

On the Feasibility of cm-Accurate Positioning via a Smartphone's Antenna and GNSS Chip

Todd E. Humphreys, Matthew Murrian
Department of Aerospace Engineering
and Engineering Mechanics
The University of Texas at Austin
Austin, Texas 78712
todd.humphreys@mail.utexas.edu
matthew.murrian@utexas.edu

Frank van Diggelen, Sergei Podshivalov
Broadcom
{fvandiggelen, spodshiv}@broadcom.com

Kenneth M. Pesyna, Jr.
Radiosense
ken@radiosense.com

Abstract—The feasibility of centimeter-accurate carrier-phase differential GNSS (CDGNSS) positioning using a smartphone's internal GNSS antenna and GNSS chip is investigated. Precise positioning on a mass-market platform would significantly influence the world economy, ushering in a host of consumer-focused applications that have so far been hampered by the several-meter-level errors in traditional GNSS positioning. Previous work has shown that GNSS signals received through a mass-market smartphone's GNSS antenna can be processed to yield a centimeter-accurate CDGNSS position solution, but this earlier work processed all GNSS signals externally to the smartphone. The question remains whether a smartphone's internal oscillator and GNSS chip can produce observables of sufficient quality to support centimeter-accurate carrier-phase-based positioning. This paper answers the question by accessing and processing the raw code- and carrier-phase observables produced by a mass-market smartphone GNSS chip—observables that have heretofore been unavailable to the research community. The phone's carrier phase measurements are shown to suffer from five anomalies compared to those from a survey-grade GNSS receiver, four of which are readily fixed in post-processing. The remaining anomaly, an error in the phase measurement that grows approximately linearly with time, currently prevents the phone's phase measurements from satisfying the conditions for CDGNSS positioning. But the phone's measurements seem otherwise fully capable of supporting cm-accurate carrier-phase differential GNSS positioning. A separate analysis of a smartphone's GNSS signal strength dependency on azimuth and elevation reveals that multipath-induced deep fading and large phase errors remain a significant challenge for centimeter-accurate smartphone positioning.

Keywords—CDGNSS, smartphone, precise positioning, multipath, antenna pattern

I. INTRODUCTION

Recent demonstrations have established that GNSS signals received through a mass-market smartphone's GNSS antenna and low-noise amplifier can be processed to yield a centimeter-accurate CDGNSS position solution [1], [2]. But in these demonstrations, signal digitization and processing were performed external to the smartphone. Internal (within the phone's GNSS chip) GNSS processing has come a long way over the past decade; indeed, modern mass-market GNSS chips actually outperform survey-grade chipsets in acquisition sensitivity and time to code-phase fix [3]. But mass-market chips have serious

limitations as regards CDGNSS processing. First, the low-quality temperature-compensated crystal oscillator (TCXO) used to drive a smartphone's internal GNSS chip exhibits significant low-frequency phase noise, making cycle-slip-free carrier phase recovery more challenging [4]–[6]. By contrast, external sampling and downmixing can be referenced to a high-quality oscillator, such as the oven-controlled crystal oscillator (OCXO) used in [1], [2]. A higher-quality oscillator permits carrier recovery within a narrow-bandwidth carrier tracking loop, leading to higher loop signal-to-noise-ratio (SNR). Second, the signal bandwidth, sampling frequency, and quantization resolution of mass-market chips is typically lower than that of a high-cost radio frequency front end, resulting in lower SNR. Third, the tracking loops within the phone's internal chip are optimized for robust code-phase positioning but not for continuous and reliable carrier-phase recovery, as expected for CDGNSS. Fourth, although popular smartphone GNSS chips have an internal delay lock loop and phase lock loop, which imply code phase and carrier phase measurement, the raw code- and carrier-phase observables have not typically been available to the smartphone manufacturers, much less to the user.

Of these limitations, only the fourth—the inaccessibility of the raw observables—is a definitive impediment to CDGNSS processing on a smartphone. Fortunately, for current- and future-generations of the popular Broadcom GNSS chips, this is not a hardware limitation. With proprietary software, the chips' raw output data stream can be accessed and parsed to produce all observables necessary for CDGNSS, which can be written to a file in a standard RINEX format. This capability to access the raw carrier and code phase observables from a mass-market smartphone GNSS chip, first reported and demonstrated in this paper, allows a study of whether the other three limitations mentioned above can be overcome to enable centimeter-accurate positioning on a smartphone.

References [1], [2] conclude that the primary challenge of CDGNSS processing using signals received through a smartphone antenna is not the signals' low carrier-to-noise ratio (C/N_0), but rather their significant multipath-induced phase errors, owing to the antenna's poor multipath suppres-

sion. While it is true that the antenna’s poor sensitivity—on average 11 dB worse than a survey-grade antenna —increases the thermal noise in the double-differenced CDGNSS phase measurements, the long time to integer ambiguity resolution was shown to be more strongly affected by multipath errors than by thermal noise [1]. This conclusion suggests that a CDGNSS solution likely remains possible despite another 1-2 dB drop in C/N_0 caused by the low bandwidth and coarse quantization of a mass-market GNSS chip. Similarly, although an OCXO was used for the processing in [1], [2], it was shown that reliable carrier phase recovery was possible using complex correlation products coherently integrated over only 20-to-40 ms, which is within the coherence time of some mass-market TCXOs. Finally, whereas the default phaselock loop (PLL) settings on the Broadcom chips are configured for navigation data recovery and for code-phase tracking loop aiding, they are fully adjustable and can be optimized for carrier recovery. Therefore, it appears that CDGNSS positioning may well be possible using a smartphone’s GNSS antenna and chip despite the first three limitations mentioned above. A test of this conjecture is this paper’s primary contribution.

A secondary contribution is an examination of the gain pattern of a popular smartphone’s internal antenna. Errors in the double differenced phase residuals studied in [1] exceeded a quarter wavelength, which for a small antenna with a well-defined phase center can only be the case when a multipath signal is received with greater amplitude than the direct-path signal. These large errors thus suggest that the smartphone antenna’s gain pattern is highly irregular, exhibiting pockets of unusually low gain. To investigate this conjecture, this paper carries out a deliberate study of a smartphone’s received signal power relative to a reference antenna with a well-defined gain pattern.

