

## GNSS-BASED POSITIONING



# Deep-Urban Unaided Precise Global Navigation Satellite System Vehicle Positioning

**Todd E. Humphreys and Matthew J. Murrian**

*Is with the Department of Aerospace Engineering  
and Engineering Mechanics,  
the University of Texas at Austin.*

*Email: todd.humphreys@utexas.edu;  
matthew.murrian@utexas.edu*

**Lakshay Narula**

*Is with the Department of Electrical  
and Computer Engineering,  
the University of Texas at Austin.*

*Email: lakshay.narula@utexas.edu*

**Abstract**—This article presents the most thorough study to date of vehicular carrier-phase differential Global Navigation Satellite System (CDGNSS) positioning performance in a deep urban setting unaided by complementary sensors. Using data captured during approximately 2 h of driving in and around the dense urban center of Austin, Texas, a CDGNSS system is demonstrated to achieve 17-cm-accurate 3D urban positioning (95% probability) with a solution availability greater than 87%. The results are achieved without the aid of inertial, electro-optical, or odometry sensors. The development and evaluation of the unaided, GNSS-based precise positioning system is a

Digital Object Identifier 10.1109/MITS.2020.2994121  
Date of current version: 11 June 2020

key milestone toward the overall goal of combining precise GNSS, vision, radar, and inertial sensing for all-weather, high-integrity, high-absolute-accuracy positioning for automated and connected vehicles. The system described and evaluated herein is composed of a densely spaced reference network, a software-defined GNSS receiver, and a real-time kinematic (RTK) positioning engine. A performance-sensitivity analysis reveals that navigation data wipeoff for fully modulated GNSS signals and a dense reference network are key to high-performance urban RTK positioning. A comparison with existing unaided systems for urban GNSS processing indicates that the proposed system has significantly greater availability or accuracy.

Future vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) connectivity will permit vehicles to relay their positions and velocities to each other with millisecond latency, enabling tight, coordinated platooning and efficient intersection management. More ambitiously, broadband V2V and V2I enabled by 5G wireless networks will permit vehicles to share unprocessed or lightly processed sensor data. Ad hoc networks of vehicles and infrastructure will then function as a single sensing organism. The risk of collisions, especially with pedestrians and cyclists—notoriously unpredictable and much harder to sense reliably than vehicles—will be significantly reduced as vehicles and infrastructure contribute sensor data from multiple vantage points to build a blind spot-free model of their surroundings.

Such collaborative sensing and traffic coordination requires vehicles to know and share their own position. How accurately? The proposed Dedicated Short Range Communications basic safety message, a first step in V2V coordination, does not yet define a position accuracy requirement, effectively accepting whatever accuracy a standard Global Navigation Satellite System (GNSS) receiver provides [1]. But automated intersection management [2], tight-formation platooning, and the unified processing of sensor data—all involving vehicles of different makes that may not share a common map—will be greatly facilitated by globally referenced positioning with sub-30-cm accuracy.

Poor weather also motivates high-accuracy absolute positioning. Every automated vehicle initiative of which the authors are aware depends crucially on lidar or cameras for fine-grained positioning within their local environment. But, these sensing modalities perform poorly in low-visibility conditions such as a snowy whiteout, dense fog, or heavy rain. Moreover, high-definition 3D maps cre-

ated with lidar and camera data, maps that have proven crucial to recent progress in reliable vehicle automation, can be rendered dangerously obsolete by a single snowstorm, leaving vehicles relying on such maps for positioning no option but to fall back on GNSS and radar to navigate a snow-covered roadway in low-visibility conditions. When, as is often the case on rural roads, such snowy surroundings offer few radar-reflective landmarks, radar also becomes useless. GNSS receivers operate well in all weather conditions, but only a highly accurate GNSS solution, e.g., one whose absolute errors remain under 30 cm 95% of the time, could prevent a vehicle's drifting onto a snow-covered road's soft shoulder. Code- and Doppler-based GNSS solutions can be asymptotically accurate (averaged over many sessions) to better than 50 cm, which may be adequate for digital mapping [3], but they will find it challenging to meet a 30-cm 95% standalone requirement, even with modernized GNSS offering wideband signals at multiple frequencies.

Carrier-phase-based GNSS positioning—also referred to as *precise GNSS positioning* even though it actually offers absolute accuracy, not just precision (repeatability)—can meet the most demanding accuracy requirements envisioned for automated and connected vehicles but has historically been either too expensive or too fragile, except in open areas with a clear view of the overhead satellites, for widespread adoption. Coupling a carrier-phase differential GNSS (CDGNSS) receiver with a tactical-grade inertial sensor, as in [4]–[7], enables robust high-accuracy positioning even during the extended signal outages common in dense urban areas. But GNSS-inertial systems with tactical-grade inertial measurement units (IMUs) cost tens of thousands of U.S. dollars and have proven stubbornly resistant to commoditization. Coupling a GNSS receiver with automotive- or industrial-grade IMUs is much more economical and significantly improves performance, as shown in [8]. But such coupling allows only approximately 5 s of complete GNSS signal blockage before the IMU no longer offers a useful constraint for so-called integer-ambiguity resolution [9], which underpins the fastest, most accurate, and most robust CDGNSS techniques, namely, single-baseline real-time kinematic (RTK), network RTK, and Precise Point Positioning RTK (PPP-RTK) [10], [11].

Previous research has suggested an inexpensive technique for robustifying RTK positioning: tightly coupling carrier-phase-based GNSS positioning with inertial sensing and vision [12], [13]. Such coupling takes advantage of the remarkable progress in high-resolution, low-cost cameras within the intensely competitive smartphone market. We are currently engaged in developing a high-integrity RTK-vision system for high-accuracy vehicular positioning in rural and urban environments. Further coupling with radar will make the system robust to low-visibility conditions. As a step toward this goal, it is of interest to evaluate the performance of standalone RTK techniques—those

unaided by IMUs, odometry, or vision—in urban environments. This type of study will reveal why and when aiding is necessary and how an RTK positioning system might behave if aiding were somehow impaired or unavailable, whether due to sensor faults or, in the case of exclusive visual aiding, poor visibility conditions.

Little prior work has explored unaided vehicular RTK performance in urban environments, no doubt because performance results have historically been dismal.

Short-baseline RTK experiments between two vehicles in [14] revealed that multifrequency (L1-L2) GPS and GLONASS RTK yielded poor results in residential and urban environments. Only along a mountain highway with a relatively clear view of the sky was availability greater than 90% and accuracy better than 30 cm. RTK positioning in downtown Calgary was disastrous, with less than 60% solution availability and root-mean-square errors exceeding 9 m.

More recently, Li et al. [8] have shown that, with the benefit of greater signal availability, unaided, professional-grade dual-frequency GPS + BDS + GLONASS RTK can achieve correct integer-fixing rates of 76.7% on a 1-h drive along an urban route in Wuhan, China, where BDS refers to the BeiDou Navigation Satellite System. But Li et al. do not provide data on the incorrect fixing rate, nor a full error distribution, so the significance of their results is difficult to assess.

Recent urban RTK testing by Jackson et al. [15] indicates that no low-to-mid-range consumer RTK solution offers greater than 35% fixed (integer-resolved) solution and availability in urban areas, despite a dense reference network and dual-frequency capability. A key failing of existing receivers appears to be their slow recovery after passing under bridges or overpasses.

This article describes and evaluates an unaided RTK positioning system that has been designed for vehicular operation in both rural and urban environments. Preliminary performance results were published in a conference version of this article [16]; however, this article improves upon the conference version in four ways: 1) the test route is both more challenging and more comprehensive, 2) a proper independent ground-truth trajectory is used as the basis of error evaluation, 3) data-modulation wipeoff for improved carrier-tracking robustness is applied not only on GPS L1 C/A signals, as previously, but now also on satellite-based augmentation system (SBAS) L1 signals, and 4) the performance benefit of vehicle GNSS antenna calibration is assessed.

This article's primary contributions are 1) a demonstration of the performance that can be achieved with a

Carrier-phase-based GNSS positioning—also referred to as *precise GNSS positioning* even though it actually offers absolute accuracy, not just precision (repeatability)—can meet the most demanding accuracy requirements envisioned for automated and connected vehicles but has historically been either too expensive or too fragile, except in open areas with a clear view of the overhead satellites, for widespread adoption.

low-cost, software-defined unaided RTK GNSS platform in a dense urban environment and 2) an evaluation of the relative importance of various factors (e.g., data bit wipeoff, age of reference data, rover antenna calibration, and reference network density) to the overall system performance.

