Fixing the Ambiguities

Are You Sure They're Right?

Peter Joosten and Christian Tiberius Delft University of Technology

Fast and precise relative satellite positioning demands resolution of the integer cycle ambiguities. Only then will the corresponding carrier-phase measurements act as if they were high-precision range measurements, thereby allowing the receiver coordinates to be estimated with comparable high precision.

Researchers have studied the GPS ambiguity problem for the past 20 years and have proposed a wide variety of methods to resolve ambiguities. So far, most of these methods have concentrated on the estimation *of the ambiguities. The problem of assessing the* correctness *of the integer numbers obtained, often referred to as "ambiguity validation," has received considerably less attention.*

The "mission" of this article is to point out that ambiguity resolution is not strictly a matter of computing integer values for the ambiguities. Before really fixing or constraining the ambiguities to the computed integers in a final baseline computation, we should assess their accuracy. In other words, we should ask ourselves "How sure am I that these values are correct?" In this month's contribution, we will look at how we might answer this question and discuss some new developments in dealing with the stochastic properties of the integer ambiguity estimator. The ambiguity success rate is presented as a tool for determining the probability of correct integer estimation.

Our authors are Peter Joosten, who holds an M.Sc. degree from the Delft University of Technology, and Christian Tiberius, who holds M.Sc. and Ph.D. degrees from that institution. Both are employed at Delft University of Technology's Department of Mathematical Geodesy and Positioning. This department is directed by Professor Peter Teunissen, who authored the LAMBDA method for ambiguity resolution in 1993. This method has found widespread use around the world. Recently, research has been extended to the stochastic properties of the integer ambiguity estimator, the topic of this month's column.

Global Navigation Satellite System (GNSS) ambiguity resolution is the process of effectively accounting for the integer property of the unknown initial cycle ambiguities of carrier-phase data, usually in the form of double differences. It applies to a great variety of GNSS data-processing models. This holds true not only for the current Global Positioning System (GPS), but also for GLONASS, the future modernized GPS, and the proposed European Galileo system. The GNSS models range from single-baseline models used for kinematic positioning to multibaseline models used as a tool for monitoring and studying geophysical phenomena such as plate tectonics and ionospheric behavior. The models may have the relative receiver–satellite geometry included (referred to as geometry-based) or excluded (referred to as geometry-free).

The geometry is included through the unit direction vectors in the model's design matrix. When the geometry is excluded, the receiver baseline components are not involved as unknowns in the model, but rather the receiver–satellite ranges themselves. The models may also be discriminated as to whether the remote receivers are in motion or not. When the receivers are moving, we solve for one or more trajectories, because with the receiver–satellite geometry included, we will have new coordinate unknowns for each new epoch. We may also discriminate as to whether the differential atmospheric delays are included as unknowns or not. In case of sufficiently short baselines, these delays are often neglected.

Despite the differences in application of the various GNSS models, their ambiguity-resolution problems are intrinsically the same. In all cases, the aim is to incorporate the integer property of the ambiguities into the least-squares adjustment of the data so as to improve the precision of the results. Once the integer ambiguities are known, the corresponding carrier-phase measurements will act as if they are high-precision pseudorange measurements, thereby allowing the remaining parameters, such as receiver coordinates or baseline components, to be estimated with a comparable high precision.

The improvement, obtained by exploiting the fact that the ambiguities are integers, is illustrated in Figure 1. This figure is based on 1,200 single-epoch experiments, each separated by 3 seconds. The graph on the top presents a position scatter plot of the so-called float solutions. Note the meter-level scale. Each dot represents a computed position, based on observations at a single epoch. For this graph, the ambiguities were estimated as real (floating point) variates. The graph on the bottom shows the fixed solution based on exactly the same observations, but using the fact that the ambiguities are integers and constraining the solution to these integer values. Note the scale difference. Comparing the two graphs shows that the integer-ambiguity information greatly strengthens the data-processing model and yields a much more precise solution.

