

Low-Cost Inertial Aiding for Deep-Urban Tightly Coupled Multi-Antenna Precise GNSS

James E. Yoder | Todd E. Humphreys

Department of Aerospace Engineering
and Engineering Mechanics, The
University of Texas at Austin, Austin,
Texas

Correspondence

James E. Yoder, Department of Aerospace
Engineering and Engineering Mechanics,
The University of Texas at Austin, Austin,
TX 78712.

Email: jyoder@utexas.edu

Summary

A vehicular pose estimation technique is presented that tightly couples multi-antenna carrier-phase differential GNSS (CDGNSS) with a low-cost MEMS inertial sensor and vehicle dynamic constraints. This work is the first to explore the use of consumer-grade inertial sensors for tightly coupled urban CDGNSS, and first to explore the tightly coupled combination of multi-antenna CDGNSS and inertial sensing (of any quality) for urban navigation. An unscented linearization permits ambiguity resolution using traditional integer least-squares while both implicitly enforcing known-baseline-length constraints and exploiting the multi-baseline problem's inter-baseline correlations. A novel false fix detection and recovery technique is developed to mitigate the effect of conditioning the filter state on incorrect integers. When evaluated on the publicly available TEX-CUP urban positioning data set, the proposed technique achieves, with consumer- and industrial-grade inertial sensors, respectively, a 96.6% and 97.5% integer fix availability, and a 12.0-cm and 10.1-cm overall (fix and float) 95th percentile horizontal positioning error.

Keywords

CDGNSS, low-cost RTK positioning, urban vehicular positioning

1 | INTRODUCTION

The rise of connected and automated vehicles has created a need for robust globally referenced positioning with lane-level (e.g., sub-30-cm) accuracy (Reid et al., 2019). Much automated ground vehicle (AGV) research focuses on the use of lidar and cameras for navigation, but these sensing modalities often perform poorly in low-illumination conditions or during adverse weather such as heavy fog or snowy white-outs. By contrast, positioning techniques based on radio waves, such as automotive radar or GNSS, are robust to poor weather and lighting conditions (Narula et al., 2022). Recent work has found that fusing measurements from low-cost automotive radars with inertial sensing can provide lane-level accuracy in urban environments (Narula et al., 2022). But radar-based positioning in a global coordinate frame requires the production and maintenance of radar maps, which is a time-consuming and costly endeavor.

GNSS signals provide a source of high-accuracy all-weather absolute positioning that does not require expensive investments in systems for map production, storage, maintenance, and dissemination. If the so-called integer ambiguities associated with the carrier-phase measurements can be correctly resolved, carrier-phase-based GNSS positioning offers exquisite accuracy. However, GNSS signal blockage, diffraction, and multipath effects make this family of techniques extremely challenging to use in urban areas. Carrier-phase differential GNSS (CDGNSS), whose real-time variant for mobile platforms is commonly known as real-time kinematic (RTK) GNSS, is a centimeter-accurate positioning technique that differences a receiver's GNSS observables with those from a nearby fixed reference station to eliminate most sources of measurement error (Teunissen & Montenbruck, 2017). Previous work by this paper's authors probed the limits of unaided CDGNSS in the deep urban environment and found that the combination of a GNSS measurement engine optimized for urban positioning and robust estimation techniques for outlier exclusion make CDGNSS feasible in deep urban environments (Humphreys et al., 2020b). But the unaided CDGNSS system described in Humphreys et al. (2020b) suffers from availability gaps of up to 90 seconds in duration, making it insufficient to serve as the sole navigation sensor for an AGV.

A natural method to bridge such availability gaps is to incorporate measurements from an inertial measurement unit (IMU). These measurements are uniquely valuable due to their invulnerability to environmental effects such as radio interference and weather. Combined GNSS and inertial navigation systems that incorporate only GNSS position solutions as measurements for a downstream navigation filter are termed *loosely coupled*, whereas *tightly coupled* systems directly incorporate raw GNSS observables (pseudorange, Doppler, or carrier-phase; Teunissen & Montenbruck, 2017). While both loosely and tightly coupled aiding can bridge availability gaps, tightly coupled aiding additionally reduces these gaps' frequency and duration; the inertial sensor provides probabilistic constraints between GNSS measurement epochs that increase the success rate of carrier-phase integer ambiguity resolution. These constraints additionally make the navigation solution observable with fewer GNSS measurements.

AGV navigation filter performance can be further improved by tightly coupling with so-called vehicle dynamics constraints (VDCs). One such technique exploits the natural motion constraints of four-wheeled ground vehicles, commonly referred to as non-holonomic constraints (NHCs). A second VDC technique infers a lack of vehicle motion by monitoring, for example, wheel odometry ticks, or by detecting a lack of road vibration. This constraint is then enforced as a strong zero-velocity pseudo-measurement called a zero velocity update (ZUPT) in the literature.

This paper extends the navigation filter component of the CDGNSS system described in Humphreys et al. (2020b) by tightly coupling with an inertial sensor and with vehicle dynamics constraints, and by incorporating measurements from multiple vehicle-mounted GNSS antennas. It also develops a novel robust estimation technique to mitigate the effects of multipath and allow for graceful recovery from incorrect integer fixes.

1.1 | Related Work

This subsection reviews relevant existing literature on urban GNSS positioning, inertial aiding, vehicle motion constraints, and multi-antenna CDGNSS.

1.1.1 | *Unaided Urban CDGNSS*

Performance of CDGNSS unaided by inertial sensing in urban environments has historically been poor. Experiments in Ong et al. (2009) suffered from poor availability (<60%) and large positioning errors (>9 m root-mean-square [RMS]) in suburban and urban environments. A 2018 assessment of commercial CDGNSS receivers found that no low-cost solution offered greater than a 35% fixed-integer solution availability in urban environments (Jackson et al., 2018). Li et al. (2018) achieved a 76.7% unaided correct integer fixing rate in urban Wuhan, China, using dual-frequency CDGNSS with a professional-grade receiver. In 2019, Humphreys et al. (2020b) achieved an unaided correct integer fix rate of 84.8% in the urban core of Austin, Texas.

1.1.2 | *Inertial Aiding*

Tightly coupled inertial aiding has long been employed as a method to increase CDGNSS solution availability and robustness. Early systems built around highly accurate but expensive tactical-grade IMUs were capable of providing robust positioning in dense urban areas (Kennedy et al., 2006; Petovello et al., 2004; Scherzinger, 2006; Zhang, 2006). The recent emergence of inexpensive consumer- and industrial-grade micro-electromechanical system (MEMS) inertial sensors has led to a new chapter of research in low-cost inertial aiding for urban CDGNSS.

Li et al. (2018) demonstrated that tight coupling of single-antenna professional-grade GNSS measurements with an industrial-grade MEMS IMU increased the integer fix availability of single-frequency CDGNSS from 44.7% to 86.1% on a test route in urban Wuhan, China. However, the authors did not provide the GNSS data set, information on the incorrect integer fix rate, or a full error distribution, making these results difficult to assess.

This paper, in contrast, is the first to demonstrate an increased CDGNSS integer fix rate in an urban environment via tightly coupling with a consumer-grade inertial sensor. Furthermore, it incorporates vehicle dynamics constraints and multiple vehicular GNSS baselines. The system's performance is evaluated on a publicly available urban positioning data set, allowing for a head-to-head comparison of techniques by the urban positioning research community.

1.1.3 | *Tightly Coupled Urban PPP*

One disadvantage of CDGNSS is that it requires observations from a nearby base station to eliminate modeling errors (e.g., for atmospheric delays or satellite clocks and orbits) common to both the base station (the reference) and the vehicle (the rover). Short-baseline CDGNSS, which offers the greatest robustness against urban multipath (Murrian et al., 2016), is limited to reference-rover baseline lengths below approximately 10 km (Odijk, 2002). To avoid the requirement for a nearby base station, attention has recently focused on extending precise point positioning (PPP), which is based on precise orbit, clock, and atmospheric corrections, to urban areas by tightly coupling with inertial sensors.