To stimulate further innovation in urban precise positioning, all of the data from this article's urban driving campaign have been posted at <http://radionavlab.ae.utexas.edu> under “Public Datasets: ATX Urban Positioning Challenge Dataset” including wideband (10-MHz) intermediate-frequency (IF) samples from both the reference and rover antennas, RINEX-formatted rover and reference observables, and the ground-truth trajectory.

### Challenges of Mobile Precise Positioning in Urban Environments

The mobile urban satellite-to-user channel is distinguished by rapid channel evolution. As the vehicle travels along streets closely lined with tall buildings, only glimpses of power are available from signals arriving from directions roughly perpendicular to the roadway. A GNSS receiver designed to provide phase-locked carrier measurements for RTK positioning in such environments must simultaneously prevent frequency unlock during the deep fades caused by building occlusions and exploit momentary signal availability by immediately acquiring full-cycle phase lock and indicating this to downstream processing.

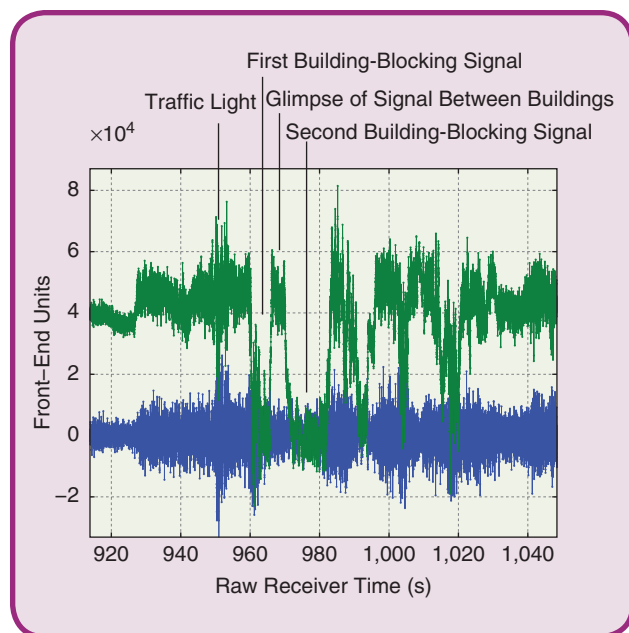
Tracking in the mobile urban channel is unlike indoor or weak-signal tracking, such as that which is explored in [17] and [18], in that the urban fading environment is substantially binary: either the line-of-sight signal is present at a fairly healthy carrier-to-noise ratio,  $C/N_0$ , or it is hopelessly attenuated after passing through entire buildings constructed of concrete, steel, and glass. The traditional weak-signal-tracking technique of extending the signal-integration time and lowering the tracking-loop bandwidths can be useful to slow the rate of frequency unlock during such fading but not for actually recovering a weak signal from the noise, as there is simply no signal to recover.



Figure 1 illustrates this fact. The initial disturbance at 950 s is due to an overhead traffic light. This is followed in rapid succession by a complete signal blockage due to a tall building on the south side of the east-west street, a brief (4-s) interval of clear satellite availability as the receiver catches a glimpse of the signal between two buildings, and another signal eclipse by a second building.

A GNSS receiver designed for urban tracking will make full use of such between-building glimpses. This requires the immediate (within approximately 100 ms) recovery of full-cycle phase lock, which is only possible on suppressed-carrier signals like GPS L1 C/A if the receiver can accurately predict the modulating data symbols. Downstream RTK processing must also be poised to exploit signal glimpses by identifying and rejecting observables from blocked or otherwise compromised signals and by immediately re-evaluating the corresponding integer ambiguities when signals reappear. A multistage cycle-slip detection and recovery technique, such as the one proposed in [19], is too slow for urban positioning.

A related hallmark of the urban mobile channel is the wide and rapid variation of the number of signals available for RTK positioning. The number  $N_{DD}$  of double-difference (DD) signals (each one providing a DD pseudorange and a DD carrier phase observable) varies widely whenever the vehicle is moving. The implication for RTK processing is that integer-ambiguity continuity will often be lost, requiring the rapid and continuous re-estimation of ambiguities.



**FIG 1** In-phase (green, top) and quadrature (blue, bottom) 10-ms complex-correlation products for a GPS L1 C/A signal at 35° elevation arriving from the south to a vehicle traveling west on an urban roadway. The 20-ms LNAV navigation data bits have been wiped off to allow for full carrier cycle recovery. Rapid fading—and rapid recovery—occur as buildings intermittently block the signal.

## System Description

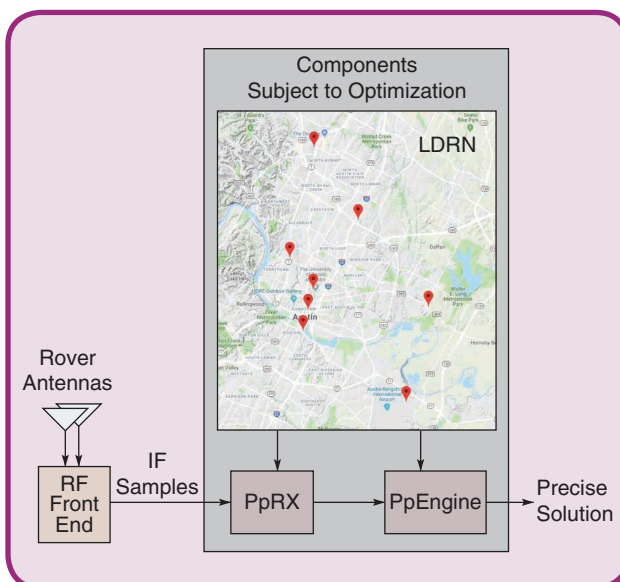
### Overview

The GNSS components of this article’s precise positioning system are shown in Figure 2. The subcomponents enclosed in the gray box are the target of the present work’s optimization efforts for good performance in urban environments.

Two rover antennas feed analog signals to a radio-frequency (RF) front end, which downmixes and digitizes the signals, producing a stream of IF samples. The RF front end used in the present work produces samples at 10 MHz for two antennas and two frequencies: a band centered at GPS L1 and one centered at GPS L2. The (single-sided) analog bandwidth of each band is 4 MHz—wide enough to capture more than 90% of the power in the GPS L1 C/A, Galileo E1 BOC(1,1), and GPS L2C signals.

Four IF sample streams, one for each antenna and band, are fed to PpRx, an embeddable, multifrequency software-defined GNSS receiver developed primarily at the University of Texas at Austin [20]–[22]. PpRx draws ephemeris data, GPS LNAV, and SBAS [Wide Area Augmentation System (WAAS)] data bit estimates from the Longhorn Dense Reference Network (LDRN), a set of eight GNSS reference stations deployed in Austin, Texas. Each reference station in the LDRN runs a strict real-time variant of PpRx and sends its data to a central network server, from which any compatible receiver can draw assistance data and network observables.

PpRx feeds code and carrier observables—and other useful signal information—to an RTK engine called PpEngine. For the results presented in this article, PpEngine draws observables and ephemeris data from a single LDRN reference station at a time, the traditional RTK topology. The precise solution produced by PpEngine is a fixed



**FIG 2** The University of Texas’ precise positioning system. RF: radio frequency; LDRN: Longhorn Dense Reference Network.

(integer-resolved) or float solution depending on the results of an integer aperture test [23].

### Performance Metrics

The performance of precise positioning systems in safety-of-life applications is assessed in terms of integrity, accuracy, and availability [24], [25]. For several emerging applications of practical interest, such as automated and connected vehicles, no regulatory body has set clear positioning performance requirements. An industry consensus appears to be emerging, one that calls for a 95% accuracy requirement of 30 cm, but it is not clear what the associated integrity risk or continuity requirements should be. It is likely that the U.S. National Highway Traffic Safety Administration, and other regulatory bodies worldwide, will eventually issue positioning performance requirements for connected and automated vehicles.

This article focuses on four related performance metrics: 1)  $d_{95}$ , the 95th-percentile error magnitude for fixed solutions; 2)  $P_V$ , the probability (or rate in continuous trials) that a validated (fixed) solution is available at each epoch, as opposed to a fallback float solution; 3)  $P_S$ , the probability of correctly (successfully) resolving the full integer set at each epoch; and 4)  $P_F$ , the probability that one or more integer ambiguities failed to resolve correctly at each epoch [24].  $P_V$ ,  $P_S$ , and  $P_F$  are related by  $P_V = P_S + P_F$ . A fourth probability,  $P_U = 1 - P_F - P_S = 1 - P_V$ , that of the undecided event, is the probability that a float solution, or no solution at all, is produced, due to an aperture test failure or failure of some other validation test.

