

Simultaneous Position Estimation & Ambiguity Resolution (SPEAR) for High-integrity Carrier-phase Navigation with Robustness to Tracking Loss

Jason H. Rife
Tufts University

>> Accepted Article <<

CITATION:

J. H. Rife, "Simultaneous position estimation & ambiguity resolution (SPEAR) for high-integrity carrier-phase navigation with robustness to tracking loss," 2014 *IEEE/ION Position, Location and Navigation Symposium - PLANS 2014*, Monterey, CA, 2014, pp. 692-700.
doi: 10.1109/PLANS.2014.6851434

CORRESPONDING AUTHOR:

Jason Rife
jason.rife@tufts.edu

COPYRIGHT:

© 2014 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

FINANCIAL SUPPORT:

NASA Safe and Efficient Surface Operations (Contract 11-1-A3.01-8393 ARC)

Simultaneous Position Estimation & Ambiguity Resolution (SPEAR) for High-Integrity Carrier-Phase Navigation with Robustness to Tracking Loss

Dr. Jason H. Rife

Dept. of Mechanical Engineering
Tufts University
Medford, MA
jason.rife@tufts.edu

Abstract—This paper introduces a GNSS processing algorithm called simultaneous position estimation and ambiguity resolution (SPEAR), with the goal of delivering high-accuracy, high-integrity navigation with robustness to carrier-tracking interruptions. The algorithm operates by continuously applying integer constraints to estimate the navigation solution recursively as a histogram of possible state values. The proposed approach is distinctive in that it decomposes the carrier-wave measurement into two components: an apparent integer difference, used for Bayesian prediction, and a fractional phase, used for Bayesian correction. The algorithm has potential application for high-accuracy, high-integrity safety of life operations near buildings or other obstructions of the sky, as may occur during aircraft surface movement or automobile navigation in urban canyons.

Keywords- Bayesian Estimation, Integer Ambiguity Resolution, Integrity, Aircraft Surface Movement

I. INTRODUCTION

Attaining a high accuracy, high integrity GNSS solution in an environment with frequent satellite-tracking interruptions remains a significant challenge. For certain applications such as aircraft navigation on the airport surface, all of these issues must be resolved simultaneously. Ideally, future GNSS capabilities would enable operations in the ramp area of an airport (also called the apron), even in zero visibility conditions. Ramp-area operations require high accuracy, since aircraft must navigate near buildings and close to the gate. High integrity is also essential, in the sense that the probability of large errors must be strictly bounded (e.g., less than one time in ten million operations) in order to avoid collisions that might damage aircraft or harm personnel. The ability to tolerate brief tracking outages is also important, since tall buildings and other aircraft may occlude portions of the sky, blocking out signals from one or more GNSS satellites [1]. Similar requirements also apply to automated driving in urban canyons [2].

Existing algorithms typically handle only two of the above three requirements. For example, most high-precision GNSS algorithms based on carrier-phase processing provide

centimeter-level accuracy through a process called integer ambiguity resolution [3]-[5] with no concern for integrity. The ambiguity resolution process (also called *integer fixing*) achieves accuracy by estimating the integer number of carrier wavelengths in the path length (or differential path length) to each satellite at a given instant in time. In many applications, integer ambiguity resolution can be performed using a short data set or even a snapshot algorithm that processes data at a single epoch [6]-[8]. Thus, continuous tracking is not a necessity. These algorithms do not make guarantees for the integrity risk associated with an incorrect integer fix, however. Hence these algorithms are not intended for high-integrity, safety-of-life applications.

By contrast, carrier-phase algorithms that provide high-accuracy with high-integrity often require long duration filtering before ambiguity resolution can occur [9]-[11]. For safety-critical applications like airborne refueling or aircraft landing on an aircraft carrier, five to fifteen minutes of continuous prefiltering may be required. These algorithms are thus vulnerable to signal interruptions caused, for instance, when a receiver loses line-of-site to one or more satellites. Although tracking interruptions are rare in open environments (e.g. in many farming applications), tracking loss presents a greater risk for safety-critical operations near large buildings or vehicles that can block the sky.

This paper investigates a class of navigation algorithms that can obtain high accuracy and high integrity without a requirement for continuous carrier-phase tracking. The algorithm is dubbed Simultaneous Position Estimation and Ambiguity Resolution (SPEAR). The term SPEAR implies an integration of ambiguity resolution directly into the navigation solution. SPEAR never truly fixes integers, but rather uses integer equations as constraints on allowable values of the position solution. Accordingly, SPEAR does not require continuous carrier-phase tracking.

The theoretical basis for unifying ambiguity resolution and position estimation is the following. If true position and clock bias are known for a receiver, then the number of integer

carrier wavelengths along the path between the receiver and satellite should be known. In other words, the integer vector \mathbf{N} at a particular time is an algebraic function of the true receiver state vector \mathbf{x} :

$$\mathbf{N} = f(\mathbf{x}). \quad (1)$$

Hence, at least in theory, it should not be necessary to “resolve” integer ambiguities, as (1) implies that integers are a deterministic function of the receiver state estimate $\hat{\mathbf{x}}$.

This paper implements SPEAR using a Bayesian estimation framework to maintain a distribution of possible state values $\hat{\mathbf{x}}$, where each state estimate is algebraically associated with an integer. A critical characteristic of the proposed algorithm is that it decomposes the carrier-phase measurement into two components, an *apparent integer difference* component used to constrain Bayesian state propagation and a *fractional phase* component used in the Bayesian measurement update.

The remainder of this paper describes the new algorithm in detail. First, carrier-wave measurement models are provided, both for conventional accumulated Doppler measurements and for the derived measurements by SPEAR (namely the apparent integer difference and the fractional phase). As background, the next section briefly describes a least-squares-based floating-point integer estimation algorithm, a process which is often employed as a pre-filter for high-integrity integer ambiguity resolution. Subsequently, the Bayesian state-estimation algorithm that is the main focus of this paper is presented. To assess performance, an analysis section discusses simulations that apply the conventional least-squares and novel Bayesian algorithms to synthetic data. In particular, this simulation-based analysis considers aircraft surface movement, the application which inspired the SPEAR concept. A brief summary concludes the paper.

