

A Phase-Reconstruction Technique for Low-Power Centimeter-Accurate Mobile Positioning

Kenneth M. Pesyna, Jr., *Student Member, IEEE*, Zaher M. Kassas, *Senior Member, IEEE*,
Robert W. Heath, Jr., *Fellow, IEEE*, and Todd E. Humphreys, *Member, IEEE*

Abstract—A carrier phase reconstruction technique is presented as an enabler for low-power centimeter-accurate mobile positioning. Reliable carrier phase reconstruction permits the duty cycling of a Global Navigation Satellite System (GNSS) receiver whose outputs are used for precise carrier-phase differential GNSS (CDGNSS) positioning. Existing CDGNSS techniques are power intensive because they require continuous tracking of each GNSS signal’s carrier phase. By contrast, the less-precise code-ranging technique that is commonly used in mobile devices for 3-to-10-meter-accurate positioning allows for aggressive measurement duty-cycling, which enables low-power implementations. The technique proposed in this paper relaxes the CDGNSS continuous phase tracking requirement by solving a mixed real and integer estimation problem to reconstruct a continuous carrier phase time history from intermittent phase measurement intervals each having an ambiguous initial phase. Theoretical bounds on the probability of successful phase reconstruction, corroborated by Monte-Carlo-type simulation, are used to investigate the sensitivity of the proposed technique to various system parameters, including the time period between successive phase measurement intervals, the duration of each interval, the carrier-to-noise ratio, and the line-of-sight acceleration uncertainty. A demonstration on real data indicates that coupling a GNSS receiver with a consumer-grade inertial measurement unit enables reliable phase reconstruction with phase measurement duty cycles as low as 5%.

I. INTRODUCTION

Existing commodity Global Navigation Satellite System (GNSS) solutions perform positioning based on code ranging [1]. This technique lends itself well to aggressive duty cycling, which permits a low-power implementation, but its positioning accuracy is limited to approximately 3 to 10 meters [2]. By contrast, the carrier-phase differential GNSS (CDGNSS) technique developed by the surveying and precise GNSS communities during the 1990s offers exquisite (centimeter-level) relative accuracy [3]–[5] but has not been seen as amenable to a duty-cycled low-power implementation because of a supposed continuous carrier phase tracking requirement [6], [7]. This, among other reasons, has prevented adoption of CDGNSS in power-constrained mobile devices. Note that CDGNSS based on a single measurement epoch, as in [5] and

[8], would not require continuous carrier tracking, but this technique is only reliable in favorable signal and multipath environments [9], [10] such as would be uncommon for mobile devices. Accordingly, this paper assumes that successful CDGNSS will require multiple measurement epochs. Relaxing the continuous phase tracking requirement of multi-epoch CDGNSS by first reconstructing a continuous phase time history from intermittent carrier phase measurements prior to performing a CDGNSS solution, thus permitting power-saving duty cycling, is the focus of this paper. While previous work has investigated multi-epoch CDGNSS-based attitude determination on the basis of intermittent carrier phase measurements [11], insofar as this paper’s authors are aware, no prior work has investigated performing a more-general multi-epoch CDGNSS position solution on the basis of such measurements.

A continuous phase time history can be reconstructed from intermittent (duty-cycled) phase measurements if so-called phase ambiguities can be resolved. These phase ambiguities are unknown integer- or fractional-cycle offsets from the true phase that arise at the beginning of each measurement interval for two reasons: (1) the receiver does not track the carrier phase evolution between measurement intervals, and (2) the receiver’s phase discriminator is not capable of measuring absolute phase.

The problem of reconstructing a continuous phase time history from intermittent, ambiguous phase measurements can be posed as a mixed real and integer estimation problem where the real parameter is the time-varying continuous phase and the integer parameters are the phase ambiguities. Prior work in mixed real and integer estimation has led to the development of a general Kalman-filter- [12] and smoother-based [13] framework which has been implemented, not for carrier-phase reconstruction, but for CDGNSS ambiguity resolution [14]. In prior work by the current authors, this framework was modified to construct a continuous carrier phase from time division multiple access (TDMA) Iridium satellite communication signals, enabling their use in navigation [15]. Other authors have constructed a similar framework for fixed-baseline CDGNSS attitude determination [11]. The integer ambiguities in these problems and in the current work bear a strong resemblance to similar integer ambiguities in CDGNSS positioning and in sphere decoding, the resolution of which has been the subject of much research in the GNSS community [5], [16], [17] and in the communications community [18]–[20], respectively.

This paper makes two contributions. First, a technique is

Copyright © 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

This work was funded by the Department of Defense through the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program and by the National Science Foundation under Grant No. 1218338.

The authors are with The University of Texas at Austin, Austin, TX 78712 USA (email: kpesyna@utexas.edu, zkassas@ieee.org, rheath@utexas.edu, todd.humphreys@mail.utexas.edu).

developed for continuous carrier phase reconstruction from duty-cycled phase measurements. The technique builds on the Kalman-filter-and-smoother-based solution to the mixed real and integer estimation problem introduced in [13], but modifies this to incorporate measurement models characterizing carrier phase errors such as the receiver front-end noise, propagation-induced phase errors, and line-of-sight range-errors. The modified solution also incorporates carrier phase structure parameters such as the measurement burst length, the measurement burst period, and the unknown whole- or fractional-cycle phase ambiguities. Moreover, the technique presented in this paper improves upon the framework in [15] by differencing phase measurements made by two GNSS receivers to substantially cancel errors that would otherwise make it more difficult to resolve the phase ambiguities and reconstruct a carrier phase time history. A second improvement is the addition of an inertial measurement unit (IMU) into the reconstruction framework, which enables more accurate between-burst receiver motion modeling, leading to improved phase reconstruction.

The second contribution of this paper is a performance evaluation of the proposed reconstruction technique. This evaluation extends the Monte-Carlo-simulation-based sensitivity study of [15] to include theoretical upper and lower bounds on the probability of successful integer ambiguity resolution. These bounds allow for a quick measure of the reconstruction technique's reliability as a function of the previously-mentioned signal error and structure parameters.

The next section sets up the system model, motivating the need for and the effects of carrier-phase reconstruction. Subsequent sections describe in detail the reconstruction technique and evaluate its performance.

II. CARRIER PHASE MODELS

Three GNSS carrier phase models are introduced in this section; one applies before, one during, and one after carrier phase reconstruction.

A. Before Reconstruction: Undifferenced Residual Carrier Phase Model

Let the undifferenced residual carrier phase $\phi_r(t)$ be defined as the measured phase after duty-cycled downmixing and correlation with the local signal replica. The term "residual" refers to this phase being the difference between the received carrier phase and the receiver's best prediction of the received carrier phase. The undifferenced residual carrier phase for a GNSS signal can be modeled by the following adaptation of the GPS carrier phase measurement model given in [21]:

$$\begin{aligned} \phi_r(t) &\triangleq \text{received carrier phase} - \text{predicted carrier phase} \\ &= \begin{cases} \tilde{\phi}_r(t) + \frac{1}{M}\eta(t) & \text{for } t_{bi} \leq t < t_{bi} + T_b, \\ & i = 0, 1, \dots, N_b - 1 \\ \text{undefined} & \text{otherwise} \end{cases} \quad (1) \\ \tilde{\phi}_r(t) &= \frac{1}{\lambda}r_e(t) + \frac{c}{\lambda}[\delta t_{RX}(t) - \delta t_{TX}(t)] + \gamma_0 - \psi_0 \\ &\quad + \epsilon_p(t) + v_\phi(t) \end{aligned}$$

with the following definitions:

$\tilde{\phi}_r(t)$	the continuous, ambiguity-free residual carrier phase, in cycles.
t_{bi}	the start time of the i^{th} discrete phase measurement interval, or burst, in seconds.
T_b	the burst duration, in seconds.
N_b	the number of bursts.
M	the ambiguity factor used to depict whole-cycle phase ambiguities ($M=1$) or fractional-cycle phase ambiguities ($M > 1$), whichever is appropriate for the receiver setup. If the broadcast binary phase-shift keying (BPSK) GNSS navigation data symbols are provided to the receiver and the receiver knows its position and time to a fraction of a data symbol interval such that it can align the data symbols to the incoming signal to perform data symbol wipeoff, or if the receiver is tracking a data-symbol-free pilot signal, then $M = 1$; otherwise $M = 2$ due to the necessary usage of a squaring-type phase detector [22], that is, a detector which is insensitive to half-cycle phase changes induced by the data symbols.
$\eta(t)$	an integer that remains constant during each discrete phase measurement interval; i.e., $\eta(t) = n_i$ for $t_{bi} \leq t < t_{bi} + T_b$. When scaled by $\frac{1}{M}$ this represents the offset of the signal's measured phase from that of the unambiguous phase at the beginning of each burst. In this paper, n_i will be referred to as the integer ambiguity over the i^{th} burst.
λ	the carrier wavelength, in meters.
$r_e(t)$	the error in the predicted range between the receiver and transmitter, in meters. This term includes errors due to the receiver's inertial measurement unit noise, as discussed briefly in the next paragraph and in detail in section V-B1.
c	the speed of light, in meters per second.
$\delta t_{RX}(t)$	the difference between the predicted and actual receiver clock offset from true time, in seconds.
$\delta t_{TX}(t)$	the difference between the predicted and actual transmitter clock offset from true time, in seconds.
γ_0	the initial replica carrier phase at receiver clock time 0, in cycles.
ψ_0	the initial transmitted carrier phase at satellite clock time 0.
$\epsilon_p(t)$	the carrier phase deviation due to unmodeled propagation and multipath effects, in cycles.
$v_\phi(t)$	the measurement noise introduced by the receiver front-end, in cycles.

This model captures all the significant effects that cause the received carrier phase to be different from what the receiver would predict on the basis of its own clock, its assumed position, and its internal models for propagation and multipath effects, satellite motion, and satellite clock offset.

It is important to point out that to facilitate reliable reconstruction, a dynamic receiver needs to have a rough estimate of its motion. Unmodeled or poorly modeled receiver motion may result in large variations in the receiver-satellite range error $r_e(t)$, which, as will be seen later, decreases the prob-

ability of successful reconstruction. Fortunately, a simple 3-axis inertial measurement unit (IMU) can be used to measure the receiver's 3-dimensional acceleration and angular velocity. These measurements, in conjunction with an initial position and orientation, can be integrated to predict the receiver's position changes and substantially eliminate variations in $r_e(t)$. Of course, inertially-aided motion prediction is imperfect: noise in the IMU measurements will still produce residual variations in $r_e(t)$, which must be accurately characterized to enable optimal phase reconstruction. Further details on IMU-aided phase reconstruction will be presented in Sections V-B and VII.

B. During Reconstruction: Double-Differenced Residual Carrier Phase Model

It is possible to reconstruct a continuous time phase history from measurements of the undifferenced residual carrier phase, as done in [15]. But some error sources modeled in (1), such as errors in the predicted transmitter and receiver clock offsets $\delta t_{TX}(t)$ and $\delta t_{RX}(t)$ and the propagation errors $\epsilon_p(t)$ can often be too unstable to support reliable reconstruction. By implementing a technique known as double-differencing, where measurements between two GNSS satellites and two GNSS receivers (a rover and a reference) are differenced, many of these error sources can be entirely or substantially canceled, increasing the probability of successful reconstruction. To enable this, carrier phase measurements made by the rover and reference receivers can be passed off to a cloud server using a cellular or other wireless connection, where double-differencing (and subsequent reconstruction) can be performed. A secondary benefit of double-differencing is that the resulting reconstructed carrier phase will be in the proper form for CDGNSS, which, as mentioned earlier, is a commonly-used technique that achieves cm-accurate relative positioning by taking advantage of double differencing, not for phase reconstruction, as suggested here, but for precise positioning [3]–[5]. The CDGNSS positioning solution can similarly be performed in the cloud subsequent to reconstruction.