Rabbou et al. in 2015 explored tight coupling of PPP with a tactical-grade inertial sensor in mostly open-sky conditions with simulated GNSS outages, achieving centimeter-level accuracy (Abd Rabbou & El-Rabbany, 2015). Gao et al. (2017) and Vana (2021) extended tightly coupled PPP to industrial-grade MEMS inertial

sensors in highway and suburban environments. More recently, Elmezayen and El-Rabbany (2021) demonstrated tightly coupled PPP using both a geodetic-grade and a low-cost GNSS receiver and an industrial-grade MEMS sensor along an urban route in downtown Toronto, Canada, but only achieved meter-level accuracy when using the low-cost GNSS receiver. A drawback of PPP-based positioning is that the aforementioned results all required a roughly 10-minute convergence period before producing an accurate navigation solution. Short-baseline CDGNSS positioning with a modern multi-frequency, multi-constellation receiver, by contrast, typically yields instantaneous initialization.

1.1.4 | Vehicle Dynamics Constraints

Recent research has also explored the tight coupling of CDGNSS measurements with vehicle dynamics constraints. Nagai et al. (2021) found, in a simulation study using a realistic 3D map of an urban environment, that a tightly coupled CDGNSS system using GPS only could feasibly provide high-integrity decimeter-level positioning when aided with vehicle-dynamics constraints, a tactical-grade IMU, and odometry based on wheel-speed sensors. Yang et al. (2021) explored tightly coupled single-antenna CDGNSS with non-holonomic constraints and a tactical-grade fiber-optic IMU, but only evaluated their system under open-sky GNSS conditions with simulated GNSS degradations.

1.1.5 | Multi-Antenna CDGNSS

The use of multiple GNSS antennas on the vehicle for CDGNSS offers four advantages. First, the full six-degrees-of-freedom vehicle pose (position and orientation) becomes instantaneously observable when CDGNSS measurements are combined with the gravity vector as measured by an inertial sensor. With a single GNSS antenna, the vehicle yaw is observable only over multiple epochs, and only if the vehicle accelerates during the observations (Hong et al., 2005). Second, the shared reference antenna creates redundancy in the measurement model that allows for better ambiguity resolution performance than any CDGNSS baseline taken individually (Medina et al., 2020). Third, the additional set of GNSS measurements at the second antenna provides reduced position estimation error. Fourth, a highly effective method for GNSS spoofing detection, the multi-antenna defense (Psiaki et al., 2014), can readily be implemented.

Multi-antenna GNSS has long been used for attitude-determination applications with snapshot estimation methods such as C-LAMBDA (Teunissen, 2006) and MC-LAMBDA (Giorgi & Teunissen, 2010), which provide globally optimal single-epoch maximum-likelihood solutions to the full nonlinear GNSS attitude determination problem, and have been successfully extended to the pose estimation case (Wu et al., 2020). Other work has incorporated special cases of a priori attitude information into the nonlinear solution process (Henkel & Günther, 2012). These snapshot methods, however, are computationally demanding, and their extension to recursive estimation for tight coupling with other sensors is not straightforward and remains unexplored.

Fan et al. (2019) found that a hard constraint using an a priori known vehicle attitude to combine CDGNSS observations from multiple vehicle antennas can increase ambiguity resolution and urban CDGNSS performance. However, this method requires a highly accurate independent source of attitude information,

such as from an expensive gyrocompass-capable tactical-grade IMU following an initial static alignment period.

Medina et al. (2020) proposed pose estimation based on multiple vehicle antennas for inland waterway navigation. This work sidestepped the complexity of C-LAMBDA or MC-LAMBDA by linearizing the attitude model in an extended Kalman filter (EKF) update and propagating the state with a simple motion model. This formulation was found to increase ambiguity resolution performance over either the positioning or attitude determination problems taken independently. However, the authors made no attempt to incorporate an inertial sensor or additional motion constraints.

Hirokawa and Ebinuma (2009) developed a multi-antenna GNSS system for aircraft pose estimation that tightly coupled with a MEMS inertial sensor, but only used CDGNSS for attitude measurements, relying on standard pseudorange measurements for the estimator's position component.

Henkel et al. (2020) tightly coupled triple-antenna CDGNSS with an industrial-grade inertial sensor for a micro air vehicle navigation application, but only evaluated the system's performance over a single, short test flight in open-sky conditions, and did not compare against a ground truth reference.

Previous work (Humphreys et al., 2020a; Yoder et al., 2020) by this paper's authors explored a suboptimal federated filtering approach to the tightly coupled multi-antenna CDGNSS + inertial problem, additionally incorporating monocular vision measurements from Yoder et al. (2020). But the approach did not properly model the multi-antenna CDGNSS measurement update, instead resolving the position and attitude baselines separately.

1.2 | Contributions

This paper makes five contributions:

1. An estimation technique that tightly couples multi-antenna CDGNSS with vehicle dynamics constraints and inertial measurements; to the best of the authors' knowledge, this paper is the first in the open literature to explore the tightly coupled combination of multi-antenna CDGNSS and inertial sensing for navigation in urban environments. Furthermore, it is the first to explore the use of *consumer-grade* inertial sensors for tightly coupled deep urban CDGNSS (Sections 3 and 4).
2. A novel application of the unscented transform for the multi-baseline CDGNSS integer ambiguity resolution and measurement update step, which widens the operating regime of the filter to allow for significantly greater attitude uncertainty without suffering from excessive integer least-squares (ILS) failures seen by existing EKF approaches (Section 3)
3. A novel false fix detection and recovery technique that limits the degree to which an incorrectly resolved integer ambiguity can corrupt the tightly coupled CDGNSS estimator's state (Section 4.4)
4. Demonstration of state-of-the-art deep urban CDGNSS performance, achieving, by tightly coupling with consumer-grade and industrial-grade inertial sensors, respectively, a 96.6% and 97.5% integer fix availability, and 12.0-cm and 10.1-cm overall (fix and float) 95th percentile horizontal positioning error on the publicly available TEX-CUP urban positioning data set (Narula et al., 2020; Sections 5.1 to 5.3).

5. A detailed evaluation and breakdown of the positioning and ambiguity resolution performance contribution of various sensors and algorithmic components (Section 5.5)

2 | COORDINATE AND NOTATION CONVENTIONS

2.1 | Vector Notation, Sensor Platform, and Coordinate Frames

Superscripts indicate the coordinate frames associated with vectors and rotation matrices. For example, \mathbf{r}^w denotes a vector \mathbf{r} expressed in the \mathbf{w} frame, and \mathbf{R}^{wb} denotes a rotation matrix that converts vectors from their representation in the \mathbf{b} frame to their representation in the \mathbf{w} frame (i.e., $\mathbf{r}^w = \mathbf{R}^{wb}\mathbf{r}^b$).

The sensor platform described in this paper and used in the evaluation in Section 5 is the University of Texas Sensorium (Narula et al., 2020), a roof-mounted vehicular perception platform incorporating multiple grades of inertial sensors, two GNSS antennas (denoted *primary* and *secondary*), stereo cameras, and three automotive radars. Only the inertial sensors and GNSS antennas are used in this work.

Several coordinate frames are referenced in this paper:

- u:** The *IMU frame* is centered at and aligned with the IMU accelerometer triad.
- b:** The *body frame* has its origin at the phase center of the Sensorium's primary GNSS antenna. Its x -axis points towards the phase center of the secondary antenna, its y -axis is aligned with the boresight vector of the primary antenna, and its z -axis completes the right-handed triad.
- v:** The *vehicle frame* is a body-fixed frame, centered at the vehicle's center of rotation as determined by an offline calibration using GNSS and IMU data. Its x -axis points in the direction of vehicle travel with no steering angle deflection, its z -axis points upwards, and its y -axis completes the right-handed triad.
- w:** The *world frame* is a fixed geographic east-north-up (ENU) frame with its origin at the phase center of the reference GNSS antenna, which is located at a fixed base station with known coordinates.

Figure 1 shows the relationships between these frames.

