Centimeter Positioning with a Smartphone-Quality GNSS Antenna

Kenneth M. Pesyna, Jr., Robert W. Heath, Jr., and Todd E. Humphreys
The University of Texas at Austin

BIOGRAPHIES

Kenneth M. Pesyna Jr. is a Ph.D. candidate in the Department of Electrical and Computer Engineering at The University of Texas at Austin. He is a member of the University of Texas Radionavigation Laboratory and the Wireless Networking and Communications Group. His research interests are in estimation and GNSS signal processing with a focus on precise and low-power mobile positioning.

Robert W. Heath Jr. is a Cullen Trust Endowed Professor in the department of Electrical and Computer Engineering at The University of Texas at Austin, and Director of the Wireless Networking and Communications Group. He received his Ph.D. in Electrical Engineering from Stanford University. He specializes in signal processing for communications with recent emphasis on millimeter wave communication systems.

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.

ABSTRACT

This paper demonstrates for the first time that centimeter-accurate positioning is possible based on data sampled from a smartphone-quality Global Navigation Satellite System (GNSS) antenna. Centimeter-accurate smartphone positioning will enable a host of new applications such as globally-registered fiduciary-marker-free augmented reality and location-based contextual advertising, both of which have been hampered by the several-meter-level errors in traditional GNSS positioning. An empirical analysis of data collected from a smartphone-grade GNSS antenna reveals the antenna to be the primary impediment to fast and reliable resolution of the integer ambiguities which arise when solving for a centimeter-accurate carrier-phase differential position. The antenna's poor multipath suppression and irregular gain pattern result in large time-correlated phase errors which significantly increase the time to integer ambiguity resolution as compared to even a low-quality stand-alone patch antenna. The time to integer resolution—and to a centimeter-accurate fix—is significantly reduced when more GNSS signals are tracked or when the smartphone experiences gentle wavelength-scale random motion.

I. Introduction

GNSS chipsets are now ubiquitous in smartphones and tablets. Yet the underlying positioning accuracy of these consumer-grade GNSS receivers has stagnated over the past decade. The latest clock, orbit, and atmospheric models have improved ranging accuracy to a meter or so [1], leaving receiver-dependent multipath- and front-end-noise-induced variations as the dominant sources of error in current consumer devices [2]. Under good multipath conditions, 2-to-3-meter-accurate positioning is typical; under adverse multipath, accuracy degrades to 10 meters or worse.

Yet outside the mainstream of consumer GNSS receivers, centimeter—even millimeter—accurate GNSS receivers can be found. These high-precision receivers are used routinely in geodesy, agriculture, and surveying. Their exquisite accuracy results from replacing standard code-phase positioning techniques with carrier-phase differential GNSS (CDGNSS) techniques [3], [4]. Currently, the primary impediment to performing CDGNSS positioning on smartphones lies not in the commodity GNSS chipset, which actually outperforms survey-grade chipsets in some respects [5], but in the antenna, whose chief failing is its poor multipath suppression. Multipath, caused by direct signals reflecting off the ground and nearby objects, induces centimeter-level phase measurement errors, which, for static receivers, have decorrelation times of hundreds of seconds. The large size and strong time correlation of these errors significantly increases the initialization period—the so-called time-to-ambiguity-resolution (TAR)—of GNSS receivers employing CDGNSS to obtain centimeter-level positioning accuracy [6], [7].

Prior work on centimeter-accurate positioning with low-cost mobile devices has focused on external devices, or "pucks," which contain a GNSS antenna and chipset. These devices interface with the smartphone via Bluetooth or a wired connection [8], [9], [10]. Such solutions, which enjoy the better sensitivity and multipath suppression offered by their comparatively large, high-quality GNSS

antennas, do not provide insight into the feasibility of CDGNSS on a stand-alone smartphone platform.

This paper makes three primary contributions. First, it demonstrates for the first time that centimeter-accurate CDGNSS positioning is indeed possible based on data sampled from a smartphone-quality GNSS antenna. This result has far-reaching significance for precise mass market positioning. Second, it offers an empirical analysis of the average gain and carrier phase multipath error susceptibility of smartphone-grade GNSS antennas. Third, it demonstrates that, for low-quality GNSS antennas such as those in smartphones, wavelength-scale random antenna motion substantially improves the time to integer ambiguity resolution.

This paper focuses on single-frequency CDGNSS rather than multiple-frequency CDGNSS or other carrier-phase-based techniques, such as precise-point positioning (PPP), for three reasons. First, virtually all smartphones are equipped with single-frequency GNSS antennas tuned to the L1 band centered at 1575.42 MHz, and single-frequency CDGNSS will likely forever remain the cheapest option [11]. Second, as compared to PPP, CDGNSS converges much faster to centimeter accuracy [12], which will be important for impatient smartphone users. Finally, as centimeter-accurate GNSS moves into the mass market, GNSS reference stations will proliferate so that the vast majority of users can expect to be within a few kilometers of one [13]. In this so-called short baseline regime, the differential ionospheric delay between the reference and mobile receivers becomes insignificant, obviating differential delay estimation via multi-frequency measurements [14]. Of course, the additional signal measurements produced by multiple-frequency receivers would lead to faster convergence times and improved robustness, but for many applications, single-frequency measurements will be adequate.

II. Test Architecture

This section describes the test architecture used to (1) collect data from a smartphone-grade antenna and higher-quality antennas, (2) process these data through a software-defined GNSS receiver, and (3) compute a CDGNSS solution on the basis of the carrier phase measurements output by the GNSS receiver.

