

Low-cost Precise Vehicular Positioning in Urban Environments

Todd E. Humphreys
Aerospace Engineering
The University of Texas at Austin
todd.humphreys@mail.utexas.edu

Matthew Murrian
Aerospace Engineering
The University of Texas at Austin
matthew.murrian@utexas.edu

Lakshay Narula
Electrical and Computer Engineering
The University of Texas at Austin
lakshay.narula@utexas.edu

Abstract—A system developed at The University of Texas for low-cost precise urban vehicular positioning is demonstrated to achieve a probability of correct integer fixing greater than 96.5% for a probability of incorrect integer fixing surely less than 2.3% and likely less than 1%. This is demonstrated using data captured during 3.4 hours of driving on a repeating urban test route over three separate days. The results are achieved without any aiding by inertial or electro-optical sensors. Development and evaluation of the unaided GNSS-based precise positioning system is a key milestone toward the overall goal of combining precise GNSS, vision, radar, and inertial sensing for all-weather high-integrity precise 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 whose processing can be executed on general-purpose commodity hardware, and a real-time kinematic (RTK) positioning engine. All components have been tailored in their design to yield competent sub-decimeter positioning in the mobile urban environment. A performance sensitivity analysis reveals that navigation data bit prediction on the GPS L1 C/A signals is key to high-performance urban RTK positioning.

Keywords—urban vehicular positioning; CDGNSS; low-cost RTK positioning.

I. INTRODUCTION

Few of the leading self-driving car projects exploit carrier-phase differential GNSS (CDGNSS) positioning for vehicle localization. Waymo's GNSS needs are reportedly satisfied by a pair of ublox M8 receivers performing a standard code- and Doppler-based navigation solution. Likewise, a recent teardown of Tesla's Hardware 2.0 Autopilot module revealed nothing more than a single ublox M8L receiver [1], which is not capable of precise CDGNSS-type positioning. But as automated vehicles become increasingly connected, and as they enter markets beyond the sunny confines of Silicon Valley and Chandler, Arizona, they will need some way of determining their globally-referenced position to better than 30 cm in all weather conditions.

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 DSRC basic safety message, a first step in V2V coordination, does not yet define a position accuracy requirement, effectively accepting whatever accuracy a standard GNSS receiver provides [2]. But automated intersection management [3], tight-formation platooning, and unified processing of sensor data—all involving vehicles of different makes who 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 present 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 created 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 who rely 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 but few radar-reflective landmarks, radar too becomes useless. GNSS receivers operate well in all weather conditions, but only a GNSS receiver whose errors remain under 30 cm 95% of the time could avoid drifting onto a snow-covered road's soft shoulder. Code- and Doppler-based GNSS solutions will find it challenging to meet this requirement, even with modernized GNSS offering wideband signals at multiple frequencies.

Carrier-phase-based GNSS positioning 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 CDGNSS receiver with a tactical grade inertial sensor, as in [4]–[7] enables robust precise positioning even during the extended signal outages common in dense urban areas. Apollo 2.0, Baidu’s open self-driving platform, includes such a coupled positioning system, in particular, the NovAtel SPAN system, among the core sensors of its reference hardware¹.

But GNSS-inertial systems with tactical-grade inertial measurement units (IMUs) cost tens of thousands of 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 only allows approximately 5 seconds of complete GNSS signal blockage before the IMU no longer offers a useful constraint for so-called integer ambiguity resolution, which underpins the fastest, most accurate, and most robust CDGNSS techniques, namely, single-baseline RTK, network RTK, and PPP-RTK [9]–[11].

University of Texas researchers have recently proposed a novel and 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 markets. The current authors are engaged in developing a high-integrity RTK-vision system for precise 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 stand-alone RTK techniques—those unaided by IMUs or vision—in urban environments. Such a 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 stand-alone RTK performance in urban environments, likely because the performance results tend to be somewhat poor. Short-baseline RTK experiments between two vehicles in [14] revealed that multi-frequency (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 centimeters. RTK positioning in downtown Calgary was disastrous, with less than 60% solution availability and RMS errors exceeding 9 meters.

More recently, Li et al. [8] show that, with the benefit of greater signal availability, stand-alone dual-frequency GPS + BDS + GLONASS RTK can achieve correct integer fixing rates of 76.7% on a 1-hour drive along an urban route in Wuhan, China. But Li et al. do not provide data on the incorrect fixing rate, nor do they assess the accuracy of their ground truth trajectory, so the significance of their correct fixing rate 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 solution availability in urban areas, despite a dense reference network and dual-frequency capability. A key failing of existing receivers is their slow recovery after passing under bridges or overpasses. Jackson et al. show that the Piksi Multi and the Eclipse P307 dual-frequency receivers require from 25 (Eclipse) to 40 (Piksi) seconds to recover a fixed position after passing under an overpass.

Professional-grade receivers appear to handle momentary obstructions better, but the primary markets of professional-grade receivers have traditionally been surveying, machine control, and precision agriculture—applications that typically enjoy a relatively unobstructed view of overhead satellites compared to urban vehicular positioning. Consider that in 2018 a professional-grade multi-constellation multi-frequency Trimble receiver requires 4 seconds to produce an RTK fix on a short baseline with a clear sky view (see Fig. 15 in [16]). This may be an acceptably short time-to-fix for traditional RTK applications, but not for urban vehicular positioning.

This paper describes and evaluates an unaided RTK positioning system developed at The University of Texas that has been designed for vehicular operation in both rural and urban environments. Solution availability and accuracy have been markedly improved since publication of preliminary results in [17]. This paper’s primary contributions are (1) a demonstration of the remarkably good performance that can be achieved with a low-cost software-defined stand-alone RTK GNSS platform in an urban environment, and (2) a thorough evaluation of the system’s sensitivity to various impairments.

II. 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 (1) prevent frequency unlock during the deep fades caused by building occlusions, and (2) 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 explored in [18], [19], 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. There is simply no signal to recover.

¹<http://apollo.auto/platform/hardware.html>

Fig. 1 illustrates this fact. The initial disturbance at 950 seconds 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 (four-second) 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 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 (1) identifying and rejecting observables from blocked or otherwise compromised signals, and (2) immediately re-evaluating the corresponding integer ambiguities when signals reappear. A multi-stage cycle slip detection and recovery technique, such as proposed in [20], is too slow for urban positioning.




Fig. 1. In-phase (green, top) and quadrature (blue, bottom) 10-ms complex correlation products for a GPS L1 C/A signal at 35 degrees 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 full carrier cycle recovery. Rapid fading—and rapid recovery—occur as buildings intermittently block the signal.

A related hallmark of the urban mobile channel is the wide and rapid variation of the number of signals available for RTK positioning. Fig. 2 shows the number N_{DD} of double-difference (DD) signals (each one providing a DD pseudorange and a DD carrier phase observable) over a 600-second segment of urban driving. N_{DD} remains constant only when the vehicle is stopped. The implication for RTK processing is that integer ambiguity continuity will often be lost, requiring rapid and continuous re-estimation of ambiguities.

The histogram shown in 3 indicates that, although volatile, N_{DD} remains above 16 for more than half the measurement epochs, which implies that this particular urban environment is not an especially challenging one. Even still, because




Fig. 2. The number of double difference signals, N_{DD} , available for RTK positioning over a 600-second interval of urban driving.

single-epoch integer ambiguity resolution becomes unlikely for $N_{DD} < 14$ [21], one should expect unaided RTK to struggle.