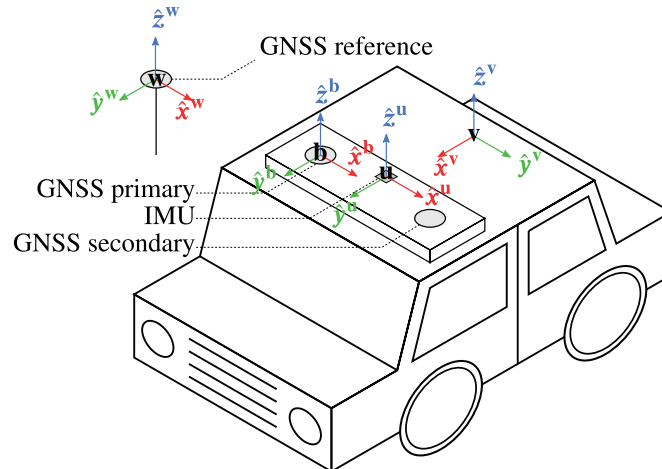


FIGURE 1 Diagram of relevant University of Texas Sensorium coordinate frames

2.2 | State Representation and Error-State Filtering

The tightly coupled navigation estimator described in this work is an unscented Kalman filter (UKF) that recursively fuses inertial measurements, double-difference (DD) GNSS pseudorange and carrier-phase measurements, and vehicle dynamics pseudomeasurements. The estimator's state at epoch k is given by the ordered set

$$\mathbf{x}_k = (\mathbf{r}_k^{\mathbf{w}}, \mathbf{v}_k^{\mathbf{w}}, \mathbf{R}_k^{\mathbf{wb}}, \mathbf{b}_{ak}^{\mathbf{u}}, \mathbf{b}_{gk}^{\mathbf{u}})$$

where $\mathbf{r}^{\mathbf{w}} \in \mathbb{R}^3$ is the position of the \mathbf{u} frame origin in the \mathbf{w} frame; $\mathbf{v}^{\mathbf{w}} \in \mathbb{R}^3$ is the velocity of the \mathbf{u} frame origin relative to the \mathbf{w} frame, expressed in the \mathbf{w} frame; $\mathbf{R}^{\mathbf{wb}} \in \text{SO}(3)$ is the attitude of the \mathbf{b} frame relative to the \mathbf{w} frame; and $\mathbf{b}_a^{\mathbf{u}}, \mathbf{b}_g^{\mathbf{u}} \in \mathbb{R}^3$ are the IMU's accelerometer and gyro biases, respectively.

This paper adopts the conventions and notation of Solà et al. (2018), which appeals to Lie theory to unify and generalize various error-state filtering methods so that the filtering equations are agnostic to the specific choice of attitude parameterization. The filter state, \mathbf{x}_k , is a point on the composite manifold $\mathcal{X} \triangleq \mathbb{R}^3 \times \mathbb{R}^3 \times \text{SO}(3) \times \mathbb{R}^3 \times \mathbb{R}^3$, which has $N_x = 15$ independent degrees of freedom. A state increment, $\delta \mathbf{x}_k$, is defined on the tangent space of \mathcal{X} at \mathbf{x}_k , which can be parameterized with the vector space \mathbb{R}^{N_x} . These spaces are related using the operators $\oplus: \mathcal{X} \times \mathbb{R}^{N_x} \rightarrow \mathcal{X}$ and $\ominus: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{N_x}$, which correspond to standard addition and subtraction for vector-valued components of \mathbf{x}_k and to more complex operations for the attitude component:

$$\mathbf{x}_k \oplus \delta \mathbf{x}_k \triangleq \begin{bmatrix} \mathbf{r}_k^{\mathbf{w}} + \delta \mathbf{r}_k^{\mathbf{w}} \\ \mathbf{v}_k^{\mathbf{w}} + \delta \mathbf{v}_k^{\mathbf{w}} \\ \mathbf{R}_k^{\mathbf{wb}} \circ \text{Exp}(\delta \mathbf{R}_k^{\mathbf{wb}}) \\ \mathbf{b}_{ak}^{\mathbf{u}} + \delta \mathbf{b}_{ak}^{\mathbf{u}} \\ \mathbf{b}_{gk}^{\mathbf{u}} + \delta \mathbf{b}_{gk}^{\mathbf{u}} \end{bmatrix} \in \mathcal{X}$$

$$\mathbf{x}_k \ominus \mathbf{x}_j \triangleq \begin{bmatrix} \mathbf{r}_k^{\mathbf{w}} - \mathbf{r}_j^{\mathbf{w}} \\ \mathbf{v}_k^{\mathbf{w}} - \mathbf{v}_j^{\mathbf{w}} \\ \text{Log}(\mathbf{R}_k^{\mathbf{wb}} \circ \mathbf{R}_j^{\mathbf{wb}-1}) \\ \mathbf{b}_{ak}^{\mathbf{u}} - \mathbf{b}_{aj}^{\mathbf{u}} \\ \mathbf{b}_{gk}^{\mathbf{u}} - \mathbf{b}_{gj}^{\mathbf{u}} \end{bmatrix} \in \mathbb{R}^{N_x}$$

where \circ denotes rotation composition, $\text{Exp}: \mathbb{R}^3 \rightarrow \text{SO}(3)$, and $\text{Log}: \text{SO}(3) \rightarrow \mathbb{R}^3$. With \mathbf{x}_k denoting the true system state at k , the filter's a priori estimate is denoted $\bar{\mathbf{x}}_k$ with error covariance, $\bar{\mathbf{P}}_k$. The a posteriori state is $\hat{\mathbf{x}}_k$ with error covariance, $\hat{\mathbf{P}}_k$.

3 | AN UNSCENTED MULTI-BASELINE CDGNSS MEASUREMENT UPDATE

3.1 | CDGNSS Measurement Model

At each GNSS measurement epoch, the estimator ingests N_k pairs of double-difference (DD) GNSS observables, each pair composed of a pseudo-range and a carrier-phase measurement, across all baselines. The baselines and

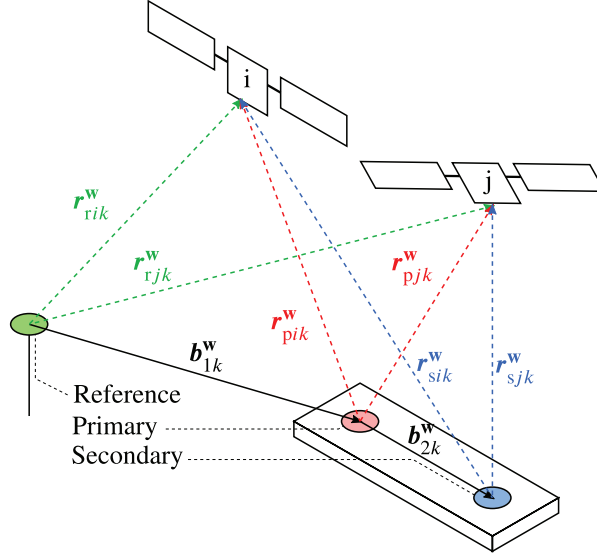


FIGURE 2 Baseline and antenna-to-satellite vectors of the multi-antenna CDGNSS measurement model for pivot satellite i and non-pivot satellite j ; \mathbf{r}_{rik}^w , \mathbf{r}_{rjk}^w , and \mathbf{r}_{snk}^w refer to vectors pointing from the reference, primary, and secondary GNSS antennas, respectively, to the antenna phase center of GNSS satellite n at epoch k .

relevant relative position vectors are shown in Figure 2. The measurement vector at epoch k is

$$\mathbf{z}_{gk} \triangleq [\rho_{1k}^T, \phi_{1k}^T, \rho_{2k}^T, \phi_{2k}^T]^T \in \mathbb{R}^{2N_k}$$

where ρ_{mk} and ϕ_{mk} are vectors of double-difference pseudorange and carrier-phase measurements, both in meters, for baseline $m \in \{1, 2\}$ at epoch k . The measurement \mathbf{z}_{gk} is a function of the state, \mathbf{x}_k , the integer ambiguity vector, $\mathbf{n}_k \in \mathbb{Z}^{N_k}$, and zero-mean white Gaussian measurement noise, ϵ_{gk} (Psiaki & Mohiuddin, 2007):

$$\mathbf{z}_{gk} = \mathbf{h}_{gk}(\mathbf{b}(\mathbf{x}_k), \mathbf{n}_k) + \epsilon_{gk}, \epsilon_{gk} \sim \mathcal{N}(0, \Sigma_{gk}) \quad (1)$$

Σ_{gk} contains off-diagonal elements due to the effect of the shared pivot satellite. The function $\mathbf{b}(\mathbf{x}_k)$ relates the baseline vectors, \mathbf{b}_{1k}^w and \mathbf{b}_{2k}^w , shown in Figure 2 to the position and attitude components of the state:

$$\mathbf{b}(\mathbf{x}_k) \triangleq \begin{bmatrix} \mathbf{b}_{1k}^w \\ \mathbf{b}_{2k}^w \end{bmatrix} = \begin{bmatrix} \mathbf{r}_k^w + \mathbf{R}_k^{wb}(\mathbf{r}_p^b - \mathbf{r}_u^b) \\ \mathbf{R}_k^{wb}(\mathbf{r}_s^b - \mathbf{r}_p^b) \end{bmatrix} \quad (2)$$

Here, \mathbf{r}_u^b , \mathbf{r}_p^b , and \mathbf{r}_s^b are the body-frame positions of the IMU, primary GNSS antenna, and secondary GNSS antenna, respectively. Under this formulation, the known length of \mathbf{b}_{2k}^w serves as an implicit constraint on integer ambiguity resolution due to the parameterization of \mathbf{b}_{2k}^w as a function solely of attitude.