II. CARRIER-WAVE MEASUREMENT MODELS

Most current day GNSS receivers track carrier phase using a phase-locked loop (PLL). PLL tracking precision is on the order of 10° one-sigma, which corresponds to a tracking error on the order of 1 cm for the 19 cm wavelength of the L1 frequency [12]. This precision is nearly two orders of magnitude better than that for conventional GPS code tracking, which delivers closer to 1 m accuracy (one sigma, for differential GNSS).

The caveat for carrier-phase processing is that only the *phase* along the wavelength is truly measured; the *integer* number of wavelengths along the line-of-site between the receiver and any GNSS satellite is completely ambiguous. In this sense, the carrier-phase measurement is precise but not accurate. Resolving the integer ambiguity is possible but time consuming, because ambiguity resolution requires measurement filtering over many epochs. To support high-accuracy, high-integrity landing on an aircraft carrier or airborne refueling, for example, five or more minutes of filtering may be necessary [9]-[11].

This section discusses models for measurements obtained from carrier-phase tracking. Collectively, these measurements

are labeled *carrier-wave* measurements. The most widely studied carrier-wave measurement is obtained by integrating carrier phase changes over time. This integral is sometimes called the *accumulated Doppler* measurement.

In this paper, the variable ϕ will be used to identify the vector of accumulated Doppler measurements for all tracked satellites. The accumulated Doppler ϕ is a function of the receiver clock bias b , the vector range to each satellite \mathbf{r} , the integer ambiguity vector for all satellites \mathbf{N} , and the measurement error vector $\boldsymbol{\varepsilon}_\phi$ [12]. Here the error vector $\boldsymbol{\varepsilon}_\phi$ combines random receiver noise and multipath with systematic errors, such as troposphere and ionosphere effects. By convention, the receiver clock bias b is defined in units of meters (where time is converted to length through multiplication with the speed of light).

$$\phi = \mathbf{r}(\mathbf{p}) + b + \lambda \mathbf{N}(\mathbf{p}) + \boldsymbol{\varepsilon}_\phi \quad (2)$$

The unknown receiver position \mathbf{p} is expressly included in the above equation to emphasize an important fact: that both the range \mathbf{r} and the integer ambiguity \mathbf{N} are dependent on the unknown receiver position \mathbf{p} (at least at an initial time step). Although the dependence of the integer ambiguities \mathbf{N} on receiver position \mathbf{p} is typically neglected to simplify the position solution algorithm, this paper will explicitly consider the constraints relating the range and integer vectors, \mathbf{r} and \mathbf{N} .

In this paper, the accumulated Doppler signal is decomposed to obtain two related carrier-wave measurements: *fractional phase*, denoted by the vector $\boldsymbol{\psi}$, and the *apparent integer*, denoted by the vector \mathbf{L} . These signals are related to the accumulated Doppler signal at any epoch k by the following equation:

$$\phi_k = \boldsymbol{\psi}_k + \lambda \mathbf{L}_k. \quad (3)$$

The fractional phase $\boldsymbol{\psi}$ is the instantaneous phase of the arriving carrier signal. This phase might be defined as an angle (between 0 and 2π radians); for convenience, however, this paper defines the fractional phase with units of length (between 0 and the carrier wavelength λ). A model for the fractional phase can be obtained by applying the modulus function to calculate the remainder of the accumulated Doppler measurement after dividing by wavelength.

$$\begin{aligned} \boldsymbol{\psi}_k &= \text{mod}(\phi_k, \lambda) \\ &= \text{mod}(\mathbf{r}(\mathbf{p}_k) + b_k + \boldsymbol{\varepsilon}_{\phi,k}, \lambda) \end{aligned} \quad (4)$$

By projecting onto the range $[0, \lambda)$, the above modulus function removes the integer, such that the fractional phase is inherently unambiguous.

The apparent integer is the complement to fractional phase. More specifically, the apparent integer is the number of phase boundary transitions that have occurred since carrier tracking began for a particular satellite.

$$\mathbf{L}_k = \text{floor} \left(\frac{\phi_k}{\lambda} \right) \quad (5)$$

By definition, \mathbf{L}_k must always be zero at the start of tracking. Later, a negative apparent integer implies the zero boundary has been passed at least once; a positive apparent integer implies the one-wavelength boundary (i.e. the 2π radian boundary) has been passed at least once. In the absence of cycle slips, the apparent integer is related to the receiver states by the following equivalent model:

$$\mathbf{L}_k = \text{floor} \left(\frac{\mathbf{r}(\mathbf{p}_k) + b_k + \boldsymbol{\varepsilon}_{\phi,k}}{\lambda} \right) - \mathbf{N}_0. \quad (6)$$

Here \mathbf{N}_0 is a vector of the (unknown) initial integers for each satellite, at the beginning of tracking. Unlike the initial integer vector \mathbf{N}_0 or the prompt integer vector \mathbf{N} , which are deterministic, the vector \mathbf{L} consists of random integers that depends on the noise term $\boldsymbol{\varepsilon}_{\phi,k}$. Even a small random noise contribution can shift the value of the integer \mathbf{L} if the argument of the *floor* function is near an integer boundary.

