position error distributions that would be governed by a certain number of selected geometries. In the latter case, the geometries can be selected independently of the data collection, which provides a larger degree of freedom in the geometries that can be studied. The main tradeoff that has to be made is the one between the number of geometries used in the check versus the amount of processing time one is prepared to invest.

6. TESTING FOR OVERBOUNDING

Overbounding is not related to a single quantile of an error distribution, but rather to the whole distribution function: all quantiles of the actual error distribution should be smaller than the corresponding quantiles of the overbounding model distribution. As the actual error distribution is unknown, all inferences should be based on its sample equivalent that is usually referred to as the empirical distribution function, here written as $\hat{F}_{actual}(L)$, which is defined as the fraction of samples that remain below L:

$$\hat{F}_{actual}(L) = \frac{number\ of\ samples \leq L}{N}. \tag{13}$$

One way to check whether overbounding is present[1], is to compare the observed and the model distribution at some grid of selected quantiles of interest [Barrett03. For simplicity, consider the case that two points L_i and L_j are checked and that confidence values $P_{Conf,i}$ and $P_{Conf,j}$ are obtained in the statement that the model is indeed conservative at these points. The question then becomes, how much confidence the two separate checks together provide in the statement that overbounding is present at *both* quantiles. When the tests would be independent, the answer to this question would be easily obtained: the probability of overbounding in both checks equals the product $P_{Conf,i}$ $P_{Conf,j}$ - the confidence is thus reduced due to the fact that each of the checks can fail: the more checks are done, the higher the probability of finding at least one quantile that is not modelled conservatively. In reality, the checks will be highly dependent, making the effect of a reduced confidence level much less severe than sketched above. Unfortunately, the problem of exactly computing the influence of dependence is unsolved [Dufour97], which means that this fact cannot be exploited easily.

It is conceptually simpler to use a test based on a global measure of the distance between the error distribution and the overbounding model. One of the standard tests that could be considered is the Kolmogorov-Smirnov test. The standard Kolmogorov statistic KS is defined as the largest difference between the empirical (data-based) distribution $\hat{F}_{actual}(L)$ of the normalised absolute error and the normal model:

$$KS = \max_{L \geq 0}\left(2\bar{\Phi}(L) - \hat{\bar{F}}_{actual}(L)\right) \tag{14}$$

KS can be used to assess whether the actual distribution significantly differs from the model distribution: a large value indicates that at some point, the model lies much above the actual distribution [Conover80]. However, to validate overbounding, it needs to be checked whether the model overbounds *everywhere*, which rather calls for using the *smallest* vertical margin between the model and the empirical distribution (which should be sufficiently large to provide sufficient confidence). This gives the candidate test statistic KS', see also Figure 3:

$$KS' = \min_{L \geq 0}\left(2\bar{\Phi}(L) - \hat{\bar{F}}_{actual}(L)\right) \tag{15}$$

[1] As the overbounding requirement has exactly the same form in the range and the position domain, the discussion here refers to both domains

It can be seen that intuitively, large values of *KS'* indicate that overbounding is likely, while values near or below zero make overbounding unlikely. Note however, that the test statistic as defined above will never be smaller than zero, because:

- near zero, the actual distribution and the model distribution necessarily coincide, that is, $\hat{\bar{F}}_{actual}(0) = 2\bar{\Phi}(0) = 1$
- at infinity the model and the empirical distribution again coincide, that is, $\hat{\bar{F}}_{actual}(\infty) = 2\bar{\Phi}(\infty) = 0$

In practice, the test statistic should therefore be adapted to exclude the regions around zero and infinity to read:

$$KS' = \min_{L_{Min} \le L \le L_{Tail}} \left(2\bar{\Phi}(L) - \hat{\bar{F}}_{actual}(L) \right) \qquad (16)$$

The new test statistic now only tests for overbounding on the interval $L_{Min} \le L \le L_{Tail}$, in which L_{Tail} is the start of the tail of the distribution. This can be justified on basis of the following observations:

- for the near-zero errors below some small threshold L_{Min} non-overbounding is acceptable as these will not degrade the performance significantly;
- for $L > L_{Tail}$ a separate test for overbounding will be devised shortly.