Importantly, the off-diagonal blocks of the measurement noise covariance matrix, Σ_{gk} , are nonzero because baselines one and two share a GNSS antenna. By consuming the GNSS measurements for all baselines in a single update, this correlation is exploited and typically yields a higher integer fix success rate than

either baseline taken individually (Medina et al., 2020). Additionally, the vehicle attitude $\mathbf{R}_k^{\mathbf{wb}}$ is often known a priori to sub-degree precision, providing a tight constraint on all three degrees of freedom of $\mathbf{b}_{2k}^{\mathbf{w}}$. This constraint both strengthens the combined integer model (Fan et al., 2019) and increases positioning accuracy, since the off-diagonal blocks of Σ_{gk} encode the sensitivity of GNSS measurements on the secondary vehicle antenna to $\mathbf{r}_k^{\mathbf{w}}$.

3.2 | Linearization

The function $\mathbf{h}_{gk}(\mathbf{b}, \mathbf{n}_k)$ is accurately modeled as linear due to the extreme distance to GNSS satellites relative to the CDGNSS baseline lengths:

$$\mathbf{h}_{gk}(\mathbf{b}(\mathbf{x}_k), \mathbf{n}_k) = \begin{bmatrix} \rho_{1k} \\ \phi_{1k} \\ \rho_{2k} \\ \phi_{2k} \end{bmatrix} = \begin{bmatrix} \mathbf{G}_{1k} \mathbf{b}_1^{\mathbf{w}} \\ \mathbf{G}_{1k} \mathbf{b}_1^{\mathbf{w}} + \Lambda_1 \mathbf{n}_{1k} \\ \mathbf{G}_{2k} \mathbf{b}_2^{\mathbf{w}} \\ \mathbf{G}_{2k} \mathbf{b}_2^{\mathbf{w}} + \Lambda_2 \mathbf{n}_{2k} \end{bmatrix}$$

Here, $\mathbf{n}_{mk} \in \mathbb{Z}^{N_{mk}}$ is the vector of carrier-phase integer ambiguities for baseline m , $\mathbf{n}_k = [\mathbf{n}_{1k}^T, \mathbf{n}_{2k}^T]^T$, and Λ_m is a diagonal matrix composed of the wavelengths in meters of each DD carrier-phase measurement. The geometry matrix, \mathbf{G}_{mk} , is defined as

$$\mathbf{G}_{mk} \triangleq \begin{bmatrix} (\hat{\mathbf{r}}_{ik}^{\mathbf{w}} - \hat{\mathbf{r}}_{1k}^{\mathbf{w}})^T \\ \vdots \\ (\hat{\mathbf{r}}_{ik}^{\mathbf{w}} - \hat{\mathbf{r}}_{N_{mk}k}^{\mathbf{w}})^T \end{bmatrix}$$

for pivot satellite i and non-pivot satellites one to N_{mk} , where $\hat{\mathbf{r}}_{jk}^{\mathbf{w}}$ denotes a unit vector in the \mathbf{w} frame directed from a GNSS antenna to GNSS satellite j of baseline m at epoch k . Under the small-angle approximation, these unit vectors are assumed to be approximately equal for all receiver antennas involved:

$$\hat{\mathbf{r}}_{ik}^{\mathbf{w}} \approx \frac{\mathbf{r}_{rik}^{\mathbf{w}}}{\|\mathbf{r}_{rik}^{\mathbf{w}}\|} \approx \frac{\mathbf{r}_{pik}^{\mathbf{w}}}{\|\mathbf{r}_{pik}^{\mathbf{w}}\|} \approx \frac{\mathbf{r}_{sik}^{\mathbf{w}}}{\|\mathbf{r}_{sik}^{\mathbf{w}}\|}$$

The nonlinearity of Equation (2) due to the manifold structure of vehicle attitude $\mathbf{R}_k^{\mathbf{wb}}$ is nontrivial. Optimal snapshot estimators for the nonlinear GNSS attitude problem, such as C-LAMBDA (Teunissen, 2006) and MC-LAMBDA (Giorgi & Teunissen, 2010), have been studied, but these estimators are computationally expensive and their extension to recursive filtering does not seem straightforward. Instead, extant Kalman-filtering-based multi-baseline CDGNSS estimators (Henkel et al., 2020; Medina et al., 2020) typically linearize the baseline measurement model about the current state estimate with a simple first-order Taylor expansion in order to perform integer ambiguity resolution with a standard ILS solver such as the well-known LAMBDA method (Teunissen, 1995). This is essentially an extension of the LC-LAMBDA method described by Teunissen and Giorgi (2009)

to the multi-baseline recursive estimation case. As demonstrated in Teunissen & Giorgi (2009), this method performs poorly for ultra-short (length $\lesssim 1\text{m}$) baselines, for which the a priori attitude estimate and the pseudorange measurements cannot, together, offer a sufficiently accurate estimate about which to linearize.

3.2.1 | Alternatives

Recursive Bayesian estimation of \mathbf{x}_k requires finding the distribution of $\mathbf{b}(\mathbf{x}_k)$ given a Gaussian prior for \mathbf{x}_k with mean $\bar{\mathbf{x}}_k$ and covariance $\bar{\mathbf{P}}_k$. While the true distribution of $\mathbf{b}(\mathbf{x}_k)$ is nontrivial, it is desirable to approximate it as Gaussian to enable ambiguity resolution with standard ILS techniques, which are computationally efficient and well understood. Dropping the k subscripts for notational clarity, such an approximation produces a joint Gaussian distribution over \mathbf{x} and $\mathbf{b}(\mathbf{x})$:

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{b}(\mathbf{x}) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \bar{\mathbf{x}} \\ \bar{\mathbf{b}} \end{bmatrix}, \begin{bmatrix} \bar{\mathbf{P}} & \mathbf{P}_{xb} \\ \mathbf{P}_{xb}^T & \mathbf{P}_{bb} \end{bmatrix}\right)$$

This can be parameterized in terms of mean, $\bar{\mathbf{b}}$, Jacobian, \mathbf{H}_b , and additional baseline uncertainty, Σ_b^* , that accounts for errors due to linearization:

$$\begin{aligned} \mathbf{P}_{bb} &\triangleq \mathbf{H}_b \bar{\mathbf{P}} \mathbf{H}_b^T + \Sigma_b^* \\ \mathbf{P}_{xb} &\triangleq \bar{\mathbf{P}} \mathbf{H}_b^T \end{aligned} \quad (3)$$

Extant multi-baseline CDGNSS Kalman filters obtain these parameters with a first-order Taylor expansion scheme, which is shown in Algorithm 1.

An alternative approach is to approximate the distribution of $\mathbf{b}(\mathbf{x})$ using a deterministic sampling technique such as the unscented transform (UT), which is used in the UKF (Julier & Uhlmann, 2004). The UT infers the probability distribution of a transformed Gaussian by transforming a set of weighted sigma points through the nonlinearity and evaluating the statistics of the transformed points. The UT implementation used by this paper's CDGNSS measurement update is shown in Algorithm 2.

A simplified two-dimensional example of the two linearization schemes for a single constrained baseline is shown in Figure 3. The EKF distribution is infinitely thin, as EKF linearizations only consider baseline vectors on a plane tangent to the true spherical distribution at the a priori estimate, causing ILS failures when the a priori uncertainty is large enough that the sphere significantly diverges from the tangent plane. The UT, by contrast, yields an approximate Gaussian distribution over the baseline vector that more closely matches the true mean and covariance of $\mathbf{b}(\mathbf{x})$. For this reason, ILS-based ambiguity resolution has a higher success rate when linearization is based on `linearizeUkf` rather than `linearizeEkf` under large a priori attitude uncertainty, as will be demonstrated in Section 3.4.