The initial integer ambiguity \mathbf{N}_0 can be eliminated by defining a related signal, the *apparent integer difference*, which is a difference between \mathbf{L} values at different epochs (as measured by the same receiver for the same satellites). The following equation defines the apparent integer difference \mathbf{M} .

$$\mathbf{M}_k = \mathbf{L}_k - \mathbf{L}_{k-1} \quad (7)$$

The apparent integer difference \mathbf{M} can be related to the receiver states by substituting (6) into the above definition. The result is unambiguous in the sense that the unknown initial integer vector \mathbf{N}_0 cancels:

$$\begin{aligned} \mathbf{M}_k = & \text{floor} \left(\frac{\mathbf{r}(\mathbf{p}_k) + b_k + \boldsymbol{\varepsilon}_{\phi,k}}{\lambda} \right) \\ & - \text{floor} \left(\frac{\mathbf{r}(\mathbf{p}_{k-1}) + b_{k-1} + \boldsymbol{\varepsilon}_{\phi,k-1}}{\lambda} \right). \end{aligned} \quad (8)$$

Instead of processing accumulated Doppler, SPEAR instead processes two derived carrier-wave measurements: the fractional phase and the apparent integer difference. The fact that both of these derived carrier-wave measurements are unambiguous simplifies algorithm implementation.

A final detail relevant to modeling carrier-wave measurements is that SPEAR assumes differential corrections (relative to a fixed base station) are used to remove systematic biases from the fractional phase and apparent integer difference measurements. For clarity, the notation δ will be introduced to indicate a differentially corrected quantity. For instance, differentially corrected accumulated Doppler measurements are labeled $\delta\mathbf{p}$.

In the case of differential corrections over reasonably short baselines (order tens of kilometers), it is reasonable to assume the differential range between the reference station and the user receiver $\delta\mathbf{r}$ can be linearized in terms of a geometry matrix \mathbf{G} [12].

$$\delta\mathbf{r} = \mathbf{G}\delta\mathbf{x} \quad (9)$$

Here the differential state vector concatenates the relative user-to-reference-station position vector and clock bias:

$\delta\mathbf{x} = [\delta\mathbf{p}^T \quad \delta b]^T$. Applying the above linearization to (4), the differential fractional phase $\delta\boldsymbol{\psi}$ can be related directly to the relative state vector $\delta\mathbf{x}$.

$$\delta\boldsymbol{\psi}_k = \text{mod} \left(\mathbf{G}_k \delta\mathbf{x}_k + \boldsymbol{\varepsilon}_{\delta\phi,k}, \lambda \right) \quad (10)$$

Similarly, the differential apparent integer difference can be modeled as

$$\delta\mathbf{M}_k = \delta\mathbf{L}_k - \delta\mathbf{L}_{k-1}, \quad (11)$$

where the apparent integer for the differentially corrected measurements is

$$\delta\mathbf{L}_k = \text{floor} \left(\frac{\mathbf{G}_k \delta\mathbf{x}_k + \boldsymbol{\varepsilon}_{\delta\phi,k}}{\lambda} \right). \quad (12)$$

Here the differentially-corrected error vector $\boldsymbol{\varepsilon}_{\delta\phi}$ is modeled as zero mean, reflecting the removal of ionospheric, tropospheric and other systematic biases.

III. THE FLOAT SOLUTION

To provide a point of comparison for SPEAR, it is useful to first review the so-called float solution, which is a conventional least-squares-based process for estimating a real valued approximation of the integer vector \mathbf{N} . The designation *float solution* emphasizes that the estimates are not truly integers (but rather are floating point values used by digital computers to represent real numbers). Before integer-fixing possible, typical high-integrity ambiguity resolution algorithms require that the float solution converge within less than a wavelength.

The float solution operates by fusing together differential code and accumulated Doppler measurements, labeled $\delta\mathbf{p}$ and $\delta\phi$, to obtain a differential receiver-state vector $\delta\mathbf{x}_k$ and the relative integer vector $\delta\mathbf{N}$. Thus, in some sense, the float solution might be considered to be primarily a position estimation algorithm, which just happens to produce real-valued integer estimates as a byproduct. These integer estimates are also useful, however, in pre-filtering prior to integer fixing. Eventually, if the linear estimator combines enough data points, the real-valued integer estimates become sufficiently accurate that integer fixing is possible with a high degree of confidence.

Though the float solution is typically implemented as a recursive least-squares (or a Kalman Filtering) algorithm, for illustration purposes it is useful to model the float solution as an equivalent batch solution, one which processes all measurements simultaneously.

$$\begin{bmatrix} \delta \mathbf{p}_0 \\ \delta \phi_0 \\ \vdots \\ \delta \mathbf{p}_k \\ \delta \phi_k \end{bmatrix} = \begin{bmatrix} \mathbf{G}_0 & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{G}_0 & \cdots & \mathbf{0} & \lambda \mathbf{I} \\ \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \cdots & \mathbf{G}_k & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{G}_k & \lambda \mathbf{I} \end{bmatrix} \begin{bmatrix} \delta \mathbf{x}_0 \\ \vdots \\ \delta \mathbf{x}_k \\ \delta \mathbf{N}_k \end{bmatrix} + \begin{bmatrix} \boldsymbol{\varepsilon}_{\delta p,0} \\ \boldsymbol{\varepsilon}_{\delta \phi,0} \\ \vdots \\ \boldsymbol{\varepsilon}_{\delta p,k} \\ \boldsymbol{\varepsilon}_{\delta \phi,k} \end{bmatrix} \quad (13)$$

So long as continuous tracking is maintained, the differential integer vector is a constant. Hence, the estimator can effectively difference the code and carrier measurements at each epoch and thereby obtain a new integer observable, somewhat corrupted by carrier and code measurement noise vectors, $\boldsymbol{\varepsilon}_{\delta \phi}$ and $\boldsymbol{\varepsilon}_{\delta p}$. The key to obtaining this observable is that the differential code measurement $\delta \mathbf{p}$, modeled below, is unambiguous.

$$\delta \mathbf{p}_k = \mathbf{G}_k \delta \mathbf{x}_k + \boldsymbol{\varepsilon}_{\delta p,k}. \quad (14)$$

The code measurement errors of $\boldsymbol{\varepsilon}_{\delta p}$ are large (meter-level). However, by averaging over time, the algorithm attenuates measurement noise and provides an integer estimate that is increasingly accurate.

