DOES THE HPL BOUND THE HPE?

Christian Tiberius and Dennis Odijk

Delft Institute of Earth Observation and Space Systems (DEOS)
Delft University of Technology
Kluyverweg 1, 2629 HS Delft, The Netherlands
Email: c.c.j.m.tiberius@tudelft.nl, d.odijk@tudelft.nl

ABSTRACT

The Horizontal Position Error (HPE) must be no greater than the confidence bound, known as the Horizontal Protection Level (HPL), beyond the specified probability (integrity risk). The definition of protection levels is one of the keys to integrity in Satellite Based Augmentation Systems (SBAS) as the Wide Area Augmentation System (WAAS) and the European Geostationary Navigation Overlay Service (EGNOS), with the main application of providing guidance to aircraft during approach. In this contribution we show and explain, also by means of geometrical interpretation, that of the two defined HPLs, theoretically — with the HPE as the radial error in the two dimensional horizontal plane — one provides a lowerbound (for the probability, and hence is safe), but that the other one leads to an upperbound (hence *over*estimating the probability contained within the HPL-area). These effects are noticeable in the size of the integrity risk. We present, as an alternative, an *exact* evaluation of the circular probability, and for completeness we include the Chebyshev inequality, which provides a lowerbound for the probability regardless the actual distribution of the position error. The measures are mutually compared, and assumptions and relations are pointed out.

INTRODUCTION

Is the Horizontal Position Error (HPE) bounded by the Horizontal Protection Level (HPL), with the required and specified probability? That is the central question of this contribution. In operational circumstances the actual Position Error (PE) can not be assessed, only in controlled experiments an accurate ground-truth for the position or trajectory can be available. In practice a statistical measure on the position error has instead to represent (in part) navigational performance.

In satellite navigation, and with a Satellite Based Augmentation System (SBAS) in particular, the use of the Protection Level (PL) as a statistical bound to the Postion Error (PE) is common place. In one dimension, setting the interval length that bounds the error in one component of the position, given the probability, is usually straightforward, once the distribution of the position error is available. In two (or more) dimensions, this may get complicated by statistical dependence between the position components (and by different possibilities for the geometric area or volume to bound the error).

In aviation typically the (one-dimensional) vertical component of the user position error is of primary concern; the horizontal components are of secondary concern. This situation may be just the other way around in other applications as positioning and navigation on land and at sea. This formed the incentive to the underlying study.

In general we would like to capture the two-dimensional position error in a measure or statement that relates the *radial* error to probability, preferably by exact means. The adjective 'radial' translates into a circle as the area that contains the position error in two dimensions. There is no distinction between the two components. When an exact assessment is not possible, one should underbound the probability or conversely overbound the radius, in order to be 'on the safe side'. The performance measure shall be conservative, as in the end, it is all about ensuring user safety.

There exists a large amount of literature on scalar measures to evaluate or bound multi-dimensional errors, see e.g. [1] for the two-dimensional case, and the topic is also addressed in [2]. In this contribution we review specifically two measures that have been outlined explicitly in earlier versions of the Minimum Operational Performance Standards (MOPS) for SBAS [3]. First, the probability contained in the ellipse just enclosed by the circle provides a lowerbound for the circular probability (and hence is on the safe side). Second, evaluation of the probability in only the worst case direction (and neglecting the other) leads to upperbounding the desired circular probability (and thereby does *not* represent guaranteed navigational performance). Then we present an *exact* measure of probability on the horizontal position error radius. To the authors' knowledge this measure has not been in use yet in the context of SBAS. The basic assumption for all three measures so far is a normally distributed position error. To put this assumption into perspective also the multi-variate Chebyshev inequality is used, which provides a lowerbound for the circular probability, irrespective of the position error distribution. In the last section all four measures are systematically compared in five numerical examples.

This contribution is *not* about error modelling, as overbounding individual error contributions to the total pseudorange error, nor about Fault Detection and Exclusion (FDE), we consider nominal conditions, the so-called 'fault or bias free case' H_0. This contribution *is* about translating the error distribution of a two dimensional position into a single scalar probabilistic measure that describes or bounds the error.