3.2.2 | UKF Measurement Update

A vector \mathbf{z}_{gk} of DD GNSS observables is ingested by the estimator at epoch k . Linearizing about the a priori state estimate, $\bar{\mathbf{x}}_k$:

$$\begin{bmatrix} \bar{\mathbf{b}}_k, \mathbf{H}_{bk}, \Sigma_{bk}^* \end{bmatrix} = \text{linearizeUkf}(\bar{\mathbf{x}}_k, \bar{\mathbf{P}}_k)$$

ALGORITHM 1

linearizeEKF

Input: $\bar{\mathbf{x}}, \bar{\mathbf{P}}$ **Output:** $\bar{\mathbf{b}}, H_b, \Sigma_b^*$

- 1 $\bar{\mathbf{b}} = \begin{bmatrix} \mathbf{b}_1^w(\bar{\mathbf{x}}) \\ \mathbf{b}_2^w(\bar{\mathbf{x}}) \end{bmatrix}$
- 2 $H_b = \left[\begin{array}{c} \frac{\partial \mathbf{b}_1^w(\mathbf{x})}{\partial \mathbf{x}} \\ \frac{\partial \mathbf{b}_2^w(\mathbf{x})}{\partial \mathbf{x}} \end{array} \right]_{\bar{\mathbf{x}}}$
- 3 $\Sigma_b^* = \mathbf{0}$

ALGORITHM 2

linearizeUKF

Input: $\bar{\mathbf{x}}, \bar{\mathbf{P}}$ **Output:** $\bar{\mathbf{b}}, H_b, \Sigma_b^*$

- 1 $\alpha = 0.001, \kappa = 0$
- 2 $\lambda = \alpha^2 (N_x + \kappa) - N_x$
- 3 $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_{N_x}] = \text{chol}(\bar{\mathbf{P}})^T$
- 4 $\bar{\mathbf{x}}^{(0)} = \bar{\mathbf{x}}$
- 5 $\bar{\mathbf{b}}^{(0)} = \begin{bmatrix} \mathbf{b}_1^w(\bar{\mathbf{x}}^{(0)}) \\ \mathbf{b}_2^w(\bar{\mathbf{x}}^{(0)}) \end{bmatrix}$
- 6 $w_m^{(0)} = \frac{\lambda}{N_x + \lambda}$
- 7 $w_c^{(0)} = \frac{\lambda}{N_x + \lambda} + 1 - \alpha^2 + \beta$
- 8 **for** $i \in [1, 2N_x]$ **do**
- 9 $\bar{\mathbf{x}}^{(i)} = \begin{cases} \bar{\mathbf{x}}^{(0)} \oplus \sqrt{N_x + \lambda} \mathbf{s}_i & i \in [1, N_x] \\ \bar{\mathbf{x}}^{(0)} \oplus -\sqrt{N_x - \lambda} \mathbf{s}_i & i \in [N_x + 1, 2N_x] \end{cases}$
- 10 $\bar{\mathbf{b}}^{(i)} = \begin{bmatrix} \mathbf{b}_1^w(\bar{\mathbf{x}}^{(i)}) \\ \mathbf{b}_2^w(\bar{\mathbf{x}}^{(i)}) \end{bmatrix}$
- 11 $w_m^{(i)} = w_c^{(i)} = \frac{1}{2(N_x + \lambda)}$
- 12 **end**
- 13 $\bar{\mathbf{b}} = \sum_{i=0}^{2N_x} w_m^{(i)} \bar{\mathbf{b}}^{(i)}$
- 14 $\begin{bmatrix} \bar{\mathbf{P}}_{xx} & \bar{\mathbf{P}}_{xb} \\ \bar{\mathbf{P}}_{xb}^T & \bar{\mathbf{P}}_{bb} \end{bmatrix} = \sum_{i=0}^{2N_x} w_c^{(i)} \begin{bmatrix} \bar{\mathbf{x}}^{(i)} \ominus \bar{\mathbf{x}}^{(0)} \\ \bar{\mathbf{b}}^{(i)} - \bar{\mathbf{b}} \end{bmatrix} \begin{bmatrix} \bar{\mathbf{x}}^{(i)} \ominus \bar{\mathbf{x}}^{(0)} \\ \bar{\mathbf{b}}^{(i)} - \bar{\mathbf{b}} \end{bmatrix}^T$
- 15 $H_b = (\bar{\mathbf{P}}_{xx}^{-1} \bar{\mathbf{P}}_{xb})^T$
- 16 $\Sigma_b^* = \bar{\mathbf{P}}_{bb} - (H_b \bar{\mathbf{P}}_{xx} H_b^T)$

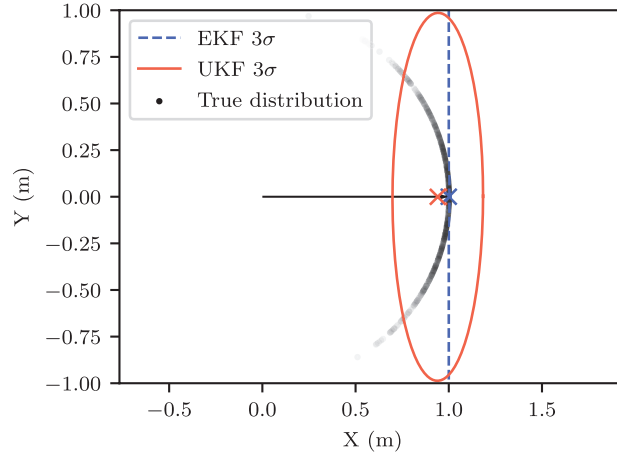


FIGURE 3 Simplified two-dimensional example of Gaussian distributions produced by `linearizeEKF` and `linearizeUKF` for a single constrained baseline of a length of 1 m with a Gaussian prior over the attitude angle; the lines represent the cost contours in the position domain associated with each Gaussian distribution at the 3σ likelihood; the \times symbols represent the means.

yields the following approximation of the measurement model for innovations vector $\mathbf{v}_{gk} \triangleq \mathbf{z}_{gk} - \mathbf{h}_{gk}(\bar{\mathbf{b}}_k, \mathbf{0})$:

$$\begin{aligned} \mathbf{v}_{gk} &= \underbrace{\begin{bmatrix} \mathbf{G}_{1k} & \mathbf{0} \\ \mathbf{G}_{1k} & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_{2k} \\ \mathbf{0} & \mathbf{G}_{2k} \end{bmatrix}}_{\mathbf{H}_{rk}} \mathbf{H}_{bk} \delta \mathbf{x}_k + \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \Lambda_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Lambda_2 \end{bmatrix}}_{\mathbf{H}_{nk}} \underbrace{\begin{bmatrix} \mathbf{n}_{1k} \\ \mathbf{n}_{2k} \end{bmatrix}}_{\mathbf{n}_k} + \epsilon_{gk} + \epsilon_{gk}^* \\ &= \mathbf{H}_{rk} \delta \mathbf{x}_k + \mathbf{H}_{nk} \mathbf{n}_k + \epsilon_{gk} + \epsilon_{gk}^* \\ \epsilon_{gk}^* &\sim \mathcal{N}(\mathbf{0}, \mathbf{H}_{bk} \Sigma_{bk}^* \mathbf{H}_{bk}^T) \end{aligned}$$

Here, the state estimate error vector, $\delta \mathbf{x}_k \triangleq \bar{\mathbf{x}}_k \ominus \mathbf{x}_k$, is expressed in the tangent space of \mathcal{X} at $\bar{\mathbf{x}}_k$, and ϵ_{gk}^* represents additional measurement error caused by approximation of $\mathbf{b}(\mathbf{x}_k)$.