II. ACCESSING AN ANDROID SMARTPHONE’S GNSS OBSERVABLES

GNSS code and carrier phase measurements have long been produced by assisted GNSS processing in smartphones, but these measurements have not been available to application developers. Yet, as this section details, the right proprietary software unlocks access to all standard GNSS observables.

This paper treats only smartphones employing the Android operating system. The Android software stack is shown in Fig. 1, with the GPS components and data flow highlighted. Up to and including Android release “L,” the raw GNSS observables were available only up to the Hardware Abstraction Layer. In Release “M,” observables are available to the Application Framework. But, as shown in Fig. 2, the framework itself does not provide the necessary Application Programming Interface (API) to expose the observables to applications. As a result, no Android application can yet access the GNSS observables.

Fig. 3 shows the GNSS satellite data that is currently available to the Application Layer; these include

- 1) The satellite identifier (e.g., PRN code), from which one can infer the constellation (GPS, GLONASS, etc).
- 2) Azimuth and elevation angles to each satellite.

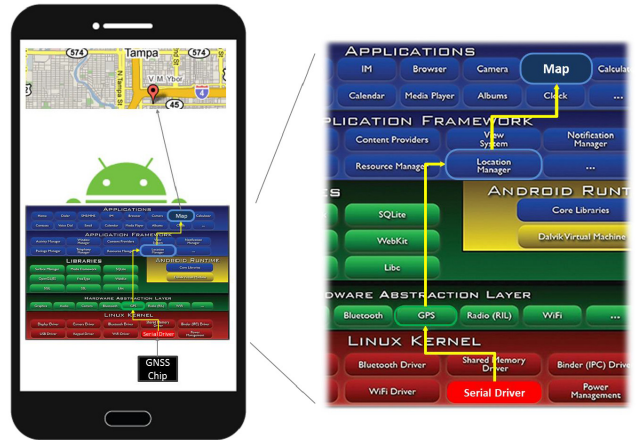


Fig. 1. A smartphone’s Android software stack with the GNSS components and data flow highlighted.

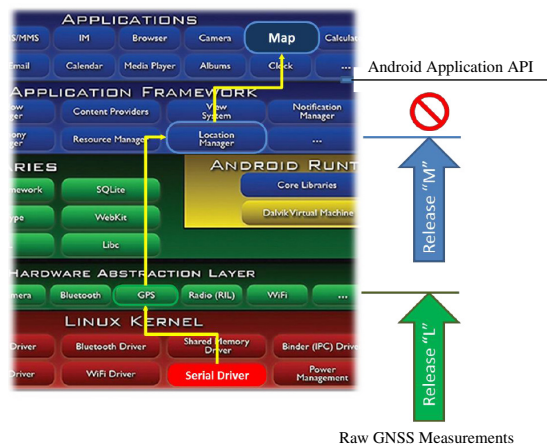


Fig. 2. Android release “L” made GNSS observables available only to the Hardware Abstraction Layer. Release “M” makes them available to the Application Framework, but does not provide an API to expose the measurements to the Applications Layer.

- 3) The carrier to noise density ratio, C/N_0 .
- 4) Flags indicating whether the receiver has received (1) almanac, and (2) ephemeris data for a particular satellite, and whether (3) the satellite’s measurements are being used in the GNSS fix.

Fig. 4 shows the PVT (position velocity and time) data currently available to the Application Layer; these include

- 1) Latitude, longitude, and altitude.
- 2) Time, speed, and “bearing” (horizontal direction of travel).
- 3) Ancillary information such as solution accuracy.

Any commercially-available GNSS application currently used on an Android phone builds up its displays and logs only from these data, which are exposed by the Android Application API. So how do the authors produce RINEX files with raw GNSS observables? Broadcom provided a test phone with customized software, including a modified GPS library that records RINEX files to the phone’s storage card. This GPS

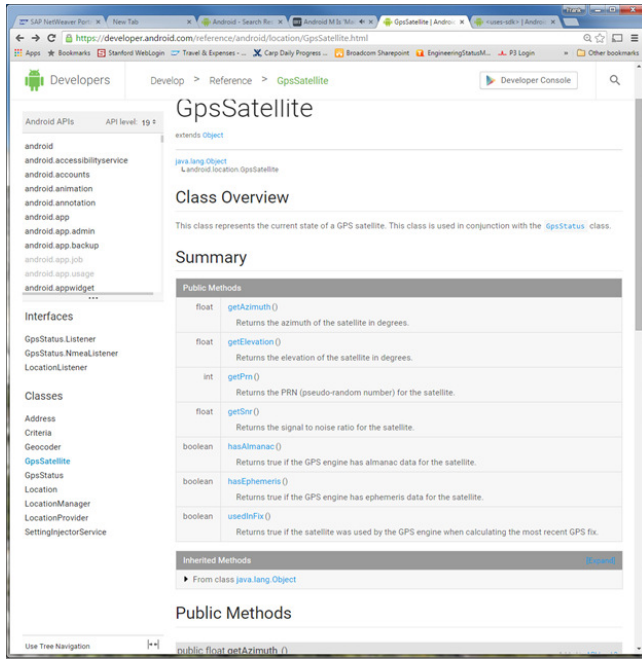


Fig. 3. GNSS satellite information currently available to Android applications.

library is part of the Hardware Abstraction Layer shown in Figs. 1 and 2, and it has all the standard GNSS observables available to it.

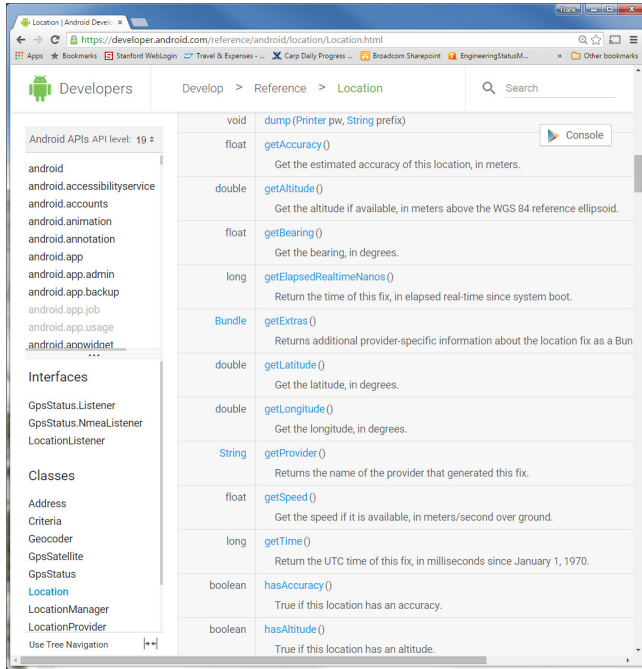


Fig. 4. Position, velocity, and time measurements currently available to Android applications.