An unavoidable tradeoff between  $P_S$  and  $P_F$  exists such that any widening of the integer aperture region to increase  $P_S$  comes at the expense of an increase in  $P_F$  (not necessarily of the same amount) [26]. Therefore, an optimization problem can be stated in terms of  $P_S$  and  $P_F$  as follows: maximize  $P_S$  for  $P_F \leq \bar{P}_F$ , where  $\bar{P}_F$  is a fixed tolerable probability of failed fixing. Integer aperture bootstrapping techniques such as those presented in [26] and its generalization to partial ambiguity resolution in [24] analytically determine thresholds for the integer aperture test to ensure  $P_F \leq \bar{P}_F$ . For the optimal integer least squares (ILS) approach adopted in this article, it is not possible to calculate an analytical aperture threshold, but an approximate one can be obtained via simulation such that  $P_F \leq \bar{P}_F$  is satisfied almost surely [27]. A value of  $\bar{P}_F = 0.001$  was adopted for this article, meaning that a fixing failure rate of less than one in 1,000 epochs was deemed acceptable. However, multipath, GNSS signal passage through foliage and other signal impairments common in urban areas cause the empirical  $P_F$  to significantly exceed  $\bar{P}_F$  when the aperture threshold is chosen accord-

An unavoidable tradeoff between  $P_S$  and  $P_F$  exists such that any widening of the integer aperture region to increase  $P_S$  comes at the expense of an increase in  $P_F$ .

ing to the Gaussian error assumptions ubiquitous in the integer aperture literature. Thus, a looser empirical upper bound,  $\bar{P}_F$ , must be chosen. The optimization problem is then to maximize  $P_S$  subject to the empirical  $P_F$  respecting the bound  $P_F \leq \bar{P}_F$ .

### Design Philosophy

With origins in scintillation-resistant carrier tracking [28], [29], PpRx was designed from the beginning for robust carrier recovery. Likewise, from its inception, PpEngine was targeted for the harsh urban environment. Over the past few years, the development of PpRx, PpEngine, and the LDRN has proceeded as a parallel evolution, with each subsystem benefiting from the improvements in the others.

The overriding design philosophy of this development has been to adapt, rebuild, and reconfigure all three subsystems, separately and in parallel, with the goal of minimizing  $d_{95}$  while maximizing  $P_V$ , or, relatedly, maximizing  $P_S$  subject to  $P_F \leq \bar{P}_F$ . This approach benefits greatly from a purely software-based approach to GNSS signal processing (as opposed to processing that exploits dedicated silicon or field-programmable gate arrays) for two reasons. First, a software-defined approach is almost infinitely flexible: all processing downstream from the RF front end can be reconsidered, rebuilt, and re-evaluated in a rapid iterative process using an efficient and common high-level programming language. Second, software-defined receivers can exploit multiple cores to run faster than real time on recorded IF samples [21]. The PpRx-PpEngine pipeline runs at 10-times real time on a six-core Intel Xeon 2.27-GHz processor, enabling rapid iteration cycles for quickly probing the optimization landscape.

### Carrier and Code Tracking

GNSS carrier and code tracking in an urban environment must be opportunistic, taking advantage of short, clear glimpses of overhead satellites as they present themselves. PpRx's code- and carrier-tracking architecture, illustrated in Figure 3, has been designed for the immediate (within approximately 100 ms) recovery of full-cycle phase lock after a blockage, and, importantly, for prompt lock indication. The following sections describe the essential elements of PpRx's tracking strategy, calling out parameters whose values significantly affect urban RTK performance.

### Correlation and Accumulation

Correlation and accumulation is performed on a sequence of noisy IF samples  $x(\tau_j)$ ,  $j = 0, 1, \dots$ , where  $\tau_j$  denotes the time of the  $j$ th sample according to the receiver's clock. Within the correlation and accumulation block, a complex, local replica signal is formed with code and carrier-phase estimates  $\hat{t}_s(\tau_j)$  and  $\hat{\theta}(\tau_j)$  provided by the code and carrier-phase-tracking loops. The outputs of the correlation and accumulation block are prompt, early, and late complex-correlation products  $S_k$ ,  $S_{e,k}$ , and  $S_{l,k}$  of the form  $S_k = I_k + jQ_k$ , where  $I_k$  and  $Q_k$  are the in-phase and quadrature accumulations. (The green and blue traces in Figure 1 correspond to  $I_k$  and  $Q_k$ , respectively.) The accumulation interval,  $T_a$ , is an important configuration parameter for urban RTK.

### Navigation Data Bit Wipeoff

The GPS L1 C/A and SBAS L1 signals have no dedicated pilot component. The phase ambiguity introduced by their full, suppressed-carrier binary modulation makes it challenging to recover an accurate carrier-phase measurement in an urban environment. However, the reference network can provide low-latency estimates  $\hat{d}_k$  with which the incoming modulation can be "wiped off," allowing full-cycle carrier recovery.

GPS L1 C/A data bit wipeoff has been employed for years to improve weak-signal acquisition in smartphones [30] but, as far as we are aware, it has not been previously applied in the context of CDGNSS positioning. SBAS L1 data

wipeoff, a novel technique introduced in this article, is even more valuable on a per-signal basis than GPS L1 C/A data wipeoff because the short, 2-ms binary SBAS symbol period otherwise renders SBAS signals of little use for urban-precise positioning.

### Lock Statistic Calculation

Also key to robust urban RTK is the ability to exclude corrupt or otherwise inaccurate carrier-phase measurements. However, due to poor signal availability, an urban RTK engine cannot afford to be overly conservative: it must minimize the number of adequate-quality measurements that get falsely labeled as corrupt. An important indicator for this wheat-from-tares separation is the lock statistic  $s_\theta$ . Let  $I$  and  $Q$  be coherent sums of  $I_k$  and  $Q_k$  over  $N_L$  accumulation intervals. Then  $s_\theta$  is calculated as [31]

$$s_\theta = \frac{I^2 - Q^2}{I^2 + Q^2}.$$

The goal of the carrier-tracking loop is to adjust its phase estimate  $\hat{\theta}_k$  to shift signal power from  $Q_k$  to  $I_k$ . Thus, for a loop in lock,  $I^2 \gg Q^2$  and  $s_\theta$  is near unity.

A new lock statistic is produced every  $N_L$  accumulations.  $N_L$  must be chosen large enough to suppress thermal noise in  $I_k$  and  $Q_k$ , but small enough to provide a prompt indicator of phase lock to all dependent processing. PpEngine relies crucially on  $s_\theta$  to screen out bad measurements. Note from Figure 3 that  $s_\theta$  is also fed to the code-tracking loop and to PpRx's central state estimator: each one adapts

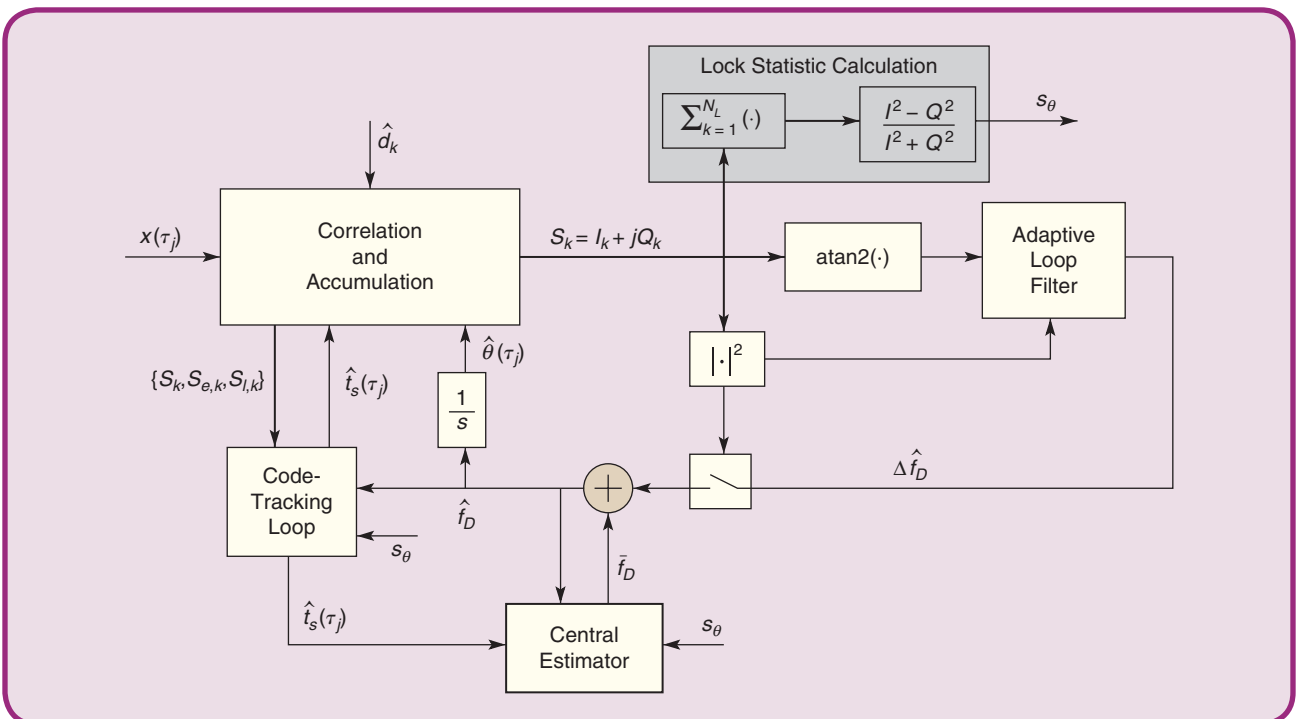


FIG 3 The PpRx's carrier- and code-tracking architecture.

its behavior to rely less on Doppler measurements when  $s_\theta$  is low.