Fig. 3. Histogram of N_{DD} over a 1-hour interval of urban driving.

III. SYSTEM DESCRIPTION

A. Overview

GNSS components of the University of Texas precise positioning system are shown in Fig. 4. The sub-components 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 down-mixes and digitizes the signals, producing a stream of intermediate frequency (IF) samples. The RF front end used in the present work produces samples at 9.8 Msps 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 over 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 multi-frequency software-defined GNSS receiver developed primarily at the University of Texas [22]–[24]. PpRx draws ephemeris data and GPS LNAV data bit estimates from the Longhorn Dense Reference Network (LDRN), a set of 8 GNSS reference stations deployed in Austin, TX. 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 in Austin can draw assistance data and network observables. The master and alternate master reference stations produce observables at 5 Hz. The other stations, which produce data at a slower rate,

are used primarily to estimate ionospheric and tropospheric corrections according to simple linear model, as described in [17], [25].

PpRx feeds code and carrier observables, and other useful signal information, to PpEngine, an RTK engine developed by members of the UT Radionavigation Laboratory. For the results presented in this paper, 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 (integer-resolved) or float solution depending on the results of an integer aperture test [26].

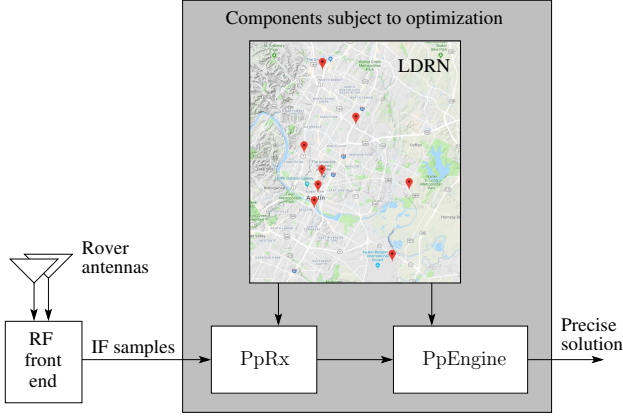


Fig. 4. The University of Texas precise positioning system.

B. Performance Metrics

The performance of precise positioning systems in safety-of-life applications is assessed in terms of integrity, accuracy, and availability [27], [28]. 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 which 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.

Assuming that a CDGNSS solution whose integer ambiguities are correctly fixed will be more than sufficiently accurate for all urban RTK applications of practical interest, this paper focuses on two performance metrics: P_S , the probability (or rate in repeated trials) of correctly (successfully) resolving the full integer set, and P_F , the probability that one or more integer ambiguities failed to resolve correctly [27]. A third probability, $P_U = 1 - P_F - P_S$, 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 increase in P_S comes at the expense of an increase

in P_F (not necessarily of the same amount) [29]. 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 [29] and its generalization to partial ambiguity resolution in [27] 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 paper, 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 [30]. A value of $\bar{P}_F = 0.001$ was adopted for the present paper, meaning that a fixing failure rate less than 1 in 1000 epochs was deemed acceptable. However, as will be illustrated later on, multipath, GNSS signal passage through foliage, and other signals impairments common in urban areas cause the empirical P_F to significantly exceed \bar{P}_F when the aperture threshold is chosen according 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 $P_F \leq \bar{P}_F$.

C. Design Philosophy

With origins in scintillation-resistant carrier tracking [31], [32], 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, development of PpRx, PpEngine, and the LDRN has proceeded as a parallel evolution, with each subsystem benefiting from 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 a singular goal, namely, to maximize 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 FPGAs), 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 [23]. The PpRx-PpEngine pipeline runs at 10x real time on a 6-core Intel Xeon 2.27 GHz processor, enabling rapid iteration cycles for quickly probing the optimization landscape.

Rapid iteration is especially important when the GNSS signal environment is difficult to characterize and highly variable, as is the case for urban GNSS. Signal tracking and integer fixing strategies optimized for traditional RTK applications such as surveying and precision agriculture tend to perform poorly when confronted with bridges, tall reflective buildings, and overhead foliage in an urban environment [15]. In fact, it is doubtful whether robust urban RTK can be grounded so strongly on fundamental signal processing

theory as weak signal GNSS tracking [18], [19] and benign-environment mixed real/integer filtering [33]. To be sure, optimal approaches such as those in [33] offer an excellent starting point. But the current authors have been surprised to see some reasonable adaptations of these strategies fail, and counterintuitive strategies succeed, in urban RTK (examples will be offered in subsequent sections). This has engendered a certain epistemological humility and renewed respect for rapid-cycle experimentation.

D. Carrier and code tracking

GNSS carrier and code tracking in an urban environment must be opportunistic, taking advantage of short clear glimpses to overhead satellites as they present themselves. PpRx's code and carrier tracking architecture, illustrated in Fig. 5 has been designed for immediate (within approximately 100 ms) recovery of full-cycle phase lock after a blockage, and, importantly, for prompt lock indication. The following subsections describe the essential elements of PpRx's tracking strategy, calling out parameters whose values significantly affect urban RTK performance.

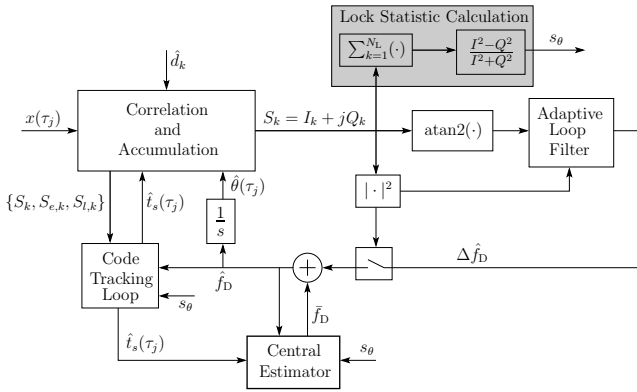


Fig. 5. PpRx's carrier and code tracking architecture.

1) *Correlation and accumulation*: Correlation and accumulation is performed on a sequence of noisy IF samples $x(\tau_j)$, $j = 0, 1, \dots$, where τ_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 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 Fig. 1 correspond to I_k and Q_k , respectively.) The accumulation interval, T_a , is an important configuration parameter for urban RTK.

2) *LNAV data bit prediction*: For the GPS L1 C/A signal, which has no dedicated pilot component, estimates \hat{d}_k of the LNAV data bits, whose structure is detailed in [34], are provided to the local signal replica generator by a prediction engine within PpRx. Extensive testing has shown that

subframes 1-3, which contain ephemeris data, can be nearly-perfectly predicted after the first 18 seconds following a 2-hour GPS time boundary. Almanac data in subframes 4 and 5 are also highly predictable, as they remain constant between approximately daily update events. However, it is worth noting that almanac data can be updated asynchronously—at times not aligned with a 2-hour boundary. For all subframes, reserved bits of the TLM word, and the TLM word's parity bits, must be treated with caution as they can be updated asynchronously with unpredictable content.

The PpRx LNAV prediction engine labels each data bit with one of three certainty levels. The 30-bit HOW word, non-reserved and non-parity bits of the TLM word, and all ephemeris words that have passed parity are marked *certain*. Parity-checked almanac data and reserved and parity bits of the TLM word are marked *almost certain*. Incomplete words, or those that fail parity checking, are marked *uncertain*. The source of each data bit in the local signal replica depends on the prediction engine's corresponding certainty. Prediction engine bits marked *uncertain* are never applied; instead, the observed bit from the received signal is used in the signal replica. This amounts to decision-directed suppressed carrier recovery [35]. Predicted bits marked *almost certain* are used if the signal to noise ratio (SNR) of the prompt complex correlation product S_k falls below a threshold SNR_{pull} ; otherwise the observed bit is used. Bits marked *certain* are used unconditionally in the local replica.