PRELIMINARIES

In this section we give a concise review of the mathematical procedure that turns the pseudorange measurements (of a single epoch) into the receiver position solution. Next, we introduce the ellipse of concentration, and review the eigenvalue decomposition of the variance matrix and finally we summarize the computation of the Horizontal Protection Level (HPL).

Measurement Model

We start from the following model $\underline{y} \sim N(Ax, Q_{yy})$, where vector y contains the SBAS-corrected pseudorange measurements (the underscore denotes the stochastic nature) and vector x denotes the (unknown) parameter vector, with typically three coordinates and the receiver clock offset. The originally non-linear functional relations between the measurements and the parameters have been approximated by linear ones, and design matrix A (full rank) represents the linearized relations. The expectation (or mean) of $\underline{y}$ is $E(\underline{y}) = Ax$. Matrix Q_{yy} is the pseudoranges' variance matrix (positive definite, symmetric), typically a diagonal matrix; the variances for the SBAS-corrected pseudoranges are specified in [3]. Stochastic vector $\underline{y}$ is (assumed to be) normally distributed, indicated by the capital N.

With $\underline{y} = Ax + \underline{e}$, the pseudorange errors $\underline{e}$ have zero mean and are normally distributed: $\underline{e} \sim N(0, Q_{yy})$, and this meets the assumption with the Protection Level equation in earlier versions of [3].

The results in [4], based on data from several experiments, indicate that the Gaussian distribution provides an adequate model for the error in pseudorange code and carrier phase measurements. It shows that in good circumstances, with decent equipment and an appropriate functional model, the measurement noise is normally distributed.

Parameter Estimation and Error Propagation

Based on the functional relation $E(\underline{y}) = Ax$ and an assumed weight matrix W, the Weighted Least-Squares Estimator (WLSE) for the unknown parameters can be obtained $\hat{\underline{x}} = (A^T WA)^{-1} A^T W\underline{y}$ (estimator indicated by the hat-symbol). The estimator is unbiased $E(\hat{\underline{x}}) = x$ and the variance matrix is $Q_{\hat{x}\hat{x}} = (A^T WA)^{-1}\, A^T WQ_{yy}WA(A^T WA)^{-1}$.

When the weight matrix W is taken as the inverse of the measurements' variance matrix, $W = Q_{yy}^{-1}$, as prescribed in [3] for precision approach, the estimator achieves minimum mean squared error (minimum variance) among all linear and unbiased estimators, and is the Best Linear Unbiased Estimator (BLUE).

As the estimator is a linear function of the measurements, the normal distribution is propagated, and we arrive at $\hat{\underline{x}} \sim N(x, Q_{\hat{x}\hat{x}})$.

Notation

In the sequel we will consider the (estimation) error in the position: $\hat{\underline{x}} - x$, for which we have $\hat{\underline{x}} - x \sim N(0, Q_{\hat{x}\hat{x}})$. We will skip the unknown true values x, hence $\hat{\underline{x}} \sim N(0, Q_{\hat{x}\hat{x}})$. For notational convenience we skip the hat-symbol. And by default we assume the position coordinates to be expressed in a local topocentric system East (E), North (N) and Up (U), of which we consider the horizontal coordinates, hence $\underline{x} = \begin{pmatrix} \underline{E} \\ \underline{N} \end{pmatrix}$ and variance matrix $Q_{xx} = \begin{pmatrix} \sigma_E^2 & \sigma_{EN} \\ \sigma_{NE} & \sigma_N^2 \end{pmatrix}$.

The norm of vector x with respect to the metric by matrix M, is given by $\|x\|_M = \sqrt{x^T Mx}$. With the standard metric, matrix $M = I$ equals the identity matrix and is notionally omitted, the set $\{x \in R^2 \mid \|x\| \leq r\}$ describes the area of a circle, centered at the origin, with radius r.

Ellipse of Concentration

When the matrix M is taken as the inverse of the variance matrix $M = Q_{xx}^{-1}$, the set

$$\{x \in R^2 \mid \|x\|^2_{Q_{xx}^{-1}} = \tilde{r}^2\} \tag{1}$$

describes the border-line of an ellipse. This ellipse is called the ellipse of concentration. This set contains points x that all have the same probability density (the PDF is set equal to a constant c, hence $f_{\underline{x}}(x) = c$), as illustrated in Fig. 1.