For ambiguity resolution to be successful, the ambiguities need to be estimated at their correct integer values, as incorrect integers generally bias the receiver coordinates considerably. However, the integer values are determined from noisy observations, and the noise or uncertainty in the observations propagates into the integer values for the ambiguities, making them uncertain. Absolute certainty about an ambiguity's integer value is therefore not possible. But if the uncertainty is too large, there is a serious risk of ending up with severely offset receiver coordinates. It is thus desirable to achieve as high a degree of certainty as possible. And to effectively control this uncertainty, it is desirable to have a mechanism for assessing the probability of correctly estimating the integer ambiguities. After a more detailed discussion of the resolution problem, we will introduce a diagnostic tool that will enable one to rigorously assess the reliability of ambiguity resolution.

INTEGER AMBIGUITY ESTIMATION

For the purpose of ambiguity resolution, GNSS data processing is usually carried out in three sequential steps. In the first step, no distinction is made between the nature of the ambiguities and the other estimated parameters, like receiver coordinates and atmospheric delays. The parameter-estimation problem is solved without taking into account the special integer characteristic of the ambiguities. The result so obtained is often referred to as the float solution

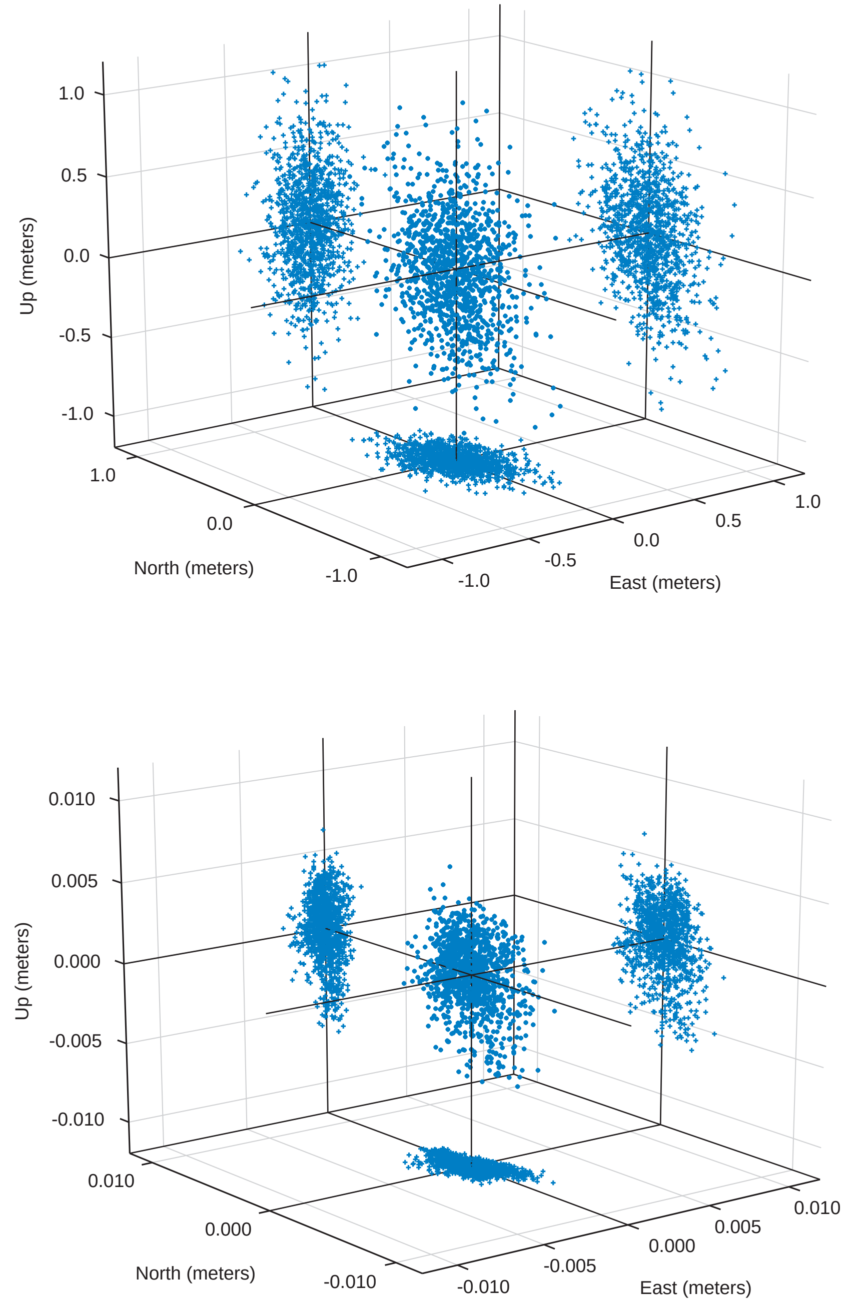

Figure 1. Relative positioning results on a short baseline are expressed in east, north, and up components. Shown are 1,200 single-epoch solutions for the case with ambiguities real-valued (top) and fixed (bottom). For the fixed solution, the ambiguities are resolved correctly in all cases. Fixing the integer ambiguities incorrectly would generally shift the position solution by a decimeter or more. After successful fixing, the precision of the coordinates is below the 1-centimeter level.