Fig. 1 illustrates the test architecture as configured for an in situ study of a smartphone-grade GNSS antenna. The architecture has been designed such that the antenna is left undisturbed within the phone; data are collected by tapping off the analog signal immediately after the phone's internal bandpass filter and low-noise amplifier. This analog signal is directed to an external radio frequency (RF) front-end and GNSS receiver. Use of an external receiver permits well-defined GNSS signal processing unencumbered


Fig. 1. Test architecture designed for an in situ study of a smartphone-grade GNSS antenna. The analog GNSS signal is tapped off after the phone's internal bandpass filter and low-noise amplifier and is directed to a dedicated RF front-end for downconversion and digitization. Data are stored to file for subsequent post-processing by a software GNSS receiver and CDGNSS filter.

by the limitations of the phone's internal chipset and clock.

The clock attached to the external front-end was an oven-controlled crystal oscillator (OCXO), which has much greater stability than the low-cost oscillators used to drive GNSS signal sampling within smartphones. However, it was found that reliable cycle-slip-free GNSS carrier tracking only required a 40-ms coherent integration (pre-detection) interval, which is within the coherence time of a low-cost temperature-compensated crystal oscillator (TCXO) at the GPS L1 frequency [15].

Although only a single model of smartphone was tested using this architecture—a popular mass-market phone—the results are assumed representative of all smartphones from the same manufacturer.

Using this architecture, many hours of raw high-rate (∼6 MHz) digitized intermediate frequency samples were collected and stored to disk for post processing. Also stored to disk were high-rate data from a survey-grade antenna, which served as the reference antenna for CDGNSS processing. An in-house software-defined GNSS receiver, known as GRID [16], [17], [18], was used to generate, from these samples, high-quality carrier phase measurements. GRID is a flexible receiver that can be easily adapted to maintain carrier lock despite severe fading. Complex baseband accumulations output from GRID allowed detailed analysis of the signal and tracking loop behavior to ensure that no cycle slips occurred. The generated carrier phase measurements were subsequently passed to a CDGNSS filter, a model for which is described in the next section.

III. CDGNSS Processing

The CDGNSS filter described in this section ingests double-differenced carrier phase measurements output

from GRID and processes them to produce (1) the centimeter-accurate trajectory estimate of the mobile antenna, (2) a time history of phase residuals, (3) the carrier-phase integer ambiguities, (4) theoretical integer ambiguity resolution success bounds, and (5) empirical integer ambiguity resolution success rates. These outputs are used to analyze the performance of the smartphone-grade antenna and compare its performance to higher-quality antennas.

A. CDGNSS Filter Model

A.1 State

The filter's state has a real-valued component that models the relative antenna position and velocity, and an integer-valued component that models phase ambiguities. Such integer ambiguities are inherent to carrier phase differential positioning techniques; their resolution has been the topic of much past research [3], [19] and is required to produce a CDGNSS positioning solution.

The filter's real-valued state component at time t_k is denoted $\mathbf{x}_k$, where $t_k = kT$ and T is the time between consecutive filter updates. This component can be expressed as

$$\mathbf{x}_k = [\mathbf{r}_o, \mathbf{r}_k, \dot{\mathbf{r}}_k]^\mathsf{T} \tag{1}$$

with the following definitions:
$\mathbf{r}_o$ a 3×1 vector, modeled as constant, describing the relative position between the reference antenna and the center of motion of the mobile antenna.
$\mathbf{r}_k$ a 3×1 vector describing the relative position between the center of motion and the mobile antenna at time t_k.
$\dot{\mathbf{r}}_k$ the rate of change of $\mathbf{r}_k$ at time t_k.

The filter's integer-valued state component $\mathbf{n}_k$ at time t_k can be expressed as

$$\mathbf{n}_k = [\mathbf{N}_1, \mathbf{N}_2, \dots, \mathbf{N}_{N_{SV}-1}]^\mathsf{T} \tag{2}$$

with the following definitions:
$\mathbf{n}_k$ a $(N_{\text{SV}-1}) \times 1$ vector of integer phase ambiguities, one for each satellite pair, where N_{SV} is the total number of satellites tracked. A reference satellite is common between all satellite pairs.

A.2 Dynamics Model

The real-valued state component $\mathbf{x}_k$ is assumed to evolve as a mean-reverting second-order Gauss-Markov process. This process models the time-correlated and mean-reverting motion a smartphone experiences when moved in the extended hand of an otherwise stationary user. The integer-valued state component $\mathbf{n}_k$ is modeled as constant, since the phase ambiguities remain fixed so long as the receiver retains phase lock on each signal.

The real-valued state evolves as

$$\mathbf{x}_{k+1} = \mathbf{\Phi}\mathbf{x}_k + \mathbf{\Gamma}\mathbf{w}_k \tag{3}$$

with the following definitions:
$\mathbf{\Phi}$ the state transition matrix.
$\mathbf{\Gamma}$ the process noise influence matrix.
$\mathbf{w}_k$ the process noise at time t_k, modeled as a discrete-time zero-mean white Gaussian random sequence with covariance matrix $\mathbf{Q}$, i.e., $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$, $\mathbf{E}\left[\mathbf{w}_k \mathbf{w}_j^\mathsf{T}\right] = \mathbf{Q}\delta_{kj}$.