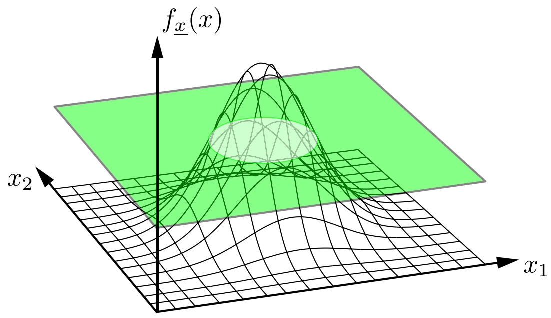

Figure 1: Horizontal cross-section of the two dimensional normal Probability Density Function (PDF) results in an ellipse of concentration.

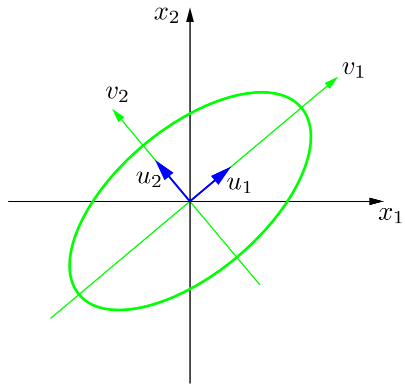

Figure 2: The eigenvectors u_1 and u_2 give the directions of the axes of the ellipse of concentration.

The ellipse represents a contour line of the PDF. The shape and orientation of the ellipse are determined by the variance matrix Q_{xx}.

Eigenvalue Decomposition

The positive definite and symmetric matrix Q_{xx} can be decomposed as follows

$$Q_{xx} = U \Lambda U^T \tag{2}$$

where diagonal matrix Λ contains the (positive) eigenvalues of Q_{xx} (we assume $\lambda_1 \geq \lambda_2$), and (orthogonal) matrix U contains the corresponding eigenvectors u_1 and u_2 as columns (the vectors are orthonormal, i.e. $\|u_1\| = \|u_2\| = 1$ and $u_1^T u_2 = 0$).

The eigenvectors dictate the directions of the (principal) axes of the ellipse of concentration, as shown in Fig. 2. The length of the semi major and semi minor axis are $\sqrt{\lambda_1}\tilde{r}$ and $\sqrt{\lambda_2}\tilde{r}$ respectively.

When we define vector $\underline{v}$ to be the linear combination $\underline{v} = U^T \underline{x}$ (which can be interpreted as a rotation of the coordinate system), error propagation leads to $\underline{v} \sim N(0, \Lambda)$; the two (principal) components $\underline{v}_1$ and $\underline{v}_2$ are uncorrelated (and independent in fact, with the normal distribution). The ellipse of concentration (1) can also be written in terms of v:

$$\|x\|^2_{Q_{xx}^{-1}} = x^T Q_{xx}^{-1} x = v^T \Lambda^{-1} v = \tilde{r}^2 \tag{3}$$

Horizontal Protection Level (HPL)

The equations for computing the protection levels are prescribed in the MOPS [3].

$$HPL_{SBAS} = \begin{cases} K_{H,NPA}\, d_{major} \\ K_{H,PA}\, d_{major} \end{cases} \tag{4}$$

with

$$d_{major} = \sqrt{\frac{\sigma_E^2 + \sigma_N^2}{2} + \sqrt{\left(\frac{\sigma_E^2 - \sigma_N^2}{2}\right)^2 + \sigma_{EN}^2}} \tag{5}$$

computed using the elements of variance matrix Q_{xx}. We have that $d_{major} = \sqrt{\lambda_1}$. Thereby $\sigma_{v_1} = d_{major}$ (regardless of the distribution), and d_{major} is the standard deviation of the one-dimensional position error in the direction of the semi-major axis.

[3] states that 'd_{major} corresponds to the error uncertainty along the semi-major axis of the *error ellipse*', which suggests an elliptically contoured distribution for the position error. Specifically with the normal distribution, d_{major} is the length of the semi-major axis of the standard ellipse of concentration ($\tilde{r} = 1$).