because the data-processing software estimates the ambiguities as floating-point numbers. The parameters are usually estimated using a least-squares algorithm, which is commonly accepted as the standard approach to deal with the inconsistencies in the data due to measurement noise.

Two additional steps are necessary to exploit the inherent integer nature of the ambiguities. In the second step, the ambiguity float solu-

The LAMBDA Method

To determine a fixed ambiguity solution from a float ambiguity vector (with n ambiguities), a mapping from the n-dimensional space of reals R^n onto the n-dimensional space of integers Z^n is necessary. Although several approaches are available to achieve this, we recommend applying the least-squares criterion that leads to the integer least-squares estimator for the ambiguities. The solution is optimal in the sense that it maximizes the probability that indeed the correct vector of integer ambiguities is found. Considering the solution geometrically, it minimizes the distance between the float ambiguity vector ($\hat{\mathbf{a}}$) and the integer ambiguity vector ($\check{\mathbf{a}}$), where this distance is measured in the metric of the ambiguity variance-covariance matrix of the float solution. If the variance-covariance matrix of the ambiguities were to be diagonal, the float ambiguity estimates would be uncorrelated and the fixed solution could be obtained by a simple rounding of the float estimates. In general however, the estimates will be correlated, and the fixed solution has to be identified by a discrete search over a subspace of Z^n, specifically, the ambiguity search ellipsoid (see figure below).

In GNSS applications, the ambiguity search space is highly elongated because of a usually high correlation between the ambiguity estimates. It also stretches over a considerable range of wavelengths or cycles as a result of the individual estimates' usual low precision. To improve the computational efficiency of the discrete search, the LAMBDA method employs a decorrelating Z-transformation prior to the search. This Z-transformation yields ambiguities that are less correlated and have improved precision, while retaining the integer character of the minimization problem. The corresponding transformed sphere-like search space allows a relatively efficient identification of the optimal integer least-squares solution.

In summary, the LAMBDA method largely decorrelates the ambiguities, whereafter it carries out a search procedure to efficiently obtain the integer ambiguity vector that has shortest distance to the float ambiguity, thereby maximizing the probability of identifying the correct integer vector. For more details about the method, please consult <http://www.geo.tudelft.nl/mgp/>. Available on request are a FORTRAN and a MATLAB implementation of the LAMBDA method. Directly available for download is an extensive description of the method and its implementation in the report, "The LAMBDA Method for Integer Ambiguity Estimation: Implementation Aspects." The MATLAB implementation comes with a separate guide and also includes a user-friendly demonstration application, which can be used for solving small problems interactively.