It will be shown in a later section that LNAV data bit prediction is critically important for robust urban RTK.

3) *Lock statistic calculation*: Also key to robust urban RTK is the ability to exclude corrupt or otherwise inaccurate carrier phase measurements. But, 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 process is the lock statistic s_θ . Let I and Q be coherent sums of I_k and Q_k over N_L accumulation intervals. Then s_θ is calculated as [36]

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

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/ggQ^2 and s_θ 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_θ to screen out bad measurements. Note from Fig. 5 that s_θ is also fed to the code tracking loop and to PpRx's central state estimator; each one adapts its behavior to rely less on Doppler measurements when s_θ is low.

4) *Carrier tracking*: As illustrated in Fig. 5, PpRx employs a vector signal tracking architecture wherein a central estimator, implemented as a Kalman filter with a nearly-constant-velocity polynomial-type dynamics model, receives

observables from all tracking channels and drives local replica generation for each channel [37]. More particularly, PpRx employs a hybrid strategy in which, for each channel, a local phase tracking loop is closed around a modeled Doppler value \bar{f}_D provided by the central estimator. The local loop's residual Doppler frequency estimate $\Delta\hat{f}_D$ is added to \bar{f}_D to produce the full estimate \hat{f}_D used in replica generation.

A four-quadrant arctangent phase discriminator, $\text{atan2}(Q_k, I_k)$, 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 [38]. The filter adapts its bandwidth B_θ 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_θ so to maintain a constant loop SNR 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 [18]. However, this reasonable approach was found to yield reduced urban RTK performance. More effective is a three-tiered schedule that reduces B_θ when $|S_k|$ falls below a fairly low threshold γ_1 , and sets B_θ to zero if $|S_k|$ falls below another threshold $\gamma_0 < \gamma_1$. 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 \bar{f}_D . The lock statistic s_θ continues to be calculated. If s_θ is sufficiently close to unity, the central estimator, the code tracking loop, and the RTK engine continue to treat $\hat{\theta}(\tau_j)$ as a valid measurement. But this is a rare occurrence; s_θ is typically far from unity in open-loop mode.

Open loop tracking has been found to be useful for preventing frequency unlock during intervals when signals are entirely blocked, e.g., by buildings or bridges, and for enabling fast re-acquisition of carrier lock immediately following the blockage.

5) *Code tracking*: PpRx's code tracking loop, which is aided by the Doppler estimate \hat{f}_D , is implemented as a 1st-order loop that toggles between a non-coherent (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_θ) has been detected; otherwise, the non-coherent discriminator is applied. A flag attached to each code phase measurement $\hat{t}_s(\tau_j)$ indicates whether it was produced under coherent or non-coherent tracking.

As with carrier tracking, the code tracking loop filter's bandwidth, B_{t_s} , is adaptive. But rather than responding to $|S_k|$ as the carrier loop's bandwidth does, B_{t_s} takes on a different value for each of four code tracking modes: (1) pre-phase lock, (2) first post-lock transient, (3) second post-lock transient, and (4) steady-state. These modes are designed to ensure rapid convergence of the code phase estimate $\hat{t}_s(\tau_j)$ after initial signal acquisition, or in the aftermath of phase

unlock.

6) *Discussion on Optimization*: The tracking architecture diagrammed in Fig. 5 and described in the foregoing subsections yields remarkably good urban RTK performance, but it is not claimed to be optimal for this application. It is likely that further vectorization, in which the code tracking loop's code phase estimate is also driven by a central-estimator-provided model value $\bar{t}_s(\tau_j)$, would yield better performance. Or perhaps entirely replacing the local code and carrier loops with a local batch estimator that takes in an interval of $\{S_k, S_{e,k}, S_{l,k}\}$ values generated using central-estimator-provided $\bar{t}_s(\tau_j)$ and $\bar{f}_D(\tau_j)$ time histories over the interval, and computes from this batch of correlation products residual code phase, carrier phase, and Doppler estimates, would prove to be even more robust to the frequent signal outages and severe multipath in urban tracking.

The larger point is that the usefulness of reasoning *a priori* about optimal architectures for urban RTK is somewhat limited. Likewise, it is hard to say *a priori* what values of the various parameters in the current tracking architecture are optimal in the sense of maximizing P_S for $P_F \leq \bar{P}_F$. Consider the parameters, listed below with default values for the GPS L1 C/A channels, that have been shown to yield good urban RTK tracking for a land vehicle. Default parameter values for the other PpRx channel types are similar. Note too that different parameter values are adopted for the PpRx instances embedded in the LDRN network, as the LDRN stations are static and positioned with a clear view of the sky.

T_a	accumulation interval, 10 ms
N_L	phase lock calculation averaging interval, 2
\bar{s}_θ	phase lock statistic threshold for the code tracking loop and for the central estimator, 0.4
B_θ	default carrier tracking loop bandwidth, 25 Hz
γ_0	SNR of S_k below which carrier tracking operates in open-loop mode, 14.8 dB
γ_1	SNR of S_k above which carrier tracking operates closed-loop mode with bandwidth B_θ , 20 dB
B_θ^1	carrier tracking loop bandwidth that applies when S_k 's SNR lies between γ_0 and γ_1 , 5 Hz
$B_{t_s}^0$	pre-phase-lock code tracking bandwidth, 3 Hz
$B_{t_s}^1$	first transient code tracking bandwidth, 3 Hz
$B_{t_s}^2$	second transient code tracking bandwidth, 0.5 Hz
B_{t_s}	steady-state code tracking bandwidth, 0.5 Hz

Several other tracking parameters, such as the timing for the staged code phase tracking modes and the interval over which $s_\theta \geq \bar{s}_\theta$ must hold before a transition to coherent code phase tracking, have been omitted for brevity. Moreover, the many parameters that govern the behavior of the central estimator, which, due to the vector tracking architecture shown in Fig. 6, influences signal tracking, have not been mentioned at all. Clearly, there are many degrees of freedom over which signal tracking may be optimized for urban RTK.

Are the values listed above optimal for urban RTK? Surely not. Could better values be selected *a priori* based on signal processing theory? Not likely. Is it reasonable to adjust parameter values one at a time to maximize P_S for $P_F \leq \bar{P}_F$? Yes;

such was the approach taken for the current paper. But it is quite possible that a combinatorial approach such as a genetic algorithm would yield better results.

E. Precise positioning

PpRx and the LDRN send carrier and code phase observables, together with signal quality indicators s_θ and C/N_0 , and various other meta-data, 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 paper, PpEngine was invoked only in its simplest single-antenna mode, producing a precise 3-dimensional 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 [8].

1) *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 double differences (DDs) from the rover and reference data, 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 [33], but limits growth of the number of integer states by either (1) marginalizing at each epoch over float-valued integer ambiguity states modeled as Gaussian-distributed, or (2) 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 so-called float solution, can be thought of as a special case of the sub-optimal filter in [33] with a window length $i = 1$. The conditioning option, which yields the so-called fixed solution, is invoked only if the integer estimates, found by integer least squares (ILS) [39], 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 vicious 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.