Let the double-differenced residual carrier phase $\nabla\Delta\phi_{rAB}^{ij}(t)$ be defined as the difference of the undifferenced residual carrier phases made between satellites i and j and receivers A and B :

$$\nabla\Delta\phi_{rAB}^{ij}(t) \triangleq [\phi_{r_A}^i(t) - \phi_{r_A}^j(t)] - [\phi_{r_B}^i(t) - \phi_{r_B}^j(t)]. \quad (2)$$

In this model, receiver B differences its undifferenced residual carrier phase measurements made by tracking satellites i and j . This difference is then subtracted from the difference made at receiver A between the same two satellites. Performing the subtractions in (2) and dropping the sub- and superscripts for clarity yields

$$\nabla\Delta\phi_r(t) = \begin{cases} \nabla\Delta\tilde{\phi}_r(t) + \frac{1}{M}\nabla\Delta\eta(t) & \text{for } t_{bi} \leq t < t_{bi} + T_b, \\ & i = 0, 1, \dots, N_b - 1 \\ \text{undefined} & \text{otherwise} \end{cases} \quad (3)$$

$$\nabla\Delta\tilde{\phi}_r(t) = \frac{1}{\lambda}\nabla\Delta r_e(t) + \nabla\Delta\epsilon_p(t) + \nabla\Delta v_\phi(t)$$




Fig. 1. Illustration of the double differenced residual carrier phase measurements formed during each burst. The solid gray trace represents $\nabla\Delta\tilde{\phi}_r(t)$, the continuous but unmeasurable ambiguity-free phase. To represent the measurable phase $\nabla\Delta\phi_r(t)$, $\nabla\Delta\tilde{\phi}_r(t)$ is structured into periodic bursts and aliased between 0 and $\frac{1}{M}$ cycles, forming the intermittent dark traces. The aliasing leads to a phase ambiguity for each burst and occurs due to the insensitivity of the receiver's phase detector to $\frac{1}{M}$ -cycle phase offsets. T_p represents the burst period and T_b represents the burst duration.

with the following new definitions:

$\nabla\Delta\tilde{\phi}_r(t)$ the continuous, ambiguity-free double-differenced residual carrier phase, in cycles.

$\nabla\Delta\eta(t)$ an integer, constant over each measurement burst and measured in cycles, that represents the double difference of the integer $\eta(t)$ term from (1). In addition, $\nabla\Delta\eta(t)$ incorporates the double-difference of the initial transmitter and receiver replica carrier phases ψ_0 and γ_0 between the two satellites and two receivers. For properly designed GNSS receivers this latter double difference is an integer and remains constant during the entire dataset [21].

$\nabla\Delta r_e(t)$ the error in the double-differenced predicted range between the two receivers and two satellites, in meters. This term includes errors due to the receiver's inertial measurement unit noise, as discussed later on in Sec. V-B1.

$\nabla\Delta\epsilon_p(t)$ the double-differenced carrier phase deviation due to unmodeled propagation and multipath effects, in cycles.

$\nabla\Delta v_\phi(t)$ the double-differenced measurement noise induced by the receivers' front-ends, in cycles.

Note that the double-differencing operation has canceled the error terms $\delta t_{TX}(t)$ and $\delta t_{RX}(t)$ introduced in (1). Because of these cancellations and substantial reductions in the variations of other terms, it is more effective to apply phase reconstruction to the double-differenced residual carrier phase rather than to the undifferenced residual carrier phase. Accordingly, (3) will be taken to model the received carrier phase *during* reconstruction.

Fig. 1 illustrates the formation of the phase ambiguities $\frac{1}{M}\nabla\Delta\eta(t)$ modeled in (3). The upper gray trace represents the continuous and ambiguity-free double-differenced residual carrier phase $\nabla\Delta\tilde{\phi}_r(t)$, which could be measured if both the reference and rover receivers were continually tracking GNSS signals. Instead, due to the measurement duty-cycling by one

or both receivers, the measurable phase becomes periodic and phase-aliased as illustrated by the lower black trace. Aliasing is caused by the insensitivity of the receiver's phase discriminator to whole- or fractional-cycle phase drifts between bursts and leads to the formation of the phase ambiguities.

C. After Reconstruction: Reconstructed Double-Differenced Carrier Phase Model

One final model is presented here to characterize the double-differenced carrier phase after reconstruction. Although this model is not used during reconstruction, it nicely illustrates the effects of reconstruction errors and relates them to the so-called "ideal" or error-free reconstructed carrier phase. In this model, the reconstructed double-differenced carrier phase $\nabla\Delta\phi_R(t)$ is given by

$$\nabla\Delta\phi_R(t) = \nabla\Delta\phi_{\text{ideal}}(t) + \beta(t) + \frac{1}{M} [\nabla\Delta\eta(t) - \nabla\Delta\hat{\eta}(t)]. \quad (4)$$

Here, the following new definitions apply:

$\nabla\Delta\phi_{\text{ideal}}(t)$ the ideal double-differenced residual carrier phase. This term represents the double-differenced carrier phase as it would appear if it were perfectly reconstructed, i.e., if the receivers involved in the double-differencing were continuously tracking the GNSS signals and there was no measurement noise.

$\beta(t)$ the non-ambiguity related reconstruction errors, measured in cycles. This term encompasses all non-ambiguity-related deviations of $\nabla\Delta\phi_R(t)$ from $\nabla\Delta\phi_{\text{ideal}}(t)$.

$\nabla\Delta\hat{\eta}(t)$ the reconstruction technique's best estimate of the time-varying double-differenced integer ambiguity term $\nabla\Delta\eta(t)$, measured in cycles. The difference $\frac{1}{M}[\nabla\Delta\eta(t) - \nabla\Delta\hat{\eta}(t)]$ is the time-varying reconstruction error that arises during ambiguity resolution.

$\nabla\Delta\phi_R(t)$ is the reconstruction technique's best estimate of $\nabla\Delta\phi_{\text{ideal}}(t)$, the continuous, noise-free, and ambiguity-free double-differenced residual carrier phase. Errors in phase reconstruction cause $\nabla\Delta\phi_R(t)$ to deviate, sometimes significantly, from $\nabla\Delta\phi_{\text{ideal}}(t)$. This deviation is modeled by $\beta(t)$ and $\frac{1}{M}[\nabla\Delta\eta(t) - \nabla\Delta\hat{\eta}(t)]$, the non-ambiguity and ambiguity-related reconstructed errors, respectively. Typically, the second term dominates, as errors in ambiguity resolution tend to be much larger than non-ambiguity errors.

Because the receiver only has access to the intermittent ambiguous phase $\nabla\Delta\phi_r(t)$, as represented by the lower dark trace in Fig. 1, the reconstruction algorithm must determine in which whole-cycle vertical partition (or fractional-cycle partition if $M > 1$) each solid black curve would reside if $\nabla\Delta\phi_r(t)$ were instead unambiguous. That is, it must determine the time-varying integer-valued phase-ambiguity term $\nabla\Delta\eta(t)$. Fig. 2 helps to illustrate this challenge. The horizontal dashed lines illustrate vertical partitions in which the reconstructed phase $\nabla\Delta\phi_R(t)$ could lie. Here $M = 1$, so each partition is 1 cycle in height. These partitions repeat infinitely in each direction along the vertical axis. This leads to an infinite number of




Fig. 2. Illustration of possible reconstructed phase trajectories, only 1 of which corresponds to the true trajectory. It is the job of the reconstruction algorithm to reconstruct the true trajectory using measurements of the periodic and aliased double-differenced residual carrier phase $\nabla\Delta\phi_r(t)$, which, in this particular illustration, has 1 cycle phase ambiguities, i.e., $M = 1$.

possible phase time histories, or trajectories, 16 of which are depicted in the figure. However, just one of these trajectories accurately depicts the continuous, ambiguity-free phase time history $\nabla\Delta\phi_r(t)$. It becomes the task of the reconstruction algorithm to use past, present, and future measurements of $\nabla\Delta\phi_r(t)$ to resolve the phase ambiguities and attempt to reconstruct $\nabla\Delta\phi_r(t)$. If one or more ambiguities are resolved incorrectly, an incorrect reconstructed phase trajectory would be chosen, leading $\nabla\Delta\phi_R(t)$ to deviate significantly ($\geq \frac{1}{M}$ cycles) from $\nabla\Delta\phi_r(t)$. Such errors will degrade the utility of the reconstructed phase time history in the context of a CDGNSS solution, as they will lead to a positioning solution that is no longer cm-accurate (see Sec. VII). Accordingly, it becomes useful to examine the probability of correctly resolving the phase ambiguities; this will be done both theoretically and empirically in subsequent sections.

III. RECONSTRUCTION TECHNIQUE

This section presents the proposed technique for reconstructing a continuous carrier phase time history from intermittent phase measurements made by the receiver. By intermittent, we mean non-continuous measurements. The measurements are assumed in this paper to be periodic, but the reconstruction technique would need to be modified only slightly to accommodate aperiodic measurements. The technique takes the double-differenced residual carrier phase measurements $\nabla\Delta\phi_r(t)$ [modeled in (3)] and forms the smoothed reconstructed double-differenced carrier phase time history $\nabla\Delta\phi_R(t)$ [modeled in (4)]. It resolves phase ambiguities with an integer least-squares solver and "stitches" discrete phase measurements together with a Kalman filter and smoother. Square-root information implementations of the filter and smoother ensure that phase reconstruction is performed in an accurate and computationally-efficient manner [13], [23]. For a tutorial on square-root information filtering and its relationship to traditional Kalman filtering, see [24].

A. Estimation State and Dynamics and Measurement Model

This section describes the reconstruction filter and smoother's state as well as its dynamics and measurement models.

1) *State*: The state has a real-valued component that models the noise- and ambiguity-free double-differenced residual carrier phase, and an integer-valued component that models the phase ambiguities. The real-valued state component at time t_k is denoted \mathbf{x}_k , where $t_k = kT$ and $T \leq T_b$ is the time between consecutive filter and smoother updates. This component can be expressed as

$$\mathbf{x}_k = [\nabla\Delta\phi_{\text{ideal},k}, \omega_k]^\top \quad (5)$$

with the following definitions:

$\nabla\Delta\phi_{\text{ideal},k}$ the discrete-time noise- and ambiguity-free ideal double-differenced residual carrier phase at time t_k , in cycles, i.e., $\nabla\Delta\phi_{\text{ideal},k} = \nabla\Delta\phi_{\text{ideal}}(t_k)$, where $\nabla\Delta\phi_{\text{ideal}}(t)$ was defined in (4).

ω_k the rate of change of $\nabla\Delta\phi_{\text{ideal}}(t)$ at time t_k , in Hz. The integer-valued state component \mathbf{n}_k at time t_k can be expressed as

$$\mathbf{n}_k = [n_1, n_2, \dots, n_{i_k}]^\top \quad (6)$$

with the following definitions:

- \mathbf{n}_k an $i_k \times 1$ vector of integers, one for each measurement burst that began between time 0 and t_k .
- n_{i_k} the integer corresponding to the i_k^{th} measurement burst.
- i_k a counter representing the number of measurement bursts that begin at or before time t_k .

2) *Dynamics Model*: The real and integer components of the state evolve separately; thus their dynamics will be modeled separately. The real-valued state component \mathbf{x}_k is assumed to evolve as a first-order Gauss-Markov process with process noise representing the variations due to $\nabla\Delta r_e(t)$ and $\nabla\Delta\epsilon_p(t)$ from (3). The integer-valued state component \mathbf{n}_k evolves under the assumption that a new ambiguity is introduced with each measurement burst.

The following models describe the time evolution of the real- and integer-valued state components:

$$\mathbf{x}_{k+1} = \Phi\mathbf{x}_k + \Gamma\mathbf{w}_k \quad (7)$$

$$\mathbf{n}_{k+1} = \begin{cases} \begin{bmatrix} \mathbf{n}_k \\ n_{i_{k+1}} \end{bmatrix} & \text{if a new burst began within the interval } (t_k, t_{k+1}] \\ \begin{bmatrix} \mathbf{n}_k \end{bmatrix} & \text{otherwise} \end{cases} \quad (8)$$

with the following definitions:

- Φ the state transition matrix.
- Γ the process noise influence matrix.
- \mathbf{w}_k the process noise at time t_k , modeled as a discrete-time zero-mean \mathbf{Q} covariance Gaussian random vector, i.e., $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$.
- \mathbf{Q} the process noise covariance matrix.