3.3 | Square-Root Formulation

The CDGNSS measurement update can be cast in square-root form for greater numerical robustness and algorithmic clarity (Psiaki & Mohiuddin, 2005). Given \mathbf{v}_{gk} , \mathbf{H}_{rk} , \mathbf{H}_{nk} , $\bar{\mathbf{x}}_k$, and $\bar{\mathbf{P}}_k$, the measurement update can be defined as finding $\delta \mathbf{x}_k$ and \mathbf{n}_k to minimize the cost function:

$$J_k(\delta \mathbf{x}_k, \mathbf{n}_k) = \left\| \mathbf{v}_{gk} - \mathbf{H}_{rk} \delta \mathbf{x}_k - \mathbf{H}_{nk} \mathbf{n}_k \right\|_{\Sigma_k^{-1}}^2 + \left\| \delta \mathbf{x}_k \right\|_{\bar{\mathbf{P}}_k^{-1}}^2$$

where $\Sigma_k = \Sigma_{gk} + \mathbf{H}_b \Sigma_{bk}^* \mathbf{H}_b^T$. The vector cost components can be normalized by left-multiplying with square-root information matrices based on Cholesky factorization $\mathbf{R}_{gk} = \text{chol}(\Sigma_k^{-1})$, $\bar{\mathbf{R}}_{xxk} = \text{chol}(\bar{\mathbf{P}}_k^{-1})$:

$$J_k(\delta \mathbf{x}_k, \mathbf{n}_k) = \left\| \begin{bmatrix} 0 \\ \mathbf{R}_{gk} \mathbf{v}_{gk} \end{bmatrix} - \begin{bmatrix} \bar{\mathbf{R}}_{xxk} \\ \mathbf{R}_{gk} \mathbf{H}_{rk} \end{bmatrix} \delta \mathbf{x}_k - \begin{bmatrix} 0 \\ \mathbf{R}_{gk} \mathbf{H}_{rk} \end{bmatrix} \mathbf{n}_k \right\|^2 = \left\| \mathbf{v}'_k - \begin{bmatrix} \mathbf{H}'_{rk} \\ \mathbf{H}'_{nk} \end{bmatrix} \begin{bmatrix} \delta \mathbf{x}_k \\ \mathbf{n}_k \end{bmatrix} \right\|^2$$

The cost, J_k , can be decomposed via QR factorization:

$$[\mathbf{Q}_k, \mathbf{R}_k] = \text{qr} \left(\begin{bmatrix} \mathbf{H}'_{rk} \\ \mathbf{H}'_{nk} \end{bmatrix} \right)$$

where matrix \mathbf{Q}_k is orthogonal and \mathbf{R}_k is upper triangular. Because \mathbf{Q}_k is orthogonal, the components of J_k inside the norm can be left-multiplied by \mathbf{Q}_k^T without changing the cost, and J_k can be decomposed into three terms:

$$\begin{aligned} J_k(\delta \mathbf{x}_k, \mathbf{n}_k) &= \left\| \mathbf{Q}_k^T \mathbf{v}'_k - \mathbf{R}_k \begin{bmatrix} \delta \mathbf{x}_k \\ \mathbf{n}_k \end{bmatrix} \right\|^2 = \left\| \begin{bmatrix} \mathbf{v}''_{1k} \\ \mathbf{v}''_{2k} \\ \mathbf{v}''_{3k} \end{bmatrix} - \begin{bmatrix} \mathbf{R}_{xxk} & \mathbf{R}_{xnk} \\ \mathbf{0} & \mathbf{R}_{nnk} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \delta \mathbf{x}_k \\ \mathbf{n}_k \end{bmatrix} \right\|^2 \\ &= \underbrace{\left\| \mathbf{v}''_{1k} - \mathbf{R}_{xxk} \delta \mathbf{x}_k - \mathbf{R}_{xnk} \mathbf{n}_k \right\|^2}_{J_{1k}(\delta \mathbf{x}_k, \mathbf{n}_k)} + \underbrace{\left\| \mathbf{v}''_{2k} - \mathbf{R}_{nnk} \mathbf{n}_k \right\|^2}_{J_{2k}(\mathbf{n}_k)} + \underbrace{\left\| \mathbf{v}''_{3k} \right\|^2}_{J_{3k}} \end{aligned} \quad (4)$$

If both the measurement model and $\bar{\mathbf{R}}_{xxk}$ are not ill-conditioned, then \mathbf{R}_{xxk} and \mathbf{R}_{nnk} are invertible. J_{3k} is the irreducible cost, and, under a single-epoch ambiguity resolution scheme, can be shown to be equal to the normalized innovations squared (NIS) associated with the double-difference pseudorange measurements. J_{2k} is the extra cost incurred by enforcing the integer constraint on \mathbf{n}_k . If \mathbf{n}_k is allowed to take any real value (the float solution), J_{2k} can be zeroed due to the invertibility of \mathbf{R}_{nnk} . Similarly, J_{1k} can be zeroed for any value of \mathbf{n}_k due to the invertibility of \mathbf{R}_{xxk} . The float solution, $\{\delta \tilde{\mathbf{x}}_k, \tilde{\mathbf{n}}_k\}$, can therefore be formed by choosing $\delta \tilde{\mathbf{x}}_k$ and $\tilde{\mathbf{n}}_k$ to zero J_{1k} and J_{2k} . Because \mathbf{R}_k is upper triangular, these values can be found by efficient backsubstitution. The fixed solution, $\{\delta \check{\mathbf{x}}_k, \check{\mathbf{n}}_k\}$, is found via an ILS solver, yielding

$$\begin{aligned} \check{\mathbf{n}}_k &= \arg \min_{\mathbf{n}_k \in \mathbb{Z}^n} J_{2k}(\mathbf{n}_k) \\ \delta \check{\mathbf{x}}_k &= \mathbf{R}_{xxk}^{-1} (\mathbf{v}''_{1k} - \mathbf{R}_{xnk} \check{\mathbf{n}}_k) \end{aligned} \quad (5)$$

\mathbf{R}_{xxk} is the a posteriori state vector square-root information matrix conditioned on $\mathbf{n}_k = \check{\mathbf{n}}_k$. Therefore, if the fixed solution is accepted (having passed validation via an integer aperture test), the a posteriori state and covariance are

$$\begin{aligned} \hat{\mathbf{x}}_k &= \bar{\mathbf{x}}_k \oplus \delta \check{\mathbf{x}}_k \\ \hat{\mathbf{P}}_k &= (\mathbf{R}_{xxk}^T \mathbf{R}_{xxk})^{-1} \end{aligned} \quad (6)$$

If, instead, the float solution is accepted, the a posteriori state and covariance are found by marginalizing over the distribution of \mathbf{n}_k :

$$\begin{aligned} \hat{\mathbf{x}}_k &= \bar{\mathbf{x}}_k \oplus \delta \tilde{\mathbf{x}}_k \\ \hat{\mathbf{P}}_k &= (\mathbf{R}_k^T \mathbf{R}_k)^{-1}_{[1:N_x, 1:N_x]} \end{aligned} \quad (7)$$

Here, $[1 : n, 1 : n]$ denotes taking the first n rows and columns of the matrix.

3.4 | Evaluation of Unscented Multi-Antenna Update

The UKF linearization of $\mathbf{b}(\mathbf{x}_k)$ yields a Gaussian prior in the position domain that better captures the true mean and covariance of the constrained attitude baseline than the EKF linearization. Consequently, the UKF linearization achieves a higher ILS success rate than the EKF linearization when attitude uncertainty is large.

Figure 4 demonstrates this effect by evaluating integer aperture success, failure, and float rates for a multi-baseline CDGNSS measurement update via Monte-Carlo simulation with 10^6 samples. The vehicle pitch and roll angles were assumed to be known to 2° (1σ), and the a priori yaw angle uncertainty was varied from 0° to 90° (1σ). The simulation assumed that GPS L1 C/A signals were visible from the equator with a representative satellite constellation and a 15° elevation mask angle. Two vehicle GNSS antennas were simulated with a baseline length of 1.0668 m (equivalent to that of the Sensorium). Other GNSS measurement model parameters were selected as in Table 2. The threshold function approximation to the fixed-failure rate distance test (Wang & Verhagen, 2015) was used with fixed-failure rate $\bar{P}_f = 0.01$.

For a dual-antenna platform similar to that of the University of Texas Sensorium with loose yaw knowledge, this effect becomes apparent as 1σ yaw uncertainty exceeds approximately 8° , whereupon the ILS success rate with the EKF linearization begins a rapid decline with increasing uncertainty, eventually falling below even that of the unconstrained multi-antenna CDGNSS snapshot estimator. In contrast, the integrity of the integer aperture test is maintained for the UKF case due to the approximate linearization error term, Σ_b^* , provided by the UKF linearization.