How often is this cycle entered? For an aperture test configured for a fixed failure rate of \bar{P}_F , it occurs with probability as high as \bar{P}_F even in the ideal case where code and carrier measurement errors are mutually independent, independent in time, and Gaussian-distributed. In fact, all three of these conditions are violated in significant measure for urban RTK, leading to false fixing rates several times larger than \bar{P}_F .

Of course, if the P_F experienced in practice could somehow be bounded below an extremely small value, the risk associated with carrying forward integer-conditioned real-valued states could be made tolerably low. To this end, one might expect that, for good geometry and an adequate number of signals (a strong *a priori* model in the language of [28]), and for some arbitrarily small bound \bar{P}_F , there ought to exist an aperture test that, when configured to respect a sufficiently small \bar{P}_F , would yield an empirical $P_F \leq \bar{P}_F$ while still maintaining a reasonably high P_S . But this does not appear to be the case. Urban multipath and blockage can at times conspire to generate an extremely self-consistent yet erroneous measurement set whose associated integer solution can pass all but the strictest of aperture tests. To prevent such erroneous integer estimates from slipping through and corrupting the real-valued states requires an aperture test so strict that it also excludes almost all correctly-fixed integer estimates, despite a strong underlying model.

Given the above considerations, it would seem folly to pursue a strategy of carrying forward integer-conditioned real-valued states. Recognizing this, the authors are developing a generalization of PpEngine that can manage growth in the number of integer state elements using a variant of the suboptimal approach of [33]. Nonetheless, as will be shown, the current approach is remarkably effective in a moderate urban environment. It also has the virtue of simplicity, making it a good choice for initial development, and of computational efficiency, permitting rapid experimental iteration.

2) *Dynamics Model:* Because this paper'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 available inertial measurements and instead rely on a simple nearly-constant-velocity dynamics model for state propagation between measurements [40]. The dynamics model assumes roughly equivalent process noise in the along-track and cross-track directions, but much smaller process noise (by a factor of 10) in the vertical direction, in keeping with a land vehicle operating in a relatively flat urban environment.

3) *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 [41].

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 [42], [43], 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 multi-level exclusion process, depicted in Fig. 6, 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_θ , and elevation angle θ_{el} . Signals whose values fall below user-selected thresholds for these quantities are excluded from all DD combinations.

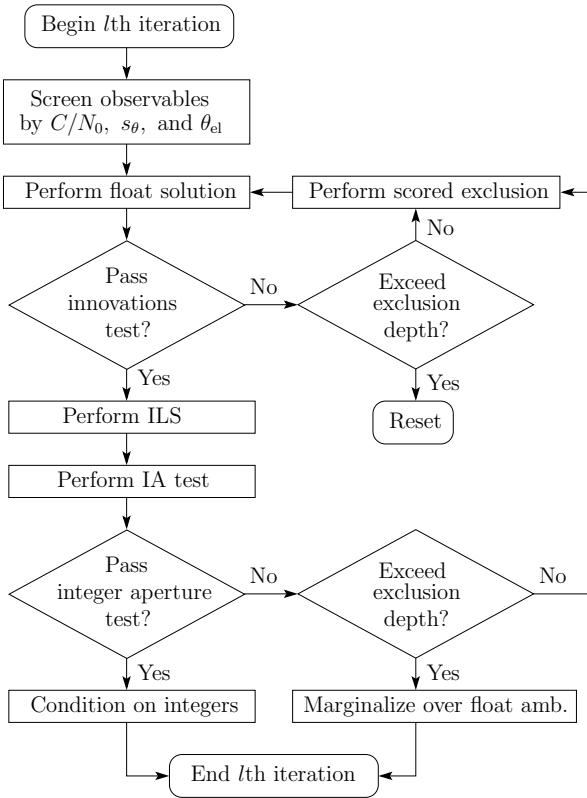


Fig. 6. Flow diagram for the PpEngine exclusion and fixing logic.

A second level of exclusion occurs as part of the float solution. A χ^2 -type test is applied to all DD measurement innovations [40], with exclusion triggered if the normalized innovations squared statistic exceeds a chosen threshold. For the current implementation of PpEngine, this test is only effective at excluding anomalous DD code phase (pseudorange) measurements, since 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_θ , and θ_{el} . 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 on successful validation by some type of aperture test; otherwise, the float solution is accepted [27]. 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 [43].

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.

IV. EXPERIMENTAL SETUP

The precise positioning system was evaluated experimentally using data collected on December 18 and 21, 2017, and January 15, 2018 in a moderate urban environment north of the University of Texas campus in Austin, TX. A total of 3.4 hours of data were collected in 4 sessions over the three days.

A. Rover platform

The rover GNSS receiver is one among several sensors housed in an integrated perception platform called the Sensorium, pictured in Fig. 7. Designed for connected and automated vehicle research, the Sensorium is a self-contained sensor housing that can be mounted atop any standard passenger vehicle. Although hardly visible in Fig. 7, two Antcom 53G1215A-XT dual-frequency patch antennas are flush-mounted in the cross-track direction on the Sensorium’s

upper plate, separated by just over 1 meter. The antennas' signals are routed to a unified RF front end whose output IF samples are processed in real time (to within less than 10 ms latency) by the Sensorium's onboard computer. The samples are also stored to disk for post-processing.

Data from both the primary (driver's side) and secondary (passenger's side) antennas were used to reconstruct the ground truth trajectory, as described below. But only data from the primary antenna were used in the urban RTK performance evaluation. No other Sensorium sensors were involved in the current paper's results.



Fig. 7. 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, dual-frequency software-defined GNSS receiver, 4G cellular connectivity, and a powerful internal computer.

B. Test route

The test route was primarily composed of two nested loops, as shown in Fig. 8, although other roadways near this area were also explored. The route includes a variety of light-to-moderate urban conditions, from open-sky to overhanging trees to built-up areas. The Dean Keeton corridor, pictured in Figs. 9 and 10, is the most challenging stretch along the test route. It passes below two pedestrian bridges and is flanked on both sides by buildings ranging from 30 to 65 meters tall set back 28 meters from the center of the roadway.

A repeating route is not ideal for evaluating urban RTK performance; a route with greater variety of urban locations would be better. The route in Fig. 8 was chosen to support other research in visual 3D mapping for which repeated sessions through the same corridors was necessary to explore the map's accuracy convergence. Future work in urban RTK will explore routes with greater variety. Meanwhile, the current route remains useful for urban RTK performance evaluation because satellite movement causes multipath and signal blockage conditions to differ significantly from lap to lap and day to day.

V. GROUND TRUTH TRAJECTORY

A primary challenge in urban RTK performance evaluation is obtaining a ground truth trajectory against which to compare the reported trajectory of the system under test. Ideally, the ground truth should be complete and provably accurate to better than 5 cm. Under such conditions, the incorrect fixing rate



Fig. 8. Overview of the test route. The smaller triangular loop was driven 30 times, the larger loop 19 times, both clockwise. The area shown lies on the north side of The University of Texas campus in Austin, TX.



Fig. 9. A 3D overview of the Dean Keeton corridor, spanned by two pedestrian bridges and flanked by buildings on both sides.

P_F can be measured by declaring an incorrect fix whenever the reported trajectory deviates by more than 15 cm from the ground truth. A more straightforward comparison of resolved integers between the reference and test systems, as in [4], is not generally possible in urban areas because the systems may not track the same set of satellites at each epoch throughout the test.