The state transition matrix for the real-valued state is that of a discrete-time mean-reverting second-order Gauss-Markov process from t_k to t_{k+1}:

$$\mathbf{\Phi} = \begin{bmatrix} \mathbf{I}_{3\times3} & \mathbf{0}_{3\times3} & \mathbf{0}_{3\times3} \\ \mathbf{0}_{3\times3} & f \cdot \mathbf{I}_{3\times3} & \frac{1-f^2}{\sqrt{1+f^2}} \cdot \mathbf{I}_{3\times3} \\ \mathbf{0}_{3\times3} & \mathbf{0}_{3\times3} & f \cdot \mathbf{I}_{3\times3} \end{bmatrix} \tag{4}$$

where

$$f = e^{-T/\tau_0} \tag{5}$$

is the correlation coefficient for the state elements $\mathbf{r}_k$ and $\dot{\mathbf{r}}_k$, with τ_0 being the average decorrelation time of the antenna motion. The process noise influence matrix is defined as

$$\mathbf{\Gamma} = \begin{bmatrix} \mathbf{0}_{3\times3} & \mathbf{0}_{3\times3} & \sqrt{1-f^2} \cdot \mathbf{I}_{3\times3} \end{bmatrix} \tag{6}$$

and the process noise covariance matrix is defined as

$$\mathbf{Q} = \sigma^2 \left[\mathbf{I}_{3\times3}\right] \tag{7}$$

where σ represents the per-dimension standard deviation of the time-dependent antenna variations modeled by the real-valued state element $\mathbf{r}_k$ introduced in (1).

To adapt the dynamics model to a static antenna constraint, one sets the decorrelation time of the antenna motion to infinity, i.e., $\tau_0 = \infty$, and sets the process noise to zero, i.e., $\mathbf{Q} = \mathbf{0}_{3\times3}$.

A.3 Measurement Model

The filter ingests measurement vectors $\mathbf{y}_k$ for $k = 1, ..., K$, each populated with a single epoch of double-differenced carrier phase measurements. The filter's measurement model relates $\mathbf{y}_k$ to the real- and integer-valued state components $\mathbf{x}_k$ and $\mathbf{n}_k$ through the following adaptation of the linearized GNSS carrier phase measurement model given in [4]:

$$\mathbf{y}_k \triangleq \begin{bmatrix} \phi_{\text{AB},k}^{21} \\ \phi_{\text{AB},k}^{31} \\ \vdots \\ \phi_{\text{AB},k}^{N_{\text{SV}}1} \end{bmatrix}$$
$$= \mathbf{r}_{\mathbf{x}k} + \mathbf{H}_{\mathbf{x}k}\left(\mathbf{x}_k - \bar{\mathbf{x}}_k\right) + \mathbf{H}_{\mathbf{n}k}\mathbf{n}_k + \mathbf{v}_k \tag{8}$$

with the following definitions:

$\phi_{\mathrm{AB},k}^{i1}$ the double-differenced phase measurement between the reference receiver A, the mobile receiver B, satellite i, and satellite 1, the reference satellite, at time t_k.

$\mathbf{r}_{\mathrm{x}k}$ the vector of double-differenced modeled ranges based on the filter's real-valued state prior $\bar{\mathbf{x}}_k$.

$\mathbf{H}_{\mathrm{x}}$ the measurement sensitivity matrix for the real-valued state components.

$\mathbf{H}_{\mathrm{n}}$ the measurement sensitivity matrix for the integer-valued state components at time t_k.

$\mathbf{v}_k$ the discrete-time double-differenced measurement noise vector. Each element of $\mathbf{v}_k$ is modeled as a zero-mean discrete-time Gaussian white noise process, i.e., $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$, $\mathbf{E}\left[\mathbf{v}_k\mathbf{v}_j^{\mathsf{T}}\right] = \mathbf{R}\delta_{kj}$, where $\mathbf{R}$ is the $N_{\mathrm{SV}-1} \times N_{\mathrm{SV}-1}$ measurement noise covariance matrix.

The measurement noise covariance matrix is

$$\mathbf{R} = \sigma_\phi^2 \begin{bmatrix} 4 & 2 & \ldots & 2 \\ 2 & 4 & & \vdots \\ \vdots & & \ddots & \\ 2 & \ldots & & 4 \end{bmatrix} \tag{9}$$

where σ_ϕ^2 is the standard deviation of the undifferenced phase measurement noise, which, for simplicity, is modeled as equivalent for all antenna-to-satellite pairs. The measurement sensitivity matrices are

$$\mathbf{H}_{\mathrm{x}k} = \begin{bmatrix} \rho_{\mathrm{B},k}^{21} & \rho_{\mathrm{B},k}^{21} & \mathbf{0}_{1\times3} \\ \rho_{\mathrm{B},k}^{31} & \rho_{\mathrm{B},k}^{31} & \mathbf{0}_{1\times3} \\ \vdots & \vdots & \vdots \\ \rho_{\mathrm{B},k}^{N_{\mathrm{SV}}1} & \rho_{\mathrm{B},k}^{N_{\mathrm{SV}}1} & \mathbf{0}_{1\times3} \end{bmatrix} \tag{10}$$

$$\tag{11}$$

$$\mathbf{H}_{\mathrm{n}} = \begin{bmatrix} \lambda & 0 & \ldots & 0 \\ 0 & \lambda & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \lambda \end{bmatrix} \tag{12}$$

where $\rho_{\mathrm{B},k}^{i1}$ is the 1×3 single difference of the unit position vectors between the mobile receiver's antenna position estimate, satellite i, and satellite 1, the reference satellite at time t_k. λ is the GNSS signal wavelength.