In the following two sections we will address the choice for the two K-factors in (4), the $K_{H,NPA}$ and $K_{H,PA}$.

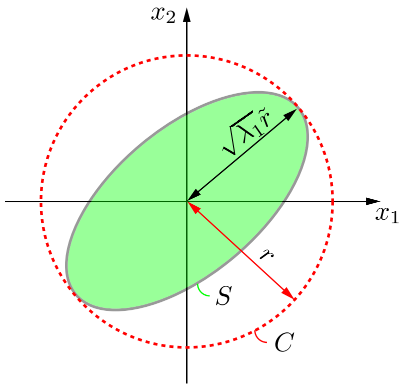

Figure 3: Circle C, with radius r, which just encloses ellipse S. The semi-major axis of the ellipse has length $\sqrt{\lambda_1}\tilde{r}$.

ENCLOSED ELLIPSE

The probability that the position solution lies inside its own so-called ellipse of concentration, is easily obtained. Scalar random variable $\underline{z} = \|\underline{x}\|^2_{Q_{xx}^{-1}} = \underline{x}^T Q_{xx}^{-1} \underline{x}$ has a central Chi-square distribution with two degrees of freedom, hence $\underline{z} \sim \chi^2(2)$.

The probability $P(\|\underline{x}\|^2_{Q_{xx}^{-1}} \leq \tilde{r}^2)$ can thus easily be evaluated through $P(\underline{z} \leq \tilde{r}^2)$, which can be obtained from the central Chi-square Cumulative Distribution Function (CDF) (two degrees of freedom).

$$P(\underline{z} = \underline{x}^T Q_{xx}^{-1} \underline{x} \leq \tilde{r}^2) = \int_0^{\tilde{r}^2} f_{\underline{z}}(z) dz \tag{6}$$

The probability that the position solution lies inside the ellipse S is a (guaranteed) lowerbound for the probability that the solution lies inside the circle C, which just encloses the ellipse. We have $P(\underline{x} \in S) \leq P(\underline{x} \in C)$, as $S \subset C$, as shown in Fig. 3. The probability inside the circle requires integration of the PDF $f_{\underline{x}}(x)$ over a larger region, and hence gives a *larger* probability. The equality holds when $\lambda_1 = \lambda_2$.

intermezzo: Rayleigh Distribution

When we take $\underline{\tilde{z}} = \sqrt{\lambda_1 \underline{z}} = \sqrt{\underline{v}_1^2 + \frac{\lambda_1}{\lambda_2}\underline{v}_2^2}$, then the term under the square root is the sum of squares of two normally distributed variables ($\underline{v}_1$ and $\sqrt{\frac{\lambda_1}{\lambda_2}}\underline{v}_2$) both with zero mean and variance equal to λ_1. Hence $\underline{\tilde{z}}$ has a Rayleigh distribution with parameter $\sqrt{\lambda_1}$, i.e. $\underline{\tilde{z}} \sim Ray(\sqrt{\lambda_1})$.

The above probability $P(\underline{z} \leq \tilde{r}^2)$ can also be obtained as $P(\underline{\tilde{z}} \leq \sqrt{\lambda_1}\tilde{r})$. In practice the standard Rayleigh distribution is used ($Ray(1)$, with parameter equal to 1) and with an integrity risk of $5 \cdot 10^{-9}$ (per independent sample, as specified in an earlier version of [3]), $K_{H,NPA} = 6.18$ is obtained, e.g. in Matlab as `raylinv(1-5e-9,1)`.

The occurrence of the Rayleigh distribution is no surprise as when $\underline{t} \sim Ray(1)$ then $\underline{s} = \underline{t}^2$ is $\underline{s} \sim \chi^2(2)$ Chi-square distributed. Loosely speaking, the Rayleigh distribution is the square root of the Chi-square distribution with two degrees of freedom.

When $\lambda_1 = \lambda_2$ we have $\|\underline{x}\| \sim Ray(\sqrt{\lambda_1})$.