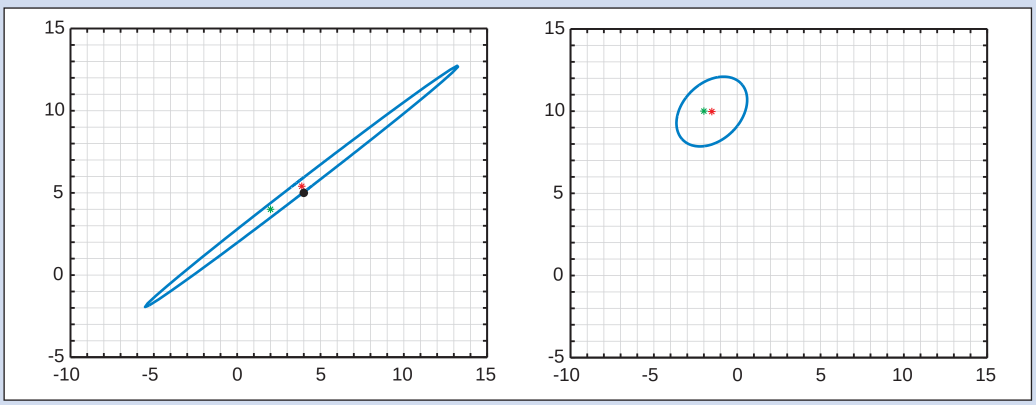

This figure illustrates an example of the search ellipse for two ambiguities showing the search space shape before (left) and after (right) decorrelating the ambiguities by means of the Z-transformation. The search space of the decorrelated ambiguities is clearly less elongated, which allows for an efficient identification of the integer ambiguity solution. Note that the volume of the search space is preserved by the transformation. The red star indicates the float solution, ($\hat{\mathbf{a}}$= [3.875, 5.400], $\hat{\mathbf{z}}$ = [-1.525, 9.975]), and the green star the fixed solution ($\check{\mathbf{a}}$ = [2, 4], $\check{\mathbf{z}}$ = [–2, 10]). The Z-transformation matrix is $\mathbf{Z}^T$ = [1, –1; –3, 4]. Note that rounding the original ambiguities would give a wrong result ($\hat{\mathbf{a}}$= [4, 5]), indicated by a blue circle.

tion is used to estimate the integer ambiguity values. Here we could choose from a wide variety of integer estimation methods. These methods range from simple rounding schemes to more advanced methods based on integer searches. One popular approach is the LAMBDA (Least-squares Ambiguity Decorrelation Adjustment) method, developed at the Delft University of Technology. With this method, the ambiguities are estimated by means of integer least squares using a very efficient search procedure (see the sidebar entitled "The LAMBDA Method" for further details).

Finally in the third step, the computed integer ambiguities are used to improve the first-step solution for the remaining parameters. These parameters are recomputed, but this time with the ambiguities constrained to the integer values obtained from the second step. This final result is referred to as the fixed solution, and it generally inherits a much higher precision than the previously obtained float solution, as was demonstrated in Figure 1.

AMBIGUITIES ARE STOCHASTIC

When computing the fixed solution, the integer ambiguities are usually assumed to be known with certainty. But how sure can we be? After all, the integer ambiguities are determined from noisy data. Only in the hypothetical case of perfect observations, without any noise or other errors, would the float solution always yield the correct integer ambiguity values. In reality, however, this is not the case. Any uncertainty (noise) in the observations will propagate and manifest itself as uncertainty in the integer ambiguities.

A single-frequency example based on the geometry-free GNSS model is shown in Figure 2. The figure illustrates empirically how uncertainty in the data (top left) propagates into the ambiguity float estimate (top right) and finally into the integer ambiguity estimate (lower left). The correct integer for the ambiguity is known to be four in this case, but as one can see from the graph at the lower left, other integer values are frequently obtained.

To capture the integer-ambiguity uncertainty, we have to treat the estimated integer ambiguities as stochastic (random) variates. This is not too different from standard adjustment practice. In standard adjustments, where all parameters are real-valued, we also propagate the observational uncertainty to obtain the