B. Phase Residuals

After processing data through the CDGNSS filter, the filter outputs, in addition to a time history of centimeter-accurate position estimates, a time history of phase residuals $\tilde{\mathbf{y}}_k$, which can be thought of as departures of each double-differenced phase measurement from phase alignment at the phase center of the antenna. These residuals can be modeled as

$$\tilde{\mathbf{y}}_k = \mathbf{y}_k - \mathbf{r}_{\mathrm{x}k} - \mathbf{H}_{\mathrm{n}}\hat{\mathbf{n}}_K \tag{13}$$

TABLE I
ANTENNA PROPERTIES

Antenna Class	Axial Ratio	Polarization	Relative Loss
Survey-Grade [20]	1 dB	Circular	0 dB
High-quality Patch [21]	2 dB	Circular	0 - 0.5 dB
Low-quality Patch [22]	3 dB	Circular	0.6 dB
Smartphone-Grade	10+ dB	Linear	11 dB

TABLE II
ANTENNAS UNDER TEST

Survey-Grade	High-Quality Patch	Low-Quality Patch	Smartphone-Grade

where $\mathbf{r}_{\mathrm{x}k}$ is now based on the filter's real-valued state estimate $\hat{\mathbf{x}}_k$ at time t_k, and $\hat{\mathbf{n}}_K$ represents the filter's estimate of the integer ambiguities based on all measurements up to t_K.

Phase residuals have been produced for batches of data collected from four different grades of antennas, as described next. These residuals will be used to analyze the suitability of each antenna for CDGNSS positioning.

IV. Antenna Performance Analysis

This section describes four antennas from which data were captured and processed using the test architecture and CDGNSS filter described previously. It also quantifies the characteristics that make low-quality smartphone-grade antennas poorly suited to CDGNSS.

Table I describes a range of antenna grades of decreasing quality, noting properties relevant to CDGNSS. The loss numbers in the rightmost column represent the average loss in gain relative to a survey-grade antenna, where the average is taken over elevation angles above 15 degrees. Table II shows four antennas, one of each grade, from which many hours of data have been collected using the test architecture. Survey-grade antennas, whose properties are described in the first row of Table I, have a uniform quasi-hemispherical gain pattern, right-hand circular polarization, a stable phase center, and a low axial ratio. These are all desirable properties for CDGNSS. Unfortunately, these properties inhere in the antennas' large size; the laws of physics dictate that smaller antennas will typically be worse in each property. Also listed in Table I are properties for three other antenna grades. The second and third rows list properties for high- and low-quality patch antennas. These antennas have similar properties to a survey-grade antenna and lose, on average, less than 0.5 dB and 1 dB respectively in sensitivity as compared to the survey-grade antenna [21], [22].
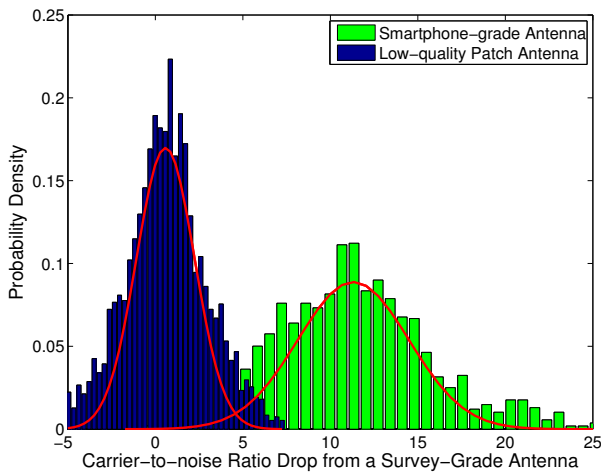
Fig. 2. Normalized histograms displaying the drop in carrier-to-noise ratio between a survey-grade antenna and a smartphone-grade (right) and low-quality patch (left) antenna. Each histogram was generated from 2 hours of data and 9 tracked satellites ranging in elevation from 15 to 90 degrees. The antennas remained stationary. The red traces represent Gaussian distribution models fit to each histogram.

The last row of Table I lists the properties for a smartphone-grade antenna. As shown subsequently, this antenna loses between 5 and 15 dB in sensitivity as compared to the survey-grade antenna. Such a loss makes it difficult to retain lock on GNSS signals. In addition, this antenna's linear polarization leads to extremely poor multipath suppression.

A. Antenna Gain Analysis

Fig. 2 quantifies one of the obvious drawbacks of a smartphone-grade antenna, namely, its low gain. The rightmost histogram, in green, shows that the decrease in carrier to noise ratio as compared to a survey-grade antenna is on average 11 dB, such that the smartphone-grade antenna only captures approximately 8% of the signal power as compared its survey-grade counterpart. For comparison, shown on the left, in blue, is a histogram of the decrease in carrier-to-noise ratio for the low-quality patch antenna. This antenna only suffers about a 0.6 dB drop in power on average relative to the survey-grade antenna. Each histogram was generated from 2 hours of data with 9 tracked satellites ranging in elevation from 15 to 90 degrees. The antennas remained stationary. The variation in signal power around the means is due to the multipath-induced power variations in the signal as well as to the different gain patterns between each antenna and the survey-grade antenna.

B. Phase Residual Analysis

Shown in Figs. 3, 4, and 5 are 2000-second segments of double-differenced phase residual time histories for data collected from a survey-grade, a low-quality patch, and a smartphone-grade antenna, respectively. To produce


Fig. 3. Time histories of double-differenced phase residuals for a 2000-second batch of data captured from a survey-grade antenna. Each trace represents a residual for a different satellite pair. The ensemble average standard deviation of the residuals is 3.4 millimeters.