WORST CASE DIRECTION

When we evaluate the probability in only one direction, the worst case direction, and neglect the other direction, we arrive at an upperbound for the circular probability. The position error has largest standard deviation in the direction of the first eigenvector u_1. The position error in this direction is distributed as $\underline{v}_1 \sim N(0, \lambda_1)$. The probability reads

$$P(|\underline{v}_1| \leq \sqrt{\lambda_1}\tilde{r}) \tag{7}$$

The problem has been reduced to the evaluation of the probability of a one dimensional normal distribution. In the context of [3], $d_{major} = \sqrt{\lambda_1}$ and $K_{H,PA} = \tilde{r}$, and with an integrity risk of $2 \cdot 10^{-9}$ (per independent sample, as specified in an earlier version of [3]), $K_{H,PA} = 6.0$ is obtained, e.g. in Matlab as `norminv(1-1e-9,0,1)`.

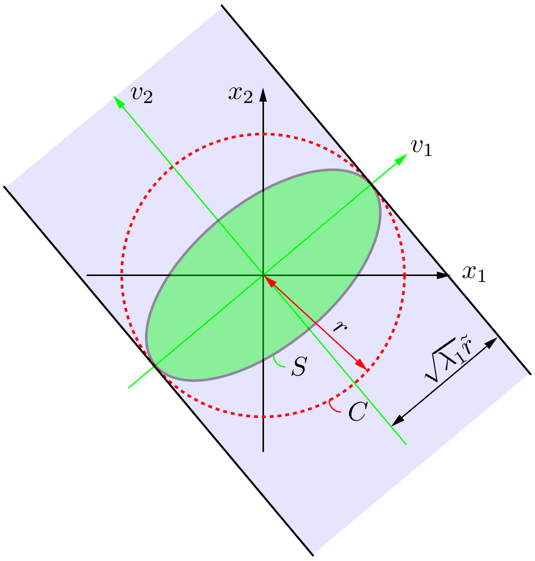

Figure 4: Circle C, with radius r is just enclosed by band B. The one-sided bandwidth is $r = \sqrt{\lambda_1}\tilde{r}$.

The probability computed with (7) represents the probability that the two dimensional position solution is lying in the band B, as shown in Fig. 4. This band has a (one sided) width of $r = \sqrt{\lambda_1}\tilde{r}$ in the direction of u_1, and is unbounded in the other. We have $P(\underline{x} \in B) > P(\underline{x} \in C)$. Or, $P(\underline{x} \notin B) < P(\underline{x} \notin C)$; all of the very small accepted probability (of the position being too far off), the integrity risk, is assigned to only one component, leaving absolutely nothing for the other direction.

Discussion

The MOPS [3] do mention that the evaluation (7) concerns only one dimension/direction (as the along-track direction is of far less relevance), but the general definition of the HPL in the same document, may suggest otherwise. 'The Horizontal Protection Level (HPL_{SBAS}) is the *radius of a circle* in the horizontal plane (...), with its center being at the true position, that describes the region assured to contain the indicated horizontal position.' The Horizontal Alert Limit (HAL), which is the maximum acceptable HPL, is defined with the same words, except for '...the region that is required to contain the indicated horizontal position with the *required probability*.'

One might be inclined to use the HPL as a bound for the HPE, the latter straightforwardly defined as $\|x\| = \sqrt{E^2 + N^2}$, for instance in SBAS verification activities, in particular through the Stanford diagram [5] and [6], where under controlled circumstances the actual HPE is assessed and compared with the given HPL, as e.g. in [7] and [8]. With the HPL too optimistic, one would experience more integrity failures (PE>PL and PE>AL>PL), than one would expect to see on the basis of the K-factor. Though at the required probabilities empirical evidence is rare.

The belief of the PL bounding the PE is corroborated by the concept of overbounding the pseudorange error, see e.g. [6] and [9], which is then propagated into the position domain. The current version of the MOPS [3] points out that it is the responsability of SBAS service providers to broadcast parameter values (for components that build up the pseudorange error variance) such that the HPL_{SBAS} and VPL_{SBAS} bound their respective errors with target probabilities.