Unfortunately, the distribution of *KS'* will depend on the unknown distribution F_{actual}. However, one can derive a conservative bound that does not depend on this distribution [Conover80] and use it to derive the probability $\hat{P}_{Over}$ that overbounding is achieved when a value of KS'_{obs} is observed:

$$P\left(overbounding \mid KS'_{obs}\right) \ge \hat{P}_{Over} \qquad (17)$$

Note that the 'hat notation' is used to indicate that the confidence level/overbounding probability is estimated from the data. Alternatively, resampling procedures could be used to assess the distribution of *KS'* and the value of $\hat{P}_{Over}$ using only the data themselves, see for example [Politis99] and [Lahiri03].

6.1 Testing the tails of the distribution

One important issue has not yet been addressed in the discussion above. SBAS validation requires assessment of the extremely small probability that the integrity of the system is compromised. It will therefore be required to obtain insight into the likelihood of errors that have a near-zero probability to occur. Such errors will generally not be observed in the data, as they are just too rare. To enable the extrapolation of the error distributions into the tails, it is proposed to use extreme-value theory (EVT) rather than imposing strong and unverifiable assumptions on the tails (such as the often used assumption of normally distributed errors). EVT is a recently developed but already well-established and mature field in statistics that provides

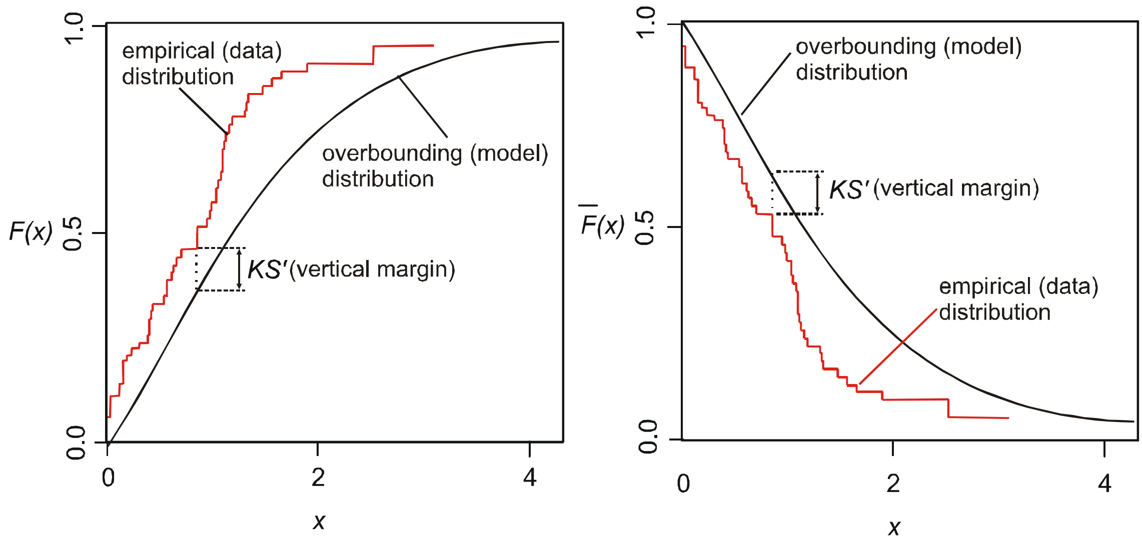

Figure 3. The Kolmogorov statistic KS' indicates whether the empirical distribution function always exceeds the model distribution with sufficient margin. KS' has a negative sign here, corresponding to a favourable margin. As noted in the text, the margin will be zero around 0 and at infinity – the start and the tail of the distribution are therefore to be disregarded to arrive at a practical margin. On the left hand side the original cdfs are depicted, on the right hand side the corresponding tail probability.