The UKF linearization expands the operating regime of the CDGNSS navigation estimator to greater levels of attitude uncertainty than with the EKF linearization. This effect is relevant for low-cost urban CDGNSS; while even a low-cost accelerometer can provide pitch and roll angles with degree-level accuracy, it is desirable to tolerate large yaw uncertainty, as may occur, for example, following a long GNSS




FIGURE 4 Integer aperture (IA) success (validated correct fix), failure (validated false fix), and float (failed IA validation) rates found via Section 3.4's Monte Carlo simulation. UKF and EKF denote the ILS success rates for models derived from `linearizeUkf` and `linearizeEkf`, respectively. Note that $P_u = 1 - P_f - P_s$ (Green & Humphreys, 2018).

outage in a parking garage. Upon emerging from such an outage, a system using an EKF linearization may require re-initialization using a snapshot attitude estimator. This scheme also allows initialization of the estimator with loose attitude knowledge, as may be provided, for example, by a magnetometer whose heading measurement may be uncertain in the presence of nearby buildings and vehicles.

4 | TIGHTLY COUPLED NAVIGATION FILTER

This section presents the remaining development of the full tightly coupled multi-antenna CDGNSS recursive estimator with vehicle dynamics constraints and false integer fix mitigation.

4.1 | Propagation Step

State propagation is based on a model replacement approach in which IMU measurements supplant a vehicle dynamics and kinematics model. This is implemented as a standard UKF time update step, and is described in detail in this paper's supplemental material.

4.2 | Vehicle Dynamics Constraints

This paper adopts the VDC scheme of Narula et al. (2022) with minor modifications to the NHC sideslip model and ZUPT detection mechanism. Details of the VDC scheme are provided in this paper's supplemental material.

4.3 | Outlier Rejection Using Pseudorange Innovations

Let $v_{\rho k} = [v_{\rho k1}, \dots, v_{\rho kn}]^T$ contain the elements of the GNSS innovation vector, v_{gk} , that correspond to pseudorange measurements. Its covariance matrix $v_{\rho k}$ is

$$\bar{P}_{\rho k} = \begin{bmatrix} G_{1k} & \mathbf{0} \\ \mathbf{0} & G_{2k} \end{bmatrix} P_{bbk} \begin{bmatrix} G_{1k} & \mathbf{0} \\ \mathbf{0} & G_{2k} \end{bmatrix}^T + \Sigma_{\rho k} \quad (8)$$

where P_{bbk} is formed as in Equation (3), and $\Sigma_{\rho k}$ is formed by selecting the rows and columns of Σ_{gk} that correspond to pseudorange measurements. If the estimator is consistent and the pseudorange measurements are not corrupted by large multipath errors, then $v_{\rho kn} \sim \mathcal{N}(0, (\bar{P}_{\rho k})_{nn})$. A test can be used to detect outlier pseudorange measurements using the detection statistic:

$$q_{kn} = v_{\rho k,n}^2 / (\bar{P}_{\rho k})_{nn}$$

with a threshold, γ^2 , selected to yield a sufficiently low false-positive rate. If a pseudorange outlier is detected for some index n (i.e., $q_{kn} > \gamma^2$), then the corresponding DD pseudorange measurement is assumed to be corrupted by multipath. Because the effect of multipath on GNSS signals is primarily a function of the line-of-sight vector to the transmitter, all pseudorange and carrier-phase measurements

associated with the offending non-pivot satellite on all frequencies and baselines are assumed to be corrupt. They are removed from the measurement vector, \mathbf{z}_{gk} , before proceeding with the measurement update.

4.4 | False-Fix Detection and Recovery

4.4.1 | Single-Epoch Ambiguity Resolution

An optimal CDGNSS filter (in the maximum a posteriori sense) must append a new carrier-phase integer ambiguity to its state each time a cycle slip is detected in a carrier tracking loop (Psiaki, 2010). This causes the update and ILS solve operations to quickly become computationally intractable when applied to an urban environment where cycle slips are common and detection of discrete cycle slips is often impossible.

A scheme could be imagined whereby cycle slips are modeled to occur with some probability, P_{CS} , at each epoch, and a suboptimal dynamic multiple-model estimator such as the interacting multiple-model or generalized pseudo-Bayesian estimator (Bar-Shalom et al., 2001) is used to handle the resulting multiple hypotheses over past cycle slips. However, for large P_{CS} , as is the case for urban CDGNSS, these methods would require measurement update and ILS solve operations for a quickly growing number of alternate hypotheses. Such a method would be computationally prohibitive without aggressive hypothesis pruning.

This paper's estimator adopts a posture of maximum pessimism regarding cycle slips: Each carrier tracking loop in the GNSS receiver is assumed to slip cycles between each pair of measurement epochs (i.e., $P_{CS} = 1$). Ambiguity state growth is curtailed by discarding all ambiguity states at every GNSS measurement epoch, either by conditioning the state vector on the candidate fixed solution (if the candidate fix is validated), as in Equations (5) and (6), or by accepting the float solution and marginalizing over the ambiguities, as in Equation (7). This single-epoch ambiguity resolution scheme is suboptimal because it discards what continuity may be present in the integer ambiguities from epoch to epoch, weakening the integer model strength. But it renders the navigation filter entirely insensitive to cycle slips, which may be a practical necessity for urban CDGNSS, and is computationally efficient. Moreover, the conditioning operation, when applied, greatly increases the integer model strength for subsequent epochs, allowing the filter to hold on to fixes by virtue of the tight epoch-to-epoch position domain constraint provided by the inertial sensor and the vehicle dynamics pseudo-measurements.

4.4.2 | False Integer Fixes

Even in the absence of modeling errors and measurement outliers, the integer fixing procedure of Section 3.3 occasionally yields an incorrect integer vector, $\hat{\mathbf{n}}_k$, that passes ambiguity validation tests, a situation referred to as a *false integer fix* or an *ambiguity resolution failure* (Green & Humphreys, 2018). This occurs because validation tests can only provide probabilistic guarantees of integer correctness (Teunissen & Verhagen, 2009). In an urban environment, measurement outliers due to multipath and diffraction cause the true false integer fix rate, P_f , to greatly exceed the specified rate, \bar{P}_f , for the validation test (Humphreys et al., 2020b). Thus, conditioning the filter state on fixed and validated integers at each epoch, as is done under a single-epoch ambiguity resolution scheme, is perilous because false

integer fixes eventually corrupt the filter state with incorrect but highly confident priors. This causes future GNSS measurement updates to accept a similarly incorrect fix with high probability, repeatedly conditioning the filter state on incorrect ambiguities. While the simple rectilinear motion model of the unaided estimator in Humphreys et al. (2020b) contains sufficient process noise that the filter eventually re-fixes to the correct ambiguities, the tight epoch-to-epoch constraints of the present paper's tightly coupled estimator can cause these cycles of false fixes to persist indefinitely.

To mitigate the effect of conditioning on incorrect integer ambiguities, this paper's estimator employs a fault detection and exclusion technique based on solution separation. A float-only filter, configured to never attempt fixing integer ambiguities, is operated in parallel to the primary navigation filter. Under the single-epoch ambiguity resolution scheme, the float-only filter's behavior is equivalent to accepting only pseudorange measurements, discarding carrier-phase measurements entirely.

4.4.3 | *Carrier-Phase Residuals Testing*

Recall from Equation (4) that $J_{2k}(\mathbf{n}_k)$ is the residual cost associated with the carrier phase in the square-root formulation of the CDGNSS measurement update. When evaluated at $\check{\mathbf{n}}_k$, the fixed solution for the integer vector \mathbf{n}_k found via an ILS solver, this cost becomes the fixed-ambiguity residual cost:

$$\epsilon_{\phi k} = J_{2k}(\check{\mathbf{n}}_k)$$

This quantity is small whenever the carrier-phase measurements are consistent with the prior state estimate, the pseudorange measurements, and with the assumption of integer-valued carrier-phase ambiguities. It is one of several acceptance test statistics used to decide whether the fixed solution $\check{\mathbf{n}}_k$ is correct with high probability (Teunissen & Montenbruck, 2017).