Fig. 4. Time histories of double-differenced phase residuals for a 2000-second batch of data captured from a low-quality patch antenna. Each trace represents a residual for a different satellite pair. The ensemble average standard deviation of the residuals is 5.5 millimeters.

these residuals, the antenna position was locked to its estimated value within the CDGNSS filter. The residuals represent departures of the carrier phase measurements from perfect alignment at the average phase center of the antenna. Each different colored trace corresponds to a different satellite pair. While the data segments were not captured at the same time of day, they were captured at the same location, and thus the multipath environment was similar.

The ensemble average residual standard deviations increase with decreasing antenna quality. The residuals for the survey-grade, low-quality patch, and smartphone-grade antennas have ensemble average standard deviations

Fig. 5. Time histories of double-differenced phase residuals for a 2000-second batch of data captured from a smartphone-grade antenna. Each trace represents a residual for a different satellite pair. The ensemble average standard deviation of the residuals is 11.4 millimeters.


Fig. 6. A successful CDGNSS solution using data collected from the antenna of a smartphone. The cluster of red near the lower left-hand corner of the phone represents 500 CDGNSS solutions over an 8-minute interval, superimposed on the photo and properly scaled.

of 3.4, 5.5, and 11.4 millimeters, respectively. This increase is due to the lower gain and less effective multipath suppression of the lower-quality antennas.

Fig. 5 shows the presence of outlier residuals in the data collected from the smartphone-grade antenna. These outliers, one of which persists for over 1000 seconds, are likely caused by either large and irregular azimuth- and elevation-dependent antenna phase center variations or a combination of poor antenna gain in the direction of the non-reference satellite coupled with ample gain in the direction of a multipath signal such that the multipath signal is received with more power than the direct-path signal. Obvious outliers such as these can be automatically excluded by the CDGNSS filter via an innovations test. However, the standard deviation of the remaining residuals still remains large compared to that of the other antennas; the ensemble average standard deviation decreases from 11.4 to 8.6 millimeters upon exclusion of the two large outliers.

For antennas with a large ensemble average standard deviation in their double-differenced phase errors, the time correlation in the phase errors becomes more important. This time correlation, which persists for 100-200 seconds, is a well-studied phenomenon caused by slowly-varying carrier phase multipath [23], [6]. While correlation is present in the residuals of all antenna types, and manifests approximately the same decorrelation time, its effect is more of a problem for low-quality antennas because the phase errors are larger. Such correlation, coupled with a large deviation, ultimately leads to a longer time to ambiguity resolution, as will be shown subsequently.

Given a smartphone antenna's extremely poor gain and multipath suppression as compared to even a low-quality stand-alone patch antenna, one might question the wisdom of attempting a CDGNSS solution using such an antenna. However, the next section reveals that it is indeed possible to achieve a centimeter-accurate positioning solution using a smartphone GNSS antenna despite its poor properties.

V. CDGNSS Performance using a Smartphone Antenna

This section discusses the results of performing a CDGNSS solution using data collected from a smartphone-grade antenna and presents two strategies for improving the performance of CDGNSS on smartphones.

Fig. 6 shows the result of an early attempt to compute a CDGNSS solution using data collected from the GNSS antenna of a smartphone. The cluster of red near the lower left-hand corner of the phone represents 500 CDGNSS position estimates over an 8-minute interval, superimposed on the photo and properly scaled. The integer ambiguities were resolved correctly, as verified through analysis of the phase residuals and by physical measurement. Although early experiments were done with the large conductive backplane shown here, later experiments revealed the backplane to be unnecessary. Also, whereas the phone was oriented face down in early testing, it was later discovered that the phone's irregular gain pattern is better oriented when the phone is face up. Accordingly, all other smartphone results presented in this paper are for data collected with the phone oriented face up and supported only by a plastic box as shown in Fig. 1. Furthermore, although this scenario enjoyed a very short baseline to a reference antenna (less than 10 meters), similar ambiguity resolution performance is to be expected for baselines

shorter than approximately 5 kilometers, as differential ionospheric and tropospheric delays are negligible in this short-baseline regime [14].

The possibility of CDGNSS-enabled centimeter positioning using a smartphone antenna has been previously conjectured [24], but—to the authors' knowledge—Fig. 6 represents the first published demonstration that this is indeed possible. This significant result portends a vast expansion of centimeter-accurate positioning into the mass market. However, serious challenges must be overcome before mass-market CDGNSS can become practical, as described below.

A. CDGNSS Performance in a Static Scenario

Fig. 7 shows the empirical probability of successful ambiguity resolution for data collected from four antennas, one of each of the different grades discussed earlier. For each antenna, 7 satellites were tracked at approximately the same location and time of day. Each trace was computed from 12 batches of double-differenced carrier phase data. Code-phase (pseudorange) measurements were assumed to be distilled into a single *a priori* position estimate modeled as a Gaussian vector with a 70-meter deviation along each axis. With such a highly uncertain prior, filter performance is dominated by the carrier phase measurements, which were modeled as having undifferenced deviations of $\sigma_\phi = 2$ cm. Each separate batch of data was treated as a Monte Carlo run with the prior position estimate randomly generated as modeled. For the traces corresponding to the low-quality patch, high-quality patch, and survey-grade antennas, there were at least 100 seconds separating the start of each batch with no overlap between batches. For the trace derived from batches collected from the smartphone-grade antenna, due to the difficulty in recording long segments of data, there were only 70 seconds separating the start of each batch, resulting in significant overlap between batches.