### Carrier Tracking

As shown in Figure 3, PpRx employs a vector signal-tracking architecture wherein a central estimator, implemented as a Kalman filter with a nearly constant-velocity dynamics model, receives observables from all tracking channels and drives local replica generation for each channel [32]. More particularly, PpRx employs a hybrid strategy in which, for each channel, a local phase-tracking loop is closed around a modeled Doppler value  $\hat{f}_D$  provided by the central estimator. The local loop's residual Doppler frequency estimate  $\Delta\hat{f}_D$  is added to  $\hat{f}_D$  to produce the full estimate  $\hat{f}_D$  used in replica generation.

A four-quadrant arctangent phase discriminator,  $a \tan 2(Q_k, I_k)$ , which is nearly optimal for decision-directed carrier recovery and optimal for data-free signals, or when data bit wipeoff is error-free, feeds a phase-error measurement at every accumulation interval to the carrier-tracking loop filter. PpRx's carrier loop filter is designed according to the controlled-root formulation of [33]. The filter adapts its bandwidth  $B_\theta$  at every accumulation interval according to the value of  $|S_k|$ . The adaptation schedule has a significant effect on RTK performance.

One might expect that adapting  $B_\theta$  to maintain a constant-loop signal-to-noise ratio as  $|S_k|$  varies would yield the best results. This is effectively the adaptation schedule that gets applied in Kalman filter-based weak-signal tracking [17]. However, this reasonable approach was found to yield reduced urban RTK performance. More effective is a three-tiered schedule that reduces  $B_\theta$  when  $|S_k|$  falls below a fairly low threshold,  $\gamma_1$ , and sets  $B_\theta$  to zero if  $|S_k|$  falls below another threshold, that is,  $\gamma_0 < \gamma_1$ . Within this lowest tier,  $\Delta\hat{f}_D$  is also driven to zero over a few accumulation intervals, thereby breaking the local feedback loop. In this open-loop mode, the local replica's phase estimate is driven entirely by the model Doppler  $\hat{f}_D$ . The lock statistic  $s_\theta$  continues to be calculated. If  $s_\theta$  is sufficiently close to unity, the central estimator, code-tracking loop, and RTK engine continue to treat  $\hat{\theta}(\tau_j)$  as a valid measurement. But this is a rare occurrence;  $s_\theta$  is typically far from unity in open-loop mode.

Such open-loop tracking is useful for preventing frequency unlock during intervals when signals are entirely blocked, e.g., by buildings or bridges, and for enabling fast reacquisition of carrier lock immediately following the blockage.

### Code Tracking

PpRx's code-tracking loop, which is aided by the Doppler estimate  $\hat{f}_D$ , is implemented as a first-order loop that toggles between a noncoherent (dot product) discriminator and a coherent discriminator. The coherent discriminator

is applied when the channel is phase locked and no recent phase trauma (indicated by  $s_\theta$ ) has been detected; otherwise, the noncoherent discriminator is applied. A flag attached to each code-phase measurement  $\hat{I}_s(\tau_j)$  indicates to downstream processes whether it was produced under coherent or noncoherent tracking.

As with carrier tracking, the code-tracking loop filter's bandwidth,  $B_c$ , is adaptive. But rather than responding to  $|S_k|$  as the carrier loop's bandwidth does,  $B_c$  takes on a different value for each of four code-tracking modes: 1) prephase lock, 2) first postlock transient, 3) second postlock transient, and 4) steady state. These modes are designed to ensure the rapid convergence of the code-phase estimate  $\hat{I}_s(\tau_j)$  after initial signal acquisition or in the aftermath of phase unlock.

### Precise Positioning

PpRx and the LDRN send carrier and code phase observables, together with signal quality indicators  $s_\theta$  and  $C/N_0$ , and various other metadata to PpEngine for processing. PpEngine is capable of processing observables from both rover antennas simultaneously, exploiting the known distance between these. But for the results presented in this article, PpEngine was invoked only in its simplest single-antenna mode, producing a precise 3D baseline between the primary rover antenna and a selected reference station antenna in the LDRN. This simple, single-baseline RTK mode was chosen so that the precise positioning system's performance could be evaluated in a familiar configuration and easily compared with other single-baseline RTK evaluations such as those in [8].

### Treatment of Real- and Integer-Valued States

The current embodiment of PpEngine adopts a straightforward approach to RTK. It first forms code and carrier measurement DDs from the rover and reference data and then sends these to a mixed real/integer extended Kalman filter for processing. The filter is implemented as a square-root information filter, as in [34], but limits the growth of the number of integer states by either marginalizing at each epoch over float-valued integer-ambiguity states modeled as Gaussian-distributed or conditioning on the estimated integer values. Thus, PpEngine's current approach is to discard all integer states, by marginalization or by conditioning, after each measurement epoch. The marginalization option, which yields the float solution, can be thought of as a special case of the suboptimal filter in [34] with window length  $i = 1$ . The conditioning option, which yields the fixed solution, is invoked only if the integer estimates, found by ILS [35], are validated by an aperture test.

Conditioning the real-valued states on the lowest-cost integer estimates yields a maximum a posteriori 3D baseline estimate. After each measurement update, the real-valued states are propagated to the next

measurement epoch, whereupon a new set of integer estimates are formed and conditioning or marginalization occurs yet again. Importantly, if the integer states are validated at the  $l$ th measurement epoch, it is the integer-conditioned real-valued states that are propagated to the  $(l + 1)$ th measurement epoch. Thus, although all integer states are discarded between measurement updates, correct integer resolution is highly likely at the  $(l + 1)$ th epoch if integer ambiguities were correctly resolved at the  $l$ th epoch because the real-valued states carry forward a decimeter-accurate position estimate.

Carrying forward integer-conditioned real-valued states is perilous because eventually an erroneous integer estimate passes the aperture test, whereupon the integer-conditioned real-valued states are corrupted by conditioning on the incorrect fix. What is more, the associated square-root information matrices indicate high confidence in the erroneous real-valued state, raising the chances that the next integer estimates, which are constrained by the prior real-valued states, will also be incorrectly fixed. This cycle, which can persist for several seconds, is eventually broken

by an aperture test failure prompted by signal loss, large measurement errors, or the persistent lack of consistency between the incoming observables and the current state.

In view of this peril, we are developing a generalization of PpEngine that can manage growth in the number of integer state elements using a variant of the suboptimal approach depicted in [34]. Meanwhile, PpEngine's single-epoch integer resolution has the virtue of being insensitive to the cycle slips that occur between measurement epochs, which are common in the urban environment.

#### Dynamics Model

Because this article's focus is on RTK unaided by any non-GNSS sensors, the mixed real- and integer-valued state estimator within PpEngine was configured to ignore all of the available inertial measurements and instead rely on a simple, nearly constant-velocity dynamics model for state propagation between measurements. The dynamics model assumes roughly equivalent process noise variance in the along- and cross-track directions, but smaller variance (by a factor of 100) in the vertical direction, in keeping with a land vehicle operating in a relatively flat urban environment.

#### Robust Measurement Update

Urban multipath and diffraction cause code and carrier observables to exhibit large errors with a much higher probability than even a conservative Gaussian model would predict. Dealing with measurement-error processes such as these, which have thick-tailed distributions, requires robust estimation techniques; that is, techniques with reduced sensitivity to measurement outliers.

Outliers are especially problematic for integer fixing in RTK positioning. By action of the decorrelation adjustment preceding ILS, a single bad measurement can contaminate multiple measurements in the decorrelated domain, rendering resolution of the associated integers impossible. Partial ambiguity resolution, as in [36] and [37], offers little relief in such cases because contamination caused by outliers is not necessarily limited to an identifiable subset of integers. It is more effective to exclude questionable measurements before the decorrelation adjustment.