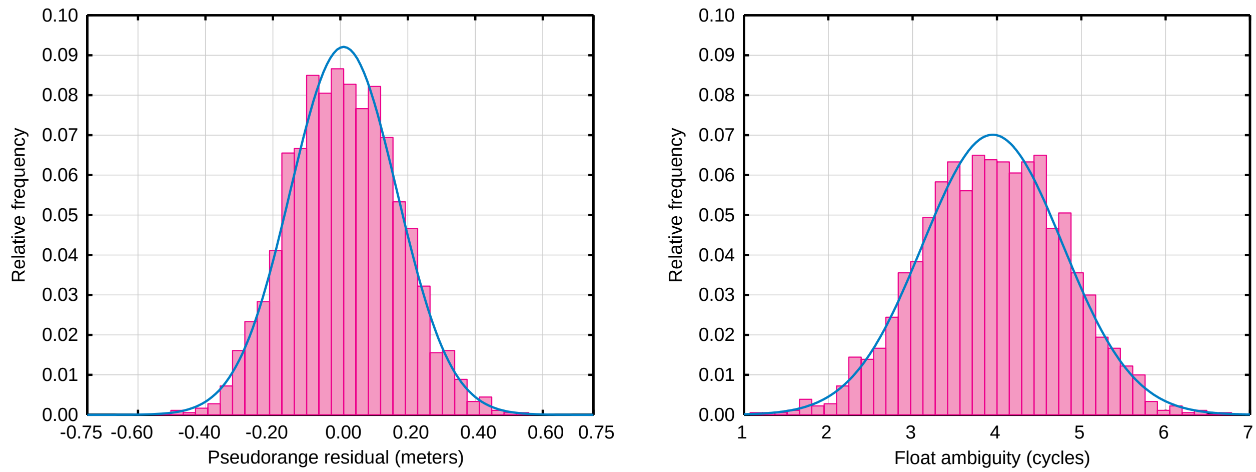

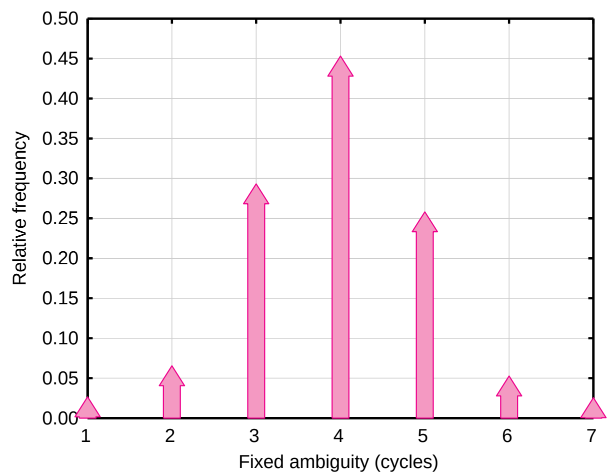

Figure 2. Using single-frequency pseudorange and carrier-phase data, the phase ambiguity of the geometry-free GPS model is estimated in 1,800 single-epoch experiments at a 1-second interval. The histogram at the top left shows the residuals of the (double-difference) pseudorange measurements; the noise is at the decimeter level. The histogram at the top right describes the float ambiguity. It is primarily the noise in the pseudorange that is reflected in the noise of the float ambiguity, and as the L1 wavelength is about 2 decimeters, the corresponding uncertainty in the float ambiguity is at the one-cycle level. The top graphs also both show the formal Gaussian probability distribution. Finally, the integer ambiguity was computed for each experiment, yielding the histogram at the lower left. In this case, the integer ambiguity is estimated correctly (four cycles) in only 43 percent of the experiments (indicated by base of arrowhead).

uncertainty of the estimated parameters. This uncertainty is then described by the probability distribution of these parameters. The real difference between a standard and an integer adjustment lies in the type of probability distribution. In the standard case, the distribution will be continuous, whereas in the integer case it will be discrete as shown in the lower left plot of Figure 2. That is, the distribution of the estimated integer ambiguities will be a probability mass function. Such a distribution is also obtained from other discrete phenomena such as throwing a pair of dice.

Without any knowledge of the probability mass function of the integer ambiguities, we have no way of knowing how often to expect the computed ambiguity solution to coincide with the correct but unknown integers. Is this nine times out of ten, 99 times out of 100, or an even higher percentage? In the example shown in Figure 2, it is actually less than 45 percent. This implies that when carrying out an experiment according to the assumption made in the example, one has about a 55-percent chance of computing a wrong integer ambiguity.

AMBIGUITY SUCCESS RATE

If we treat the computed integer ambiguities as deterministic quantities, as we usually do in practice, we will have to ensure that their uncertainty is sufficiently small to be indeed neglected. This is the case when the frequency with which estimated integer ambiguity values coincide with the correct but unknown values is sufficiently large. This concept is formalized in a probabilistic measure, referred to as the *ambiguity success rate*. The success rate is a number between 0 and 1, or 0 and 100 percent, and it expresses the chance, or probability, that the integer ambiguities are correctly estimated.

The ambiguity success rate depends on three contributing factors: the observation equations (functional model), the precision of the observables (the stochastic model), and the chosen method of integer estimation. Changes in any one of these will affect the success rate. The first two contributing factors reflect the data model's strength and they are given once the measurement set-up is known. As to the method of integer estimation, one has a variety of options available. However, because different methods of integer estimation will generally result in different success rates, we might wish to use the method that maximizes the success rate. It has recently been proven that the integer least-squares estimator has the largest success rate of all admissible integer estimators. The success rate of the LAMBDA method is therefore larger than, or at least as large as, any other integer ambiguity estimator.