If, for precision approach, the HPE is no longer $\|x\|$, but instead taken to be equal to the projection of the position error x onto the worst-case direction given by vector u_1 (with unit length, and obtained from the variance matrix), the HPE can be properly compared to the above HPL. The HPE is then given by $v_1 = u_1^T x$.

EXACT CIRCULAR PROBABILITY

In this section an exact evaluation is presented of the probability that the position solution lies inside a circle with a given radius. The probability is given by, see [10],

$$P(\|\underline{x}\|^2 \leq r^2) = P(\underline{\chi}^2_{2,r_2^2} \leq r_1^2) - P(\underline{\chi}^2_{2,r_1^2} \leq r_2^2) \tag{8}$$

where two non-central Chi-square distributions are involved, both with two degrees of freedom. The Chi-square random variables are denoted by $\underline{\chi}^2$. The non-centrality parameters r_2^2 and r_1^2 respectively, can be computed from the eigenvalues

λ_1 and λ_2 of the variance matrix Q_{xx}.

$$r_1 = \frac{1}{2}(\frac{r}{\sqrt{\lambda_1}} + \frac{r}{\sqrt{\lambda_2}}) \text{ and } r_2 = \frac{1}{2}(\frac{r}{\sqrt{\lambda_2}} - \frac{r}{\sqrt{\lambda_1}})$$

The eigenvalues should be ordered such that $\lambda_1 \geq \lambda_2 \geq 0$.

Equation (8) gives in principle the Cumulative Distribution Function (CDF) of a quadratic form, the squared radius, of normally distributed variables.

It should be emphasized that (8) provides an *exact* evaluation of the probability of $\underline{x}$ lying inside the circle C of Fig. 3, and that it is based on the same information as the previous two measures, namely the mean (assumed to be zero in our discussion), the variance matrix Q_{xx}, and the normal distribution for $\underline{x}$. Tables and algorithms for the non-central Chi-square distribution are available (the numerical evaluation is in principle based on an infinite sum, but in practice the probabilities in (8) can be accurately computed). This implies a clear recommendation for the use of (8), over the previous two measures which are (only) approximations.

intermezzo: Lateral and Vertical Error

By default it has been assumed in this discussion that the two dimensional position error x consisted of the local East and North components. In aviation, in particular in approach mode, as mentioned before, the vertical, or Up component is most important, if not critical. Next the lateral, or cross-track component is relevant, whereas the along-track component is of less importance.

If instead of East, North and Up components, Along-track (L), Cross-track (X), and Up (U) are used (Up-component is maintained), the two-dimensional vector could contain the Cross-track and Up-components $\underline{x} = \begin{pmatrix} \underline{X} \\ \underline{U} \end{pmatrix}$. The circular probability of Cross-track and Up together being within a certain radius, is readily obtained through (8), and can be incorporated in the visualization of the desired flight path by means of a tunnel-in-the-sky, at the current epoch, and possibly predicted a short time span ahead as well. The information of the pair of navigation error probability and radius can be overlaid with the vertical cross-section of the tunnel tube.

CHEBYSHEV INEQUALITY

The last measure does *not* need, as opposed to the previous three, knowledge on the position error distribution. The multivariate Chebyshev inequality requires, next to the mean, just the variance matrix. It provides a (guaranteed) lowerbound for the circular probability. It reads, see e.g. [11],

$$P(\|\underline{x}\|^2 \geq r^2) \leq \frac{\text{trace}(Q_{xx})}{r^2}$$

for every $r > 0$. The probability that $\underline{x}$ resides outside the circle, centered at the origin, and with radius r, is always smaller than (or at most equal to) the sum of the variances divided by r^2. This inequality can be rewritten into

$$P(\|\underline{x}\|^2 < r^2) \geq 1 - \frac{\text{trace}(Q_{xx})}{r^2} \quad (9)$$

EXAMPLES

In this section we employ the four measures to five different variance matrices. The performance is compared in Fig. 5.