PpEngine implements a multilevel exclusion process, depicted in Figure 4, to mitigate the effects of measurement outliers. At each measurement epoch, measurements are first screened on the basis of three quality indicators: carrier-to-noise ratio  $C/N_0$ , phase lock statistic  $s_\theta$ , and elevation angle  $\theta_{el}$ . Signals whose values fall below user-selected thresholds for these quantities are excluded from all DD combinations.

A second level of exclusion occurs as part of the float solution. A  $\chi^2$ -type test is applied to all DD measurement innovations [38], with exclusion triggered if the normalized innovations squared statistic exceeds a chosen threshold. For the current implementation of PpEngine, this test is only

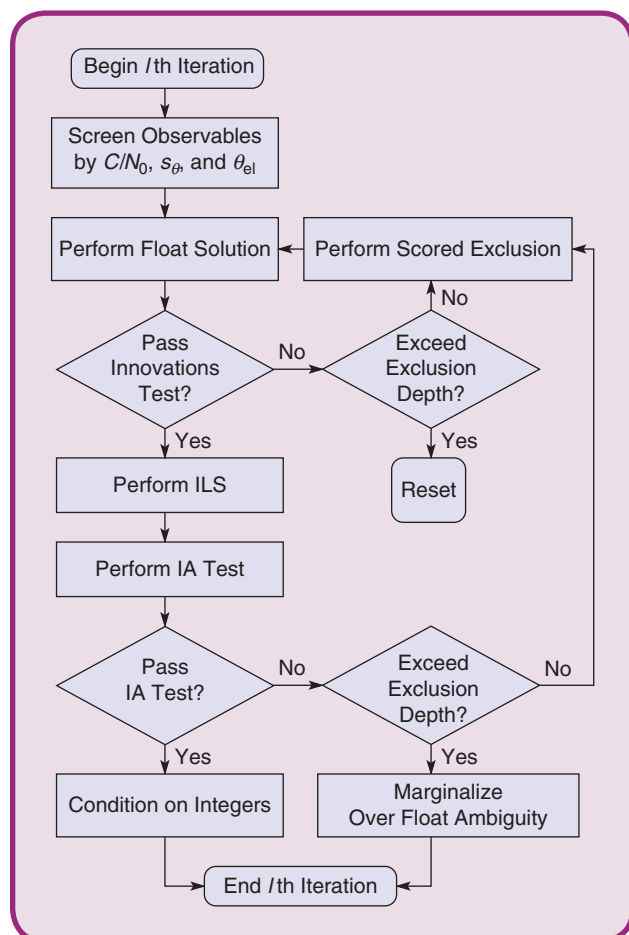


FIG 4 The flow diagram for the PpEngine exclusion and fixing logic. IA: integer aperture.



effective at excluding anomalous DD code-phase (pseudorange) measurements because the float states are discarded, and thus unconstrained, from epoch to epoch. Note that innovations testing benefits strongly from a correctly integer-constrained state because the exclusion threshold can be made tighter. However, with an incorrectly integer-constrained state, innovations testing may end up excluding the very measurements necessary to correct the state.

If a set of innovations fails the innovations test, DD measurements (both code and carrier for a particular DD combination) are excluded one at a time (with replacement). Exclusion is ordered such that the next DD combination removed is the one with the next-lowest quality score that has not yet been removed. A quality score is formed for each DD combination via a linear combination of scores based on  $C/N_0$ ,  $s_e$ , and  $\theta_{ei}$ . If such  $N$ -choose-1 elimination fails to create a subset of DD measurements that passes the innovations test, exclusion can proceed to  $N$ -choose- $m$  elimination, with  $m > 1$ . If a user-configurable exclusion depth is exceeded, the estimator state is reset.

The third level of exclusion is based on the integer aperture test following integer estimation via ILS. This is the standard data-driven integer-fixing process whereby the integer-fixed solution is selected only upon successful validation by some type of aperture test; otherwise, the float solution is accepted [24]. The aperture test is configured for a fixed failure rate (under independent Gaussian errors) of  $\bar{P}_F$ . If the integer aperture test fails,  $N$ -choose-1 exclusion (with replacement) is attempted, starting with the lowest-scoring DD combinations and working up through higher-scoring combinations.  $N$ -choose- $m$  exclusion, with  $m > 1$ , is currently not attempted at this layer of exclusion because testing a large number of subsets is eventually “doomed to succeed” at passing the aperture test, causing  $P_F$  to significantly exceed  $\bar{P}_F$  even under benign conditions [37].

If the aperture test is passed before the permissible exclusion depth is exceeded, the solution is conditioned on the integers and the integer states are dropped; otherwise, the integer state elements are marginalized out as float values. In either case, the state is propagated to the next measurement epoch via the dynamics model and the process repeats.

## Experimental Setup

The precise positioning system was evaluated experimentally using data collected on 1 August 2018 during approximately 2 h of driving in and around the dense urban center of Austin. The rover GNSS receiver is one among several sensors housed in an integrated perception platform called the *Sensorium*, which is displayed in Figure 5. Although hardly visible in Figure 5, two Antcom G8 triple-frequency patch antennas are flush mounted in the cross-track direction on the Sensorium’s upper plate, separated by slightly more than 1 m. The antennas’ signals are routed to a unified RF front end whose output IF samples are processed in

real time (to within less than a 10-ms latency) by the Sensorium’s onboard computer. The samples are also stored to disk for postprocessing.

Data from both the driver- and passenger-side antennas were used to produce the PpRx standard navigation solution, but only data from the driver-side antenna were used in the urban RTK performance evaluation. No other Sensorium sensors were involved in the current article’s results.

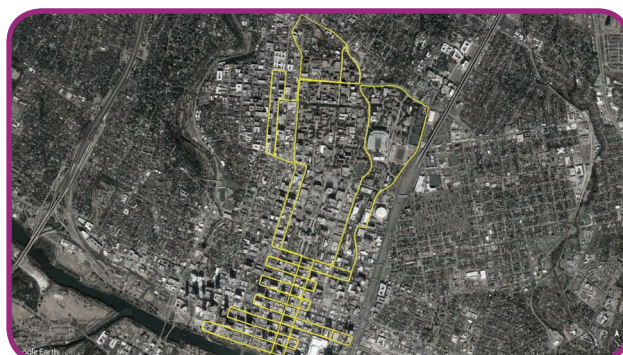
The test route, presented in Figure 6, runs the gamut of light-to-dense urban conditions, from open sky to narrow streets with overhanging trees to the high-rise urban city center. A time history of route coordinates, in the form of a Google Earth KML file, is packaged with the other campaign data so that readers can explore the route. The route begins with a 10-min stationary interval and ends with a 4-min stationary interval in open-sky conditions to allow confident bookending for the ground-truth system. The number  $N_{DD}$  of DD signals available to PpEngine at each epoch over the 2-h test interval ranged from one to 18, with an average of 12.5.

## Ground-Truth Trajectory

A trustworthy ground-truth trajectory against which to compare the reported trajectory of the system under test



**FIG 5** The University of Texas’ Sensorium is a platform for automated and connected vehicle perception research. It includes stereo visible-light cameras, an industrial-grade IMU, an automotive radar unit, a dual-antenna, a dual-frequency software-defined GNSS receiver, 4G cellular connectivity, and a powerful internal computer.



**FIG 6** An overview of the test route through the urban core of Austin, Texas. (Source: Google Maps.)

The positioning performance appears excellent, with 95% of horizontal and vertical errors below 14 and 8 cm, respectively. That the vertical errors are smaller than the horizontal errors is explained by the vehicle motion's greater vertical predictability.

bound  $\bar{P}_F$ . This section discusses the baseline system's configuration and performance, while the following section compares the baseline system against several alternative configurations of the PpRx-PpEngine pipeline.

### Configuration

PpRx's carrier- and code-tracking loops were configured as detailed in [16]. PpRx was set to track the following signal types:

GPS L1 C/A, GPS L2C (combined M + L tracking), Galileo E1 BOC(1,1) (combined B + C tracking), and SBAS (WAAS) on L1. It was configured to output observables at 5 Hz.

is indispensable for urban positioning evaluation. The ground-truth generation process described in [16] was unsatisfactory for two reasons. First, it lacked independence, as it drew, in part, on the same underlying precise solutions that were to be evaluated. Second, it was not possible to create a complete ground-truth trajectory even for the moderate urban test route discussed in [16]. Gaps in the ground truth prevented an accurate determination of  $P_F$ .