A two-dimensional example will show us how to determine the success rate of the integer least-squares ambiguities. Based on the measurement precision and the assumed relationship between observations and unknown parameters, we can obtain a probabilistic description of the uncertainty in the float ambiguities. In this example, the uncertainty is given

FURTHER READING

For a thorough discussion of the carrier-phase ambiguity, see

■ "GPS Carrier Phase Ambiguity Fixing Concepts" by P.J.G. Teunissen, Chapter 8 in *GPS for Geodesy*, 2nd edition, edited by P.J.G. Teunissen and A. Kleusberg, Springer-Verlag, Berlin, 1998.

For an introduction to the LAMBDA ambiguity-fixing approach, see

■ "A New Way to Fix Carrier-Phase Ambiguities" by P.J.G. Teunissen, P.J. de Jonge, and C.C.J.M. Tiberius in *GPS World*, Vol. 6, No. 4, April 1995, pp. 58–61.

More information about the LAMBDA approach can be found on the Department of Mathematical Geodesy and Positioning, Delft University of Technology, Website:

■ <http://www.geo.tudelft.nl/mgp/>.

For further details about the probabilistic theory of ambiguity fixing and its consequences, see

■ "The Probability Distribution of the GPS Baseline for a Class of Integer Ambiguity Estimators" by P.J.G. Teunissen, in the *Journal of Geodesy*, Vol. 73, 1999, pp. 275–284.

■ "An Optimality Property of the Integer Least-squares Estimator" by P.J.G. Teunissen, in the *Journal of Geodesy*, Vol. 73, 1999, pp. 587–593.

For a discussion of the ambiguity success rate, see

■ "A Probabilistic Evaluation of Correct GPS Ambiguity Resolution" by P.J.G. Teunissen, D. Odijk, and P. Joosten, in the *Proceedings of ION GPS-98*, the 11th International Technical Meeting of the Satellite Division of The Institute of Navigation, Nashville, Tennessee, September 15–18, 1999, pp. 1315–1323.

To learn more about random variables and probability distributions, consult one of the classic books on statistics, such as

■ *Introduction to Mathematical Statistics*, 5th edition, by R.V. Hogg and A.T. Craig, Prentice Hall, Englewood Cliffs, New Jersey, 1995.

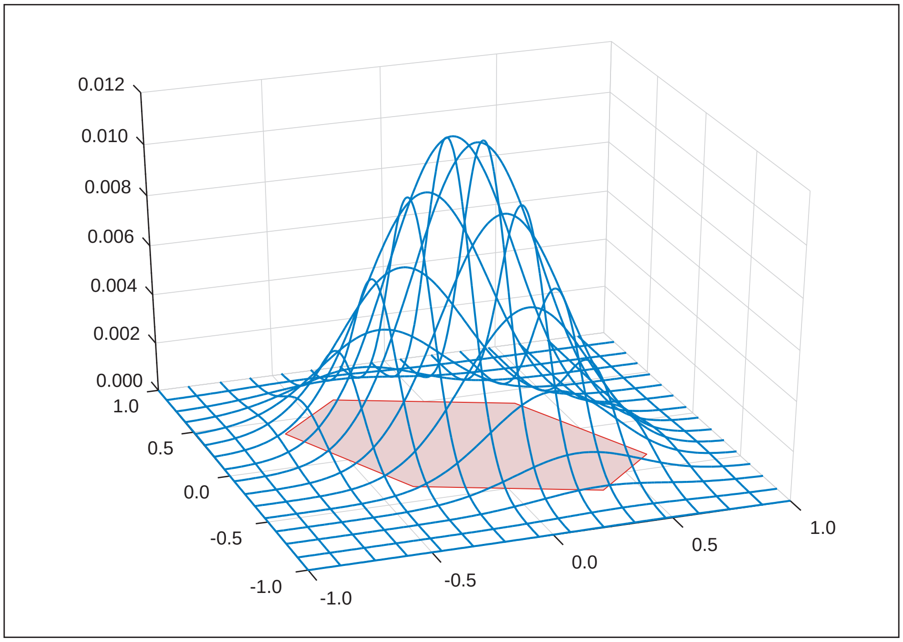

Figure 3. In this example of a joint probability density function for two float ambiguities, the ambiguity values (in cycles) are along the horizontal axes and the probability density along the vertical. The red area on the bottom of the graph indicates the pull-in region, which is the float ambiguities area that is mapped onto the correct integer ambiguity vector, in this case (0,0), using integer least-squares estimation. By taking the integral of the probability density function over the pull-in region, the success rate is obtained. It is the probability of correct integer estimation and about 85 percent in this example.