The five example variance matrices are:

1. $Q = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}$, no correlation, same variance; the variance matrix is a scaled identity matrix; eigenvalues $\lambda_1 = 2$ and $\lambda_2 = 2$
2. $Q = \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}$, no correlation, different variances; eigenvalues $\lambda_1 = 4$ and $\lambda_2 = 2$
3. $Q = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$, moderate correlation, same variance; eigenvalues $\lambda_1 = 3$ and $\lambda_2 = 1$

4. $Q = \begin{pmatrix} 2 & 1 \\ 1 & 4 \end{pmatrix}$, moderate correlation, different variances; typical practical example from satellite navigation; eigenvalues $\lambda_1 = 3 + \sqrt{2}$ and $\lambda_2 = 3 - \sqrt{2}$
5. $Q = \begin{pmatrix} 2 & 2.6 \\ 2.6 & 4 \end{pmatrix}$, large correlation ($\rho_{x_1 x_2} = 0.92$), different variances; eigenvalues $\lambda_1 = 5.79$ and $\lambda_2 = 0.21$

The worst case direction measure always overbounds the probability, the dashed line is always above the bold solid line of the exact circular probability.

The dotted line of the multi-variate Chebyshev inequality is always below the bold solid line. The inequality provides a guaranteed lowerbound to the probability, and it does not need knowledge about the error distribution, but it is a very loose bound when the contained probability is large, i.e. for large radii. The lower-right graph shows one-minus-the-probability, the integrity risk, in a logarithmic scale, for variance matrix 4. The worst case direction is slightly too optimistic on the size of the integrity risk.

CONCLUSION

The motivation for this study lies in the quest for a measure that adequately represents the horizontal position error. The goal was to find the radius (= the HPL) of the circle, centered at the true position, that contains the position solution with a pre-set probability, or the other way around, to specify the radius and find the probability. Applications lie in precise positioning and navigation on land at at sea with different levels of required probability.

In this contribution four methods have been presented. The HPL based on the enclosed ellipse gives a lowerbound for the probability, and hence is safe. The HPL based on the worst-case direction gives an upperbound for the probability. When one interprets this HPL as a measure for the HPE radius, it is too optimistic. The answer to the question in the title therefore reads a 'no', the HPL does *not* provide a guaranteed bound on the HPE. The presented exact evaluation of the circular probability is fairly straightforward and therefore recommended.

The above three measures start from a normal distribution for the position error. The extent to which this assumption is adequate and acceptable depends on the application and the circumstances. The Chebyshev inequality can dispense with knowledge about the distribution, but is hampered by (severe) overbounding in terms of radius, and practical use is questionable.

References

[1] Leenhouts, P.P., On the computation of bi-normal radial error, *Navigation*, Vol. 32, No. 1, 1985, pp. 16–28.

[2] Ober, P.B., *Integrity prediction and monitoring of navigation systems*, Integricom, Leiden, The Netherlands, 2003.

[3] SC-159 of the RTCA, *Minimum Operational Performance Standards for Global Positioning System/ Wide Area augmentation system airborne equipment*, RTCA - Washington DC, Document DO-229D, December 13, 2006.

[4] Tiberius, C. and K. Borre, Probability distribution of GPS code and phase data, *Zeitschrift für Vermessungswesen*, Vol. 124, No. 8, 1999, pp. 264-273.

[5] website on Stanford diagram, http://waas.stanford.edu/metrics.html, accessed August 2008.

[6] Walter, T., A. Hansen and P. Enge, Validation of the WAAS MOPS integrity equation, *Proceedings of ION 1999 Annual Meeting*, June 28–30, 1999, Cambridge MA, pp. 217–226.

[7] Pullen, S., T. Walter and P. Enge, System overview, recent developments, and future outlook for WAAS and LAAS, *Tokyo University of Mercantile Marine GPS symposium*, Tokyo Japan, November 11-13, 2002.

[8] Roturier, B., E. Chatre and J. Ventura-Traveset, The SBAS integrity concept standardized by ICAO: application to EGNOS, section 1.4 in *EGNOS - The European Geostationary Navigation Overlay System - a cornerstone of Galileo*, SP-1303, 2007, ESA Publications Division, ESA, Noordwijk, The Netherlands, pp. 43–53.

[9] DeCleene, B., Defining pseudorange integrity - overbounding, *Proceedings of ION GPS 2000*, September 19-22, 2000, Salt Lake City, UT, pp. 1916–1924.