The state transition matrix for the real-valued state models standard Euler integration from t_k to t_{k+1} :

$$\Phi = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}. \quad (9)$$

The process noise influence is defined as

$$\Gamma = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \quad (10)$$

and the process noise covariance matrix is defined as

$$\mathbf{Q} = S_g f_0^2 \begin{bmatrix} \frac{T^3}{3} & \frac{T^2}{2} & \frac{T^3}{8} \\ \frac{T^2}{2} & T & \frac{T^2}{6} \\ \frac{T^3}{8} & \frac{T^2}{6} & \frac{T^3}{20} \end{bmatrix} + S_f f_0^2 \begin{bmatrix} T & 0 & \frac{T}{2} \\ 0 & 0 & 0 \\ \frac{T}{2} & 0 & \frac{T}{3} \end{bmatrix}, \quad (11)$$

where f_0 is the GNSS signal's nominal carrier frequency, in Hz. The quantities S_g and S_f parameterize the combined phase instability caused by the process noise error components in (3), namely $\nabla\Delta r_e(t)$ and $\nabla\Delta\epsilon_p(t)$. The model for the evolution of the real-valued state elements in (7) with the process-noise covariance defined by (11) follows a two-state Gauss-Markov model commonly used to describe clock-error-induced phase variations (see [25], Ch. 11). This model will be discussed further in Sec. V-A. Note that \mathbf{w}_k is of dimension 3-by-1 while \mathbf{x}_k is of dimension 2-by-1. The third element in \mathbf{w}_k and, correspondingly, the third row and column in \mathbf{Q} is a standard way to model the average of the phase process noise over the interval $t_{k-1} < t \leq t_k$ [26] and will be needed in the measurement model discussed next.

3) *Measurement Model*: The filter ingests measurements y_k of the double-differenced residual carrier phase and relates these measurements to its state. Each measurement y_k represents the average of $\nabla\Delta\phi_r(t)$ over the interval $t_{k-1} < t \leq t_k$, i.e., $y_k = \frac{1}{T} \int_{t_{k-1}}^{t_k} \nabla\Delta\phi_r(t) dt$. It should be noted that filter measurement updates occur only within measurement bursts when measurements are available. The filter's measurement model relates y_k to the real- and integer-valued state components \mathbf{x}_k and \mathbf{n}_k and to the process-noise \mathbf{w}_{k-1} :

$$y_k = \begin{cases} \tilde{\mathbf{H}}_x \mathbf{x}_k + \tilde{\mathbf{H}}_{n_k} \mathbf{n}_k + \tilde{\mathbf{H}}_w \mathbf{w}_{k-1} + v_k & \text{for } t_{b_i} \leq t_k < t_{b_i} + T_b, \\ & i = 0, 1, \dots, N_b - 1 \\ \text{undefined} & \text{otherwise} \end{cases} \quad (12)$$

with the following new definitions:

- $\tilde{\mathbf{H}}_x$ the measurement sensitivity matrix for the real-valued state components.
- $\tilde{\mathbf{H}}_{n_k}$ the measurement sensitivity matrix for the integer-valued state components at time t_k .
- $\tilde{\mathbf{H}}_w$ the measurement sensitivity matrix for the process noise.
- v_k the average of the continuous-time double-differenced measurement noise over the interval $t_{k-1} < t \leq t_k$, i.e., $v_k = \frac{1}{T} \int_{t_{k-1}}^{t_k} \nabla\Delta v_\phi(t) dt$. v_k is modeled as a zero-mean discrete-time Gaussian white noise process, $v_k \sim \mathcal{N}(0, \sigma_{\phi_k}^2)$, where $\sigma_{\phi_k}^2$ has a nonlinear relationship with the mean carrier-to-noise ratio over the interval $(C/N_0)_k$, but for high $(C/N_0)_k$ converges to $\sigma_{\phi_k}^2 = \frac{1}{2T(C/N_0)_k}$ [27].

The measurement sensitivity matrices can be expanded as

$$\tilde{\mathbf{H}}_x = \begin{bmatrix} 1 & -\frac{T}{2} \end{bmatrix} \quad (13)$$

$$\tilde{\mathbf{H}}_{n_k} = \begin{bmatrix} 0 & 0 & \dots & 0 & \frac{1}{M} \end{bmatrix} \quad (14)$$

$$\tilde{\mathbf{H}}_w = \begin{bmatrix} -1 & \frac{T}{2} & 1 \end{bmatrix} \quad (15)$$

where T is the time between consecutive filter updates, as defined previously, and M is the ambiguity factor defined in (1). Two features of the $1 \times i_k$ matrix $\tilde{\mathbf{H}}_{nk}$ are noteworthy. First, the $\frac{1}{M}$ factor in its last element allows the integer-valued state \mathbf{n}_k to relate to a whole-cycle ($M = 1$) or a fractional-cycle ($M > 1$) phase ambiguity. Second, $\tilde{\mathbf{H}}_{nk}$ has 0s in all but its last element to ensure that the measurements made during burst i_k are only affected by the most recent integer ambiguity n_{i_k} of \mathbf{n}_k . Because y_k is an average, $\tilde{\mathbf{H}}_w$ is needed to model the accumulation of process noise into the measurement [26].

B. Cost Function

Optimal channel-by-channel estimates of the state components \mathbf{x}_k and \mathbf{n}_k for $1 \leq k \leq K$ can be obtained according to the maximum *a posteriori* criterion based on all measurements y_k from $k = 1$ to K by determining the state and process noise time histories that minimize a certain cost function subject to the dynamics models in (7) and (8). For numerical robustness, a square-root-information approach is adopted [13], [23]. Let the square-root information equation for the *a priori* estimate of the real-valued state component \mathbf{x}_0 at $k = 0$ be given by

$$\mathbf{z}_{x0} = \mathbf{R}_{xx0}\mathbf{x}_0 + \mathbf{v}_{x0} \quad (16)$$

with the following definitions:

- \mathbf{z}_{x0} the *a priori* nonhomogeneous term that stores information about \mathbf{x}_0 .
- \mathbf{R}_{x0k} the *a priori* square-root information matrix for \mathbf{x}_0 .
- \mathbf{v}_{x0} the error corresponding to \mathbf{x}_0 , a sample from a discrete-time zero-mean, unity covariance Gaussian white noise process, i.e., $\mathbf{v}_{x0} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

No *a priori* information is assumed to be available for the integer-valued state component \mathbf{n} . Let the square-root information equation for the *a priori* estimate of the process noise \mathbf{w}_k at each time index k be given by

$$\mathbf{z}_{wk} = 0 = \mathbf{R}_{ww}\mathbf{w}_k + \mathbf{v}_{wk} \quad (17)$$

with the following definitions:

- \mathbf{z}_{wk} the *a priori* nonhomogeneous term that stores information about \mathbf{w}_k .
- \mathbf{R}_{ww} the *a priori* square-root information matrix for \mathbf{w}_k , defined as $\mathbf{R}_{ww} = \mathbf{Q}^{-\frac{1}{2}}$, where \mathbf{Q} is defined in (11).
- \mathbf{v}_{wk} the error corresponding to \mathbf{w}_k , a sample from a discrete-time zero-mean, unity covariance Gaussian white noise process, i.e., $\mathbf{v}_{wk} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

The equation in (17) is set equal to zero because the process noise is assumed to be zero-mean and thus solving (17) for the *a priori* process noise estimate $\hat{\mathbf{w}}_k$ should yield $\hat{\mathbf{w}}_k = \mathbf{0}$. Now, let the measurement model in (12) be normalized by multiplying both sides by $\sigma_{\phi_k}^{-1}$. This normalized measurement model, now in standard square-root equation form, is written

$$z_k = \mathbf{H}_{xk}\mathbf{x}_k + \mathbf{H}_{nk}\mathbf{n}_k + \mathbf{H}_{wk}\mathbf{w}_{k-1} + v_{zk} \quad (18)$$

with the following definitions:

- z_k the normalized nonhomogeneous term defined as $z_k = \sigma_{\phi_k}^{-1}y_k$.

- \mathbf{H}_{xk} the normalized measurement sensitivity matrix for the real-valued state components at time t_k , defined as $\mathbf{H}_{xk} = \sigma_{\phi_k}^{-1}\tilde{\mathbf{H}}_x$
- \mathbf{H}_{nk} the normalized measurement sensitivity matrix for the integer-valued state components at time t_k , defined as $\mathbf{H}_{nk} = \sigma_{\phi_k}^{-1}\tilde{\mathbf{H}}_{nk}$.
- \mathbf{H}_{wk} the normalized measurement sensitivity matrix for the process-noise at time t_k , defined as $\mathbf{H}_{wk} = \sigma_{\phi_k}^{-1}\tilde{\mathbf{H}}_w$.
- v_{zk} the normalized measurement noise at time t_k , modeled as a zero-mean, unit variance, discrete-time, Gaussian white noise process, $v_{zk} \sim \mathcal{N}(0, 1)$.

Like y_k in (12), z_k is undefined between bursts.

Given (16), (17), and (18), the phase reconstruction problem can be posed as follows:

$$\begin{aligned} \underset{\mathbf{x}_i, \mathbf{n}_i \{i: 0 < i \leq K\}}{\underset{\mathbf{w}_i \{i: 0 \leq i < K\}}{\text{minimize}}} \quad J = & \underbrace{\|\mathbf{R}_{xx0}\mathbf{x}_0 - \mathbf{z}_{x0}\|^2}_{\text{A priori information}} + \underbrace{\sum_{k=0}^{K-1} \|\mathbf{R}_{ww}\mathbf{w}_k\|^2}_{\text{Process noise}} \\ & + \underbrace{\sum_{k=1}^K \|\mathbf{H}_{xk}\mathbf{x}_k + \mathbf{H}_{nk}\mathbf{n}_k + \mathbf{H}_{wk}\mathbf{w}_{k-1} - z_k\|^2}_{\text{Measurements}} \end{aligned} \quad (19)$$

subject to the state dynamics models in (7) and (8).

A solution to (19) can be found by breaking the reconstruction process into three stages: filtering, ambiguity resolution, and smoothing.

C. Filtering

Filtering is the first stage in the reconstruction process. Filter estimates of the state at each time index k will be produced by making optimal use of the measurement at time t_k (if within a measurement burst) and all measurements prior to t_k . It can be shown through a series of orthogonal transformations on (19) that at each time index k , the filter's best estimate of the real- and integer-valued state elements can be found by choosing \mathbf{x}_k and \mathbf{n}_k to minimize the partial cost functional [28]

$$\begin{aligned} J_k(\mathbf{x}_k, \mathbf{n}_k) = & \underbrace{\|\mathbf{R}_{xxk}\mathbf{x}_k + \mathbf{R}_{xnk}\mathbf{n}_k - \mathbf{z}_{xk}\|^2}_{\text{Term involving the integer- and real-valued states}} \\ & + \underbrace{\|\mathbf{R}_{nnk}\mathbf{n}_k - \mathbf{z}_{nk}\|^2}_{\text{Term involving only the integer-valued state}} + \underbrace{\sum_{i=1}^k \|z_{ri}\|^2}_{\text{Residual term}} \end{aligned} \quad (20)$$

with the following definitions:

- \mathbf{z}_{xk} the nonhomogeneous term corresponding to the real-valued state component at time t_k .
- \mathbf{z}_{nk} the nonhomogeneous term corresponding to the integer-valued state component at time t_k .
- z_{ri} the residual nonhomogeneous term at time t_{bi} , $1 \leq i \leq k$.
- \mathbf{R}_{xxk} the square-root information matrix corresponding to \mathbf{x}_k and \mathbf{z}_{xk} at time t_k .
- \mathbf{R}_{xnk} the square-root information matrix corresponding to \mathbf{n}_k and \mathbf{z}_{xk} at time t_k .