by the Gaussian two-dimensional or bivariate probability density function as shown in Figure 3. The standard deviations of the two ambiguities are about 0.3 cycle. The corresponding success rate follows, then, as the integral of the probability density function over the area shown in red. This area is referred to as the *ambiguity pull-in region*. It contains all locations of the float ambiguities which get pulled to the correct integer solution when using the integer least-squares principle. If we denote the probability density function of the float ambiguities as $p_a(\mathbf{x})$, and the pull-in region of the correct integer ambiguity vector as R_a, the ambiguity success rate can be written in equation form as

$$\text{Success rate} = \int_{R_a} p_a(\mathbf{x})d\mathbf{x}.$$

Some easy ways of computing or approximating this multiple integral are discussed in the sidebar entitled "How to Compute Ambiguity Success Rate."

The ambiguity success rate can be evaluated once the GNSS functional and stochastic models are known. Similar to the usage of dilution of precision (DOP) measures, it can be computed without having the actual measurements available, that is, before actual field operations. By means of the success rate, the user is given a rigorous way of assessing how often he or she can expect ambiguity resolution to be successful. Only when the success rate is close enough to 1 is one allowed to proceed as if the estimated integer ambiguities are nonstochastic. How close to 1 does the success rate need to be? This depends on the particular GNSS application and the potential impact of incorrectly fixing the ambiguities on the parameters being estimated. A smaller success rate can be accepted in the instances where the effect is small.

The success rate depends of course, as any other formal reliability measure does, on the correctness of the assumptions that underlie the model used. Incorrect specifications in the model may lead to unrealistic values for the success rate. For instance, even with a high enough success rate, fixing to the wrong integer ambiguities is still possible when one or more observations are grossly erroneous — so-called *outliers*. A success rate close enough to 1 therefore does not release us from the obligation of performing statistical tests for model validation. It does however make it much easier to perform such tests. The higher the success rate, the sooner one is allowed to apply the classical theory of statistical hypothesis testing and use, for instance, the common F-test to spot any anomaly in the observations.

CONCLUSION

In this article, we have shown that although double-difference GNSS carrier-phase ambiguities are known to be integers, they are still

How to Compute Ambiguity Success Rate

The ambiguity success rate is defined as the probability of correct integer ambiguity estimation, $P(\check{\mathbf{a}}=\mathbf{a})$. It equals the integral of the probability density function of the float ambiguities $p_{\hat{a}}(\mathbf{x})$, over the pull-in region R_a. The pull-in region R_a equals the region in which all float solutions are pulled by the integer least-squares criterion to the correct integer ambiguity solution. The success rate is given as:

$$P(\check{\mathbf{a}}=\mathbf{a}) = \int_{R_a} p_{\hat{a}}(\mathbf{x})d\mathbf{x}.$$

In the absence of any biases in the underlying observations, the success rate corresponds to the central and largest probability mass of the ambiguity probability mass function. The figure below shows, for a two-dimensional example, the probability density function of the float ambiguities on the left and the corresponding discrete distribution of the integer least-squares ambiguities on the right.

Various ways of computing or approximating the success rate exist, two of which will be given here. One way of obtaining the success rate is by simulation. Using a random number generator, we can obtain a large number of real-valued ambiguity vectors from the origin-centered probability distribution $p_{\hat{a}}(\mathbf{x})$ of the float solution. For each of these generated vectors, we then compute the corresponding integer least-squares solution using the LAMBDA method. The percentage of integer solutions that coincide with the origin yields the success rate. The number of generated samples must be large enough to obtain a close enough approximation to the success rate. For example, to achieve a success rate of 99.9 percent with a 0.1-percent uncertainty would require between 100,000 and 1,000,000 samples.