IV. BAYESIAN FUSION OF CODE AND CARRIER MEASUREMENTS

In this section, a novel approach is proposed for integrating GNSS carrier and code measurements. The proposed algorithm, dubbed SPEAR, is a Bayesian estimator [13], a form of estimation which explicitly models the Probability Density Function (PDF) for the state estimate. Bayesian estimators consist of two steps, which are implemented iteratively: a state propagation (or prediction) step and a measurement update (or correction) step. The state propagation maps the past state-estimate PDF to the current epoch. The measurement update then weights regions of this PDF based on how well they match sensor data. By applying such predictions and corrections iteratively through time, the PDF can be continually updated based on newly acquired measurement data.

The central concept of SPEAR is that the apparent integer difference measurement imposes constraints that can be embedded in Bayesian state propagation. Complementary measurements (code pseudoranges and carrier-wave fractional phase) are used in the Bayesian measurement update. Details are described below, following a brief discussion of nonparametric PDF modeling.

A. Histogram PDF Model

A characteristic of Bayesian estimators is that they are capable of modeling state-estimate distributions in a general manner. As such, state variables need not be assumed

Gaussian distributed, as is typically assumed by least-squares estimators. This characteristic is important for SPEAR, as modeling integer constraints results in a highly non-Gaussian error distribution.

In the Bayesian estimation community, two of the most widely studied non-parametric models used to represent general PDFs are the histogram and particle models. The histogram model uses a well-defined, regular grid to represent the region of interest of the state-estimate PDF. By contrast, the particle approach uses an irregular, randomly selected cluster of points (i.e. a Monte Carlo simulation) to model the state-estimate PDF.

Of the two approaches, the particle approach is much more widely used, as the particle approach greatly reduces computational overhead by efficiently redistributing particles toward the densest regions of the PDF (i.e. the peaks). The histogram approach has started to become practical in some recent high-precision estimation applications [14], however, in part because of the increasingly low cost of parallel computation.

For high-precision navigation applications, particle models have significant liabilities [15]. One of the most significant liabilities of particle models of state error is that particle locations are determined stochastically (based on a random number generator), such that the results of processing are not repeatable, even when particle models are applied twice to the same measurement set. Another liability of particle models is that they are ill-suited for high integrity navigation applications where capturing low-density (i.e. far tail) regions is essential, such as in safety critical aviation applications. This liability is related to the very characteristics that make particle filters computationally efficient (e.g. the fact that intermittent resampling drives particles away from low-density regions toward high-density regions of the PDF).

Given these trade-offs, a histogram model is used to implement the proposed Bayesian algorithm. This histogram model consists of probabilities distributed on a regular grid across the four-dimensional receiver state space (three positions and time). The resulting number of grid cells, and the corresponding computational cost of the algorithm, is relatively high. In the analysis described later in this paper, for example, a grid of 1.6 billion nodes was used.

B. Measurement Update Step

In a Bayesian estimator, the measurement update (or *correction*) step obtains a new (or *posterior*) model of the state PDF by weighting an old (or *prior*) PDF model based on currently available sensor data. In the initial time step, the prior model might simply be a uniform distribution over the state space (as initially no information exists to localize the user receiver). In subsequent time steps, the state PDF becomes better and better resolved, allowing for the receiver state to be estimated with progressively better precision.

A generic equation for the measurement update is obtained from Bayes Rule. Here $\delta \mathbf{y}$ is the vector of differentially corrected sensor measurements and $\delta \hat{\mathbf{x}}$ is the relative-receiver state estimate. Unlike the sensor vector, the state estimate is

not a deterministic value, but rather a distribution of possible values as defined by the PDF below.

$$\underbrace{p(\delta\hat{\mathbf{x}}_k | \delta\mathbf{y}_k)}_{\text{posterior}} = C \underbrace{p(\delta\mathbf{y}_k | \delta\hat{\mathbf{x}}_k)}_{\text{sensor error model}} \underbrace{p(\delta\hat{\mathbf{x}}_k)}_{\text{prior}} \quad (15)$$

This equation weights the prior distribution by a sensor error model, which takes into account the values of the sensor measurements \mathbf{y}_k at the current epoch k . The spatially uniform scaling factor C is selected to ensure that the resulting posterior $p(\hat{\mathbf{x}}_k | \mathbf{y}_k)$ integrates to unity probability, as required for a PDF.

The measurement update equation, equation (15), relies on the existence of a reasonable model for the sensor measurement error distribution, $p(\mathbf{y}_k | \hat{\mathbf{x}}_k)$. In this work, we assume Gaussian models for sensor measurement error:

$$p(\delta\mathbf{y}_k | \delta\hat{\mathbf{x}}_k) = \frac{1}{|2\pi\mathbf{R}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}\hat{\mathbf{e}}^T \mathbf{R}^{-1} \hat{\mathbf{e}}\right), \quad (16)$$

where \mathbf{R} is the measurement-error covariance and where $\hat{\mathbf{e}}$ is the error estimate, which is the difference between the actual measurement \mathbf{y}_k and to the expected measurement $\mathbf{h}(\hat{\mathbf{x}}_k)$ for each value of the state estimate $\hat{\mathbf{x}}_k$.

$$\hat{\mathbf{e}}_k = \delta\mathbf{y}_k - \mathbf{h}(\delta\hat{\mathbf{x}}_k) \quad (17)$$

The error estimate vector $\hat{\mathbf{e}}$ includes errors for all code-phase and carrier-wave fractional phase measurements.

$$\hat{\mathbf{e}}_k = \begin{bmatrix} \hat{\mathbf{e}}_{\delta\rho,k} \\ \hat{\mathbf{e}}_{\delta\psi,k} \end{bmatrix} \quad (18)$$

The code-phase error estimate is computed for each grid cell in the state-estimate histogram by rearranging (14).

$$\hat{\mathbf{e}}_{\delta\rho,k} = \delta\boldsymbol{\rho}_k - \mathbf{G}_k \delta\hat{\mathbf{x}}_k \quad (19)$$

Notionally, the carrier-phase error for each grid cell might be computed likewise by rearranging (10).