$\mathbf{R}_{\text{nn}k}$ the square-root information matrix corresponding to \mathbf{n}_k and \mathbf{z}_{nk} at time t_k .

$J_k(\mathbf{x}_k, \mathbf{n}_k)$ is the contribution to the overall cost that is obtained after filtering measurements z_1 to z_k . Each term on the right-hand side of (20) is produced during the filter's measurement updates during which *a priori* state estimates are combined with measurements. Minimization of (20) proceeds as follows: First, one determines the integer-valued vector state estimate $\hat{\mathbf{n}}_k$ that minimizes the second term of (20), the term involving only the integer-valued state. This estimate can be determined efficiently using the integer least-squares techniques discussed in the next section. Once determined, $\hat{\mathbf{n}}_k$ is inserted into the first term, the term involving both the integer- and real-valued states. At this point, it is possible to determine the real-valued state estimate $\hat{\mathbf{x}}_k$ that reduces the first term to zero, minimizing (20).

D. Integer Ambiguity Resolution

Integer ambiguity resolution is the second stage in the reconstruction process and must be performed before the real-valued state component can be determined. At any time index k during filtering, the cost functional of the form in (20) can be minimized to provide real-time (causal) estimates of the real- and integer-valued state components. This entails first minimizing the following cost function involving the integer-valued state

$$J_{\text{n}}(\mathbf{n}_k) = \|\mathbf{R}_{\text{nn}k}\mathbf{n}_k - \mathbf{z}_{\text{nk}}\|^2. \quad (21)$$

This minimization can be posed as an integer least-squares (ILS) problem whose solution has been shown to be NP-hard and has been studied extensively [5], [17], [18], [29]. ILS solution algorithms are optimal in the sense that out of the set of all admissible estimators, they have the largest possible probability of successful integer ambiguity resolution [30]. For the definition of an admissible estimator, see [30]. Solution algorithms accept the matrix $\mathbf{R}_{\text{nn}k}$ and the vector \mathbf{z}_{nk} from the filter and solve for the vector \mathbf{n}_k that minimizes (21); calling this minimizing vector $\hat{\mathbf{n}}_k$.

If desired, to save computational resources, the minimizing procedure to estimate \mathbf{n}_k only need be performed once, at the end of the dataset at time index K . This is because $\hat{\mathbf{n}}_k$ is a vector that contains integer estimates for all ambiguities up to through time k . Real-time requirements, however, may require $\hat{\mathbf{n}}_k$ to be determined more often, e.g., after each filter update, as the real-valued state components \mathbf{x}_k may be needed in real-time and these cannot be determined without first determining $\hat{\mathbf{n}}_k$.

The ILS solution algorithm can be interpreted geometrically as a closest point lattice search [17], where the lattice is defined by the $n \times n$ dimensional square-root information matrix $\mathbf{R}_{\text{nn}k}$ and the n -dimensional vector of integers \mathbf{n}_k . The product $\mathbf{R}_{\text{nn}k}\mathbf{n}_k$ forms an n -dimensional vector which spans the lattice. Given $\mathbf{R}_{\text{nn}k}$ and \mathbf{z}_{nk} , the ILS solution amounts to finding the closest lattice point $\mathbf{R}_{\text{nn}k}\mathbf{n}_k$ to \mathbf{z}_{nk} [18]:

$$\hat{\mathbf{n}}_k = \underset{\mathbf{n}_k \in \mathbb{Z}^{i_k}}{\operatorname{argmin}} \|\mathbf{R}_{\text{nn}k}\mathbf{n}_k - \mathbf{z}_{\text{nk}}\|^2. \quad (22)$$

The solution procedure can be broken into a reduction step and a search step. The reduction step attempts to reduce the search space; the search step searches for the lowest-cost solution. For the reduction step, the least-squares ambiguity decorrelation adjustment method (LAMBDA) [5] and the Lenstra-Lenstra-Lovász (LLL) reduction [31] are widely used in practice [16], [17]. Implementations of both the LAMBDA method [32] and the LLL method [33] were compared by the authors. They were found to offer comparable computational performance. For the search step, the solution algorithm introduced in [32] has been used for the results presented in this paper.

E. Smoothing

Smoothing is the third stage in the phase reconstruction process. Although a reconstructed carrier phase time history can be determined by solving for the time-varying real-valued state component \mathbf{x}_k after only the first two stages, smoothing is acausal and thus enables past, present, and future phase measurements to be incorporated into the estimates of \mathbf{x}_k at each time instant. To initialize the smoother, the integer ambiguity vector estimate $\hat{\mathbf{n}}_K$ after the final measurement update is determined as described previously and then incorporated, together with $\mathbf{R}_{\text{xn}K}$, to form the smoother's initial nonhomogeneous term $\mathbf{z}_{\text{x}K}^*$ and initial square-root information matrix $\mathbf{R}_{\text{xx}K}^*$ as follows:

$$\mathbf{z}_{\text{x}K}^* = \mathbf{z}_{\text{x}K} - \mathbf{R}_{\text{xn}K}\hat{\mathbf{n}}_K \quad (23)$$

$$\mathbf{R}_{\text{xx}K}^* = \mathbf{R}_{\text{xx}K}. \quad (24)$$

It should be noted that because the smoother is initialized with the already-resolved integer ambiguity vector $\hat{\mathbf{n}}_K$, a quantity determined solely from filter outputs as described in Sec. III-D, smoothing has no effect on integer ambiguity resolution. Consequently, the smoother's contribution to phase reconstruction is a minor one; smoothing acts only to remove abrupt innovation-induced dynamics from \mathbf{x}_k that do not conform to the filter's state dynamics model (see [15], Fig. 5). Furthermore, because smoothing is performed over a batch of measurements, a natural lag is introduced between when the measurements are taken and when the smoothed reconstructed double-differenced carrier phase estimates are formed. As a result, for real-time systems, smoothing may be forgone in favor of removing this lag. Computational lag due to filtering, ambiguity resolution, and subsequent CDGNSS processing, however, will still persist.

After this initialization, the smoother begins its processing. At each time index k , $0 \leq k \leq K$, the smoother ingests $\mathbf{z}_{\text{x}k}^*$ and $\mathbf{R}_{\text{xx}k}^*$ from the previous smoother update as well the process noise terms $\mathbf{z}_{\text{w},k-1}$, \mathbf{R}_{ww} , $\mathbf{R}_{\text{wx},k-1}$, and $\mathbf{R}_{\text{wn},k-1}$ from the filtering stage and outputs $\mathbf{R}_{\text{xx},k-1}^*$ and $\mathbf{z}_{\text{x},k-1}^*$. It then decrements k by 1 and repeats, working backward from index K until it reaches $k = 0$. Smoothed state estimates \mathbf{x}_k^* for $k = 0, 1, \dots, K$ can then be computed from the smoother output terms as follows:

$$\mathbf{x}_k^* = (\mathbf{R}_{\text{xx}k}^*)^{-1}\mathbf{z}_{\text{x}k}^*. \quad (25)$$

The minimum cost after smoothing can be shown to be [28]

$$J(\{\mathbf{x}_i^*\}_{i=1}^K, \hat{\mathbf{n}}_K, \{\mathbf{w}_i^*\}_{i=0}^{K-1}) = \underbrace{\|\mathbf{R}_{\text{nn}K} \hat{\mathbf{n}}_K - \mathbf{z}_{\text{nn}K}\|^2}_{\text{Integer-fit error}} + \underbrace{\sum_{i=1}^K \|z_{ri}\|^2}_{\text{Residual error}}. \quad (26)$$

IV. BOUNDS ON THE PROBABILITY OF SUCCESSFUL AMBIGUITY RESOLUTION

In Section III-D, it was shown that the reconstruction algorithm uses an integer least-squares solver to determine the vector of integer ambiguities \mathbf{n}_k which minimizes (21). However, because of noise, there is no guarantee that the minimizing \mathbf{n}_k , denoted $\hat{\mathbf{n}}_k$, equals the true integer phase ambiguities of the double-differenced residual carrier phase trajectory. This section discusses the probability of successful ambiguity resolution P_c , or the probability that $\hat{\mathbf{n}}_k$ equals \mathbf{n}_k . Only bounds on P_c are presented, as determination of the exact probability is NP-hard.

As discussed previously, minimizing (21) is equivalent to finding the closest lattice point $\mathbf{R}_{\text{nn}k} \mathbf{n}_k$ to $\mathbf{z}_{\text{nn}k}$, which, in turn, is equivalent to minimizing the ambiguity measurement noise vector $\mathbf{v}_{\text{nn}k}$ in the following ambiguity square-root information equation:

$$\mathbf{z}_{\text{nn}k} = \mathbf{R}_{\text{nn}k} \mathbf{n}_k + \mathbf{v}_{\text{nn}k}, \quad \mathbf{v}_{\text{nn}k} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (27)$$

The vector of integers $\hat{\mathbf{n}}_k$ that corresponds to the closest lattice point will be equal to the true vector of integer ambiguities \mathbf{n}_k if and only if the ambiguity measurement noise $\mathbf{v}_{\text{nn}k}$ is such that $\mathbf{z}_{\text{nn}k}$ remains closer to the lattice point $\mathbf{R}_{\text{nn}k} \mathbf{n}_k$ than any other point in the lattice. This is equivalent to $\mathbf{R}_{\text{nn}k} \mathbf{n}_k + \mathbf{v}_{\text{nn}k}$ falling within the *Voronoi cell* $\mathcal{V}_{\text{Rnn}k}$ of $\mathbf{R}_{\text{nn}k} \mathbf{n}_k$. $\mathcal{V}_{\text{Rnn}k}$ is formally defined as the collection of real-valued ℓ -dimensional points (where $\ell = i_k$) closer to $\mathbf{R}_{\text{nn}k} \mathbf{n}_k$ than any other lattice point. Under this framework, the probability of correct integer ambiguity resolution P_c can be defined as [17]:

$$P_c = \Pr\{\mathbf{R}_{\text{nn}k} \mathbf{n}_k + \mathbf{v}_{\text{nn}k} \in \mathcal{V}_{\text{Rnn}k}\} \quad \mathbf{v}_{\text{nn}k} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (28)$$

Because the lattice has a periodic structure, $\mathcal{V}_{\text{Rnn}k}$ is merely a translation of the origin's *Voronoi cell* \mathcal{V}_{0k} by $\mathbf{R}_{\text{nn}k} \mathbf{n}_k$. Thus P_c can be written equivalently as

$$P_c = \Pr\{\mathbf{v}_{\text{nn}k} \in \mathcal{V}_{0k}\} \quad \mathbf{v}_{\text{nn}k} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (29)$$

P_c is now a function of solely the Gaussian random ambiguity noise vector $\mathbf{v}_{\text{nn}k}$ and can be precisely determined by integrating the probability distribution function of $\mathbf{v}_{\text{nn}k}$ over \mathcal{V}_{0k} [34]:

$$P_c = \int_{\mathcal{V}_{0k}} \mathcal{N}(\mathbf{v}_{\text{nn}k}; \mathbf{0}, \mathbf{I}) d\mathbf{v}_{\text{nn}k} = \int_{\mathcal{V}_{0k}} \frac{1}{(2\pi)^{\frac{n}{2}}} \exp\left(-\frac{1}{2} \|\mathbf{v}_{\text{nn}k}\|^2\right) d\mathbf{v}_{\text{nn}k}. \quad (30)$$

In (30), $\mathcal{N}(\mathbf{v}_{\text{nn}k}; \mathbf{0}, \mathbf{I})$ is the multivariate normal distribution and $\|\cdot\|$ is the L_2 -norm. Unfortunately, determining \mathcal{V}_{0k} and integrating over it is a computationally intensive problem [17]. Nonetheless, it is possible to relax the structure of \mathcal{V}_{0k} and solve instead for bounds on P_c [17], [34], [35].