The present work adopts the more traditional approach of taking the forward-backward smoothed trajectory generated in after-the-fact processing by a coupled RTK-inertial system with a tactical-grade IMU as the ground truth [8], [14]. In particular, an iXblue ATLANS-C mobile mapping inertial navigation system/GNSS system, which incorporates a professional-grade Septentrio AsteRx3 RTK receiver, was used to generate the ground truth [39]. The ATLANS-C was rigidly mounted to the Sensorium and attached to the same antenna from which PpEngine drew observables. A cm-accurate lever-arm estimate from the inertial sensor to the GNSS antenna was determined. The self-reported 3D accuracy of the ATLANS-C's smoothed estimate varied between 2 and 20 cm (1-sigma) along the test route. Along the light-to-moderate urban portions of the test route, the ATLANS-C and PpEngine 3D estimates agreed to better than 5 cm (95%).

### Baseline System Performance

The baseline urban RTK system is the PpRx-PpEngine pipeline configured to maximize  $P_S$  while respecting  $P_F \leq \bar{P}_F$  for some chosen, empirical, incorrect fixing-probability

PpEngine was configured as follows. The master LDRN reference station, located within 4 km of all points on the test route, was taken as the reference receiver, producing reference observables at 5 Hz. The master station's antenna is a Trimble Zephyr II geodetic antenna. A single-baseline RTK solution with a near-zero age of data was performed between the rover's primary antenna and the reference station at a 5-Hz cadence. The following thresholds were applied in the first-level screening processing within PpEngine:  $C/N_0 \geq 37.5$ -dB Hz,  $s_\theta \geq 0.5$ , and  $\theta_{el} \geq 15^\circ$ . Signals whose values fell below any one of these thresholds were excluded from all the DD combinations. Elevation-dependent weighting was applied in the float solution. The threshold above which float-innovation statistics failed the normalized-innovation squared test was chosen to be two. Scored  $N$ -choose-1 exclusion was applied for both failed float-innovations tests and failed aperture tests. A depth of eight signals was allowed for the  $N$ -choose-1 exclusion, after which the estimator was either reset or the integers marginalized, according to the flow diagram in Figure 4. The difference test of [27], which was found to work as well in urban environments, was chosen as the integer aperture test. The test was configured for a fixed failure rate of  $\bar{P}_F = 0.001$ . The undifferenced pseudorange and phase-measurement error were taken to be  $\sigma_\rho = 0.9$  m and  $\sigma_\phi = 4$  mm, respectively. The nearly constant-velocity dynamics model was configured for a 0.4- and 0.06-m/s deviation in horizontal and vertical velocity, respectively, over a 1-s interval.

A calibration was carried out of the Sensorium antennas' phase center variation with elevation angle relative to the reference antenna. The calibration procedure is similar to the one presented in [40] except that it works with DDs instead of single differences. The calibration succeeded in reducing the standard deviation of L1 and L2 undifferenced carrier-phase residuals by 11 and 15%, respectively, in open-sky conditions.




FIG 7 The cumulative distribution function (CDF) for horizontal and vertical fixed-position error magnitudes with respect to the ground truth for the baseline system.

### Performance

Figure 7 shows the cumulative distribution function (CDF) of the horizontal and vertical positioning errors for fixed (aperture test-validated) PpEngine solutions. The positioning performance appears excellent, with 95% of horizontal and vertical errors below 14 and 8 cm, respectively. That the vertical errors are smaller than the horizontal errors is explained by the vehicle motion's greater vertical predictability.

Figure 8 depicts the CDF of availability gaps in the baseline system's fixed solution. These are intervals during which only a less-accurate float solution is available. Although the longest gap was 90 s, more than 99% of gaps are shorter than 2 s, a span that could be easily bridged by a microelectromechanical systems-quality inertial sensor with errors smaller than a few centimeters [6].

The baseline system's fixed-solution availability,  $P_V$ , was 87.2%. Fixed solutions were considered correctly resolved if their 3D positions were within 30 cm of the ground truth. This led to  $P_S = 84.8\%$  and  $P_F = 2.4\%$ . Note that  $P_F$  is a factor of 24 larger than  $\bar{P}_F = 0.1\%$  but may be tolerable for a larger system that combines standalone RTK with inertial and electro-optical sensing, as the Sensorium shown in Figure 5 is intended to do.

### Performance Degradation Analysis

This section presents a performance-degradation analysis in which features of the baseline system are removed or altered one at a time to assess their relative contribution to baseline system performance. Table 1, where  $P_V$ ,  $P_S$ , and  $P_F$  are as defined previously, summarizes the results of the analysis. Starting with scenario 2, the following discussion treats each scenario in turn.

#### Data Bit Prediction Disabled

Eliminating the baseline's system's LNAV and SBAS data bit prediction capability, described in the "Navigation Data Bit Wipeoff" section, has a devastating effect on performance. The availability of validated epochs drops by eight percentage points and  $P_F$  rises tremendously, from 2.4 to 25%. Clearly, data bit prediction is a key capability for urban RTK.

#### Scalar Tracking With Adaptive $B_\theta$

Eliminating vector tracking, as described in the "Carrier Tracking" section, in favor of scalar tracking, but retaining carrier-tracking loop-bandwidth adap-

tation, has no significant effect on  $P_V$  but increases  $P_F$  slightly, from 2.4 to 3%. Thus, vector tracking appears helpful but not critically so.

#### Scalar Tracking With Fixed $B_\theta$

Eliminating both vector tracking and carrier-tracking loop-bandwidth adaptation has little effect on availability, but  $P_F$  rises to 6.2%, indicating that loop-bandwidth adaptation is useful in preventing fixing errors.

#### GPS L2CL Tracking

For GPS L2C tracking, PpRx jointly tracks the pilot (CL) and data-bearing medium-length (CM) codes, wiping off the INAV data symbols modulating the CM code with symbol-value estimates based not on prediction, as with LNAV, but merely on observation. The rationale for this strategy is that the CL pilot renders prediction less necessary than for the GPS L1 C/A signal, which does not enjoy a pilot. Eliminating joint L2C L + M tracking in favor of pure




FIG 8 The CDF for baseline system-availability gaps.

Table 1. A summary of precise positioning results.

Scenario	Description	$P_V$ : Validated Epochs (%)	$P_S$ : Success (%)	$P_F$ : Failure (%)
1	Baseline system	87.2	84.8	2.4
2	LNAV and SBAS data bit prediction disabled	79.4	54.4	25
3	Scalar tracking with adaptive $B_\theta$	87.1	84.1	3
4	Scalar tracking with fixed $B_\theta$	86.9	80.7	6.2
5	GPS L2CL tracking	83.2	80.8	2.3
6	Age of data = 200 ms	87.4	85	2.4
7	Age of data = 400 ms	87.3	83.8	3.5
8	Age of data = 600 ms	87.2	82.1	5.1
9	Age of data = 1,000 ms	87	82	5
10	15-km baseline	86.6	78.9	7.7
11	Without SBAS	78.6	73.1	5.5
12	Without GPS L2C (L + M)	83.6	82.5	1.1
13	Without Galileo E1 (B + C)	77.4	75.9	1.4
14	No scored exclusion	78.8	75.6	3.2
15	No antenna calibration	87	82.8	4.2

PpRx-PpEngine enjoys a 16% availability advantage over the commercial RTK system, and neither the CSRS-PPP nor the ECPP solutions are close to sub-30-cm accuracy.

L2CL tracking might be considered a more reliable strategy given that no symbol wipeoff mistakes are ever made when tracking only the pilot. However, Table 1 indicates that this leads to a drop in availability with hardly any improvement in the error rate. Thus, it appears that joint CL and CM tracking is preferred.

#### Age of Data

Scenarios 6–9 explore the effect of increased age of the reference data, from the baseline age (near-zero latency relative to the rover stream) to 1 s. Little reduction occurs in  $P_V$ , but there appears a somewhat steady increase in  $P_F$  after 200 ms.

#### 15-km Baseline

The baseline system’s distance to the reference receiver, referred to as the *reference-rover baseline*, is no greater than 4 km. For scenario 10, the LDRN alternate master station, which sits 15 km from the furthest portion of the test route, was instead taken as reference. The alternate mas-




FIG 9 The CDF for horizontal and vertical fixed-position error magnitudes with respect to the ground truth for scenario 10.

Table 2. The comparison results.

System	Availability (%)	$P_F$ : Failure (%)	$d_{95\%}$ (cm)
PpRx-PpEngine	77	0.7	7.5
Commercial RTK	61	0.8	8.6
CSRS-PPP	74	N/A	154
ECPP	100	N/A	275

N/A: not applicable.

ter station has a Trimble Zephyr II antenna identical to the master station’s. A 15-km baseline might still be considered within the short-baseline regime for standard RTK [41]. Nonetheless, a slight decrease in  $P_V$  and a significant increase in  $P_F$  is observed, consistent with the argument in [42] that a dense reference network is helpful in urban settings with reduced signal avail-

ability. Figure 9 also shows a significant degradation in overall fixed-position accuracy.