Each trace represents an empirically-derived success rate computed from 12 batches of phase data as follows:

1. For a given batch, at each epoch, the filter outputs its best estimate of the integer ambiguities on the basis of the data ingested thus far.
2. The estimate from step 1 is compared against the true set of integer ambiguities which were acquired in advance by processing a much longer batch of data. If correct, a flag is set at that epoch to "1"; if incorrect, the flag is set to 0.
3. For each epoch, the flags produced in step 2 are averaged across all 12 batches to generate each trace.

As shown by the green trace in Fig. 7, the smartphone-grade antenna required 400 seconds to achieve a 90% ambiguity resolution success rate; in other words, it manifested a 400-second TAR at 90%. This would surely exceed the patience of most smartphone users. Also shown are traces
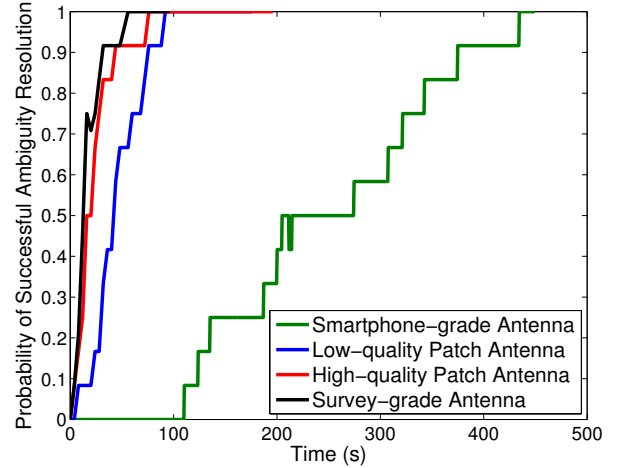

Fig. 7. Probability of successful ambiguity resolution vs. time using data collected from four antennas of varying quality. For each antenna, 7 satellites were tracked at approximately the same location and time of day. Each trace was computed from 12 batches of double-differenced carrier phase data.

for the other three antenna grades. The higher-quality antennas yield shorter TARs for a given success rate, primarily due to their superior multipath suppression.

Note that the loss in received signal power due to the smartphone antenna's poor gain turns out to be tolerable—the signals arriving from the smartphone-grade antenna can be tracked without cycle slipping. Therefore, the outstanding challenge preventing fast ambiguity resolution for data collected from smartphone-grade antennas is the severe time-correlated multipath errors in the double-differenced carrier phase data.

B. Decreasing TAR via More Signals

There are ways to mitigate the impact of multipath on the CDGNSS TAR—even the severe multipath experienced by low-quality antennas. It has been shown that the volume of the integer ambiguity search space, and thus TAR, decreases as a function of the number of double-differenced phase time histories available, which, for single-frequency CDGNSS, is one less than the number of satellites tracked [25]. Consequently, an acceptable TAR can always be achieved with enough satellites tracked.

Fig. 8 shows the reduction in TAR for an increasing number of satellites. Each trace was computed from 720 non-overlapping 2-minute batches of data taken from a survey-grade antenna over a 24-hour interval. A decreasing elevation mask angle was used to allow an increasing number of SVs to participate in the CDGNSS solution. In other words, for a given 2-minute batch of data, an elevation mask was first applied to all but the highest 5 satellites. Double-differenced phase data from these satellites were then processed by the CDGNSS filter to compute an em-

Fig. 8. Probability of successful ambiguity resolution vs. time as a function of the total number of satellite vehicles (SVs) tracked. Each trace is computed from 720 non-overlapping 2-minute batches of double-differenced carrier-phase data taken from a survey-grade antenna over a 24-hour interval. A decreasing elevation mask angle was used to allow an increasing number of SVs to participate in the CDGNSS solution.
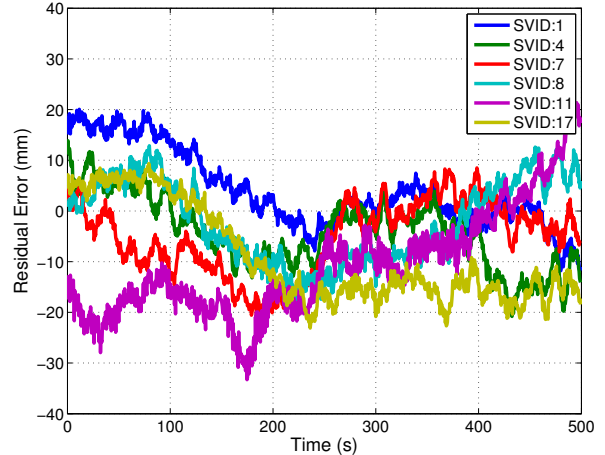

Fig. 9. Time histories of phase residuals for a batch of data captured from the smartphone-grade antenna while static. Each trace represents a double-differenced phase residual history for a different satellite pair.

pirical probability of successful integer ambiguity resolution. Next, the elevation mask was reduced until one additional satellite was in view and the process was repeated to produce all the traces shown.

As before, code-phase (pseudorange) measurements were assumed to be distilled into a single *a priori* position estimate modeled as a Gaussian vector. The prior's deviation was 100 meters along each axis and the carrier phase measurement errors were modeled as having undifferenced deviations of $\sigma_\phi = 2$ cm. Each separate 2-minute batch of data was treated as a Monte Carlo run with the prior position estimate randomly generated as modeled.