A second option for inferring the success rate is to compute a sharp lower bound of the probability of correct integer least-squares estimation. A sharp and easy-to-compute lower bound (LB) is given by:

$$LB = \prod_{i=1}^{n}\left[2\Phi\left(\frac{1}{2\sigma_{i|I}}\right) - 1\right] \le P(\check{\mathbf{a}} = \mathbf{a})$$

$$\text{with} \quad \Phi(x) = \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} dz$$

It equals a product Π of n terms (the number of ambiguities). Φ is the standard normal cumulative probability distribution and $\sigma_{i|I}$ is the standard deviation of ambiguity i, conditioned on all previous ambiguities, indicated by I. The conditional standard deviations follow directly from the triangular decomposition of the float ambiguity variance-covariance matrix $\mathbf{Q}_{\hat{a}} = \mathbf{L}^T\mathbf{DL}$ as the square root of the elements of diagonal matrix $\mathbf{D}$. This decomposition is already made in the computations for the LAMBDA method, and hence available at no extra computational cost.

For this lower bound to be sharp, it is essential that the variance–covariance matrix of the LAMBDA-transformed ambiguities be used to compute the conditional standard deviations, as they have an improved precision and decreased correlation over the original double-difference ambiguities.

This approximation of the success rate can be computed in a straightforward manner and, if it is sufficiently large, say 99 or 99.9 percent, it is guaranteed that the actual success rate of the integer least-squares method is at least equally high and thus very close to 100 percent. As it provides a lower bound, one can safely rely on this approximation.

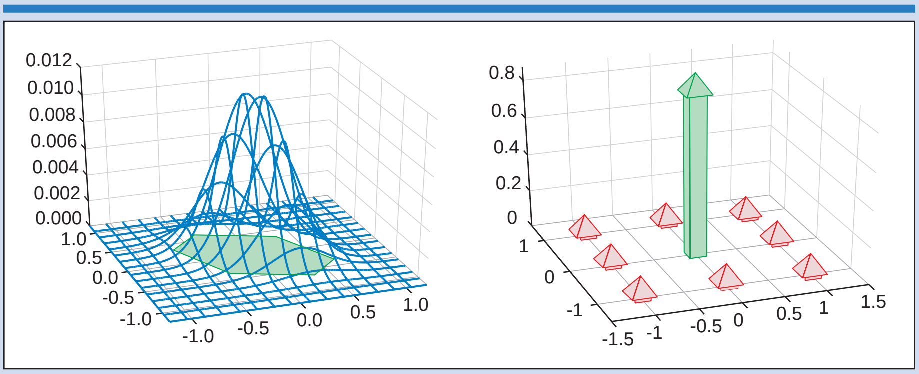

By taking the integral of the probability density function (on the left) over the pull-in region for each integer vector, the probability that this vector will result as the integer least-squares solution is obtained. The probabilities are given on the right for the integer vectors between –1 and +1. The integral over the area for the correct integer vector, in this case (0,0), gives the success rate. It is about 85 percent in this example.

stochastic variates. To safely neglect their uncertainties, one must be very sure that the integer ambiguities are indeed correctly estimated. This probability is expressed by the success rate.

Like any reliability measure, the success rate should be used to establish the probability of correct fixing. Since the success rate can be computed prior to actual measurements being made, we can ensure that the measurements are collected in such a way that successful ambiguity resolution will be feasible. In the data processing stage, the success rate is essential as well, since only when this number is sufficiently large, will it be safe to neglect remaining stochasticity; we can proceed with the integer ambiguities as fixed quantities with certainty.

The success rate is only a single number, and there exists a valuable approximation that is easy (cheap) to compute. We therefore strongly advocate its evaluation by default during any processing of GNSS measurements in which ambiguity resolution is involved. ■

MANUFACTURERS

The data used in Figures 1 and 2 were obtained from **Trimble** (Sunnyvale, California) 4000 SSi receivers.

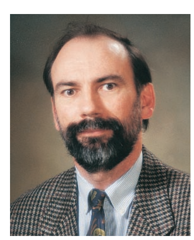

"Innovation" is a regular column featuring discussions about recent advances in GPS technology and its applications as well as the fundamentals of GPS positioning. The column is coordinated by Richard Langley of the Department of Geodesy and Geomatics Engineering at the University of New Brunswick, who appreciates receiving your comments as well as topic suggestions for future columns. To contact him, see the "Columnists" section on page 4 of this issue.