But what system is capable of provably determining the location of a mobile GNSS antenna in an urban area to within 5 cm of its true location? Prior work in urban positioning has relied on forward-backward smoothed trajectories from coupled RTK-inertial systems with a tactical-grade IMU [8], [14]. But the estimated 95% accuracy of the truth trajectory in [14] was only 55 cm in residential areas, and 59 cm in urban canyons, which is far too loose to allow confident measurement of P_F . And the authors of [8] make no attempt to assess the accuracy of their reference system.

A. Examination of DD carrier phase residuals

A defensible claim of better-than-5-cm ground truth accuracy is made in [4], where a tactical-grade IMU was tightly integrated with an RTK system on a short baseline in nearly-perfect open-sky rural conditions. The reference system did experience one brief GNSS outage as the vehicle passed under a roadway sign, but Petovello et al. assert that “an analysis of the measurement residuals and the static position at the end




Fig. 10. A street-level view of the Dean Keeton corridor just before passing below the east pedestrian bridge.

of the run confirmed that the ambiguities had indeed been resolved correctly.”

In urban areas, DD code and carrier measurement residuals likewise carry valuable information about whether the reference system’s integer ambiguities have been correctly resolved. In particular, the absence of a pronounced secular trend in the DD carrier residuals over an interval spanning 30 seconds or more during which the integer estimate remains constant is a strong indicator that the associated ambiguity has been resolved correctly over that interval. Unfortunately, these conditions do not always hold in urban testing. The DD carrier residuals shown in Fig. 11 for a particular GPS signal are small (all less than 50 mm) and appear to be free from secular trends. Moreover, the vast majority of intervals over which the associated estimated ambiguities remain constant span more than 30 seconds. For these intervals, a confident declaration of correct integer resolution can be made. But over the 1-hour interval shown, there are several constant-integer intervals that are too short to allow confident declaration of correctly-resolved integers.

B. Dual-antenna ground truth generation

Given that residuals analysis alone appears insufficient to confidently discern a truth trajectory for urban RTK, an alternative approach is adopted for this paper.

1) *Independent RTK solutions*: Independent RTK solutions are obtained for both the primary and secondary rover antennas. Because these are separated by several GNSS wavelengths on the Sensorium’s top plate, and given the wavelength-scale sensitivity of multipath phase to path length, they experience significantly uncorrelated carrier multipath except when the multipath source is to the front or rear of the vehicle, which is rare (multipath reflections come primarily from buildings to the left or right of a vehicle [44]).

2) *Strict trajectory*: From the independent RTK solutions, a *strict* trajectory is constructed. Let S_P be the set of primary antenna positions that have passed aperture test validation, and let T_P be the set of associated time points. Let S_S and T_S be equivalent sets, respectively, for the secondary antenna. The intersection set $T_\cap = T_P \cap T_S$ contains time points at which both the primary and secondary fixed solutions are available.




Fig. 11. DD carrier phase (top) and pseudorange (center) residuals for GPS PRN 26’s L1 C/A signal for the 1-hour Jan. 15, 2018 test session. The lower plot shows the satellite’s elevation angle (gray) and the pivot satellite’s elevation angle (black).

Let S_{P_\cap} and S_{S_\cap} be the sets of positions for the primary and secondary antennas associated with T_\cap . These are ordered such that the k th primary position $s_{P,k} \in S_{P_\cap}$ corresponds to the k th secondary position $s_{S,k} \in S_{S_\cap}$. The strict trajectory $S_{S_\cap}^{\text{strict}} \subseteq S_{S_\cap}$ is defined as the set of positions $s_{P,k}$ for which $|b_k - b| < \epsilon$, where

$$b_k = \|s_{P,k} - s_{S,k}\|$$

is the measured baseline length between the two antennas at the k th epoch, b is the known baseline length, and ϵ is an acceptable error threshold, taken to be 15 cm for this paper. Thus, the strict trajectory consists of all fixed primary antenna positions that can be checked against, and are consistent with, a corresponding fixed secondary antenna position. The strict trajectory for an inner loop in the test route is shown as the blue trace in Figs. 12 and 13.

3) *Filling gaps in the strict trajectory*: Approximately 95.5% of epochs in the 3.4 hours of collected data find correspondence in the strict trajectory. The remaining 4.5% of epochs are those for which either the primary or secondary antenna did not produce a fixed solution, or the magnitude of the solution difference disagreed with the known baseline b by more than $\epsilon = 15$ cm. Perhaps unsurprisingly, almost all of these missing or outlier epochs occur near the bridges along the Dean Keeton corridor. Two techniques are applied to attempt fill-in of the strict trajectory gaps.

Transfer via the baseline-constrained solution: A fixed local RTK solution between the primary and secondary rover antennas is obtained at every epoch possible. This local solution’s availability is increased by applying the known baseline constraint b , as in [45]. A missing epoch in the strict trajectory is populated via transfer from the secondary antenna under the following conditions: (i) a fixed secondary antenna position is available, (ii) a fixed baseline-constrained primary-




Fig. 12. Horizontal component of a 130-second segment of the ground truth trajectory showing strict (blue), transferred (green), interpolated (red), and unavailable epochs. The red and green points on the lower east-west segment correspond to epochs when the rover vehicle passed under the east Dean Keeton pedestrian bridge.




Fig. 13. Vertical component time history associated with the trace in Fig. 12.

to-secondary solution is available, and (iii) when the secondary and local solution vectors are added, the resulting “transferred” primary antenna location is within ϵ of the position predicted by a shape-preserving piecewise cubic interpolation across the strict trajectory gap. The interpolation is carried out via Matlab’s `interp1` function with `METHOD = 'cubic'`. Transferred epochs are shown as green points in Figs. 12 and 13.

Short-hop interpolation: After all permissible secondary-to-primary solutions have been transferred, an attempt is made to bridge the remaining gaps via cubic interpolation (again using the `interp1` function). Only spans shorter than 5 seconds are allowed to be interpolated. Interpolated epochs are shown as red points in Figs. 12 and 13.

C. Ground truth completeness and accuracy

The ground truth trajectory constructed as described above from the strict trajectory, with gaps filled where possible, encompasses approximately 97% of all epochs in the 3.4-hour data set. The missing 3% of epochs almost all coincide with areas near the bridges passing over the Dean Keeton corridor. The west bridge is especially problematic, as the rover vehicle is often stopped for an extended time near or under the bridge while waiting for the stoplight at Dean Keeton and Guadalupe.

Visual inspection of the ground truth trajectory reveals no obvious errors. However, the authors concede that better-than-5-cm accuracy of the available ground truth cannot be guaranteed. Future work will explore additional approaches for completing and verifying ground truth for urban RTK.

VI. BASELINE SYSTEM PERFORMANCE

The baseline urban RTK system is the configuration of the University of Texas precise positioning system that maximizes P_S while respecting $P_F \leq \bar{P}_F$ for some chosen empirical incorrect fixing probability bound \bar{P}_F . This section discusses the baseline system’s performance. The following section will compare the baseline system against several alternatively-configured systems.

A. Baseline configuration

1) **PpRx:** The baseline system’s carrier tracking loops were configured as detailed in Section III-D, with minor variations for the different signal types. The code tracking loops’ default bandwidth was 0.5 Hz. PpRx was configured 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. PpRx was configured to output observables at 5 Hz.