statistical methods that allow for the estimation of the probability of the occurrence of rare events regardless of the underlying error distributions [Embrechts97]. The application areas in which EVT has been successfully used are numerous and include hydrology (flood frequency analysis), finance, insurance (insurance and reinsurance risk assessment, claim size assessment, asset price analysis), environmental analysis (ozone level, pollution analysis), meteorology (wind strength and rainfall assessment), earthquake risk assessment) and many engineering areas (corrosion and fatigue prediction, telecommunication). EVT is applicable regardless of the underlying error distributions of the measurement data, relieving the need for the a priori assumption of Gaussian error distributions. The properties of the tails of the error distributions can be derived from the measurement data and some generally valid a priori considerations, allowing to meaningfully extrapolate the data into the region of misleading information, even when no (or very limited amounts of) sample values in this region are available. A goodness-of-fit check can be used to check the a priori assumptions. By finding a data-based description of the tails of the error distributions the ICAO conditions can be validated for the tails of these distributions and the actual protection level can be estimated from the measurement and/or position estimation error data.

One way to assess the probability of overbounding using EVT is to define an extrapolated version of the Kolmogorov statistic:

$$KS'_{ext} = \min_{L \geq L_{Tail}} \left(2\bar{\Phi}(L) - \hat{\bar{F}}_{actual,ext}(L) \right) \tag{18}$$

in which $\hat{F}_{actual,ext}(L)$ is no longer the 'normal' empirical error distribution as in (13), but rather its EVT-extrapolated equivalent, as indicated explicitly by the subscript 'ext' here. Overbounding will correspond to $KS'_{ext} \geq 0$ and that negative values of KS'_{ext} indicate that the model is not conservative somewhere on the distribution. By computing (18) for the whole dataset, one can thus check whether overbounding is present. For practical computations, one can use the fact that the error and model distributions are continuous and evaluate the right-hand side of (18) only for a limited number of values of L ranging from the start of the tail L_{Tail} to the largest quantile that is considered to be of interest – which should have an exceedance probability that is considerably smaller than the integrity requirement that is to be validated.

To find the probability that overbounding is achieved for the actual error distribution, one should go one step beyond this simple check and include the uncertainty in the estimation of the unknown error distribution. Because the distributional properties of KS'_{ext} are unknown, one needs to use a method like resampling (bootstrapping) to compute the probability that the model distribution indeed overbounds the actual distribution [Haan03]. This approach is based on selecting subsets of the data and checking each subset for actual overbounding. Note that unlike what was the case in deriving the properties of the bulk of the distribution, the fact that there is data dependency is not a problem when deriving properties for the tail [Haan03]. For each subset, one can compute whether overbounding is achieved by computing the 'extrapolated modified Kolmogorov-Smirnov statistic' and check whether it remains above zero (overbounding) or not (overbounding). The estimated probability $\hat{P}_{Over,tail}$ that overbounding is achieved approximately equals the fraction of subsets for which the model actually overbounds the data, that is, the fraction of subsets for which KS'_{ext} exceeds zero.

6.2 Combining the overbounding tests

When overbounding probabilities $\hat{P}_{Over}$ and $\hat{P}_{Over,tail}$ have been estimated for both the bulk and the tail of the error distribution, they can be combined into a single confidence level $\hat{P}_{Over,tot}$ for the whole distribution (see also [Schluter02] and [Dufour97]). Because the results on the tail are (asymptotically) independent of those on the rest of the distribution [Haan03], one can simply use:

$$\hat{P}_{Over,tot} \approx \hat{P}_{Over} \hat{P}_{Over,tail} \tag{19}$$

because overbounding of the whole error distribution requires both overbounding in the bulk and in the tail of the distribution.