$$\hat{\mathbf{e}}_{\delta\psi,k} = \text{mod}(\delta\boldsymbol{\psi}_k - \mathbf{G}_k \delta\hat{\mathbf{x}}_k, \lambda) \quad (20)$$

This expression is inconvenient, however, as it results in a carrier-phase error distributed on the range $[0, \lambda)$, with probability densest near the edges of the domain and sparsest in the middle. To simplify processing and maintain a unimodal distribution, it is convenient to instead reformulate the error model as follows, such that the phase-tracking error lies on the range $[-\frac{1}{2}\lambda, \frac{1}{2}\lambda)$.

$$\hat{\mathbf{e}}_{\delta\psi,k} = \delta\boldsymbol{\psi}_k - \mathbf{G}_k \delta\hat{\mathbf{x}}_k - \lambda \text{round}\left(\frac{\delta\boldsymbol{\psi}_k - \mathbf{G}_k \delta\hat{\mathbf{x}}_k}{\lambda}\right) \quad (21)$$

The resulting error distribution is unimodal and truncated over a one-wavelength range.

C. State Propagation Step

The correction step, defined by (15), assumes that a prior distribution is available for the state estimate at the current epoch k . In general, this prior must be constructed by propagating forward the state estimate from the previous time step $k-1$. This process is called the state propagation (or *prediction*) step in Bayesian estimation.

State propagation is typically performed using a physics model (e.g. a kinematic motion model) or inertial measurements (e.g. accelerometer and gyroscope data) to relate GNSS data from one time step to another. Typically, sensing or modeling errors blur the PDF during propagation. This blurring effect can be computed by the following predictor equation, which applies generally to all Bayesian estimators.

$$\underbrace{p(\hat{\mathbf{x}}_k)}_{\text{prior}} = \iiint \underbrace{p(\hat{\mathbf{x}}_k | \hat{\mathbf{x}}_{k-1})}_{\text{propagation error}} \underbrace{p(\hat{\mathbf{x}}_{k-1})}_{\text{posterior}} dV \quad (22)$$

The variable dV indicates the volumetric element of integration in the state-estimate space.

In effect, the above equation represents a convolution, where the old posterior distribution (from epoch $k-1$) is blurred through convolution with a propagation-error kernel to obtain a new prior distribution (for epoch k). In the above equation, both the prior and the posterior are conditioned on all sensor data up to the previous time step (through \mathbf{y}_{k-1}). By convention, this conditionality is omitted from the above equation for compactness.

SPEAR uses a state-propagation step that is unusual for two reasons. First, the propagation step operates not on the state directly, but rather on a function of the state: the estimated apparent integer difference $\delta\hat{\mathbf{L}}(\hat{\mathbf{x}})$. Second, the propagation step is deterministic. The update is deterministic because the apparent integer difference $\delta\mathbf{M}_k$ precisely relates the estimated apparent integers between two different epochs. That is: $\delta\mathbf{M}_k = \delta\hat{\mathbf{L}}_k - \delta\hat{\mathbf{L}}_{k-1}$. (23)

The above relationship specifies a deterministic relationship between pairs of cells in the state-value grid, tying any give cell in the state-value grid at the current epoch k to a particular grid cell at the prior epoch $k-1$. To see, this consider the structure of (8). The left-hand side of the equation is a measurement, and is hence known. The right hand side is a deterministic function of the states ($\hat{\mathbf{x}} = [\hat{\mathbf{p}}^T \ \hat{b}]^T$) specified for each grid cell. Even the estimate of measurement error $\hat{\mathbf{e}}_{\phi,k}$ is deterministic for each grid cell, and can be computed by comparing the grid-cell's state vector to the fractional phase measurement using (21). Because both sides of (8) are fully

specified for any pair of grid cells, the equation is, in effect, a constraint relating each grid cell at the prior epoch $k-1$ to a grid cell at the current epoch k . In other words, propagation is deterministic, so the state propagation model (22) can be reformulated as follows.

$$p_-(\delta\hat{\mathbf{L}}(\delta\hat{\mathbf{x}}_k)) = p_+(\delta\hat{\mathbf{L}}(\delta\hat{\mathbf{x}}_k) - \delta\mathbf{M}_k). \quad (24)$$

Note this analysis does not consider cycle slips, which would make the value of \mathbf{M} probabilistic. Instead, we will simply assume that measurements are monitored and excluded if a cycle slip appears to have been likely. Given that continuous carrier tracking is not required, such an exclusion should have minimal impact on algorithm performance.

The most challenging aspect of implementing (24) is that it is not a one-to-one mapping in state-space. Rather, multiple points in the state-space grid may map to a single integer both at the prior epoch and at the current epoch. Thus, to use (24), it is first necessary to compute the probability associated with each integer combination $\delta\hat{\mathbf{L}}$ over all associated grid points at the prior epoch (time step $k-1$). For any given integer estimate $\delta\hat{\mathbf{L}}$, the associated probability is obtained by the following integral

$$p_+(\delta\hat{\mathbf{L}}_{k-1}) = \iiint_{\Omega(\delta\hat{\mathbf{L}}_{k-1})} p(\delta\hat{\mathbf{x}}_{k-1}) d\delta V \quad (25)$$

where

$$\Omega(\delta\hat{\mathbf{L}}_k) = \left\{ \delta\hat{\mathbf{x}}_k \left\lfloor \frac{(\mathbf{G}\delta\hat{\mathbf{x}}_k + \hat{\mathbf{e}}_{\phi,k})}{\lambda} \right\rfloor = \delta\hat{\mathbf{L}}_k \right\} \quad (26)$$

and where the error estimate $\hat{\mathbf{e}}_{\phi,k}$ is computed using (21).