2) **PpEngine:** The baseline system’s RTK engine was configured as follows. The master LDRN reference station, located within 1.2 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, $s_\theta \geq 0.55$, and $\theta_{el} \geq 15$ deg. Signals whose values fell below any one of these thresholds were excluded from all 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 2. Scored N -choose-1 exclusion was applied for both failed float innovations tests and failed aperture tests. A depth of 5 signals was allowed for the N -choose-1 exclusion, after which the estimator was either reset or integers marginalized, according to the flow diagram in Fig. 6. The difference test of [30], 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 $\sigma_\rho = 1$ m and $\sigma_\phi = 3.5$ mm, respectively. The nearly-constant-velocity dynamics model was configured for a 0.2 m/s deviation in horizontal velocity, and a 0.02 m/s deviation in vertical velocity over a 1-second interval.

A rough calibration of the vertical L1-L2 phase center offset was attempted for the primary rover antenna as mounted on the Sensorium.

B. P_S and P_F for the baseline system

Because the ground truth trajectory is incomplete, empirical P_S and P_F can only be determined to within upper and lower bounds. Let P_V denote the empirical probability of successful integer validation at any given epoch, whether the estimated integers are correct or not. In other words, P_V is the availability of a fixed solution. Let P_T denote the fraction of validated epochs that are testable; that is, the fraction having a corresponding populated epoch in the ground truth trajectory. Let P_E denote the fraction of testable epochs that are assumed to have incorrect integer estimates because their 3-dimensional position differs by more than $\epsilon = 15$ cm from the ground truth position.

For the baseline system over the 3.4-hour data set taken over three days, the average values of P_V , P_T , and P_E were

$$P_V = 0.988, P_T = 0.976, P_E = 0.0048$$

With these values one can bound the probability of incorrect fix P_F as

$$P_V P_T P_E = 0.0046 \leq P_F \leq 0.0234 = P_V(1 - P_T)$$

The lower bound optimistically assumes that none of the non-testable validated epochs were in error, whereas the upper bound pessimistically assumes that all non-testable validated epochs were in error. Given that the probability of successful fix $P_S = P_V - P_F$, one can similarly bound P_S :

$$P_V P_T = 0.965 \leq P_S \leq 0.984 = P_V(1 - P_T P_E)$$

Whether the baseline system's performance is impressive depends on the actual value of P_F . Visual inspection seems to indicate that $P_F < 0.01$, which is a factor of 10 larger than \bar{P}_F but may be tolerable for a larger system that combines stand-alone RTK with inertial and electro-optical sensing, as the Sensorium of Fig. 7 is intended to do.

The current system's lower bound for P_S , 0.965, is significantly higher than the $P_S = 0.767$ reported by Li et al. [8] for dual-frequency GPS + BDS + GLONASS stand-alone RTK, but a fair comparison is complicated by the facts that (1) Li et al. do not report P_F , and (2) the Wuhan and Austin urban testing environments are different.

C. Error distribution

Fig. 14 shows the cumulative distribution function (CDF) and a time history of errors for the baseline system's testable aperture-test-validated positions for the Jan. 15, 2018 data set. Positioning performance appears excellent, with over 99.7% of testable epochs having errors smaller than 10 cm. But

one should bear in mind that errors in aperture-test-validated but non-testable solutions are not shown in Fig. 14. Also, among the few errors that do appear, two exceed 1 meter over the 1-hour interval, which would make the baseline system unacceptable as the sole positioning sensor for connected or automated vehicles.




Fig. 14. Cumulative distribution function (top) and time history (bottom) of positioning errors with respect to the ground truth for the baseline system.

D. Residuals Time Histories

It is instructive to examine the DD carrier and code residuals for urban RTK. Fig. 11, discussed previously, shows residuals for GPS PRN 26 during the 1-hour Jan. 15, 2018 session. The RMS values of the time histories correspond to undifferenced phase and pseudorange deviations of $\sigma_\phi = 4.1$ mm and $\sigma_\rho = 0.65$ m. Note that the pseudorange errors are large during the first 250 seconds, over which the vehicle was stationary. This is because PpRx was configured with a large code tracking bandwidth (0.5 Hz) and operates on fairly narrow-band signals (4 MHz) without any dedicated multipath mitigation. When the vehicle begins to move, code multipath gets averaged out non-coherently due to rapid changes in the multipath phase.

Fig. 15 is identical to Fig. 11 but for the L2C signal from GPS PRN 26. For this signal, the undifferenced deviations are slightly larger than for its L1 C/A counterpart: $\sigma_\phi = 4.6$ mm and $\sigma_\rho = 0.78$ m. This appears to be generally the case for L2C signals despite their having the same ranging code bandwidth, and nearly the same power, as L1 C/A signals.

Fig. 16 shows DD residuals for Galileo PRN 4, whose elevation time history is similar to GPS PRN 26 over the same interval. For this signal, $\sigma_\phi = 4.2$ mm and $\sigma_\rho = 0.78$ m. Finally, Fig. 17 shows DD residuals for WAAS PRN 131, for which $\sigma_\phi = 4.4$ mm and $\sigma_\rho = 0.65$ m. The WAAS signal benefits from a high (49-deg. from Austin) and constant elevation angle, but its DD pseudorange residual has a 1.5-meter bias, likely due to asymmetry in the WAAS signal's autocorrelation function. Similarly-constant biases, albeit with different values, are observed for WAAS PRNs 135 and 138.




Fig. 15. As Fig. 11 but for PRN 26's L2C signal.




Fig. 17. As Fig. 11 but for WAAS PRN 131 on L1.




Fig. 16. As Fig. 11 but for Galileo PRN 4's E1 (B+C) signal.

VII. PERFORMANCE DEGRADATION ANALYSIS

This section reports on one of this paper's most significant contributions, 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. The analysis was limited to the 1-hour data set from Jan. 15, 2018. Table I, where P_V , P_T , and P_E are as defined previously, summarizes the results of the analysis. Starting with Scenario 2, subsections below will discuss each scenario in turn.

A. LNAV data bit prediction disabled

Eliminating the baseline's system's LNAV data bit prediction capability, which was described in Section III-D2, has a devastating effect on performance. The availability of validated epochs is scarcely reduced, but P_E rises tremendously, from 0.23% to 36.4%. Fig. 18 shows that large errors persist over

many tens of seconds. It is possible that the PpRx carrier tracking strategy, or the PpEngine integer fixing strategy, could be better tailored for the case where LNAV data bit prediction is disabled, thus reducing P_E , but this would likely cause a significant drop in P_V . Clearly, LNAV data bit prediction is a key capability for urban RTK.

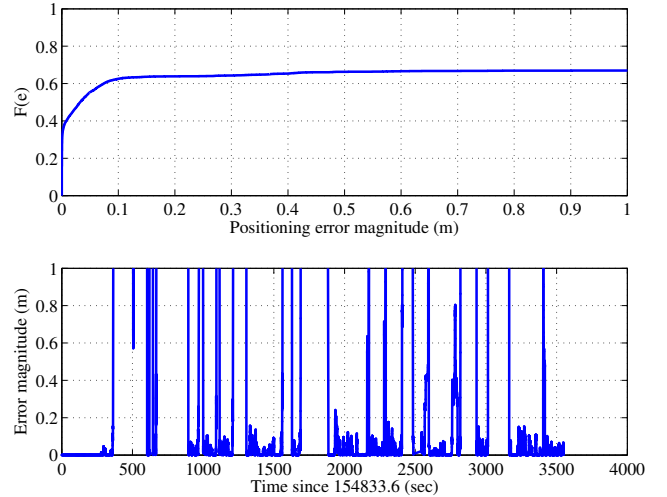


Fig. 18. As Fig. 14 but for Scenario 2.