III. CHARACTERIZATION OF OBSERVABLES

Pseudorange and carrier phase observables were extracted as described above from the internal Broadcom GNSS chip of a Samsung Galaxy S5 smartphone under two test configurations: (1) a zero-baseline configuration, which allowed the observables to be studied in isolation from multipath and antenna effects, and (2) a stand-alone configuration in which the phone rested on a metal backplane in nearly full view of the sky.

A. Measurement Models

1) *Undifferenced Observables*: The standard undifferenced GNSS pseudorange and carrier phase observables can be modeled in distance units, respectively, as [7]

$$\rho = r + c[\delta t_R - \delta t_S] + I + T + w_\rho \quad (1a)$$

$$\lambda\phi = r + c[\delta t_R - \delta t_S] - I + T + \lambda\beta + \lambda w_\phi \quad (1b)$$

where r is the geometric range between satellite and receiver, c is the speed of light in a vacuum, δt_R and δt_S are the receiver and satellite clock errors, I is the frequency-dependent ionospheric group delay, T is the neutral-atmospheric (commonly called tropospheric) delay, β is a phase offset having to do with the initial reference phase values at the satellite and receiver, λ is the carrier wavelength, and w_ρ and w_ϕ are pseudorange and carrier phase measurement errors. Note that for the undifferenced carrier phase measurement, β is not necessarily an integer number of cycles.

By convention [8], if the receiver clock is adjusted (e.g., to bring δt_R closer to zero), the adjustment must be reflected in both ρ and $\lambda\phi$. Let $t_R = t + \delta t_R$ denote receiver time, with t being the corresponding true time, and let subscripts b and a denote quantities before and after an adjustment, with $\Delta t_R \triangleq t_{R,a} - t_{R,b}$. Then a receiver time adjustment must satisfy the following conditions:

$$t_{R,a} = t_{R,b} + \Delta t_R \quad (2a)$$

$$\rho_a = \rho_b + c\Delta t_R \quad (2b)$$

$$\lambda\phi_a = \lambda\phi_b + c\Delta t_R \quad (2c)$$

2) *Double-Differenced Observables*: Differencing the observables produced by a receiver under test from those produced by a reference receiver is a convenient basis for CDGNSS positioning and offers a straightforward way to evaluate the quality of the test receiver's observables. If the receivers are driven by a common clock, then single-difference measurements are possible; otherwise, if the receivers do not share a common clock, as in the present case for evaluation of a smartphone's observables, double-difference (DD) observations more convenient.

Let A designate the reference receiver and B the receiver under test. Then for pivot satellite q , other satellite i and a single carrier wavelength λ , the pseudorange and carrier phase double differences are formed as

$$\rho_{AB}^{iq} = (\rho_A^i - \rho_A^q) - (\rho_B^i - \rho_B^q) \quad (3a)$$

$$\phi_{AB}^{iq} = (\phi_A^i - \phi_A^q) - (\phi_B^i - \phi_B^q) \quad (3b)$$

Multiplying the carrier phase by λ , the DD observations can be modeled as

$$\rho_{AB}^{iq} = r_{AB}^{iq} + I_{AB}^{iq} + T_{AB}^{iq} + w_{\rho,AB}^{iq} \quad (4a)$$

$$\lambda\phi_{AB}^{iq} = r_{AB}^{iq} - I_{AB}^{iq} + T_{AB}^{iq} + \lambda N_{AB}^{iq} + \lambda w_{\phi,AB}^{iq} \quad (4b)$$

where r_{AB}^{iq} is the DD geometric range, I_{AB}^{iq} is the DD ionospheric delay at the L1 frequency, T_{AB}^{iq} is the DD neutral atmospheric delay (commonly called the tropospheric delay), N_{AB}^{iq} is the DD integer ambiguity, and $w_{\rho,AB}^{iq}$ and $w_{\phi,AB}^{iq}$ are the pseudorange and carrier phase measurement errors.

B. Zero-Baseline Test

A zero-baseline test is a standard method for evaluating the quality of GNSS observables [9]. In this test, the observables from two receivers are compared when attached to the same antenna so that $r_{AB}^{iq} = I_{AB}^{iq} = T_{AB}^{iq} = 0$; moreover, the measurement errors $w_{\rho,AB}^{iq}$ and $w_{\phi,AB}^{iq}$ are free of multipath effects. The zero-baseline DD observables can be modeled as

$$\rho_{AB}^{iq} = \tilde{w}_{\rho,AB}^{iq} \quad (5a)$$

$$\lambda\phi_{AB}^{iq} = \lambda N_{AB}^{iq} + \lambda \tilde{w}_{\phi,AB}^{iq} \quad (5b)$$

where the pseudorange measurement noise $\tilde{w}_{\rho,AB}^{iq}$ and carrier phase measurement noise $\tilde{w}_{\phi,AB}^{iq}$, both modeled as zero-mean, are due only to (i) thermal noise within each receiver's radio frequency (RF) front end (not in the common antenna's low-noise amplifier), and (ii) differences in the measurement procedures between the reference receiver and the receiver under test.

Zero-baseline evaluation of a smartphone's GNSS observables is complicated by the lack of a radio frequency (RF) jack into which signals from a common antenna can be fed. But the setup in Fig. 5 illustrates how a zero-baseline configuration can nonetheless be realized by splitting the RF signal from a common antenna, routing one branch to a reference GNSS receiver, and re-radiating signals from the other branch to the phone within an RF enclosure. This is the zero-baseline configuration under which the Galaxy S5's observables were evaluated.