A. Upper Bound on P_c

The volume of a Voronoi cell is equal to the absolute value of the determinant of its lattice generating matrix [17]. Thus, the volume of \mathcal{V}_{0k} is $|\det \mathbf{R}_{\text{nn}k}|$. By making a simplifying assumption that \mathcal{V}_{0k} is an ℓ -dimensional hypersphere with the same volume, an upper bound on the probability of successful integer ambiguity resolution can be written as [17]

$$P_{c,\text{ub}} = \Pr\{\|\mathbf{v}_{\text{nn}k}\| < \rho\} \quad (31)$$

where ρ is the radius of the hypersphere defined as

$$\rho = \sqrt[\ell]{|\det \mathbf{R}_{\text{nn}k}| / \alpha_\ell} \quad (32)$$

where ℓ is the dimension of vector $\mathbf{v}_{\text{nn}k}$ and

$$\alpha_\ell = \pi^{\frac{\ell}{2}} / \Gamma(\ell/2 + 1) \quad \text{and} \\ \Gamma(\ell) = (\ell - 1)!.$$

Since $\mathbf{v}_{\text{nn}k}$ is an ℓ -dimensional normal random vector, $\|\mathbf{v}_{\text{nn}k}\|^2$ is equal to the sum of squares of ℓ independent normally distributed random variables [17], which has a chi-squared distribution with ℓ -degrees of freedom. As a result,

$$P_{c,\text{ub}} = F_{\chi^2}(\rho^2; \ell). \quad (33)$$

where $F_{\chi^2}(\cdot; n)$ is the cumulative distribution function of a n -degree chi-squared random variable.

B. Lower Bound on P_c

The probability of correctly resolving integer ambiguities using so-called integer bootstrapping [36] offers the sharpest known lower bound on P_c [37]. Unlike the case with ILS solvers, it is possible to compute this bootstrapping probability exactly. The bootstrapping estimator takes an approach where it rounds the float least-squares solution while taking advantage of correlation between the ambiguities into account. The bootstrapping estimator's probability of successful ambiguity resolution, which offers a lower bound on P_c , can be written as [36]

$$P_{c,\text{lb}} = \prod_{i=1}^n \left(2\Psi\left(\frac{1}{2\sigma_{\hat{n}_i|I}}\right) - 1 \right) \quad (34)$$

where $\sigma_{\hat{n}_i|I}$ are conditional variances derived from $\mathbf{R}_{\text{nn}k}$ [36], and

$$\Psi(x) = \int_{-\infty}^x \frac{1}{2\pi} \exp\left(-\frac{1}{2}y^2\right) dy. \quad (35)$$

For the sensitivity results presented later in Sec. VI, the code provided by the Ps-LAMBDA software package [38] was used to compute this lower bound.

V. SIMULATION AND TEST ENVIRONMENT

To evaluate the performance of the reconstruction technique outlined in Sec. III, a Monte-Carlo-type simulation and test environment has been designed in MATLAB. The environment performs three tasks.

First, it simulates double-differenced GNSS residual carrier phase time histories $\nabla\Delta\phi_r(t)$. Noise parameters modeling the double-differenced range error $\nabla\Delta r_e(t)$, the

double-differenced propagation- and multipath-induced effects $\nabla\Delta\epsilon_p(t)$, and the double-differenced measurement noise $\nabla\Delta v_\phi(t)$ are inputs to the simulator. Structural parameters such as the measurement burst duration T_b , the time between consecutive bursts T_p , and the ambiguity factor M are also inputs. Note that although the reconstruction technique can handle variations in T_b and T_p , i.e., a non-fixed burst duration and aperiodic bursts, for the analysis performed in this section, these quantities will be assumed fixed. From these parameters, independent time histories of $\nabla\Delta\phi_r(t)$ are generated. Note that phase-locked loop (PLL) pull-in transients need not be simulated in $\nabla\Delta\phi_r(t)$ so long as a batch estimation technique is assumed to be used by the receiver as opposed to a PLL as will be discussed in Sec. VIII-D.

Second, the reconstruction technique is applied to each generated $\nabla\Delta\phi_r(t)$ to produce smoothed reconstructed double-differenced carrier phase time histories $\nabla\Delta\phi_R(t)$.

Third, the environment evaluates the performance of the reconstruction technique by computing the empirical probability of correct integer ambiguity resolution $P_{c,\text{emp}}$ as well as the analytical upper and lower bounds $P_{c,\text{lb}}$ and $P_{c,\text{ub}}$ discussed in Sec. IV. $P_{c,\text{emp}}$ is computed as the ratio of the number of successful reconstruction attempts to the total number of attempts. A successful attempt occurs when all ambiguities are resolved successfully. In the limit, as the number of attempts approaches infinity, $P_{c,\text{emp}} \rightarrow P_c$.

A. Error Component Modeling

To create a high-fidelity simulator and to ensure near-optimal reconstruction of the simulated phase time histories, it is important to provide both the simulator and the reconstruction algorithm with accurate models for the phase variations caused by each error component of $\nabla\Delta\phi_r(t)$ detailed in (3). Some of the error components can be realistically modeled by the following flexible model:

Let $S_\phi(f)$ be the single-sided power spectral density (PSD) of some stationary phase error process $\phi(t)$. $S_\phi(f)$ can be expressed as

$$S_\phi(f) = 4 \int_0^\infty \mathcal{R}_\phi(\tau) \cos(2\pi f\tau) d\tau \quad (36)$$

where $\mathcal{R}_\phi(\tau) = \mathbb{E}[\phi(t)\phi(t+\tau)]$ is the autocorrelation function of $\phi(t)$. Let $S_\phi(f)$ be approximated by a frequency-weighted summation of five power-law parameters h_α , called h -parameters [39]:

$$S_\phi(f) = \frac{\nu_0^2}{f^2} \sum_{\alpha=-2}^2 h_\alpha f^\alpha \quad 0 < f < f_h \quad (37)$$

where ν_0 is the nominal center frequency of the phase data (e.g., the GPS L1 center frequency), in Hz, and f_h is the maximum frequency at which $S_\phi(f)$ is evaluated, typically corresponding to the Nyquist frequency of the sampled phase error process $\phi(t)$. Often only the h_{-2} component (corresponding to frequency random walk) and the h_0 component (corresponding to phase random walk) of the model are assumed to be nonzero. In this case, the five-parameter model

in (37) reduces to the two-parameter (second-order Gauss-Markov) clock error model commonly invoked in Kalman filtering [25].

Two out of the four error components of $\nabla\Delta\phi_r(t)$ can be accurately characterized by a PSD model of the form in (37): (1) double-differenced range error term $\nabla\Delta r_e(t)$ whose variations are largely induced by IMU errors, and (2) the double-differenced propagation- and multipath-induced error term $\nabla\Delta\epsilon_p(t)$. As discussed in Sec. III-A2, both of these error components are process noise and are characterized by the S_f and S_g parameters in the process noise covariance matrix \mathbf{Q} defined in (11). The relationship between S_f and S_g and the two-parameter h_{-2} and h_0 model is as follows [25]:

$$S_g = 2\pi^2 h_{-2} \quad (38)$$

$$S_f = \frac{h_0}{2}. \quad (39)$$

A third error component, the double-differenced measurement noise $\nabla\Delta v_\phi(t)$, could also be characterized by a PSD model, in particular by the h_2 parameter corresponding to white phase noise, but $\nabla\Delta v_\phi(t)$ will instead be characterized by the more-familiar carrier-to-noise ratio C/N_0 . Under this characterization, the measurement noise variance $\sigma_{\phi_k}^2$ (defined after (12)) is computed from its full nonlinear relationship to $(C/N_0)_k$, assuming a standard arctangent-type phase detector, and used to simulate the discrete-time measurement noise v_k of (12). The final term, the phase ambiguity term $\nabla\Delta\eta(t)$, need not be modeled, as ambiguities are introduced deterministically via an ‘‘ambiguity-free’’ simulation of $\nabla\Delta\phi_r(t)$ which is then aliased to between 0 and $\frac{1}{M}$ cycles.

B. Inertial Aiding

As discussed in Sec. II, an inertial measurement unit (IMU) can model the rover receiver’s changing position, enabling it to more-accurately predict its line-of-sight range to each satellite. This modeling substantially eliminates receiver-motion-induced variations from $r_e(t)$, and, consequently, from $\nabla\Delta\phi_r(t)$. It is for this reason that while the reconstruction technique can work without inertial aiding, it works much better when inertial measurements are available.

1) *Characterization of Inertial Errors:* Despite its advantages, inertially-aided motion prediction is imperfect. Noise in the IMU measurements will leave residual variations in $r_e(t)$ which enter into $\nabla\Delta\phi_r(t)$. These variations must be accurately characterized to enable optimal reconstruction. The two-parameter PSD model, discussed previously, can be used to characterize these variations. Table I lists h -parameter values that characterize the undifferenced range-error variations resulting from use of three different-quality IMUs to predict the receiver’s motion: (1) a low-end ‘‘consumer-grade’’ IMU found in consumer-electronic devices, (2) a high-end ‘‘consumer-grade’’ IMU found in commercial equipment, and (3) a ‘‘tactical-grade’’ IMU found in military equipment. These h -parameter values were determined as follows: First, white noise and bias instability values commonly used to characterize acceleration and angular velocity measurement noise were taken from the datasheet of each IMU. Second,

TABLE I
 h -PARAMETER VALUES CHARACTERIZING THE NOISE STATISTICS OF
 THREE INERTIAL MEASUREMENT UNITS

Device	h_{-2} (cycles ² -Hz)	h_0 (cycles ² /Hz)	Grade
Analog Dev. ADIS16360	5×10^{-24}	3×10^{-27}	Consumer
XSENS MTi	5×10^{-25}	3×10^{-28}	Consumer
Honeywell HG1900	5×10^{-26}	3×10^{-29}	Tactical

these values were used to simulate IMU measurement errors and, from these, IMU-specific 3-dimensional position error trajectories were generated [40]. Third, the PSD of the variations along a randomly chosen dimension was computed and a weighted least-squares solution was used to determine h -parameter-values best characterizing each PSD, as per (37). The final values listed in Table I represent average values from 20 Monte-Carlo-type simulations. While it is possible to compute the h -parameter equivalents of the white noise and bias instability parameters in isolation, it becomes much more difficult to accurately compute these parameters when the noise sources are coupled together, such as is the case in an inertial navigation system as described here.

2) *Estimation of Inertial Biases*: The foregoing two-parameter model for errors in $\nabla\Delta\phi_r(t)$ due to imperfect inertial aiding does not account for biases in the IMU's accelerometer and rate sensor measurements. Such biases could be accommodated by augmenting the real-valued state with a phase-acceleration component and, consequently, the two-parameter error model with an h_{-4} parameter. However, this turns out to be unnecessary so long as these inertial biases and the receiver orientation are periodically estimated and compensated for. Assuming that a receiver starts with a pseudorange-based initial position and an accelerometer-and-magnetometer-provided orientation, then acceleration, angular velocity, and magnetometer measurements can be integrated in an inertial navigation system (INS) [40] that approximates the receiver's change in position and orientation over an extended period of many bursts. After each extended period, e.g., roughly 20 seconds for a consumer-grade IMU, the INS-derived position must be augmented with GNSS code-phase and recently-reconstructed ambiguity-free GNSS carrier-phase measurements in a tightly-coupled INS/GNSS filter that estimates the inertial biases and receiver orientation as part of its state [41], [42]. These recent bias and orientation estimates will enable the INS to more-accurately approximate the receiver's change in position and orientation over the next extended period, allowing the reconstruction technique to accurately predict the receiver-motion-induced phase variations. This technique works so long as the inertial biases remain approximately constant over the duration of each extended period. Section VII provides a demonstration of the reconstruction technique on real data where inertial biases and receiver orientation were periodically estimated in this way.