Fig. 8 makes clear that each additional double-differenced phase time history, although corrupted by its own multipath-induced phase errors, significantly decreases the overall TAR. Note that although Fig. 8 was produced from data collected via a survey-grade antenna, a similar trend would apply for the smartphone-grade antenna. One implication of Fig. 8 is that smartphone-based CDGNSS would benefit greatly from the additional double-differenced measurements that a multi-frequency GNSS receiver could provide. For example, at the time of writing there are 14 operational GPS satellites broadcasting unencrypted civil signals at the GPS L2 frequency (1227.6 MHz), and 7 broadcasting civil signals at the GPS L5 frequency (1176.45 MHz). With some modification of the smartphone GNSS antenna and chipset, these modernized GPS signals could be exploited to reduce TAR. However, the narrow profit margins on mass-market GNSS antennas and chipsets militate against multi-frequency architectures.

C. Decreasing TAR via Random Receiver Motion

There is a second way to reduce TAR under severe multipath conditions. Unlike TAR reduction via additional signals, the theory and practice of this second technique have not been previously treated in the literature. Moreover, the technique is well-suited for smartphones, which are typically hand-held and mobile. This simple technique consists of gently moving the smartphone in a quasi-random manner within a wavelength-scale volume. The key to this technique's effectiveness is that, whereas multipath-induced phase measurement errors are typically time-correlated on the order of hundreds of seconds for a static receiving antenna [23], their spatial correlation is on the order of one wavelength, or approximately 19 centimeters at the GPS L1 frequency [26]. As a result, random wavelength-scale antenna motion transforms the phase residuals from slowly-varying when the antenna is static, as shown in Fig. 9, to quickly-varying when the antenna is dynamic, as shown in Fig. 10. Put another way, the autocorrelation time of the phase residuals decreases from hundreds of seconds when the antenna is static, as shown in Fig. 11, to less than a second when the antenna is moved even slowly (a few centimeters per second), as shown in Fig. 12. More vigorous antenna motion would be possible if the phone's inertial devices were used to aid the phase tracking loops [27].

The shorter phase error decorrelation time resulting from random antenna motion effectively increases the information content per unit time that each double-differenced phase measurement provides to the CDGNSS filter, thus decreasing the time to ambiguity resolution. Fig. 13 compares empirical success rates for three different antennas under static and dynamic scenarios. As expected, motion

Fig. 10. Time histories of phase residuals for a batch of data captured from the smartphone-grade antenna as it experienced wavelength-scale random motion of 2-5 centimeters per second. Each trace represents a double-differenced phase residual history for a different satellite pair.


Fig. 11. Autocorrelation functions corresponding to the phase residuals in Fig. 9.

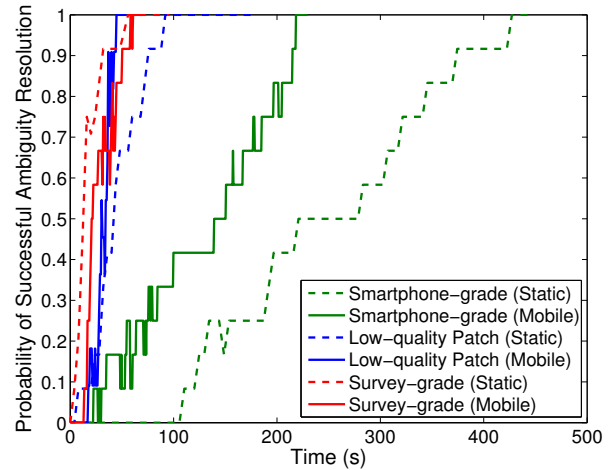Fig. 12. Autocorrelation functions corresponding to the phase residuals in Fig. 10.

Fig. 13. Probability of successful ambiguity resolution vs. time for three different antennas under static and dynamic scenarios. Data were collected at the same location and at approximately the same time of day. For each trace, 7 satellites were tracked. Prior positions and carrier phase errors were modeled, and prior positions were randomized in independent Monte Carlo runs, exactly as for Fig. 7

results in a reduced TAR for the smartphone-grade and low-quality patch antenna. But, somewhat counterintuitively, motion results in an increased TAR for the survey-grade antenna. This discrepancy reflects a tradeoff within the CDGNSS filter. While it is true that the phase measurement errors decorrelate much faster when the antenna is moving—increasing the per-epoch information provided to the filter—it is also the case that the filter can no longer employ a hard motion constraint. For the high-quality antennas, the increased information per epoch due to faster phase error decorrelation is completely counteracted by a loss in information per epoch due to uncertainty (lack of constraint) in the motion model. Also, for the high-quality antennas, multipath in the reference antenna's phase measurements is not insignificant compared to multipath in

the mobile antenna, and this reference multipath exhibits the usual 100-200 second correlation time for a static antenna. On the other hand, phase error decorrelation via random antenna motion offers the lower-quality antennas a larger net information gain because their multipath-induced phase errors are so large. Consequently, for the smartphone-grade antenna, motion substantially reduces the 90% success TAR, which drops from 400 to 215 seconds.

VI. Conclusions and Future Work

Centimeter-accurate positioning was demonstrated based on data sampled from a smartphone-quality GNSS antenna. An empirical analysis revealed that the extremely

poor multipath suppression of these antennas is the primary impediment to fast resolution of the integer ambiguities that arise in the carrier-phase differential processing used to obtain centimeter accuracy. It was shown that, for low-quality smartphone-grade GNSS antennas, wavelength-scale random antenna motion substantially reduces the ambiguity resolution time.