#### Value of Additional Signals

Scenarios 11–13 explore the degradation that occurs when all signals of a particular type are eliminated from the RTK solution. Curiously, dropping L2C and Galileo from consideration significantly reduced  $P_F$ , likely due to the misfortune that both groups were composed primarily of low-elevation satellites (below 30°) during the test run. By contrast, one notes a significant increase in  $P_F$  as the four available SBAS satellites are removed from consideration. Clearly, data-wiped SBAS signals offer significant strength to the solution.

#### Scored Exclusion

Scenario 14 removes the scored exclusion strategy described in the “Robust Measurement Update” section by setting the exclusion depth to 0. This caused a noticeable reduction in  $P_V$  and a slight increase in  $P_F$ .

#### Antenna Calibration

Scenario 15 indicates that a lack of rover antenna calibration has no discernible effect on  $P_V$  but increases  $P_F$  significantly. The effect would no doubt be larger for lower-quality rover antennas.

#### Performance Comparison

To further assess its performance, the PpRx-PpEngine system was compared against three alternative systems: 1) A high-end commercial RTK system, i.e., a Septentrio AsteRx3 receiver attached to the same GNSS antenna as the PpRx-PpEngine system, with RTK solutions produced by Septentrio’s RTK engine taking reference data from the Continually Operating Reference Station station TXAU, which sits less than 5.4 km from the furthest portion of the test route; 2) a state-of-the-art PPP solution from the CSRS-PPP service [43] (version v2.26.0 from March 2019) based on observables produced by the same receiver as in 1), processed in a “batch kinematic” mode; and 3) the so-called enhanced code-phase positioning (ECPP) solution [5] from PpRx, which draws on precise orbit and clock models from the



International GNSS Service [44], a WAAS ionospheric model, and RF signals from both the passenger- and driver-side Sensorium antennas.

Note that the receiver described in 1) and 2) is the very receiver used to generate the ground-truth trajectory, but for this comparison, its data were processed without aid from the ATLANS-C IMU. Also, to ensure a fair comparison, PpEngine exclusion and validation tests were tightened (at the expense of solution availability) until its probability of incorrect fix,  $P_F$ , was less than that of the commercial RTK system.

Table 2 lists the comparison results in terms of solution availability (equivalent to  $P_V$  for the PpRx-PpEngine and Septentrio RTK systems),  $P_F$  (which applies only to the PpRx-PpEngine and Septentrio RTK systems, as the other two do not attempt integer fixing), and  $d_{95h}$ , the horizontal 95th-percentile positioning error. Significantly, PpRx-PpEngine enjoys a 16% availability advantage over the commercial RTK system, and neither the CSRS-PPP nor the ECPP solutions are close to sub-30-cm accuracy.

## Conclusions

In this article, an RTK positioning system tailored for urban vehicular positioning was described and evaluated. To facilitate performance comparison against similar systems, the system was tested without the benefit of aid by inertial or electro-optical sensors. Over nearly 2 h of urban testing, including multiple passes through Austin's dense urban center, the system achieved an 85% probability of correct integer fix for a 2.4% probability of incorrect fix, resulting in 3D positioning errors smaller than 17 cm (95%). A performance-sensitivity analysis revealed that navigation data bit prediction on fully modulated GNSS signals is key to high-performance urban RTK positioning, and that a dense reference network, carrier-tracking bandwidth adaptation, and rover antenna calibration each offer a significant integrity benefit. A comparison with existing unaided systems for urban GNSS processing indicates that the proposed system has a significant advantage in availability and/or accuracy.

## Acknowledgments

We thank iXblue engineers Jean-Baptiste Lacambre and Tim Barford for providing the ground-truth trajectory based on the iXblue ATLANS-C system. This work has been supported by the National Science Foundation under grant 1454474 (CAREER), by the Data-Supported Transportation Operations and Planning Center, a tier 1 U.S. Department of Transportation University Transportation Center, and by the University of Texas Situation-Aware Vehicular Engineering Systems Center (<http://utsaves.org/>), an initiative of the Wireless Networking and Communications Group.

## About the Authors



**Todd E. Humphreys** ([todd.humphreys@utexas.edu](mailto:todd.humphreys@utexas.edu)) earned his B.S. and M.S. degrees in electrical and computer engineering from Utah State University, Logan, in 2000 and 2003, respectively, and his Ph.D. degree in aerospace engineering from Cornell University, Ithaca, New York, in 2008.

Currently, he is an associate professor with the Department of Aerospace Engineering and Engineering Mechanics, the University of Texas (UT) at Austin, and director of the UT Radionavigation Laboratory. He received the University of Texas Regents' Outstanding Teaching Award in 2012, the National Science Foundation CAREER Award in 2015, the Institute of Navigation (ION) Thurlow Award in 2015, and the Presidential Early Career Award for Scientists and Engineers in 2019. He specializes in the application of optimal detection and estimation techniques to problems in satellite navigation, autonomous systems, and signal processing. His recent focus has been on secure perception for autonomous systems, including navigation, timing, and collision avoidance, and on centimeter-accurate location for the mass market. He is a Member of the IEEE and a fellow of the ION.



**Matthew J. Murrian** ([matthew.murrian@utexas.edu](mailto:matthew.murrian@utexas.edu)) earned his master's degree in aerospace engineering in 2018 from the University of Texas (UT) at Austin, where he works in the UT Radionavigation Laboratory. He is currently a systems engineer at CTSi working on advanced navigation. Previously, he served as a U.S. Navy submarine officer. He is a Member of the IEEE.



**Lakshay Narula** ([lakshay.narula@utexas.edu](mailto:lakshay.narula@utexas.edu)) earned his B.Tech. degree in electronics engineering from the Indian Institute of Technology (Banaras Hindu University) and his M.S. degree in electrical and computer engineering from the University of Texas (UT) at Austin, in 2014 and 2016, respectively.

He is currently a Ph.D. student with the Department of Electrical and Computer Engineering at UT Austin and a graduate research assistant at the UT Radionavigation Laboratory. Previously, he was a visiting student at the Position Location and Navigation Group at the University of Calgary, Alberta, Canada, and a systems engineer at Accord Software & Systems, Bangalore, India. He was a recipient of the 2017 Qualcomm Innovation Fellowship. His research interests include Global Navigation Satellite System signal processing, secure perception in autonomous systems, and detection and estimation. He is a Student Member of the IEEE.