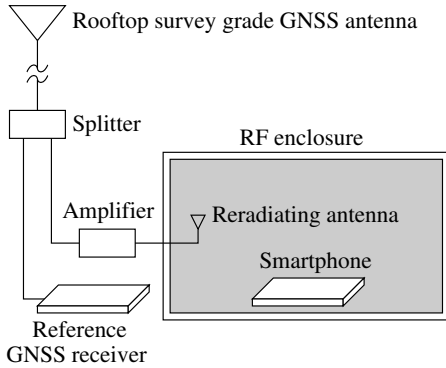


Fig. 5. Illustration of the zero-baseline measurement setup.

1) *Undifferenced Observables:* Before considering the DD observables, it is instructive to examine the smartphone's undifferenced observables under the setup in Fig. 5. The high-quality rooftop antenna and the strong re-radiated signals ensure that the smartphone's signal tracking loops operate under ideal conditions; thus, the character of the observables can be attributed to the smartphone's GNSS chip and not its built-in antenna or local environment.

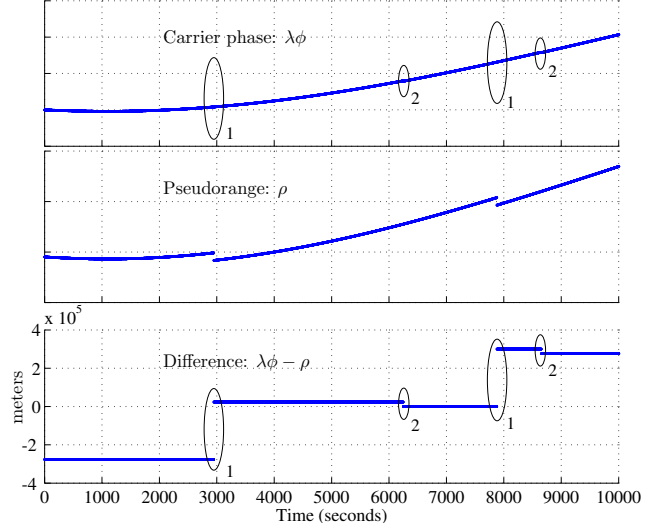


Fig. 6. Time histories of distance-equivalent carrier phase $\lambda\phi$ (top), pseudorange ρ (center), and the difference $\lambda\phi - \rho$ in meters (bottom). Type 1 and 2 carrier phase anomalies are marked.

Fig. 6 shows time histories of $\lambda\phi$, ρ , and their difference, $\lambda\phi - \rho$, for a particular GPS L1 C/A signal over a nearly 3-hour interval as output by the Galaxy S5 in the configuration of Fig. 5. The smartphone's code and carrier tracking loops apparently operated continuously over the interval, and the output carrier phase changes in the same sense as the pseudorange, pursuant to the usual convention [8]. The pseudorange measurements exhibit no unusual behavior and, as will be shown later on, yield a high-quality pseudorange-based navigation solution. However, the carrier phase in the top panel of Fig. 6 exhibits two anomalies that degrade its usefulness compared to measurements produced by a survey-grade receiver. Examination of similar time histories for other signals revealed a total of five carrier phase anomalies. Four of these are evident in the undifferenced observables, and are detailed below. The fifth is only evident in the DD observables; its discussion is delayed until the next section.

Type 1: Failure to adjust phase upon time fixup: The two discontinuities in the pseudorange time history shown in Fig. 6, both in the same direction, almost certainly reflect an adjustment of the receiver clock, as discussed in Section III-A1. But, as manifest by the large jump in the difference $\lambda\phi - \rho$, no concomitant adjustment is made to the carrier phase measurement as per the formulas in (2). Consequently, carrier phase values output by the phone after this adjustment are no longer correctly related to their time stamp. This anomaly

can readily be eliminated in post-processing by solving for the time adjustment Δt_R from the pseudorange ensemble and applying it to each carrier phase measurement as in (2).

Type 2: Small discontinuity: The carrier phase time history in the top panel of Fig. 6 exhibits two small discontinuities (marked with a 2) that are hardly visible in the $\lambda\phi$ plot but become obvious in the difference $\lambda\phi - \rho$. Such small discontinuities, amounting to approximately 122915 cycles or 23390 meters, are common throughout phase measurements produced by the Galaxy S5 and are not accompanied by any similar discontinuity in the corresponding pseudorange time history. The cause of this anomaly is unknown; it does not appear to be the result of simple cycle slipping, as it occurred for strong overhead signals received by amplified re-radiation. In any case, provided they are large enough, these discontinuities can easily be detected and eliminated in post-processing.

Type 3: Large discontinuity due to integer rollover: A phase discontinuity larger than the Type 2 anomaly is also found throughout the data. This discontinuity is exactly consistent with rollover of a 19-bit counter for the carrier phase when expressed in meters. It can be easily remedied in post-processing.

Type 4: Phase reflected: The carrier phase can suddenly become reflected so that the reflected phase's relationship to the phase without reflection is $\phi_r(t_R) = 2\phi(t_R^*) - \phi(t_R)$, where t_R^* is the receiver time at reflection, resulting in a time history that is continuous but not differentiable at t_R^* . Regardless of its cause, this phase reflection can easily be detected and fixed in post-processing.

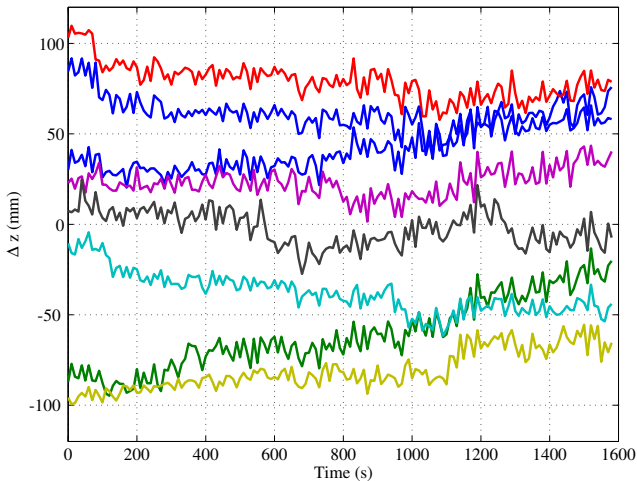


Fig. 7. Zero-baseline DD carrier phase time histories between the S5 receiver and a scientific-grade GNSS receiver.