VI. SENSITIVITY ANALYSIS

This section discusses the sensitivity of the reconstruction technique to signal structure parameters such as the burst

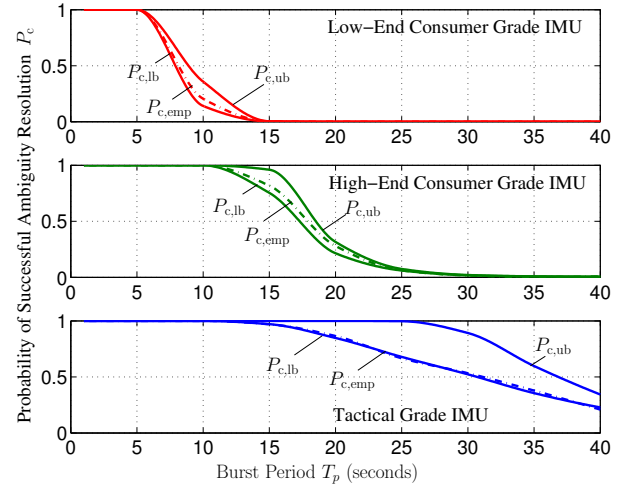


Fig. 3. Probability of successful ambiguity resolution P_c as a function of the burst period T_p and of IMU quality. The dashed traces denote the empirical estimate of P_c , $P_{c,emp}$, obtained via Monte-Carlo simulation. The solid traces denote the analytically computed upper and lower bounds $P_{c,lb}$ and $P_{c,ub}$.

period T_p and the $\frac{1}{M}$ ambiguity factor and to signal error parameters such as the carrier-to-noise ratio and the underlying IMU quality. Sensitivity is measured by computing empirical estimates and analytical bounds on the probability of correct integer ambiguity resolution P_c as a function of these parameters. The purpose of the sensitivity analysis is to discover parameter bounds beyond which the reconstruction technique will perform poorly. As it is unwieldy to test all possible combinations of parameters, testing is performed around a set of nominal parameters that model a typical low-power mobile receiver setup. In particular, during each test, sensitivity is analyzed as a function of the burst period T_p and one other parameter, namely IMU quality, the ambiguity factor $\frac{1}{M}$, or the carrier-to-noise ratio C/N_0 . During each test, the strategy will be to:

- 1) Fix the burst duration T_b to 0.05 seconds and the time duration over which reconstruction will be performed to 250 seconds.
- 2) Vary the time between bursts T_p for each test, along with one of either IMU quality, M , and C/N_0 .
- 3) Fix $M = 1$, $C/N_0 = 50$ dB-Hz, and the IMU quality to that of a low-quality consumer-grade IMU when not being varied.

This will result in three sensitivity scenarios, each of which is explored in the next three subsections.

A. Sensitivity to IMU Quality

This section illustrates the sensitivity of the reconstruction technique to the underlying IMU quality, modeled by the h_0 and h_{-2} power-law parameter values listed in Table I, and to the burst period T_p . During sensitivity testing, the IMU quality and the burst period were varied while the other important parameters were held constant at the values discussed earlier.

Fig. 3 illustrates the sensitivity results. The empirical probability of successful ambiguity resolution $P_{c,emp}$ was computed via Monte-Carlo simulation and is represented by the dashed




Fig. 4. Probability of successful ambiguity resolution P_c as a function of the burst period T_p and the ambiguity factor M . The dashed traces denote an upper bound on P_c , $P_{c,ub}$, while the solid traces denote a lower bound, $P_{c,lb}$.

trace. The lower and upper bounds $P_{c,lb}$ and $P_{c,ub}$ were computed analytically (see section IV) and are represented by the solid traces. The waterfall structure of each trace indicates a breakdown point in successful ambiguity resolution. Each subplot represents a different underlying IMU quality. It is evident that the higher the IMU quality, the larger the burst period T_p that the reconstruction technique can sustain before a breakdown occurs. A higher quality IMU allows the reconstruction technique to more accurately predict the underlying phase trajectory between bursts, making it easier to resolve the phase ambiguity at the beginning of each burst.

B. Sensitivity to the Ambiguity Factor

Fig. 4 plots $P_{c,lb}$ and $P_{c,ub}$ as a function of T_p for two different values of the ambiguity factor M . (For visual clarity, empirical results, which always lie close to $P_{c,ub}$, were not plotted.) As shown, a lower M value allows for a larger burst period T_p before a breakdown in P_c occurs. This is as might be expected: all else equal, integer-cycle ambiguities are easier to resolve than fractional-cycle ambiguities. This implies that a GNSS receiver with *a priori* knowledge of the binary navigation data symbols and an approximation of its position and time (to within a fraction of a data symbol interval) (in which case $M = 1$) has the ability to be more power efficient than a receiver with no such knowledge ($M = 2$) by extending its burst period while maintaining the same probability of successful ambiguity resolution.

C. Sensitivity to the Carrier-to-noise Ratio

Fig. 5 plots lower bounds on P_c as a function of the burst period for five different carrier-to-noise ratios. Note that for visual clarity only the lower bounds were plotted. As illustrated, a higher C/N_0 tolerates a longer burst period before a breakdown in P_c occurs. This is because for lower C/N_0 values the measurement noise variations v_k imparted by the receiver's front end become a larger share of the overall variations within y_k [see (12)]. This makes it difficult for the reconstruction technique to separate these variations from the variations due to the real- and integer-valued state components \mathbf{x}_k and \mathbf{n}_k , leading to a decrease in the probability of correct ambiguity resolution.




Fig. 5. Lower bounds on the probability of successful channel reconstruction as a function of the burst period T_p and the carrier-to-noise ratio C/N_0 (in dB-Hz).

VII. DEMONSTRATION OF A CDGNSS SOLUTION ON RECONSTRUCTED DATA

This section provides a demonstration of the reconstruction technique applied to real data collected by a reference and rover GNSS receiver. Each receiver was running a version of the GRID software [43]–[45] and is capable of capturing GNSS signal code- and carrier-phase data. Additionally, the rover contained an Xsens MTi IMU capable of providing linear acceleration measurements and attitude estimates derived from an internal filter.

A. Data Collection, Modification, and Processing

The demonstration was carried out as follows. First, GNSS signal code- and carrier-phase data were collected simultaneously and continuously by both receivers while IMU-provided acceleration measurements and attitude estimates were additionally collected by the rover. Two minutes of data were collected. During this time the rover receiver was moved about while the reference receiver remained stationary. The rover's trajectory was that of a pedestrian moving at a walking-pace holding the receiver at an approximately fixed pitch and roll angle, but allowing changes in yaw. Second, the data collected by the rover receiver were digitally modified in two ways to simulate collection by a power-constrained receiver: (1) discrete measurement intervals were selected from the continuously-recorded GNSS code- and carrier-phase data and (2) the carrier-phase data were aliased to between 0 and 1 cycle. Third, the intermittent carrier-phase time histories from the two receivers were differenced to form 7 double-differenced carrier-phase time histories from 8 GNSS satellite signals present in the recorded data. Fourth, biases in the IMU acceleration measurements were estimated once every 60 seconds via an INS/GNSS filtering technique similar to that described in Sec. V-B2. Fifth, the bias estimates, acceleration measurements, and IMU-derived attitude estimates were incorporated into an INS to approximate the receiver trajectory and remove the motion-induced variations from each double-differenced carrier-phase time history, forming double-differenced residual carrier-phase time histories. Sixth, the reconstruction technique was applied to each double-differenced residual carrier-phase time history. Finally, the reconstructed time histories (along with the intermittent code-phase measurements) were passed off to a standard CDGNSS positioning algorithm which computed a centimeter-accurate positioning solution for the rover receiver.




Fig. 6. Error in the positioning solution provided by the CDGNSS algorithm whose inputs are phase trajectories from two different reconstruction outcomes: (1) a successful attempt (lower blue trace) defined when all phase ambiguities were resolved correctly for all 7 of the reconstructed phase time histories involved in the positioning solution and (2) a failed attempt (upper red trace) defined when one or more phase ambiguities were resolved incorrectly for one or more of the reconstructed phase time histories.

B. Results

Fig. 6 illustrates the accuracy of the CDGNSS-based positioning solution under two circumstances: (1) a scenario in which all phase ambiguities from the 7 reconstructed time histories were resolved correctly, and (2) a scenario in which one or more phase ambiguities were resolved incorrectly. In the first scenario, carrier-phase measurements were provided to the reconstruction algorithm with a burst length of $T_b = 0.05$ seconds and a burst period of $T_p = 1$ seconds, which corresponds to a 5% duty cycle. For the second scenario, $T_b = 0.05$ seconds and $T_p = 2$ seconds, corresponding to a 2.5% duty cycle. The time duration of the dataset for each scenario was approximately 120 seconds. In both scenarios, the error, in meters, from the true position is computed and plotted as a function of time. The ground truth trajectory is obtained by separately computing a CDGNSS solution using the unmodified, ambiguity-free, continuous phase time histories originally captured by the receivers. As the lower trace illustrates, when the reconstruction algorithm resolves the phase ambiguities correctly, the positioning error is very small (less than 1.5 cm). The small error is primarily due to the inability of the reconstruction algorithm to perfectly reconstruct the variations in the residual carrier phase between measurement bursts. In contrast, as the upper trace illustrates, incorrectly resolved ambiguities lead to significant positioning errors well in excess of the accuracy potential of the CDGNSS algorithm. Large jumps in the positioning error denote an incorrectly resolved phase ambiguity at that time index.

In this demonstration, the availability of the ground truth trajectory enables the generation of a simple metric, i.e., baseline error, to indicate when reconstruction has failed. However, a system in the field will not have a ground truth trajectory with which the baseline error can be generated. In these cases, the system can compute the upper and lower bounds on the probability of successful reconstruction, as introduced in Sec. IV, and use one or both of these as an indicator as to when reconstruction may have failed, e.g., when $P_{c,lb}$ is not above a predetermined threshold for each double-differenced reconstructed phase time history, the resulting CDGNSS solution can be presumed inaccurate.

It should be noted that while the results shown here are promising, these results reveal the performance of the

TABLE II
GNSS CHIP POWER CONSUMPTION

Mfr.	Chip	Measurements Provided	Power (mW)
Broadcom	BCM4751	duty-cycled code phase	13
u-blox	NEO-6P	cont. code & carrier phase	117

reconstruction technique on only one set of collected data. For a more in-depth and direct analysis of the technique's performance, see Sec. VI, where reconstruction was performed on hundreds of sets of simulated data and compared against analytical performance bounds.

VIII. POWER CONSUMPTION ANALYSIS

This section provides an analysis of the power consumption of the duty-cycled measurement framework enabled by the carrier-phase reconstruction technique as compared to a framework requiring the continuous tracking of GNSS signal carrier phase.

A. Low-Power GNSS Chips

Table II lists two GNSS chips and their average power consumption. The Broadcom chip, used in many mobile devices, computes a receiver's position using only the tracked code-phase of each GNSS signal. It achieves an impressively low power draw of 13 mW by aggressively duty cycling its code-phase measurements [46]. Unlike carrier-phase measurements, code-phase measurements do not suffer from ambiguity problems when duty-cycled. Code-phase measurements, however, can only be used to compute a pseudorange-based position solution, which is much less accurate than a carrier-phase-based CDGNSS solution. The other two chips provide both code- and carrier-phase measurement outputs. To provide ambiguity-free carrier-phase measurements, these chips continuously track each GNSS signal, drastically increasing their power consumption compared to the duty-cycling Broadcom chip. The NEO-7M, a variant of the NEO-6P, has a low-power duty-cycled mode with a power draw of 14 mW [47], similar to that of the Broadcom chip. The NEO-7M, however, does not provide access its carrier-phase measurements. This paper's authors conjecture that both u-blox and Broadcom believe it is fruitless to provide carrier-phase data fraught with ambiguities and so do not provide access to these measurements in chips with duty-cycled tracking modes.

It is important to note that unlike code and carrier phase, whose measurements are duty-cycled in this low-power framework, acceleration and angular velocity must be continuously measured by an IMU such that they can be used in predicting receiver motion between phase measurement bursts (see Sec. V-B). Fortunately, there exist low-power chip-scale IMUs which consume power on the order of 10-20 mW [48], much less than the state-of-the-art u-blox chips that output continuous code and carrier phase measurements.