Future work will study the effectiveness of combining antenna motion with a motion trajectory estimate derived from non-GNSS smartphone sensors to further reduce the integer ambiguity resolution time. This technique, which is a type of synthetic aperture processing applied to the double-differenced GNSS phase measurements, effectively points antenna gain enhancements in the direction of the overhead GNSS satellites, thereby suppressing multipath arriving from other directions. Preliminary results show that this technique offers modest benefit beyond the unaided random motion technique discussed herein.

References

[1] K. Alexander, "U.S. GPS program and policy update," in *26th SBAS International Working Group*. National Coordination Office, Feb. 2014.

[2] M. Sahmoudi and R. J. Landry, "Multipath mitigation techniques using maximum-likelihood principle," *Inside GNSS*, pp. 24–29, 2008.

[3] C. C. Counselman III, R. I. Abbot, S. A. Gourevitch, R. W. King, and A. R. Paradis, "Centimeter-level positioning with GPS," *Journal of Surveying Engineering*, vol. 109, no. 2, pp. 81–89, 1983.

[4] S. Mohiuddin and M. L. Psiaki, "High-altitude satellite relative navigation using carrier-phase differential global positioning system techniques," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1628–1639, Sept.-Oct. 2007.

[5] F. van Diggelen, "Expert advice: Are we there yet? The state of the consumer industry," Mar. 2010, GPS World.

[6] C. Miller, K. OKeefe, and Y. Gao, "Operational performance of RTK positioning when accounting for the time correlated nature of GNSS phase errors," in *Proceedings of the ION GNSS Meeting*, 2010, pp. 21–24.

[7] K. O'Keefe, M. Petovello, G. Lachapelle, and M. E. Cannon, "Assessing probability of correct ambiguity resolution in the presence of time-correlated errors," *Navigation, Journal of the Institute of Navigation*, vol. 53, no. 4, pp. 269–282, 2007.

[8] L. Wirola, K. Alanen, J. Kappi, and J. Syrjarinne, "Bringing RTK to cellular terminals using a low-cost single-frequency AGPS receiver and inertial sensors," in *Proceedings of the IEEE/ION PLANS Meeting*, 2006, pp. 645–652.

[9] B. Bougard, G. Cuypers, W. De Wilde, and J.-M. Sleewaegen, "AsteRx-m – a very low-power compact GPS/GLONASS RTK receiver," in *Proceedings of the ION GNSS+ Meeting*, 2011, pp. 1429–1436.

[10] T. Takasu and A. Yasuda, "Development of the low-cost RTK-GPS receiver with an open source program package RTKLIB," in *International Symposium on GPS/GNSS*, 2009, pp. 4–6.

[11] F. van Diggelen, "Who's your daddy? Why GPS will continue to dominate consumer GNSS," *Inside GNSS*, pp. 30–41, March/April 2014.

[12] C. Rizos, V. Janssen, C. Roberts, and T. Grinter, "Precise Point Positioning: Is the era of differential GNSS positioning drawing to an end?" in *In Proceedings of FIG Working Week*, vol. 10, 2012.

[13] E. Gakstatter, "Centimeter-level RTK accuracy more and more available for less and less," 2014, GPS World.

[14] T. Takasu and A. Yasuda, "Kalman-filter-based integer ambiguity resolution strategy for long-baseline RTK with ionosphere and troposphere estimation," in *Proceedings of the ION National Technical Meeting*, 2010.

[15] K. Wesson, K. Pesyna, J. Bhatti, and T. E. Humphreys, "Opportunistic frequency stability transfer for extending the coherence time of GNSS receiver clocks," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2010.

[16] 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, pp. 326–338.

[17] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, and P. M. Kintner, Jr., "GNSS receiver implementation on a DSP: Status, challenges, and prospects," in *Proceedings of the ION GNSS Meeting*. Fort Worth, TX: Institute of Navigation, 2006, pp. 2370–2382.

[18] 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, pp. 2745–2753.

[19] P. Teunissen, P. De Jonge, and C. Tiberius, "The LAMBDA method for fast GPS surveying," in *Proceedings of International Symposium on GPS Technology Applications*, vol. 29. Citeseer, 1996, pp. 203–210.

[20] *Datasheet: GPS-702L Dual-Frequency Pinwheel GPS Antenna*, Novatel, 2014.

[21] *Datasheet: Active L1/L2 GPS Antenna, P/N: 53G1215A-XT-1*, Antcom, 2006.

[22] *Datasheet: Active L1 GPS Antenna, P/N: Dominator AA.161*, Taoglas.

[23] C. Miller, K. O'Keefe, and Y. Gao, "Time correlation in GNSS positioning over short baselines," *Journal of Surveying Engineering*, vol. 138, no. 1, pp. 17–24, 2011.

[24] T. E. Humphreys, "How to fool a GPS," Feb. 2012, http://www.ted.com/talks/todd_humphreys_how_to_fool_a_gps.

[25] P. Teunissen, "Closed form expressions for the volume of the GPS ambiguity search spaces," *Artificial Satellites- Planetary Geodesy*, vol. 32, no. 1, pp. 5–20, 1997.

[26] S. N. Sadrieh, A. Broumandan, and G. Lachapelle, "Spatial/temporal characterization of the GNSS multipath fading channels," in *Proceedings of the ION GNSS Meeting*, 2010, pp. 393–401.

[27] E. Ohlmeyer, "Analysis of an ultra-tightly coupled GPS/INS system in jamming," in *Proceedings of the IEEE/ION PLANS Meeting*. Institute of Navigation, 2006, pp. 44–53.