After the integrated probability for each integer combination is mapped to the current time step via (24), that probability must be spread over all states (i.e. all grid cells) associated with that integer. There is no a priori preference for one grid cell over the other, so probability is distributed uniformly over the associated grid cells at the current epoch k . In other words:

$$p(\delta\hat{\mathbf{x}}_k) = \frac{p_+(\delta\hat{\mathbf{L}}(\delta\hat{\mathbf{x}}_k) - \delta\mathbf{M}_k)}{\iiint_{\Omega(\delta\hat{\mathbf{L}}(\delta\hat{\mathbf{x}}_k))} dV} \quad (27)$$

The denominator describes the volume of the state space associated with each integer, over which the probability form (25) is uniformly distributed.

In summary, the Bayesian prediction equations combine equations (24), (25) and (27). This mapping takes the prior state distribution from one epoch ($k-1$), maps it briefly into the integer domain and then back into the state domain at the next

epoch (k). The mapping into the integer domain at the prior time step is performed using (25). The probabilities associated with each integer are brought forward to the new epoch using (24). Then those probabilities are converted back into a state distribution (e.g. mapped to grid cells of the histogram) using (27).

The complete SPEAR algorithm can thus be implemented by sequentially combining the prediction – equations (24), (25) and (27) – with the measurement update step of (15). The key point is that the nonlinear integer constraint relating the integers and states is explicitly captured by the proposed approach.

V. SIMULATION-BASED COMPARISON

Simulations were used to analyze and compare the performance of SPEAR. Simulations configurations were tuned to the aircraft surface-movement application. Simulation results are compared to the float solution, to show that SPEAR provides high accuracy and high integrity before the float solution has converged sufficiently to allow high-integrity ambiguity resolution.

A. Simulation Setup

In the simulation, eight GPS satellites were simulated to be visible above the horizon. All these satellites were assumed to remain visible across the entire period of the simulation. For this baseline study, continuous carrier-phase tracking was assumed, so that code and carrier-phase measurements were available at all epochs. Simulated GPS measurements were subject to randomly generated Gaussian measurement errors. The differential code-phase measurement error was assumed to be 1.0 m (one sigma) for all satellites. The differential carrier-phase error was assumed to be 2 cm (one sigma).

In order to model aircraft surface movement, two additional pieces of navigation data were introduced. First, it was assumed that the user receiver moved parallel to a ground plane at a known height, subject only to small random vertical perturbations with a standard deviation of 1 cm. Second, a vision-based sensing system capable of detecting reliable land marks (i.e. road lane markers) was assumed to be collocated with the GPS receiver, as in [16],[17]. This system was assumed to provide cross-track measurements with an error of 0.20 m (one sigma). These additional measurements were introduced not only to model a representative application, but also to compress the volume of the histogram PDF model by a factor of approximately two in the horizontal direction and by nearly two orders of magnitude in the vertical direction. This is a significant benefit for the proposed SPEAR implementation, which otherwise would require an intractably large four-dimensional histogram of possible states.

For the purposes of computing Bayesian PDFs, SPEAR used a histogram model with 10^9 voxels. Voxels were spaced with uniform separation (2 cm) in all four state-estimate dimensions (three spatial dimensions and time). The bounding dimensions of the volume were set to ± 7 m for the assumed direction of travel (along-track coordinate) and for the clock bias (b coordinate). The bounding dimension in the lateral

direction (cross-track coordinate) was set to ± 3 m; that for the vertical direction (up coordinate), to only ± 10 cm.

SPEAR results were compared to the least-squares estimator defined by equation (13). All sensors were assumed available to both algorithms.

B. Simulation Results

As might be expected, the SPEAR algorithm required vastly more computational resources than the least-squares method. Data from each epoch was processed in a fraction of a second when computing the float solution with a conventional laptop; by comparison, each epoch was processed in approximately 8 minutes when computing the Bayesian solution using the same laptop.

In exchange for burdensome processing, SPEAR provides the accuracy benefits of partial integer-fixing and high integrity before the least-squares float estimate converges enough to make integer fixing practical. Fig. 1 provides a graphical comparison of the PDFs generated by both algorithms in processing data from four epochs. The epochs were assumed to be separated over a sufficiently long interval such that all measurement errors are effectively independent. PDFs for the float solution (left column) and Bayesian approach of SPEAR (right column) are illustrated for each of four epochs (corresponding to the four rows of plots). The figure shows PDFs projected into the two-dimensional horizontal plane (along-track and cross-track directions). The illustrated PDF is a marginal probability on the horizontal plane, with probabilities integrated across the remaining two state dimensions (up coordinate and clock bias). In the illustration, dark regions indicate high probability densities, and light regions indicate low densities. Grayscale is graded logarithmically, to enhance the visibility of the PDFs across a wide range of scales. In the simulation, GPS receiver location is allowed to change from one time step to the next; the true location of the GPS receiver in each plot is identified by a circle marker.

PDF appearance gives a qualitative indication of the differences between the float solution and SPEAR. The most salient difference between the two algorithms is that the SPEAR solution is speckled (multi-peaked), whereas the float solution estimates position as a single-peaked Gaussian distribution. Each peak in the Bayesian distribution corresponds to a different integer combination.

At the first time step, both algorithms exhibit PDFs with approximately elliptical contours of constant probability. Elliptical contours are expected for the float solution, which inherently models the state-estimate PDF as Gaussian distributed. Though the Bayesian PDF need not be Gaussian distributed in general, it has this initial appearance because the dominant error source at this step is the Gaussian-distributed code measurement error.

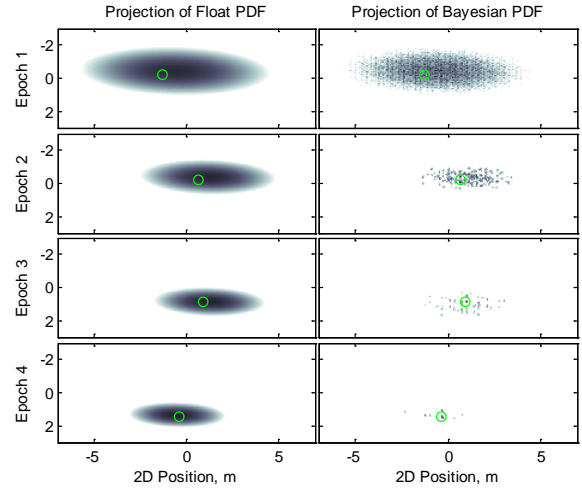


Fig. 1. Comparison of float solution (left) to Bayesian solution (right) for four sequential time steps.