2) *Double-Differenced Observables:* Fig. 7 shows the DD phase measurements taken under the zero-baseline configuration between the Galaxy S5 and the reference GNSS receiver, a scientific-grade software-defined GNSS receiver [10], [11]. For each trace, the integer ambiguity N_{AB}^{iq} has been estimated by standard integer least squares (ILS) [12], [13], assuming

a known zero baseline. The integer ambiguities were then removed so that, according to the DD phase model in (5), only zero-mean thermal noise $\tilde{w}_{\phi,AB}^{iq}$ should remain. Anomalies of type 1-4, discussed earlier, were either fixed or avoided in the time interval of data chosen.

Clearly, the traces in Fig. 7 are biased away from zero to a degree that cannot be confused with thermal noise variations. Moreover, there are obvious secular trends in several traces. Other intervals of zero-baseline DD phase measurements studies exhibited similar behavior. However, the DD pseudorange measurements invariably adhered to the model in (5).

The DD phase behavior in Fig. 7 constitutes a fifth type of carrier phase anomaly.

Type 5: Gradual accumulation of phase error: The carrier phase measurements produced by the Galaxy S5 gradually accumulate errors. This is not obvious in the undifferenced phase measurements, but becomes so in the DD phase measurements. The rate of error accumulation differs from signal to signal, but appears to be constant in time over long intervals, so that the accumulated error grows linearly with time. The phenomenon presumably afflicts all signals because the DD phase measurements typically depart significantly from their expected zero value. This anomaly's cause is unknown, but the effect is consistent with accumulated roundoff error. Of the five carrier phase anomalies enumerated, this is the most serious because it defies a simple remedy in post-processing. Moreover, the gradual accumulation of phase error prevents one from determining whether the relationship between each channel's carrier phase satisfies the conditions, detailed in [7], required to ensure the DD ambiguities N_{AB}^{iq} are integers.

3) *Zero-Baseline Navigation Solutions:* Navigation solutions based on the Galaxy S5 data were attempted in the zero-baseline configuration. First, a solution based on undifferenced GPS L1 C/A pseudorange measurements, broadcast ephemerides, broadcast ionospheric corrections, and a standard Saastamoinen tropospheric model [14] was calculated. The resulting 1-second epoch-by-epoch solutions, for the same span of data as in Fig. 7, is given by the blue scatter plot in Fig. 8. The scatter plot's mean, which is taken as the origin of the axes in Fig. 8, was displaced 1.4 meters horizontally and 2.3 meters vertically from the antenna's surveyed location in the International Terrestrial Reference Frame (ITRF). This mean is comfortably within the expected deviation even for a survey-grade GNSS receiver when operating in a single-point positioning mode, and the horizontal component standard deviations, which come to less than one meter, are as good as those obtained by scientific-grade receivers attached to the same antenna.

The time-varying biases in the zero-baseline DD carrier phase time histories shown in Fig. 7 are large enough (some exceeding a half wavelength) that no trust can be placed in a CDGNSS solution based on the Galaxy S5 phase data. Nevertheless, a float and fixed solution were calculated to demonstrate that, apart from the Type 5 anomaly, the smartphone's phase measurements can be made fully compatible with CDGNSS processing. The CDGNSS fixed and float




Fig. 8. Horizontal components of navigation solutions derived from the smartphone in the zero-baseline setup for the same data set as in Fig. 7. The coordinate axis origin is centered at the mean of the 1-second epoch pseudorange-based solutions (blue), whose horizontal error is 1.4 meters. The fixed and float solutions have respective horizontal errors of 21 and 25 cm.

solutions were based only on DD carrier phase measurements and on an *a priori* position estimate weighted as accurate to 10 meters. These solutions are shown in Fig. 8. It should be emphasized that no real significance can be ascribed to either of these solutions; in particular, any data-driven verification of the fixed solution would have failed: the cost of the ILS solution was not convincingly lower than competing solutions [15], [16].

C. Stand-Alone Test

A stand-alone test of the Galaxy S5's GNSS data was carried out with the phone placed face upward on the metal roof of the WRW building on the University of Texas campus (Fig. 9). This favorable use case, with the metal roof offering a protective backplane and with a nearly full view of the sky, was chosen to explore conditions only one step removed from the ideal conditions of the zero-baseline test. The phone collected data autonomously for several hours at a time during various sessions. GNSS observables were simultaneously captured at station TXAU of the U.S. Continuously Operating Receiver Station (CORS) Network, approximately 3 km west of the phone's location.

The undifferenced phase observables produced in the stand-alone test exhibited all the phase anomalies noted previously, but no new ones. Data dropouts were common for low-elevation signals, but signals above about 15 degrees elevation enjoyed nearly continuous tracking; indeed, cycle slips were rare in this regime.

Fig. 10 shows the DD residuals resulting from an ILS CDGNSS solution applied to the S5 and TXAU data, and Fig. 11 shows the fixed and float solutions together with the scatter plot from the pseudorange-based solution. Again, the CDGNSS fixed and float solutions were based only on




Fig. 9. The Galaxy S5 smartphone placed face-upward on a metal rooftop for a stand-alone test of its observables.

DD carrier phase measurements and on an *a priori* position estimate weighted as accurate to 10 meters.

As with the zero-baseline test, one cannot treat the float and fixed CDGNSS solutions from the stand alone test with any significance. Owing to the Type 5 phase anomaly, the residuals in Fig. 10 exhibit a significantly larger spread than those for a survey-grade GNSS receiver for the same baseline. Again, verification of the fixed solution by any of the methods in [15], [16] would have failed against a threshold set for high probability of correct integer ambiguity resolution despite the data spanning nearly half an hour and despite there being 8 participating satellites (including the pivot satellite).

Nonetheless, it is remarkable that the Galaxy S5 provided nearly continuous carrier phase data for 8 satellites during half an hour while exposed to an open outdoor environment. Evidently, the oscillator driving the S5's GNSS chip, and the S5's antenna, both of poor quality, are adequate to support CDGNSS positioning, at least under the (admittedly favorable) circumstances of the stand-alone test.