B. Scalar tracking with adaptive B_θ

Eliminating vector tracking, as described in Section III-D4, in favor of scalar tracking, but retaining carrier tracking loop bandwidth adaptation, has no significant effect on P_V but P_E increases from 0.23% to 5.1%. Fig. 19 shows that the increase in P_E is primarily due to a single 150-second interval with a persistent large error. Further inspection reveals that the error interval begins at the west bridge. Thus, vector tracking appears helpful, but not critically so.

TABLE I
SUMMARY OF PRECISE POSITIONING RESULTS

Scenario	Description	P_V : Validated Epochs (%)	P_T : Testable (%)	P_E : Errors (%)
1	Baseline system	98.9	96.0	0.23
2	LNAV data bit prediction disabled	97.6	96.2	36.4
3	Scalar tracking with adaptive B_θ	98.7	96.2	5.1
4	Scalar tracking with fixed B_θ	98.5	96.2	4.3
5	GPS L2CL tracking	94.3	96.2	1.8
6	Age of data = 200 ms	98.8	96.0	0.27
7	Age of data = 400 ms	98.4	96.4	0.27
8	Age of data = 600 ms	98.3	96.4	0.35
9	Age of data = 800 ms	98.3	96.4	0.35
10	10 km baseline	97.9	96.4	2.0
11	Sans WAAS	98.7	96.0	4.2
12	Sans GPS L2C (L+M)	97.0	96.4	2.7
13	Sans Galileo E1 (B+C)	95.9	96.5	6.8
14	No scored exclusion	96.6	96.8	4.9
15	35 deg. el. mask	92.3	96.1	6.6
16	35 deg. el. mask, 10 km baseline	81.1	95.9	16.8




Fig. 19. As Fig. 14 but for Scenario 3.

C. Scalar tracking with fixed B_θ

Eliminating both vector tracking and carrier tracking loop bandwidth adaptation does not appear to cause significant further degradation beyond eliminating vector tracking.

D. 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 L2CL tracking might be thought a more reliable strategy given that no symbol wipeoff mistakes are ever made when tracking only the pilot. But the results of Table I indicate that joint L+M is indeed valuable, as it increases P_V and decreases P_E .

E. Age of data

Scenarios 6-9 explore the effect of increased age of reference data, from the baseline age (near zero latency relative to the rover stream) to 800 ms. Very little reduction occurs in P_V , and little increase in P_E , indicating that the baseline system is not particularly sensitive to increased age of data. However, other experimentation has shown that an age of data beyond 800 ms begins to affect the WAAS carrier DDs. The cause of this degradation is the relatively poor stability of the WAAS clocks, which degrades the accuracy of carrier phase extrapolation to the rover epoch. No such effect appears for GPS or Galileo signals until an age of data beyond 10 seconds.

F. 10-km baseline

The baseline system's distance to the reference receiver, commonly referred to as the reference-rover baseline, is no greater than 1.2 km. For Scenario 10, the LDRN alternate master station, which sits 10 km from the test route, was instead taken as reference. The alternate master station has a Trimble Zephyr II antenna identical to the master station's. Note that a 10-km baseline is still considered to be within the short-baseline regime for standard RTK [46]. Nonetheless, a slight decrease in P_V and increase in P_E is observed. The CDF in the top panel of Fig. 20 also shows that the main drawback of the longer baseline is the increase in large errors, presumably due to incorrect integer fixing, rather than the slight decrease in accuracy of correctly-fixed solutions caused by the longer baseline.

G. Value of additional signals

Scenarios 11-13 explore the degradation in the baseline system that occurs when all signals of a particular type are eliminated from the RTK solution. Over the 1-hour interval, DDs based on 3 WAAS, 3 Galileo, and 4 GPS L2C signals were originally available. Table I indicates that loss of any one of the signal types degrades performance, with Galileo being the most important for the 1-hour data interval studied.




Fig. 20. As Fig. 14 but for Scenario 10.

H. No Scored Exclusion

Scenario 14 removes the scored exclusion strategy described in Section III-E3 by setting the exclusion depth to 0 caused a noticeable reduction in P_V and increase in P_E . Increasing the exclusion depth beyond 5, the baseline system's value, had no discernible effect.

I. 35 degree elevation mask angle

Scenarios 15 and 16 explore the effect on system performance of increasing the elevation mask angle from 15 to 35 degrees, simulating a denser urban environment. For these scenarios, the average number of DDs drops from above 16 (for the baseline system) to less than 11. Scenario 15 takes the usual ~ 1 -km baseline whereas Scenario 16 takes the 10-km baseline. Both scenarios exhibit significant degradation in both P_V and P_E , but the degradation is especially pronounced for the 10-km baseline, with P_V dropping to 81.1% and P_E rising to 16.8%. This is consistent with the argument in [17] that a dense reference network is especially important in urban settings with reduced signal availability.

VIII. CONCLUSIONS

A low-cost urban real-time kinematic (RTK) positioning system developed at The University of Texas for precise vehicular location has been described and evaluated. The system is unaided by inertial or electro-optical sensors. Over 3.4 hours of urban testing, the system achieves a probability of correct integer fixing greater than 96.5% for a probability of incorrect integer fixing surely less than 2.3% and likely less than 1%. Fixed integer solutions are available for 99% of measurement epochs. Of these, 96% are testable against a ground truth trajectory and are shown to be accurate to within 10 cm over 99.7% of the time. A performance sensitivity analysis revealed that navigation data bit prediction on the GPS L1 C/A signals is key to high-performance urban RTK positioning, and that various other features of the positioning system contribute in minor but cumulatively significant ways.

ACKNOWLEDGMENTS

This work has been supported by the National Science Foundation under Grant No. 1454474 (CAREER), by the Data-supported Transportation Operations and Planning Center (DSTOP), a Tier 1 USDOT University Transportation Center, and by the University of Texas Situation-Aware Vehicular Engineering Systems (SAVES) Center (<http://utsaves.org/>), an initiative of the Wireless Networking and Communications Group.