The distinction between the float and Bayesian solutions becomes more apparent in the second and subsequent time steps. Bayesian state propagation occurs for the first time between the first and second time steps. This propagation is critical to contracting the Bayesian PDF, because it is the propagation step that enforces the integer constraint.

In effect, the propagation step performs a form of integer ambiguity resolution. This “effective” ambiguity resolution appears as a reduction in the number of likely integer combinations (dark peaks) visible in Bayesian PDF in the second and subsequent epochs. Only a handful of peaks are visible at the fourth time step. By the fourth time step, the highest density region of the PDF coincides with the true position (circle marker in Fig. 1). By comparison, the float solution is still several wavelengths wide, even in the narrowest (cross-track) error direction, indicating that high-integrity integer fixing using the float solution is not yet possible.

To quantify accuracy more precisely, it is useful to plot Cumulative Distribution Functions (CDFs) for the two approaches. CDF comparisons are presented in Fig. 2 and Fig. 3. In these figures, the CDF is evaluated by integrating the planar PDF from Fig. 1 outside an elliptical region of scalable area, for which the elliptical contour is defined by a Mahalanobis distance m .

$$CDF = \int_M^{\infty} p(m) dm \quad (28)$$

The Mahalanobis distance m is the distance of a point \mathbf{x} from a center point $\boldsymbol{\mu}$ weighted by the inverse of a covariance matrix \mathbf{P} .

$$m(\mathbf{x}) = \left((\mathbf{x} - \boldsymbol{\mu})^T \mathbf{P}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right)^{\frac{1}{2}} \quad (29)$$

For the float solution, the Mahalanobis distance is computed with \mathbf{x} being the two-dimensional horizontal state and with \mathbf{P} and $\boldsymbol{\mu}$ being the covariance and mean estimate of that solution at a particular epoch. (Each value of m thus corresponds to an elliptical contour of constant probability.) For the Bayesian solution, the Mahalanobis distance is evaluated using the same parameters at each epoch, except that the mean $\boldsymbol{\mu}$ is shifted to the location of the mean of the Bayesian histogram.

At the first epoch, the CDF curve is very similar for both methods, as shown in Fig. 2. The figure depicts both CDF curves using Mahalanobis coordinates. In these coordinates, all Gaussian distributions are identical, due to normalization by mean and covariance according to (29). The float solution produces a Gaussian CDF, which tracks this characteristic curve. The SPEAR solution (labeled *Bayes* in the figure) roughly tracks the same characteristic curve; however, the CDF is nonsmooth, with steps visible corresponding to local PDF peaks and valleys (associated with particular integer combinations).

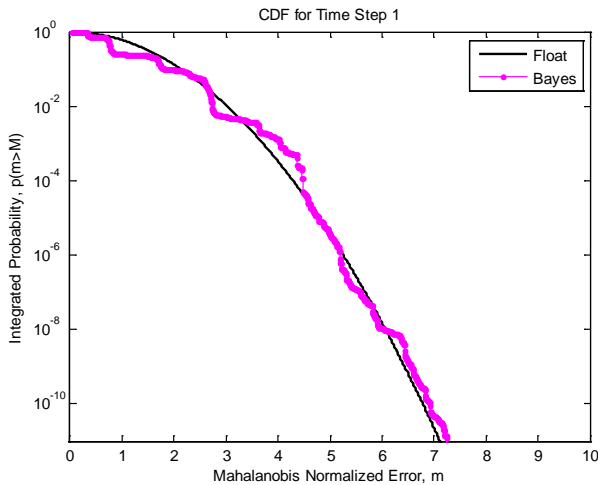


Fig. 2. Comparison of Float and Bayesian CDFs at First Epoch

The CDF curves for the float and SPEAR solutions diverge after the first epoch (see Fig. 3). The float solution consistently produces a Gaussian curve, and hence the Mahalanobis CDF remains in the same location, even though the PDF mean and covariance evolve over time. By contrast, the Bayesian solution increasingly shifts weight to particular integer combinations, such that the distribution becomes increasingly non-Gaussian.

A salient feature of Fig. 3 is that the SPEAR CDF asymptotes at a minimum probability for large values of the Mahalanobis-normalized error. This behavior is a result of the numerical implementation of the Bayesian algorithm, which clusters low-likelihood peaks together into an “unassigned” probability with the purpose of enhancing computational efficiency. Because the unassigned probability is not associated with a particular integer, it is excluded from the CDF (such that the CDF always provides a conservative bound on error).

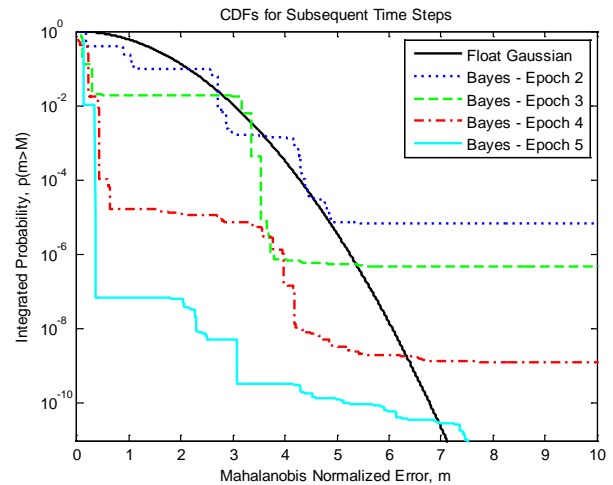


Fig. 3. Comparison of Float and Bayesian CDFs at Epochs 2 through 5.