IV. SMARTPHONE ANTENNA CHARACTERIZATION

The accuracy and reliability of carrier-phase-based precise positioning on a smartphone depends to a significant extent on the phase and amplitude pattern of its internal antenna. Over the past two decades, surveyors and geodesists have carefully studied the problem of antenna phase center variation with satellite elevation and azimuth angle [17], [18]. Geodetic- and survey-grade GNSS antennas are designed to have a highly stable phase center and a gain pattern that strongly attenuates multipath-prone low-altitude signals while presenting a fairly uniform hemispherical gain pattern to incoming signals above ~ 5 degrees elevation. As a result, the DD phase residuals produced by a CDGNSS solution using these high-quality antennas over a short baseline exhibit ensemble variations with a standard deviation less than 5 mm (Fig. 12).

By contrast, the DD phase residuals derived from smartphone-generated GNSS observables, whether extracted by external processing, as in [1], or from the phone's internal




Fig. 10. DD carrier phase residuals time histories for a fixed CDGNSS navigation solution taken between the Galaxy S5 receiver and U.S. CORS station TXAU.




Fig. 12. Time history of DD phase residuals produced by a CDGNSS solution using a survey-grade antenna, the NovAtel Pinwheel, on a baseline of approximately 10 meters. The reference antenna was a Trimble Geodetic Zephyr II. From [1].

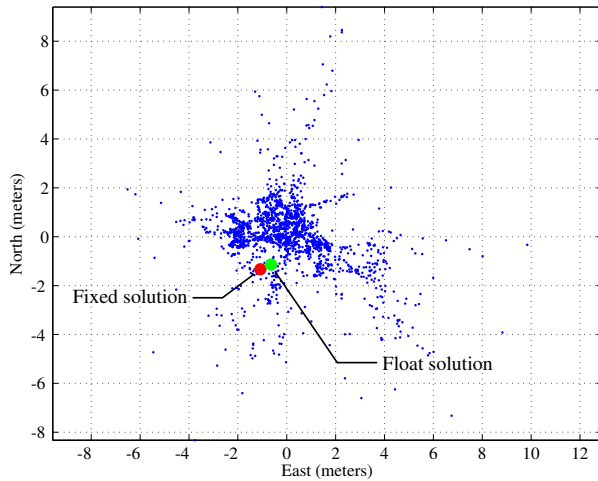


Fig. 11. Pseudorange (blue) and CDGNSS solution (fixed red, float green) for the stand-alone test. The blue scatter plot shows single-point pseudorange-based GPS L1 C/A solutions for 1750 1-second epochs. The red and green points denote attempted fixed and float CDGNSS solutions taken between the Galaxy S5 receiver and U.S. CORS station TXAU.




Fig. 13. As Fig. 12 but for a Samsung Galaxy 3 smartphone in the setup described in [1]. To ensure that the residuals represent deviations from the smartphone antenna's phase center, the location of the phase center was determined using a longer data set from which the DD phase measurements associated with large residuals were eliminated. The residuals shown are with respect to this "truth" location. From [1].

GNSS chip, as in the current paper, can exhibit standard deviations exceeding 10 mm, with temporary excursions as large as 6 cm Fig. 13. Even a low-quality and inexpensive (< 15) patch antenna, whose phase center is not so stable, nor gain pattern so uniform, as a survey-grade antenna, produces far smaller residuals than the smartphone antenna (Fig. 14).

A puzzling feature of Fig. 13 is the size and persistence of the outlier residuals. The yellow trace, which, surprisingly, corresponds to a high-elevation satellite, remains near or above 5 cm for more than 1000 seconds even though multipath for a static receiver has a time constant of only approximately 200-300 seconds. It is well known that a multipath signals whose amplitude is less than that of the direct-path signal

cannot cause phase errors larger than 1/4 wavelength, or about 4.75 cm for GPS L1 [19]. Therefore, it must be the case that the amplitude of the multipath signal associated with yellow trace exceeded the amplitude of the direct-path despite the high-elevation of the transmitting satellite. This suggests an irregular antenna gain pattern that significantly attenuates signals received in at least one "dead zone." Clearly, a study of the phone's antenna pattern is needed to better understand such behavior.




Fig. 14. As Fig. 12 but for a low-quality patch antenna. From [1].

A. Experimental Setup

A characterization of a smartphone antenna’s phase and gain pattern can be made by comparing the phase and amplitude of signals received simultaneously through the phone’s antenna and a nearby reference antenna. This technique, described in [17], yields a phase and gain antenna pattern relative to the reference antenna. If the reference antenna’s own pattern has been determined in an absolute sense, e.g., via robot calibration as in [18], then the absolute pattern of the smartphone antenna can be recovered. The setup for such a relative calibration strategy requires closely-spaced receivers in a benign multipath environment, and single-differenced carrier phase measurements, made possible by driving two GNSS receiver with the same oscillator.

Fig. 15 shows the experimental setup for the present analysis. The RF signal received through a Samsung Galaxy S3’s GNSS antenna was tapped off immediately before the phone’s internal low-noise amplifier (LNA). Note that this phone is different from the Galaxy S5 used for the results in Section III. The phone remained powered off throughout the experiment, so its internal antenna acted only as a passive element. The RF signals captured at the tap-off point were fed through an external LNA and subsequently routed to a dual-input digital recording device that downmixed the RF signals to an intermediate frequency and sampled at 10 MHz. The second input of the recording device was fed by a nearby survey-grade antenna whose signals were also sampled at 10 MHz. All downmixing and sampling was referenced to the same clock, a high-quality OCXO.

To minimize multipath into the reference antenna, a survey-grade NovAtel Pinwheel GPS-702L, it was mounted close to the surface of a metal rooftop. By contrast, the phone was cantilevered as shown in Fig. 15, with no backplane behind its internal antenna, to simulate a typical situation in which a user holds the phone by its base in her hand.



Fig. 15. Experimental setup for the Samsung Galaxy S3 antenna characterization study. The phone’s GNSS antenna and the survey-grade reference antenna, a NovAtel Pinwheel GPS-702L, are separated by approximately two meters.