7. CONCLUDING REMARKS

This paper has assessed the use of data to validate the integrity of the EGNOS system. It has defined a structured approach towards the assessment of the level of integrity that is provided. It has been shown that a system analysis in the range domain has certain advantageous above the use of the position domain; not only does it provide more insight, as the signals are essentially protected in the range domain, the strong effect of geometry on the positioning performance can be left out of the analysis. Finally, much more data will be available in the range domain. It is known that the overbounding concept can be validated in the range domain when the range error distributions obey some extra conditions (of being symmetric, unimodal and near-zero mean). However, it is hard to check these conditions. It is therefore advised to test the system for overbounding in both the range and the position domain; in which the latter is essentially used to verify that the overbounding that is hopefully found in the range domain indeed translates to the position domain. Although a test has been provided, this test is expected to be quite conservative;

unfortunately, standard tests don't seem to be available; designing a less conservative test surely has the potential to improve validation results by providing a tighter bound on the probability that overbounding is achieved and is therefore advisable.

The paper does not answer all of the many question on how to validate EGNOS in a practical campaign. One of the major issues that is yet to be fully resolved is how massive amounts of data from many different sources should be combined. Furthermore, it will not be easy to process these massive amounts of data within a reasonable amount of time – this will not only require an automated way to handle the major parts of the data analysis, but it will also be required to decide on a number of trade-offs between the accuracy of the performance assessment and the amount of computer time that will be required. Accurate decisions are envisioned to require that a substantial amount of data is processed in a number of different ways first - in order to gain insight in their consequences.

REFERENCES

[Barrett03] Barrett, G. and S.G. Donald; "Consistent Tests For Stochastic Dominance", Econometrica, 2003, vol. 71, issue 1, pages 71-104 (see also the working paper at http://economics.web.unsw.edu.au/PEOPLE/GBARRETT/papers/stocdom.pdf)
[Conover80]; Conover, W.J.; "Practical Nonparametric Statistics", John Wiley & Sons, New York, 1980
[DeCleene00] DeCleene, B.; "Defining pseudorange integrity-overbounding", Proceedings of ION-GPS, Salt Lake City, Utah, U.S.A., September 2000
[Dufour97] Dufour Jean-Marie and Olivier Torrès; "Union-Intersection and Sample-Split Methods in Econometrics with Applications to MA and SURE Models", in Handbook of Applied Economic Statistics, ed. by D. E. A. Giles, and A. Ullah, pp. 465–505. Marcel Dekker, New York, 1997
[Embrechts97] Embrechts P., C. Klueppelberg and T. Mikosch; "Modelling Extremal Events for Insurance and Finance", Springer-Verlag, Berlin, 1997
[Haan03] Haan, L. de; Personal communications, 2003
[Lahiri03] Lahiti, S.N.; "Resampling Methods for Dependent Data", Springer Verlag, August 2003
[Linton02] Linton, O., E. Maasoumi and Y-J Whang; "Consistent Testing for Stochastic Dominance: A Subsampling Approach"; Cowles Foundation Discussion Paper No. 1356; Februari 2002
[Ober04] Ober, P.B., L. de Haan and D. Li; "SBAS Integrity Validation", report for Eurocontrol, edition 1.0, March 2004
[Pincus97] Pincus R.; "A note on symmetric unimodal distributions ordered by peakedness", Preprint 27/1997, Univ. Potsdam, Inst. Math., 1997
[Roturier01] Roturier, B., E. Chatre and J. Ventura-Traveset (ESA) "The SBAS Integrity concept Standardised by ICAO-Appliation to EGNOS", Proceedings of GNSS 2001, May 2001
[SARPS] AN-WP/7556 Addendum No.1 Appendix, Draft Amendment To Annex 10 To The Convention On International Civil Aviation, Volume I – Radio Navigation Aids, International Standards And Recommended Practices, Aeronautical Telecommunications
[Schempp02] Schempp, T..R. and A.L. Rubin; "An Application of Gaussian Overbounding for the WAAS Fault Free Error Analysis", Proceedings of ION-GPS 2002
[Schluter02] Schluter, Christian and Mark Trede; "Tails of Lorenz Curves", *Journal of Econometrics*, 109: 151-166, 2002, see also http://www.iza.org/en/papers/Trede240401.pdf
[Walter03] Walter, T., P. Enge and B. DeCleene; "Integrity Lessons from the WAAS Integrity Performance Panel (WIPP)", Proceedings of the ION National Technical Meeting, 2003