B. Power Consumption of the Reconstruction Algorithm

The reconstruction technique outlined in this paper relaxes the continuous tracking requirement for GNSS chips that pro-

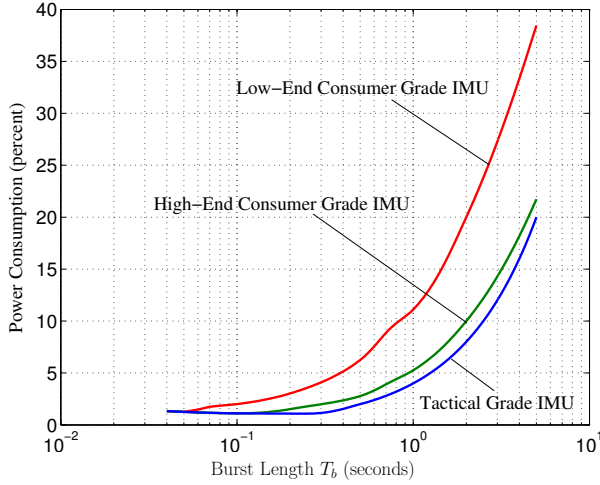


Fig. 7. Minimum achievable relative power consumption of a duty-cycled measurement and phase reconstruction framework as compared to a continuous measurement framework as function of the burst length T_b .

vide carrier-phase measurements. Although duty-cycled measurements will contain phase ambiguities, the reconstruction technique can be applied to these measurements, enabling, under favorable circumstances, an ambiguity-free continuous time phase history to be accurately reconstructed.

The minimum power consumed by a receiver duty cycling its carrier-phase measurements can be described as a percentage of the power consumed by a receiver continuously tracking the carrier phase:

$$\text{Power Consumption (\%)} = \frac{T_b}{T_p} \times 100. \quad (40)$$

This power consumption metric is a minimum as it considers only the power that will be saved by measurement duty-cycling. It ignores the overhead imposed by the reconstruction algorithm and the overhead involved in turning on and off the receiver components associated with sampling and digitizing the signal. This will be discussed later.

When evaluating the power consumption using (40), it is logical to check that the chosen combination of T_b and T_p will result in a successful reconstruction. One way to do this is to look at the lower bound of the probability of successful ambiguity resolution $P_{c,lb}$ and determine if it is above a certain threshold, e.g., 99.99%. Obviously, as $\frac{T_b}{T_p} \rightarrow 1$, $P_{c,lb} \rightarrow 1$, but the power consumption as denoted by (40) will also approach 100%, saving little power. Accordingly, there exists a tradeoff between keeping $P_{c,lb}$ close to 1 and minimizing the power consumption.

Fig. 7 provides an empirical analysis of the minimum achievable power consumption as a function of the burst length T_b for performing reconstruction on a simulated double-differenced residual GNSS signal. The signal was simulated with a single-sided PSD defined by (37) with the h_{-2} and h_0 power-law parameters varied according to the values in Table I, representing the usage of a low-end consumer-, high-end consumer-, and tactical-grade IMU. Additionally white phase noise was added to the signal to simulate front-end noise representative of a receiver $C/N_0 = 50$ dB-Hz. The

simulated dataset duration was 250 seconds. For simplicity, it was assumed that the contribution to $\nabla\Delta\phi_r(t)$ by unmodeled propagation and multipath effects was small in comparison to the other noise sources, and as such, the h -parameters characterizing the PSD of this noise source were set to 0.

In computing the power consumption values in Fig. 7, for each T_b the largest value of T_p was chosen such that $P_{c,lb}$ remained above .9999. The power consumption was then computed using (40) and plotted. From the figure, one can make two interesting observations. First, the minimal power consumption is attained when the burst length is very small. This implies that, to save power, it is beneficial for a GNSS receiver making duty-cycled phase measurements to use relatively short burst lengths and short burst periods rather than long burst lengths and long burst periods. Second, at shorter burst lengths (and burst periods), the quality of the underlying IMU has a smaller impact on reducing power consumption. This implies that a high-quality IMU can be forgone in favor of a lower-quality IMU as long as the burst length and burst period are reduced enough to achieve the desired power consumption. This is an important result as many mobile handheld devices come with consumer-grade IMUs (or separate consumer-grade accelerometers and rate sensors).

C. Power Consumption Overhead

Although measurement duty-cycling enables a large reduction in power consumption at the rover receiver, it would be negligent to assume that the added computational complexity required by the reconstruction algorithm to reconstruct a continuous phase-time history from the duty-cycled measurements does not consume any power. Much computational complexity is added by the integer least-squares ambiguity resolution algorithm, and this complexity increases exponentially with the number of integer ambiguities to resolve [18]. Fortunately, unlike carrier-phase measurement, carrier-phase reconstruction need not be executed at the rover receiver – the discrete phase measurements can be relayed to the cloud for reconstruction. Furthermore, because the CDGNSS algorithm requires a double-differencing of carrier-phase measurements from the rover and a reference station, the rover would in any case be required to offload its measurements to the network, since receiving reference station measurements from the network and computing the CDGNSS solution locally would likely consume more power than relaying local phase measurements to the network.

The additional power required to transmit data in an LTE network versus receiving it is about 400 mW per 1 Mbps [49]. The average rate at which the rover must send duty-cycled phase measurements to the cloud is $f_s \cdot \frac{T_b}{T_p}$ measurements per second, which, for aggressive duty cycling (e.g., $\frac{T_b}{T_p} = \frac{1}{10}$) and a modest sampling rate (e.g., $f_s = 50$ Hz), would result in an average code- and carrier-phase measurement rate of 5 samples per second per signal tracked. Given 10 signals tracked and 32 bits allocated per sample, the average transmission data rate is about 1.6 kbps, or approximately 0.6 mW of added power to transmit the carrier phase data rather than receive

it. Although this number does not account for the cost of transmitting the code-phase and IMU measurements, the total rover power consumption for a cloud-based CDGNSS solution will likely be far below the power consumption needed to perform a local CDGNSS solution. Thus, under the current framework, the rover will relay a batch of code- and carrier-phase measurements along with IMU measurements to the network for processing. The network will perform phase differencing, IMU bias correction, reconstruction, and CDGNSS processing and then relay back to the rover the latter's precise position time history over the batch interval.

In addition to the overhead involved in transmitting data over the network, the power consumption model in (40) also ignores the power overhead associated with switching on and off the receiver components involved in sampling and digitizing the GNSS signals. Such overhead is inversely proportional to T_p ; as T_p is decreased, there will come a point when the increase in power consumption due to this overhead will outweigh the additional reduction in power consumption from measurement duty-cycling. This "break-even" point places a lower bound on T_p (and its associated T_b), below which the power consumption will no longer decrease. For the power consumption analysis described in Sec. VIII-B whose results are displayed in Fig. 7, T_p is always larger than 1 second. Such values of T_p are assumed to be well above this break-even point. As such, the power overhead due to switching is assumed to be negligible and is not modeled by (40).

D. Avoiding Phase-Locked Loop Transients using Batch Estimation

If the rover receiver's phase-locked loop (PLL) re-synchronizes its local carrier replica with the incoming carrier phase at the beginning of each measurement burst, then this will result in short phase transients during re-synchronization [50]. These transients are manageable so long as they settle prior to the end of the measurement burst, i.e., the convergence time is less than T_b . However, only the phase measurements taken after this settling period can be used during reconstruction. Rather than forgo the information contained in these transients, which could benefit reconstruction, a different methodology avoids them altogether. Instead of attempting to track the incoming carrier-phase using a traditional PLL, the rover receiver can generate a model line of sight trajectory to each GNSS satellite tracked and employ batch estimation on raw correlation outputs to measure carrier phase, code phase, and Doppler with no phase pull-in transient. This model-trajectory methodology is standard in so-called vector tracking, where traditional tracking loops are replaced by a navigation filter that provides prior knowledge of receiver position and velocity to the local replica generators [26], [51]. The only requirement is that the model-trajectory be accurate enough that the difference between the received carrier phase and the carrier-phase predicted by the model-trajectory does not drift by more than $\frac{1}{2}$ cycle during the sampling interval $\frac{1}{f_s}$.

IX. CONCLUSIONS

A technique for reconstructing a continuous carrier-phase time history from intermittent GNSS carrier-phase measure-

ments has been developed. The technique combines an integer least-squares method for estimating the phase ambiguity that arises at the beginning of each measurement burst with a Kalman filter and smoother that correct for these ambiguities and "stitch" the bursts together.

A Monte-Carlo-type simulation and test environment has been built in MATLAB to simulate the intermittent GNSS phase measurements, implement the phase reconstruction technique, and analyze the sensitivity of the technique to determine the parameter space within which successful reconstruction is possible. Theoretical bounds predicting the probability of successful reconstruction were compared to empirical results from the Monte-Carlo simulations.

Simulation results indicate that successful carrier-phase reconstruction is strongly dependent on the burst period, the carrier-to-noise ratio, the ambiguity factor, and the quality of the underlying inertial measurement unit employed by the receiver. A demonstration on real data shows that the reconstruction technique can successfully reconstruct carrier phase measurements made at a 5% duty cycle by a GNSS receiver containing a consumer-grade IMU and receiving GNSS signals with a carrier-to-noise ratio of 50 dB-Hz. The reconstruction technique assumes the use of special batch tracking techniques to avoid PLL transients at the start of each burst. Furthermore, an analytical power analysis indicates that the reconstruction technique can permit potential power savings in excess of 95% for a receiver duty-cycling its carrier phase measurements when compared against a receiver continuously tracking the incoming carrier phase. These results suggest that the reconstruction technique could act as an enabler for high-precision positioning in energy-limited mobile devices.

REFERENCES

- [1] M. S. Braasch and A. Van Dierendonck, "GPS receiver architectures and measurements," *Proceedings of the IEEE*, vol. 87, no. 1, pp. 48–87, Jan. 1999.
- [2] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*, revised second ed. Lincoln, Massachusetts: Ganga-Jumana Press, 2012.
- [3] C. Counselman and S. Gourevitch, "Miniature interferometer terminals for earth surveying: ambiguity and multipath with global positioning system," *IEEE Transactions on Geoscience and Remote Sensing*, no. 4, pp. 244–252, 1981.
- [4] B. W. Parkinson and P. K. Enge, *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, vol. 2, ch. 1: Differential GPS, pp. 3–50.
- [5] P. Teunissen, P. De Jonge, and C. Tiberius, "Performance of the LAMBDA method for fast GPS ambiguity resolution," *Navigation, Journal of the Institute of Navigation*, vol. 44, no. 3, pp. 373–383, 1997.
- [6] M. Kirkko-Jaakkola, J. Traugott, D. Odijk, J. Collin, G. Sachs, and F. Holzapfel, "A RAIM approach to GNSS outlier and cycle slip detection using 11 carrier phase time-differences," in *Signal Processing Systems, 2009. SiPS 2009. IEEE Workshop on*. IEEE, 2009, pp. 273–278.
- [7] J. How, N. Pohlman, and C. Park, "GPS estimation algorithms for precise velocity, slip and race-track position measurements," in *SAE Conference Proceedings*. Society of Automotive Engineers, 2002, pp. 497–504.
- [8] P. Teunissen, P. de Jonge, D. Odijk, and C. Tiberius, "Fast ambiguity resolution in network mode," *Proceedings IAG Scientific Assembly*, pp. 3–9, 1997.
- [9] M. Pratt, B. Burke, and P. Misra, "Single-epoch integer ambiguity resolution with GPS-GLONASS L1 data," in *Proceedings of the 53rd Annual Meeting of The Institute of Navigation*, Albuquerque, NM, 1997, pp. 691–699.