## References

- [1] J. B. Kenney, "Dedicated short-range communications (DSRC) standards in the United States," *Proc. IEEE*, vol. 99, no. 7, pp. 1162–1182, 2011. doi: 10.1109/JPROC.2011.2152790.
- [2] D. Fajardo, T.-C. Au, S. Waller, P. Stone, and D. Yang, "Automated intersection control: Performance of future innovation versus current traffic signal control," *Transport. Res. Rec., J. Transport. Res. Board*, vol. 2259, no. 1, pp. 225–252, 2011. doi: 10.5141/2259-21.
- [3] L. Narula, J. M. Wooten, M. J. Murrian, D. M. LaChapelle, and T. E. Humphreys, "Accurate collaborative globally-referenced digital mapping with standard GNSS," *Sensors*, vol. 18, no. 8, p. 2452, 2018. doi: 10.3390/s18082452. [Online]. Available: <http://www.mdpi.com/1424-8220/18/8/2452>
- [4] M. Petovello, M. Cannon, and G. Lachapelle, "Benefits of using a tactical-grade IMU for high-accuracy positioning," *Navigation, J. Inst. Navigation*, vol. 51, no. 1, pp. 1–12, 2004. doi: 10.1002/j.2161-4296.2004.tb00357.x.
- [5] B. M. Scherzinger, "Precise robust positioning with inertially aided RTK," *Navigation*, vol. 53, no. 2, pp. 75–85, 2006. doi: 10.1002/j.2161-4296.2006.tb00374.x.
- [6] H. T. Zhang, "Performance comparison on kinematic GPS integrated with different tactical-grade IMUs," M.S. thesis, The Univ. of Calgary, Jan. 2006.
- [7] S. Kennedy, J. Hamilton, and H. Martell, "Architecture and system performance of SPAN—NovAtel's GPS/INS solution," in *Proc. 2006 IEEE Position, Location, and Navigation Symp. (IEEE/ION)*, p. 266. doi: 10.1109/PLANS.2006.1650612.
- [8] T. Li, H. Zhang, Z. Gao, Q. Chen, and X. Niu, "High-accuracy positioning in urban environments using single-frequency multi-GNSS RTK/MEMS-IMU integration," *Remote Sens.*, vol. 10, no. 2, p. 205, 2018. doi: 10.3390/rs10020205.
- [9] S. Godha, "Performance evaluation of low cost MEMS-based IMU integrated with GPS for land vehicle navigation application," M.S. thesis, The Univ. of Calgary, Jan. 2006.
- [10] P. Teunissen and A. Khodabandeh, "Review and principles of PPP-RTK methods," *J. Geodesy*, vol. 89, no. 3, pp. 217–240, 2015. doi: 10.1007/s00190-014-0771-5.
- [11] Y. Cui et al., "Feasibility analysis of low-cost GNSS receivers for achieving required positioning performance in CAV applications," in *Proc. Forum Cooperative Positioning and Service*, May 2017, pp. 555–561. doi: 10.1109/CPGPS.2017.8075154.
- [12] D. P. Shepard and T. E. Humphreys, "High-precision globally-referenced position and attitude via a fusion of visual SLAM, carrier-phase-based GPS, and inertial measurements," in *Proc. IEEE/ION Position, Location, and Navigation Symp. (PLANS)*, May 2014. doi: 10.1109/PLANS.2014.6851506.
- [13] K. M. Pesyna Jr., "Advanced techniques for centimeter-accurate GNSS positioning on low-cost mobile platforms," Ph.D. dissertation, The Univ. of Texas at Austin, Dec. 2015.
- [14] R. B. Ong, M. G. Petovello, and G. Lachapelle, "Assessment of GPS/GLONASS RTK under various operational conditions," in *Proc. ION GNSS*, 2009, pp. 5297–5308.
- [15] J. Jackson, B. Davis, and D. Gebre-Egziabher, "A performance assessment of low-cost RTK GNSS receivers," in *Proc. IEEE/ION Position, Location, and Navigation Symp. (PLANS)*, Monterey, CA, 2018, pp. 642–649. doi: 10.1109/PLANS.2018.8575458.
- [16] T. E. Humphreys, M. Murrian, and L. Narula, "Low-cost precise vehicular positioning in urban environments," in *Proc. IEEE/ION Position, Location and Navigation Symp. (PLANS)*, April 2018, pp. 456–471. doi: 10.1109/PLANS.2018.8575444.
- [17] M. L. Psiaki and H. Jung, "Extended Kalman filter methods for tracking weak GPS signals," in *Proc. ION GPS Meeting*, Portland, OR, 2002, pp. 2559–2555.
- [18] H. Niedermeier et al., "Dingpos: High sensitivity GNSS platform for deep indoor scenarios," in *Proc. 2010 Int. Conf. Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–10. doi: 10.1109/IPIN.2010.5647542.
- [19] S. Mohiuddin and M. L. Psiaki, "High-altitude satellite relative navigation using carrier-phase differential global positioning system techniques," *J. Guid., Control, Dyn.*, vol. 30, no. 5, pp. 1628–1639, Sept.–Oct. 2007. doi: 10.2514/1.27827.
- [20] 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 *Proc. ION GNSS Meeting*, Fort Worth, TX, 2006, pp. 2570–2582.
- [21] T. E. Humphreys, J. Bhatti, T. Pany, B. Ledvina, and B. O'Hanlon, "Exploiting multicore technology in software-defined GNSS receivers," in *Proc. ION GNSS Meeting*, Savannah, GA, 2009, pp. 526–538.
- [22] E. G. Lightsey, T. E. Humphreys, J. A. Bhatti, A. J. Joplin, B. W. O'Hanlon, and S. P. Powell, "Demonstration of a space capable miniature dual frequency GNSS receiver," *Navigation, J. Inst. Navigation*, vol. 61, no. 1, pp. 55–64, 2014. doi: 10.1002/navi.52.
- [23] P. Teunissen, "Integer aperture GNSS ambiguity resolution," *Artif. Satellites*, vol. 38, no. 5, pp. 79–88, 2005.
- [24] G. N. Green and T. Humphreys, "Data-driven generalized integer aperture bootstrapping for high-integrity positioning," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 55, no. 2, p. 1, 2018. doi: 10.1109/TAES.2018.2864770.
- [25] G. N. Green and T. Humphreys, "Position-domain integrity analysis for generalized integer aperture bootstrapping," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 55, no. 2, pp. 734–746, Aug. 2018. doi: 10.1109/TAES.2018.2864766.
- [26] P. Teunissen, "Integer aperture bootstrapping: a new GNSS ambiguity estimator with controllable fail-rate," *J. Geodesy*, vol. 79, no. 6–7, pp. 389–397, 2005. doi: 10.1007/s00190-005-0481-y. [Online]. Available: <http://dx.doi.org/10.1007/s00190-005-0481-y>
- [27] L. Wang and S. Verhagen, "A new ambiguity acceptance test threshold determination method with controllable failure rate," *J. Geodesy*, vol. 89, no. 4, pp. 361–375, 2015. doi: 10.1007/s00190-014-0780-2. [Online]. Available: <http://dx.doi.org/10.1007/s00190-014-0780-2>
- [28] T. E. Humphreys, M. L. Psiaki, J. C. Hinks, B. O'Hanlon, and P. M. Kintner Jr., "Simulating ionosphere-induced scintillation for testing GPS receiver phase tracking loops," *IEEE J. Sel. Topics Signal Process.*, vol. 3, no. 4, pp. 707–715, Aug. 2009. doi: 10.1109/JSTSP.2009.2024130.
- [29] T. E. Humphreys, M. L. Psiaki, B. M. Ledvina, A. P. Cerruti, and P. M. Kintner Jr., "A data-driven testbed for evaluating GPS carrier tracking loops in ionospheric scintillation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 46, no. 4, pp. 1609–1623, Oct. 2010. doi: 10.1109/TAES.2010.5595582.
- [30] F. S. T. Van Diggelen, *A-GPS: Assisted GPS, GNSS, and SBAS*. Norwood, MA: Artech House, 2009.
- [31] A. J. Van Dierendonck, "GPS receivers," in *Global Positioning System: Theory and Applications*, J. J. Spilker Jr., P. Axelrad, B. W. Parkinson, and P. Enge, Eds. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 8, pp. 329–407.
- [32] M. Lashley and D. Bevil, "What are vector tracking loops, and what are their benefits and drawbacks?" *GNSS Solutions Column, Inside GNSS*, vol. 4, no. 5, pp. 16–21, 2009.
- [33] S. A. Stephens and J. B. Thomas, "Controlled-root formulation for digital phase-locked loops," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 51, no. 1, pp. 78–95, Jan. 1995. doi: 10.1109/7.566295.
- [34] M. Psiaki, "Kalman filtering and smoothing to estimate real-valued states and integer constants," *J. Guid., Control, Dyn.*, vol. 33, no. 5, pp. 1404–1417, Sept.–Oct. 2010. doi: 10.2514/1.48567.
- [35] P. J. Teunissen, "The least-squares ambiguity decorrelation adjustment: a method for fast GPS integer ambiguity estimation," *J. Geodesy*, vol. 70, no. 1–2, pp. 65–82, 1995. doi: 10.1007/BF00865419.
- [36] A. Brack, "On reliable data-driven partial GNSS ambiguity resolution," *GPS Solutions*, vol. 19, no. 5, pp. 411–422, 2015. doi: 10.1007/s10291-014-0401-9. [Online]. Available: <http://dx.doi.org/10.1007/s10291-014-0401-9>
- [37] A. Parkins, "Increasing GNSS RTK availability with a new single-epoch batch partial ambiguity resolution algorithm," *GPS Solutions*, vol. 15, no. 4, pp. 391–402, 2011. doi: 10.1007/s10291-010-0198-0.
- [38] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: Wiley, 2001.
- [39] iXblue, "ATLANS-C mobile mapping position and orientation solution," Denver, CO, Datasheet 014-08-BR-ATLANS-C, 2014.
- [40] G. L. Mader, "GPS antenna calibration at the National Geodetic Survey," *GPS solutions*, vol. 5, no. 1, pp. 50–58, 1999. doi: 10.1007/PL00012780.
- [41] R. Odolinski, P. Teunissen, and D. Odijk, "Combined GPS+BDS for short to long baseline RTK positioning," *Measurement Sci. Technol.*, vol. 26, no. 4, p. 045801, 2015. doi: 10.1088/0957-0253/26/4/045801.
- [42] M. J. Murrian, C. W. Gonzalez, T. E. Humphreys, K. M. Pesyna Jr., D. P. Shepard, and A. J. Kerns, "Low-cost precise positioning for automated vehicles," *GPS World*, vol. 27, no. 9, pp. 52–59, Sept. 2016.
- [43] Natural Resources Canada, "Geodetic tools and applications: CSRS-PPP," Ottawa, Mar. 2019. [Online]. Available: <https://goo.gl/1Y5KGa>
- [44] International GNSS Service (IGS), "GPS satellite ephemerides/satellite & station clocks." Accessed on: Mar. 30, 2018. <http://www.igs.org/products>