[10] Johnson, N.L. and S. Kotz, *Continuous univariate distributions*, Houghton Mifflin Company, Boston, 1970.

[11] Teunissen, P.J.G., D.G. Simons and C.C.J.M. Tiberius, *Probability and observation theory*, Lecture notes AE2-E01. Faculty of Aerospace Engineering, Delft University of Technology, 2006.

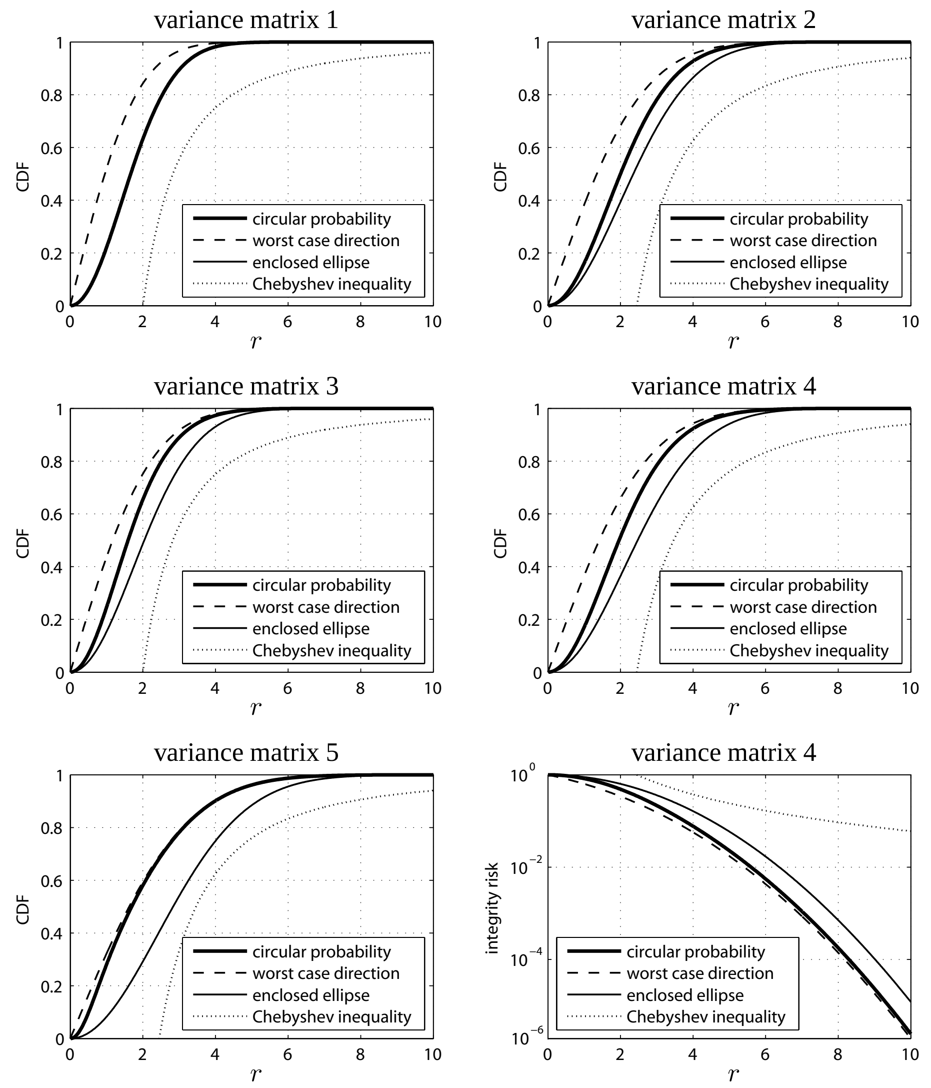

Figure 5: Circular probability for variance matrices 1 through 5, computed using the four methods, with (6) as solid line, (7) as dashed line, (8) as bold solid line, and (9) as dotted line. Probability is visualized as the Cumulative Distribution Function (CDF) with the radius of the circle along the horizontal axis (units can be thought to be in meters), and the probability along the vertical axis. The last graph (bottom right) shows one-minus-the-probability, the integrity risk, for the variance matrix of case 4, in a logarithmic scale.