- [10] S. Verhagen, P. J. Teunissen, and D. Odijk, "The future of single-frequency integer ambiguity resolution," in *VII Hotine-Marussi Symposium on Mathematical Geodesy*. Springer, 2012, pp. 33–38.
- [11] K. Q. Chiang, M. L. Psiaki, S. P. Powell, R. J. Miceli, and B. W. O'Hanlon, "GPS-based attitude determination for a spinning rocket," in *Proceedings of the ION GNSS Meeting*. Nashville, Tennessee: Institute of Navigation, 2012.
- [12] D. B. Cox and J. D. Brading, "Integration of lambda ambiguity resolution with kalman filter for relative navigation of spacecraft," *Navigation, Journal of the Institute of Navigation*, vol. 47, no. 3, pp. 205–210, 2000.
- [13] M. Psiaki, "Kalman filtering and smoothing to estimate real-valued states and integer constants," *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 5, pp. 1404–1417, Sept.–Oct. 2010.
- [14] S. Mohiuddin and M. Psiaki, "Carrier-phase differential Global Positioning System navigation filter for high-altitude spacecraft," *Journal of Guidance, Control, and Dynamics*, vol. 31, no. 4, pp. 801–814, 2008.
- [15] K. M. Pesyna, Z. M. Kassas, and T. E. Humphreys, "Constructing a continuous phase time history from TDMA signals for opportunistic navigation," in *Proceedings of the IEEE/ION PLANS Meeting*, April 2012, pp. 1209–1220.
- [16] P. De Jonge and C. Tiberius, "The LAMBDA method for integer ambiguity estimation: implementation aspects," *Publications of the Delft Computing Centre, LGR-Series*, 1996.
- [17] A. Hassibi and S. Boyd, "Integer parameter estimation in linear models with applications to GPS," *IEEE Transactions on Signal Processing*, vol. 46, no. 11, pp. 2938–2952, 1998.
- [18] B. Hassibi and H. Vikalo, "On the sphere-decoding algorithm I. Expected complexity," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 2806–2818, 2005.
- [19] J. Jaldén and B. Ottersten, "On the complexity of sphere decoding in digital communications," *IEEE Transactions on Signal Processing*, vol. 53, no. 4, pp. 1474–1484, 2005.
- [20] M. Pohst, "On the computation of lattice vectors of minimal length, successive minima and reduced bases with applications," *ACM SIGSAM Bulletin*, vol. 15, no. 1, pp. 37–44, 1981.
- [21] M. Psiaki and S. Mohiuddin, "Modeling, analysis, and simulation of GPS carrier phase for spacecraft relative navigation," *Journal of Guidance Control and Dynamics*, vol. 30, no. 6, p. 1628, 2007.
- [22] T. E. Humphreys, M. L. Psiaki, and P. M. Kintner, Jr., "Modeling the effects of ionospheric scintillation on GPS carrier phase tracking," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 4, pp. 1624–1637, Oct. 2010.
- [23] G. J. Bierman, *Factorization Methods for Discrete Sequential Estimation*. New York: Academic Press, 1977.
- [24] G. Bierman, M. Belzer, J. Vandergraft, and D. Porter, "Maximum likelihood estimation using square root information filters," *IEEE Transactions on Automatic Control*, vol. 35, no. 12, pp. 1293–1298, 1990.
- [25] R. G. Brown and P. Y. Hwang, *Introduction to Random Signals and Applied Kalman Filtering*. Wiley, 1997.
- [26] M. L. Psiaki and H. Jung, "Extended Kalman filter methods for tracking weak GPS signals," in *Proceedings of the ION GPS Meeting*. Portland, Oregon: Institute of Navigation, 2002, pp. 2539–2553.
- [27] J. T. Curran, G. Lachapelle, and C. C. Murphy, "Improving the design of frequency lock loops for GNSS receivers," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 1, pp. 850–868, 2012.
- [28] K. M. Pesyna Jr. and T. E. Humphreys, "Cost analysis of square root information filtering and smoothing with a mixed real-integer state," *Whitepaper*, 2013, <http://radionavlab.ae.utexas.edu/reconstruction/>.
- [29] E. Agrell, T. Eriksson, A. Vardy, and K. Zeger, "Closest point search in lattices," *IEEE Transactions on Information Theory*, vol. 48, no. 8, pp. 2201–2214, 2002.
- [30] P. Teunissen, "An optimality property of the integer least-squares estimator," *Journal of Geodesy*, vol. 73, no. 11, pp. 587–593, 1999.
- [31] A. Lenstra, H. Lenstra, and L. Lovász, "Factoring polynomials with rational coefficients," *Mathematische Annalen*, vol. 261, no. 4, pp. 515–534, 1982.
- [32] M. Psiaki and S. Mohiuddin, "Global positioning system integer ambiguity resolution using factorized least-squares techniques," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 2, pp. 346–356, March–April 2007.
- [33] X.-W. Chang and T. Zhou, "MILES: MATLAB package for solving Mixed Integer Least Squares problems," *GPS Solutions*, vol. 11, no. 4, pp. 289–294, 2007.
- [34] P. Xu, "Voronoi cells, probabilistic bounds, and hypothesis testing in mixed integer linear models," *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 3122–3138, 2006.
- [35] P. Teunissen, "Success probability of integer GPS ambiguity rounding and bootstrapping," *Journal of Geodesy*, vol. 72, no. 10, pp. 606–612, 1998.
- [36] —, "Gnss ambiguity bootstrapping: theory and application," in *Proceedings of International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation*, 2001, pp. 246–254.
- [37] S. Verhagen, "On the reliability of integer ambiguity resolution," *Navigation, Journal of the Institute of Navigation*, vol. 52, no. 2, pp. 99–110, 2005.
- [38] S. Verhagen, B. Li, and P. J. Teunissen, "Ps-lambda: Ambiguity success rate evaluation software for interferometric applications," *Computers & Geosciences*, vol. 54, pp. 361–376, 2013.
- [39] A. Thompson, J. Moran, and G. Swenson, *Interferometry and Synthesis in Radio Astronomy*. Wiley, 2001, ch. 9: Very-Long-Baseline Interferometry, pp. 304–382.
- [40] O. J. Woodman, "An introduction to inertial navigation," *University of Cambridge, Computer Laboratory, Tech. Rep. UCAMCL-TR-696*, 2007.
- [41] D. Shepard, "Fusion of carrier-phase differential GPS, bundle-adjustment-based visual slam, and inertial navigation for precisely and globally-registered augmented reality," Master's thesis, The University of Texas at Austin, May 2013.
- [42] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley and Sons, 2001.
- [43] T. E. Humphreys, M. L. Psiaki, P. M. Kitner, and B. M. Ledvina, "GNSS receiver implementation on a DSP: Status, challenges, and prospects," in *Proceedings of the ION GNSS Meeting*. Fort Worth, TX: The Institute of Navigation, 2006.
- [44] T. E. Humphreys, J. Bhatti, T. Pany, B. Ledvina, and B. O'Hanlon, "Exploiting multicore technology in software-defined GNSS receivers," in *Proceedings of the ION GNSS Meeting*. Savannah, GA: Institute of Navigation, 2009.
- [45] B. O'Hanlon, M. Psiaki, S. Powell, J. Bhatti, T. E. Humphreys, G. Crowley, and G. Bust, "CASES: A smart, compact GPS software receiver for space weather monitoring," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011.
- [46] *Product Brief: BCM4750 AGPS chip*, Broadcom, 2007.
- [47] *Datasheet: NEO-7 GPS/GNSS Module*, u-Blox, 2013.
- [48] *Datasheet: ST iNEMO inertial module*, STMicroelectronics, 2012.
- [49] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4g lte networks," in *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 2012, pp. 225–238.
- [50] W. Mao, H. Tsao, and F. Chang, "Intelligent GPS receiver for robust carrier phase tracking in kinematic environments," in *IEEE Proceedings, Radar, Sonar and Navigation*, vol. 151, no. 3. IET, 2004, pp. 171–180.
- [51] M. Lashley, D. M. Bevly, and J. Y. Hung, "Performance analysis of vector tracking algorithms for weak GPS signals in high dynamics," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 3, no. 4, pp. 661–673, 2009.



Kenneth M. Pesyna, Jr. (S'09) received a B.S. degree from Purdue University in 2009 and a M.S. degree from The University of Texas at Austin in 2011, both in electrical engineering. He is currently a Ph.D. candidate at The University of Texas at Austin where he is a member of The Radionavigation Laboratory, The Wireless Systems Innovation Laboratory, and The Wireless Networking and Communications Group. He won the best student paper award at IEEE/ION PLANS 2012 and a best presentation award at ION GNSS+ 2013. His research interests

are in estimation and navigation and his recent focus is in advanced techniques for precise low-power mobile positioning.



Zaher (Zak) M. Kassas (S'98-M'08-SM'011) received a B.E. with Honors in Electrical Engineering from The Lebanese American University, a M.S. in Electrical and Computer Engineering from The Ohio State University, and a M.S.E. in Aerospace Engineering from The University of Texas at Austin. He is currently a Ph.D. candidate at The University of Texas at Austin. From 2004 to 2010 he was a research and development engineer with the Control Design and Dynamical Systems Simulation Group at National Instruments Corp. From 2008 to 2011 he

was an adjunct professor at Texas State University. His research interests include estimation, navigation, autonomous vehicles, control systems, and intelligent transportation systems.



Todd E. Humphreys Todd E. Humphreys is an assistant professor in the department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, and Director of the UT Radionavigation Laboratory. He received a B.S. and M.S. in Electrical and Computer Engineering from Utah State University and a Ph.D. in Aerospace Engineering from Cornell University. He specializes in applying optimal estimation and signal processing techniques to problems in radionavigation. His recent focus is on radionavigation robustness and

security.



Robert W. Heath, Jr. (S'96 - M'01 - SM'06 - F'11) received the B.S. and M.S. degrees from the University of Virginia, Charlottesville, VA, in 1996 and 1997 respectively, and the Ph.D. from Stanford University, Stanford, CA, in 2002, all in electrical engineering. From 1998 to 2001, he was a Senior Member of the Technical Staff then a Senior Consultant at Iospan Wireless Inc, San Jose, CA where he worked on the design and implementation of the physical and link layers of the first commercial MIMO-OFDM communication system. Since January

2002, he has been with the Department of Electrical and Computer Engineering at The University of Texas at Austin where he is a Cullen Trust for Higher Education Endowed Professor, and is Director of the Wireless Networking and Communications Group. He is also President and CEO of MIMO Wireless Inc. and Chief Innovation Officer at Kuma Signals LLC. His research interests include several aspects of wireless communication and signal processing: limited feedback techniques, multihop networking, multiuser and multicell MIMO, interference alignment, adaptive video transmission, manifold signal processing, and millimeter wave communication techniques.

Dr. Heath has been an Editor for the IEEE Transactions on Communication, an Associate Editor for the IEEE Transactions on Vehicular Technology, lead guest editor for an IEEE Journal on Selected Areas in Communications special issue on limited feedback communication, and lead guest editor for an IEEE Journal on Selected Topics in Signal Processing special issue on Heterogenous Networks. He currently serves on the steering committee for the IEEE Transactions on Wireless Communications. He was a member of the Signal Processing for Communications Technical Committee in the IEEE Signal Processing Society and is a former Chair of the IEEE COMSOC Communications Technical Theory Committee. He was a technical co-chair for the 2007 Fall Vehicular Technology Conference, general chair of the 2008 Communication Theory Workshop, general co-chair, technical co-chair and co-organizer of the 2009 IEEE Signal Processing for Wireless Communications Workshop, local co-organizer for the 2009 IEEE CAMSAP Conference, technical co-chair for the 2010 IEEE International Symposium on Information Theory, the technical chair for the 2011 Asilomar Conference on Signals, Systems, and Computers, general chair for the 2013 Asilomar Conference on Signals, Systems, and Computers, founding general co-chair for the 2013 IEEE GlobalSIP conference, and is technical co-chair for the 2014 IEEE GLOBECOM conference.

Dr. Heath was a co-author of best student paper awards at IEEE VTC 2006 Spring, WPMC 2006, IEEE GLOBECOM 2006, IEEE VTC 2007 Spring, and IEEE RWS 2009, as well as co-recipient of the Grand Prize in the 2008 WinTech WinCool Demo Contest. He was co-recipient of the 2010 and 2013 EURASIP Journal on Wireless Communications and Networking best paper awards and the 2012 Signal Processing Magazine best paper award. He was a 2003 Frontiers in Education New Faculty Fellow. He is also a licensed Amateur Radio Operator and is a registered Professional Engineer in Texas.