B. Results

The lack of a backplane seriously degraded the Galaxy S3 antenna performance compared to that of the Galaxy S5, which was placed directly on the metal roof (Fig. 9). To appreciate this degradation, consider first the C/N_0 skyplot for the reference antenna, shown in Fig. 16. This skyplot is not equivalent to an antenna gain pattern, first because the C/N_0 values vary for reasons independent of the antenna gain (variation in transmit power between satellites, greater path loss for low-elevation satellites, local blockage and multipath effects), and second because the 24-hour GPS constellation geometry leaves gaps in the azimuth and elevation skyplot. Nonetheless, the evident regularity of the C/N_0 values in Fig. 16 would be an unlikely product of an antenna with an irregular gain pattern. In fact, when averaged over azimuth angle, the profile of C/N_0 dropoff from zenith closely matches the known elevation gain pattern of the the NovAtel Pinwheel antenna, and the fairly uniform azimuthal distribution is consistent with the Pinwheel’s known high azimuthal symmetry [20].

Now consider Fig. 17, which shows the *relative* C/N_0 skyplot for the Galaxy S3 with respect to the reference antenna. This plot is closer to a true gain pattern, as variations due to satellite transmit power, path loss, and large-scale local blockage are eliminated in the ratio (or dB difference). But, importantly, local multipath effects, which, as will be show, strongly affect the pattern, are not eliminated in the difference.

Immediately evident in Fig. 17 are patches of high or low relative signal strength. For example, at low elevations toward the west, the Galaxy S3’s antenna consistently outputs signals that have higher C/N_0 than those from the reference antenna—by up to 5 dB or so. Other “hot zones” where the phone’s signals exceed the reference antennas signals are evident. These zones all have a high degree of local correlation in azimuth and elevation angles, which leads one to think that they reflect the underlying gain pattern of the phone’s antenna, since multipath is unlikely to be so consistent over such a wide range of azimuth and elevation angles.

Likewise, one notices a few “dead zones” where relative




Fig. 16. C/N_0 skyplot for the reference antenna tracking GPS L1 C/A signals over a 24-hour period. The color bar has units of dB-Hz.




Fig. 17. C/N_0 skyplot for the Samsung Galaxy S3 relative to the reference antenna. The color bar has units of dB.

C/N_0 values are locally correlated but low, some reaching -25 dB. The dead zone to the southwest is striking for the strong gradient evident in all traces at about 27 degrees elevation and 215 degrees azimuth. Such gradients act as features in the relative pattern that help reveal whether the pattern is primarily driven by the antenna's inherent pattern or by its local multipath environment: if the features remain consistent even as the antenna is shifted horizontally or rotated about its yaw axis, then they reflect the antenna's gain pattern; if not, they must be caused by local multipath, despite whatever local consistency they manifest.

It is instructive to analyze the behavior of isolated C/N_0 traces to better understand the patterns evident in Fig. 17. Fig. 18 shows reference and phone traces for the same satellite through the interesting southwest gradient. As the satellite rises from a 5-degree initial elevation angle, the Galaxy S3's C/N_0 approximately tracks that of the reference antenna. But as the satellite rises above 25 degrees, the phone's C/N_0 drops steeply and consistently, then becomes somewhat more variable. Is the drop in C/N_0 , which is shared by all other proximate satellite signals, due to the Galaxy S3 antenna's pattern or to local multipath? *Prima facie*, it seems extraordinary that multipath should cause a steady drop so consistent in time and across azimuth angles.

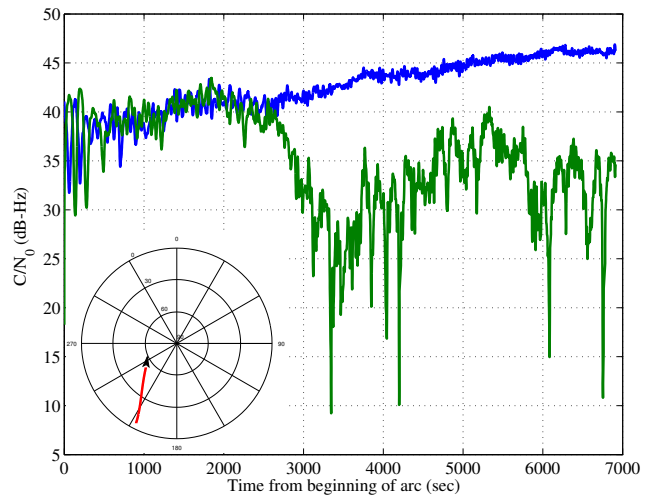


Fig. 18. Comparison of reference antenna (blue) and Galaxy S3 (green) C/N_0 values over the time span of the indicated satellite arc, which coincides with the strong southwest gradient in Fig. 17.

Fig. 19 shows C/N_0 values for another satellite arc, this one to the southeast. One notes the strikingly wide periodic oscillations in the S3's C/N_0 as the satellite rises. These variations are the hallmark of multipath: as the satellite moves and the phase of the primary multipath signal's phase becomes positively (negatively) correlated with that of the direct-path signal, constructive (destructive) interference results. Note that only a multipath-to-direct-path signal amplitude ratio above 0.8 could cause the > 20 dB peak-to-nadir variations in the green trace. The period of the variations increases as the satellite rises, starting at around 150 seconds between peaks and stretching to 500 seconds, indicating that the path length difference between the direct and multipath signals is diminishing.

Thus, it seems reasonable to conclude that rapid variations in coloration of the sky traces in Fig. 17 are caused by local multipath. It remains to be determined whether the hot and dead zones are the result of the antenna gain pattern or multipath. To investigate, the Galaxy S3 was shifted laterally by a few centimeters and another 24 hours of data were captured, resulting in the relative C/N_0 skyplot shown in Fig. 20. Remarkably, Fig. 20 appears much different from Fig.




Fig. 19. As Fig. 18 but for a different arc.

17 despite their being referenced to almost the same relative C/N_0 scale. In particular, Fig. 20 exhibits a large dead zone around 60 degrees elevation and between 330 and 0 degrees azimuth that is not present in Fig. 17. Only a faint vestige of the southwest gradient is still present in Fig. 20, whereas other fairly sharp gradients have arisen.




Fig. 20. As Fig. 17 except for a different 24-hour period and a few centimeter's lateral shift in the position of the Galaxy S3. The phone's orientation remained the same as for Fig. 17.