REFERENCES

- [1] kdday, "Inside the NVIDIA PX2 board on my HW2 AP2.0 Model S," May 2017, <https://goo.gl/wNmzh2>.
- [2] J. B. Kenney, "Dedicated short-range communications (DSRC) standards in the united states," *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1162–1182, 2011.
- [3] D. Fajardo, T.-C. Au, S. Waller, P. Stone, and D. Yang, "Automated intersection control: Performance of future innovation versus current traffic signal control," *Transportation Research Record: Journal of the Transportation Research Board*, no. 2259, pp. 223–232, 2011.
- [4] M. Petovello, M. Cannon, and G. Lachapelle, "Benefits of using a tactical-grade IMU for high-accuracy positioning," *Navigation, Journal of the Institute of Navigation*, vol. 51, no. 1, pp. 1–12, 2004.
- [5] B. M. Scherzinger, "Precise robust positioning with inertially aided RTK," *Navigation*, vol. 53, no. 2, pp. 73–83, 2006.
- [6] H. T. Zhang, "Performance comparison on kinematic GPS integrated with different tactical-grade IMUs," Master's thesis, The University of Calgary, Jan. 2006.
- [7] S. Kennedy, J. Hamilton, and H. Martell, "Architecture and system performance of SPAN—NovAtel's GPS/INS solution," in *Position, Location, And Navigation Symposium, 2006 IEEE/ION*. IEEE, 2006, p. 266.
- [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 Sensing*, vol. 10, no. 2, p. 205, 2018.
- [9] S. Godha, "Performance evaluation of low cost MEMS-based IMU integrated with GPS for land vehicle navigation application," Master's thesis, The University of Calgary, 2006.
- [10] P. Teunissen and A. Khodabandeh, "Review and principles of PPP-RTK methods," *Journal of Geodesy*, vol. 89, no. 3, pp. 217–240, 2015.
- [11] Y. Cui, X. Meng, Q. Chen, Y. Gao, C. Xu, S. Roberts, and Y. Wang, "Feasibility analysis of low-cost GNSS receivers for achieving required positioning performance in CAV applications," in *2017 Forum on Cooperative Positioning and Service*, May 2017, pp. 355–361.
- [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 *Proceedings of the IEEE/ION PLANS Meeting*, May 2014.
- [13] K. M. Pesyna, Jr., "Advanced techniques for centimeter-accurate GNSS positioning on low-cost mobile platforms," Ph.D. dissertation, The University 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. 3297–3308.
- [15] J. Jackson, B. Davis, and D. Gebre-Egziabher, "An assessment of low-cost RTK GNSS receivers," in *Proceedings of the IEEE/ION PLANS Meeting*, Monterey, CA, 2018.
- [16] S. Riley, H. Landau, V. Gomez, N. Mishukova, W. Lentz, and A. Clare, "Positioning with Android GNSS observables," *GPS World*, vol. 29, no. 1, pp. 18–34, Jan. 2018.
- [17] M. J. Murrian, C. W. Gonzalez, T. E. Humphreys, K. M. P. Jr., D. P. Shepard, and A. J. Kerns, "Low-cost precise positioning for automated vehicles," *GPS World*, vol. 27, no. 9, pp. 32–39, September 2016.
- [18] M. L. Psiaki and H. Jung, "Extended Kalman filter methods for tracking weak GPS signals," in *Proceedings of the ION GPS Meeting*. Portland, Oregon: Institute of Navigation, 2002, pp. 2539–2553.

- [19] H. Niedermeier, B. Eissfeller, J. Winkel, T. Pany, B. Riedl, T. Woz, R. Schweikert, S. Lagrasta, G. Lopez-Risueno, and D. Jimenez-Banos, "Dingpos: High sensitivity GNSS platform for deep indoor scenarios," in *Indoor Positioning and Indoor Navigation (IPIN), 2010 International Conference on*. IEEE, 2010, pp. 1–10.
- [20] S. Mohiuddin and M. L. Psiaki, "High-altitude satellite relative navigation using carrier-phase differential global positioning system techniques," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1628–1639, Sept.–Oct. 2007.
- [21] S. Verhagen, P. J. Teunissen, and D. Odijk, "The future of single-frequency integer ambiguity resolution," in *VII Hotine-Marussi Symposium on Mathematical Geodesy*. Springer, 2012, pp. 33–38.
- [22] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, and P. M. Kintner, Jr., "GNSS receiver implementation on a DSP: Status, challenges, and prospects," in *Proceedings of the ION GNSS Meeting*. Fort Worth, TX: Institute of Navigation, 2006, pp. 2370–2382.
- [23] 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.
- [24] 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, Journal of the Institute of Navigation*, vol. 61, no. 1, pp. 53–64, 2014.
- [25] M. Murrian, C. Gonzalez, T. E. Humphreys, and T. D. Novlan, "A dense reference network for mass-market centimeter-accurate positioning," in *Proceedings of the IEEE/ION PLANS Meeting*, Savannah, GA, 2016.
- [26] P. Teunissen, "Integer aperture gnss ambiguity resolution," *Artificial Satellites*, vol. 38, no. 3, pp. 79–88, 2003.
- [27] G. N. Green and T. E. Humphreys, "Data-driven generalized integer aperture bootstrapping for high-integrity positioning," *IEEE Transactions on Aerospace and Electronic Systems*, 2017, submitted for review. [Online]. Available: <http://goo.gl/pfq9EU>
- [28] —, "Position domain integrity analysis for generalized integer aperture bootstrapping," *IEEE Transactions on Aerospace and Electronic Systems*, 2017, in preparation. [Online]. Available: <http://goo.gl/qgKu9W>
- [29] P. Teunissen, "Integer aperture bootstrapping: a new GNSS ambiguity estimator with controllable fail-rate," *Journal of Geodesy*, vol. 79, no. 6-7, pp. 389–397, 2005. [Online]. Available: <http://dx.doi.org/10.1007/s00190-005-0481-y>
- [30] L. Wang and S. Verhagen, "A new ambiguity acceptance test threshold determination method with controllable failure rate," *Journal of Geodesy*, vol. 89, no. 4, pp. 361–375, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s00190-014-0780-2>
- [31] 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 Journal of Selected Topics in Signal Processing*, vol. 3, no. 4, pp. 707–715, Aug. 2009.
- [32] 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 Transactions on Aerospace and Electronic Systems*, vol. 46, no. 4, pp. 1609–1623, Oct. 2010.
- [33] M. Psiaki, "Kalman filtering and smoothing to estimate real-valued states and integer constants," *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 5, pp. 1404–1417, Sept.–Oct. 2010.
- [34] GPS Directorate, "Systems engineering and integration Interface Specification IS-GPS-200G," 2012, <http://www.gps.gov/technical/icwg/>.
- [35] T. E. Humphreys, M. L. Psiaki, and P. M. Kintner, Jr., "Modeling the effects of ionospheric scintillation on GPS carrier phase tracking," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 4, pp. 1624–1637, Oct. 2010.
- [36] A. J. Van Dierendonck, *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 8: GPS Receivers, pp. 329–407.
- [37] M. Lashley and D. Bevely, "What are vector tracking loops, and what are their benefits and drawbacks?" *GNSS Solutions Column, Inside GNSS*, vol. 4, no. 3, pp. 16–21, 2009.
- [38] S. A. Stephens and J. B. Thomas, "Controlled-root formulation for digital phase-locked loops," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 31, no. 1, pp. 78–95, Jan. 1995.
- [39] P. J. Teunissen, "The least-squares ambiguity decorrelation adjustment: a method for fast GPS integer ambiguity estimation," *Journal of Geodesy*, vol. 70, no. 1-2, pp. 65–82, 1995.
- [40] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley and Sons, 2001.
- [41] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge Univ Press, 2000, vol. 2.
- [42] A. Brack, "On reliable data-driven partial GNSS ambiguity resolution," *GPS Solutions*, vol. 19, no. 3, pp. 411–422, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s10291-014-0401-9>
- [43] 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.
- [44] E. Steingass and A. L. German, "Measuring the navigation multipath channel—a statistical analysis," in *Proceedings of the ION GNSS Meeting*, 2004.
- [45] K. Chiang, M. Psiaki, S. Powell, R. Miceli, and B. O'Hanlon, "GPS-based attitude determination for a spinning rocket," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 4, pp. 2654–2663, 2014.
- [46] R. Odolinski, P. Teunissen, and D. Odijk, "Combined GPS+BDS for short to long baseline RTK positioning," *Measurement Science and Technology*, vol. 26, no. 4, p. 045801, 2015.