VI. DISCUSSION

A significant result is that the accuracy of the SPEAR solution improves much more rapidly than that of the float solution. This phenomenon is evident in Fig. 3, which compares the Mahalanobis CDFs for epochs 2 through 5. For each successive epoch, the CDFs are normalized by the same (improving) covariance matrix. Despite the normalization, the SPEAR solution still compresses toward the left side of the figure, indicating improved accuracy (i.e., lower probabilities of large errors) relative to the float solution as well as high integrity (e.g. tight bound on rare errors). Consider the error bound that contains all except a small risk probability of 10^{-7} . At the level of 10^{-7} probability on the vertical axis of Fig. 3, the Bayesian solution has an error that is an order of magnitude smaller than the corresponding float solution error at the same risk probability. This result has important implications for the speed of integer ambiguity resolution for safety-of-life applications.

The SPEAR algorithm is robust to integer tracking loss because integers are never actually fixed; rather integers are embedded in processing as a constraint used to weight possible state values. The algorithm in no way depends on estimation of ambiguous integers. Thus, the state probability distribution obtained at each epoch by SPEAR remains valid as a prior for the next epochs, even when satellite signals are added or lost.

VII. SUMMARY

This article described a new algorithm that achieves high accuracy and high integrity with robustness to loss of carrier-phase tracking. Dubbed simultaneous position estimation and ambiguity resolution (SPEAR), the algorithm explicitly couples integer values to state estimates on a grid of possible state values. The algorithm achieves high accuracy by embedding integer constraints into the state estimation process (specifically in the state propagation step of Bayesian estimation). The algorithm achieves high integrity by evaluating all probabilities across a gridded histogram of possible state estimates. The algorithm achieves robustness to

loss of tracking by never explicitly estimating the initial integer offset at the beginning of tracking; rather, integers are only applied between time steps, so continuous carrier-phase tracking is not required. The major liability of the proposed SPEAR algorithm is its large computational overhead.

ACKNOWLEDGEMENTS

The author gratefully acknowledges NASA (Contract 11-1-A3.01-8393 ARC) for supporting this research. The opinions discussed here are those of the authors and do not necessarily represent those of NASA or other affiliated agencies.

REFERENCES

- [1] Guilloton, A., Arethens, J.-P., Escher, A.-C., Macabiau, C., Koenig, D., "Multipath Study on the Airport Surface," *Proceedings of IEEE/ION Position, Location, and Navigation Symposium (PLANS 2012)*, April, 2012.
- [2] Groves, P., Jiang, Z., Wang, L., Ziebart, M., "Intelligent urban positioning using multi-constellation GNSS with 3D mapping and NLOS signal detection," *Proceedings of ION-GNSS 2012*, pp. 458-472, 2012.
- [3] Teunissen, P., "The least-squares ambiguity decorrelation adjustment: a method for fast GPS integer ambiguity estimation," *Journal of Geodesy*, vol. 70, no. 1 2, pp. 65-82, 1995.
- [4] Teunissen, P., Odijk, D., and Joosten, P., "A probabilistic evaluation of correct GPS ambiguity resolution," *Proceedings of ION-GPS 1998*, pp. 1315-1323, 1998.
- [5] Henkel, P. and Günther, C., "Reliable integer ambiguity resolution: multi-frequency code carrier linear combinations and statistical a priori knowledge of attitude," *NAVIGATION*, Vol. 59, No. 1, Spring 2012, pp. 61-75.
- [6] Teunissen, P., De Jonge, P., and Tiberius, C., "The least-squares ambiguity decorrelation adjustment: its performance on short GPS baselines and short observation spans." *Journal of Geodesy*, vol. 71, no. 10, pp. 589-602, 1997.
- [7] Odijk, Dennis. "Instantaneous precise GPS positioning under geomagnetic storm conditions." *GPS Solutions*, vol. 5, no. 2, pp. 29-42, 2001.
- [8] Giorgi, G., Teunissen, P., and Gourlay, T. "Instantaneous global navigation satellite system (GNSS)-based attitude determination for maritime applications." *IEEE Journal of Oceanic Engineering*, vol. 37, no. 3, pp. 348-362, 2012.
- [9] Khanafseh, S., and Pervan, B., "Autonomous airborne refueling of unmanned air vehicles using the global positioning system," *AIAA Journal of Aircraft*, Vol. 44, No. 5, pp. 1670-1682, 2007.
- [10] Rife, J., Khanafseh, S., Pullen, S., De Lorenzo, D., Kim, U.-S., Koenig, M., Chiou, T. Y., Kempny, B., and Pervan, B., "Navigation, Interference Suppression, and Fault Monitoring in the Sea-Based Joint Precision Approach and Landing System," *Proceedings of the IEEE*, Vol. 96, No. 12, pp. 1958 – 1975, 2008.
- [11] Khanafseh, S., and Langel, S., "Implementation and experimental validation of cycle ambiguity resolution with position domain integrity risk constraints," *NAVIGATION*, Vol. 58, No. 1, Spring 2011, pp. 45-58.
- [12] Misra, P., and Enge, P., *Global Position System: Signals, Measurements, and Performance*. Ganga-Jamuna Press, 2006.
- [13] Thrun, S., Burgard, W., Fox, D., *Probabilistic Robotics*, The MIT Press, 2005.
- [14] Nadimi, Esmail S., Tarokh, Vahid, "Bayesian source localization in networks with heterogeneous transmission medium," *NAVIGATION*, Vol. 59, No. 3, Fall 2012, pp. 163-175.
- [15] O'Brien, K., and Rife, J. "Propagating integrity bounds in nonlinear state estimation," *Proc. ION GNSS 2009*, Savannah, GA, pp. 1807 – 1818, 2009.
- [16] J. Clanton, D. Bevly, and A. Hodel, A low-cost solution for an integrated multisensory lane departure warning system, *IEEE Trans. Intelligent Transportation Systems*, vol. 10, pp. 47-59, 2009.
- [17] Rife, J., "Collaborative vision-integrated pseudorange error removal: team-estimated differential GNSS corrections with no stationary reference receiver," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 13, No. 1, pp. 15-24, 2012.