As a final test, the Galaxy S3 was rotated 90 degrees relative to its orientation for Figs. 17 and 20, and another 24 hours of data were captured, resulting in Fig. 21, which has been rotated to facilitate comparison with Figs. 17 and 20. But as before, no obvious correlation exists in the location of hot/dead zones and gradients between this plot and the others.




Fig. 21. As Fig. 17 except for a different 24-hour period in which the S3 was rotated 90 degrees clockwise about the normal vector to the screen relative to its orientation for Fig. 17. To facilitate comparison with Fig. 17, the skyplot has been rotated so that the pattern is aligned with the phone's body as they were for Fig. 17.

C. Discussion

Given the foregoing results, it must be concluded that when the Galaxy S3 lacks a backplane but is otherwise favorably positioned in full view of the sky, the amplitudes of the RF signals exiting its internal antenna are predominantly influenced by the local multipath environment. Therefore, attempting to determine the phone's intrinsic antenna gain pattern is pointless, as local multipath is the dominant factor in determining the azimuth- and elevation-angle dependence of output C/N_0 . Similar reasoning renders pointless the determination of the phone's intrinsic phase center variation with azimuth and elevation.

Given smartphone antennas' universally poor axial ratio and poor shielding from low-elevation signals, it is likely that these results from the Galaxy S3 can be generalized to other smartphones. This implies that, under typical circumstances in which a smartphone is held in a user's hand, the effects of strong local multipath—deep unpredictable fading and large phase errors—will be a significant challenge for centimeter-accurate smartphone positioning. If, however, the phone is supplied with a backplane, as in Fig. 9, the effect of multipath is much reduced and long periods of strong signals with no cycle slipping are possible; indeed, such was the case for Fig. 10.

V. CONCLUSIONS

A smartphone with a Broadcom GNSS chip has been modified to produce standard code and carrier phase GNSS observables in RINEX format. Analysis of the phone's observables revealed that the pseudorange measurements are of high quality and conform to expected conventions, but the

carrier phase measurements suffer from five anomalies that reduce their usefulness compared to those produced by a survey-grade GNSS receiver. However, all but one of these anomalies can be remedied in post-processing. The remaining anomaly is an error in the phase measurement that appears to grow linearly with time. The error is obvious in zero-baseline double-differenced phase measurements. It may be the result of accumulated roundoff error. But the phone appears otherwise to be fully capable of supporting cm-accurate carrier-phase differential GNSS positioning, and produces solutions with plausible locations, albeit with low probability of correct integer fix.

A separate smartphone's antenna was studied to determine whether a sensible gain pattern could be inferred from differential carrier-to-noise ratio skyplots taken with respect to a reference antenna having a known gain pattern. For this test, the phone was mounted on a cantilever attached to a building so that its GNSS antenna enjoyed no external backplane, a situation meant to simulate the common hand-held smartphone use case. It was found that whatever intrinsic gain pattern the phone's antenna has was masked by strong local multipath effects. One can conclude that, under typical smartphone use, the effects of local multipath—namely, deep, unpredictable fading and large phase errors—will be a significant challenge for centimeter-accurate smartphone positioning.

ACKNOWLEDGMENTS

This work was supported in part by the Data-supported Transportation Operations and Planning Center (D-STOP), a Tier 1 USDOT University Transportation Center.

ENDNOTE

Author Todd Humphreys is Scientific Advisor to Radiosense, a startup developing technology for low-cost precise positioning.

REFERENCES

- [1] K. M. Pesyna, Jr., R. W. Heath, Jr., and T. E. Humphreys, "Accuracy in the palm of your hand: Centimeter positioning with a smartphone-quality GNSS antenna," *GPS World*, vol. 26, no. 2, pp. 16–31, Feb. 2015.
- [2] K. M. Pesyna, Jr., R. W. Heath, Jr., and T. E. Humphreys, "Centimeter positioning with a smartphone-quality GNSS antenna," in *Proceedings of the ION GNSS+ Meeting*, 2014.
- [3] F. van Diggelen, "Expert advice: Are we there yet? The state of the consumer industry," *GPS World*, Mar. 2010, GPS World.
- [4] A. J. Van Dierendonck, "How GPS receivers measure (or should measure) ionospheric scintillation and TEC and how GPS receivers are affected by the ionosphere," in *Proc. 11th International Ionospheric Effects Symposium*, Alexandria, VA, 2005.
- [5] —, *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 8: GPS Receivers, pp. 329–407.
- [6] K. D. Wesson, K. M. Pesyna, Jr., J. A. 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.
- [7] 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.
- [8] W. Gurtner and L. Estey, "RINEX-The receiver independent exchange format-version 3.00," *Astronomical Institute, University of Bern and UNAVCO, Boulder, Colorado.*, 2007.
- [9] O. Montenbruck, J. Williams, T. Wang, and G. Lightsey, "Preflight validation of the IGOR GPS receiver for TerraSAR-X," *DLR/GSOC Document No: GTN-TST-0200*, no. 1.2, p. 36, 2005.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] J. Saastamoinen, "Atmospheric correction for the troposphere and stratosphere in radio ranging of satellites," in *Geophysical Monograph 15*, S. W. Henriksen, Ed. Washington, D.C.: American Geophysical Union, 1972, pp. 247–251.
- [15] S. Verhagen, "On the reliability of integer ambiguity resolution," *Navigation, Journal of the Institute of Navigation*, vol. 52, no. 2, pp. 99–110, 2005.
- [16] S. Verhagen and P. J. Teunissen, "New global navigation satellite system ambiguity resolution method compared to existing approaches," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 4, pp. 981–991, 2006.
- [17] G. L. Mader, "GPS antenna calibration at the National Geodetic Survey," *GPS solutions*, vol. 3, no. 1, pp. 50–58, 1999.
- [18] G. Wübbena, M. Schmitz, F. Menge, V. Böder, and G. Seeber, "Automated absolute field calibration of GPS antennas in real-time," in *Proceedings of the ION GPS Meeting*, 2000, pp. 2512–2522.
- [19] J. J. Spilker, Jr., *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 14: Multipath Effects, pp. 547–568.
- [20] L. I. Basilio, J. T. Williams, D. R. Jackson, and M. A. Khayat, "A comparative study of a new GPS reduced-surface-wave antenna," *Antennas and Wireless Propagation Letters, IEEE*, vol. 4, pp. 233–236, 2005.