In Mohiuddin and Psiaki (2007), $\epsilon_{\phi k}$ was incorporated in a statistic used to detect carrier cycle slips. It can similarly be used to detect false integer fixes (as a form of integer aperture test statistic) or the lingering effects of performing the conditioning operation in Equation (6) on an incorrectly fixed integer vector at some epoch in the past. A windowed sum of $\epsilon_{\phi k}$ offers even greater sensitivity to false-fix events at the expense of a longer time to detect. The test statistic used to detect false fixes in this paper's estimator is the windowed fixed-ambiguity residual cost (WFARC), Ψ_k . This is calculated over a moving window of fixed length l of past GNSS measurement epochs. It has N_{Ψ_k} degrees of freedom and is calculated by

$$\Psi_k \triangleq \sum_{n=k-l+1}^k \epsilon_{\phi n}, \quad N_{\Psi_k} \triangleq \sum_{n=k-l+1}^k N_n$$

where N_k is the number of DD carrier-phase measurements at epoch k . If the filter is consistent and the integer ambiguities are correctly resolved, then Ψ_k should be approximately χ^2 -distributed with N_{Ψ_k} degrees of freedom. (This distribution is approximate due to the integer-folding effect; large phase residuals are not possible because of integer-cycle phase wrapping.) A statistical consistency test can be performed by choosing a desired false-alarm rate, $\bar{P}_{f,\Psi}$, and declaring a false fix if




FIGURE 5 A false-fix detection and recovery event; ! denotes a test failure ($\Psi_k > \gamma_{\Psi_k}$) at $k = 4$. The primary navigation filter's state estimate and covariance matrix are replaced with those of the float-only filter.




FIGURE 6 A re-seed event: ✓ denotes re-seed criteria being met at epoch four. The estimator replaces the state estimate and covariance of the float-only filter with a copy of that of the primary navigation filter.

TABLE 1
Re-Seed Criteria Used in the Evaluation of Section 5

Normalized residual	$\epsilon_{\phi k} / N_k$	≤ 1.0
Normalized WFARC	Ψ_k / N_{Ψ_k}	≤ 0.5
Last fix number of DD measurements	N_k	≥ 10
Time since last soft reset	$t_{sr,k}$	$\geq 2.0 \text{ s}$

Note: A re-seed operation is performed if all four conditions are met.

$\Psi_k > \gamma_{\Psi_k}$, where the threshold, γ_{Ψ_k} , is calculated by evaluating the inverse cumulative distribution function (CDF) of $\chi^2(N_{\Psi_k})$ at $\bar{P}_{f,\Psi}$.

4.4.4 | False-Fix Recovery

If a false fix is detected ($\Psi_k > \gamma_{\Psi_k}$), the estimator performs a soft reset, discarding the primary navigation filter's state estimate and covariance and replacing them with a copy of the float-only filter's state and covariance, as shown in Figure 5.

4.4.5 | Float-Only Estimator Re-Seeding

To increase the probability of a correct fix after a soft reset, the float-only filter is occasionally re-seeded with the state of the primary filter during epochs over which a set of heuristic criteria indicates that a correct fix is extremely likely. This operation carries the risk that the float-only filter could also be contaminated with information from a false integer fix in the primary filter. But in practice, heuristic criteria can be set to strictly limit this event's probability. The four criteria used in the next section's evaluation are given in Table 1.

One might argue that these criteria for re-seeding the float-only filter are redundant because an integer aperture test for validating the primary filter's integer estimate can be made arbitrarily strict, obviating additional validation. But

integer aperture theory is founded on modeling measurement error distributions as Gaussian (Teunissen, 2003), which is a poor approximation in urban environments leading to low fixed solution availability (Humphreys et al., 2020b). Teunissen's recent extension of the so-called best integer equivariant estimation to the class of elliptically contoured distributions in Teunissen (2020) may offer a means of providing a better re-seed estimate for the float-only filter in urban environments, but it has not been tested with empirical urban data. Meanwhile, application of this paper's re-seeding technique with the criteria in Table 1 will be shown in the next section to significantly increase integer fix availability while respecting a low false-fix rate.

5 | PERFORMANCE EVALUATION IN A DEEP URBAN ENVIRONMENT

5.1 | Experimental Setup

The tightly coupled CDGNSS estimator described in the foregoing sections was implemented in C++ as a new version of the PpEngine sensor fusion engine (Humphreys et al., 2020b), and was experimentally evaluated against the publicly available TEX-CUP urban positioning data set. TEX-CUP is comprised of raw GNSS intermediate-frequency (IF) samples and inertial data, which were collected on May 9, 2019, and May 12, 2019, using the University of Texas Sensorium vehicular perception research platform (Narula et al., 2020). The data set consists of a total of over 2 hours of driving in Austin, Texas, in conditions ranging from light to dense urban; routes are shown in Figure 7.

Two-bit-quantized IF samples were captured at the Sensorium and at the reference station through RadioLynx, a low-cost L1+L2 GNSS front end with a 5-MHz bandwidth at each frequency, and were processed with the PpRx software-defined




FIGURE 7 Overview of the CDGNSS reference station position and routes driven through the urban core of Austin, Texas, in the TEX-CUP urban positioning data sets; routes differ slightly from May 9 to May 12 due to road closures.

GNSS receiver (Humphreys et al., 2020b). The Sensorium RadioLynx was connected to two Antcom G8 GNSS antennas separated by 1.0668 meters in the vehicle Y direction, and the reference RadioLynx was connected to a Trimble Zephyr II geodetic-grade GNSS antenna. The number of double-difference measurements available to PpEngine over the two data sets ranged from one to 25, with an average of 11.68.

The system's performance was separately evaluated using inertial data from each of the Sensorium's two MEMS inertial sensors. The first, a LORD MicroStrain 3DM-GX5-25, is an industrial-grade sensor. The second, a Bosch BMX055, is a surface-mount consumer-grade sensor. Their relevant datasheet specifications are compared in Yoder (2021).

The system's positioning performance was evaluated by comparing with TEX-CUP's forward-backward smoothed ground-truth reference trajectory. This ground truth was generated by postprocessing the data from an iXblue ATLANS-C mobile mapping system comprised of a Septentrio AsteRx4 RTK receiver and a high-end tactical-grade IMU. The AsteRx4 was configured to use additional constellations (BDS and GLONASS) and frequencies (GPS L5, Galileo E5, GLONASS L1 and L2) that were unavailable to PpEngine. The reported accuracy of the ground truth trajectory varied between 2 and 15 cm (1σ) along the route. Due to the difficulty of directly evaluating integer ambiguity resolution performance in the presence of outlier carrier-phase measurements, the integer fixing performance in the following sections was evaluated by considering integer fixes to be correct if the 3D distance to the ground truth was below 30 cm, following Humphreys et al. (2020b).

One might naturally be concerned that the ground truth trajectory's reported accuracy, which degrades to 15 cm (1σ) at some points along the route, may not be good enough to evaluate the accuracy of the proposed system, whose errors are also at the decimeter-level. However, the errors in the TEX-CUP ground truth data set are likely to be substantially uncorrelated with the errors of the proposed system, for three reasons. First, the truth data set employed an independent receiver with access to additional GNSS constellations (BeiDou and GLONASS) and frequencies (GPS L5, Galileo E5, GLONASS L1 and L2) that were unavailable to PpEngine for the experiments reported here. Second, the high-end tactical-grade IMU used to produce the ground truth has noise characteristics much different from the low-cost IMUs of the proposed system. Third, the estimation process that produced the ground truth was non-causal forward-backward smoothing, whereas the proposed system processes the data only causally. Assuming that the ground truth errors and PpEngine's errors are not strongly correlated, it follows that the results presented in this section will, in fact, be pessimistic compared to an evaluation against the absolute truth.

Because TEX-CUP contains several minutes of no motion in an open-sky environment at the beginning and end of each capture, the estimator was run on a subset of each capture beginning approximately 10 seconds before the first motion and ending 10 seconds after the last motion. In the May 9 data set, the estimator was run and evaluated from GPS time of week (TOW) 411,003 s to 415,029 s, and in the May 12 data set, from TOW 63,770 s to 67,972 s.

5.2 | Baseline Configuration

The performance of the tightly coupled estimator with all proposed features and signals enabled (the baseline configuration) was evaluated using each IMU on the May 9 and May 12 TEX-CUP data sets. PpRx was configured for dual-frequency operation as described in Humphreys et al. (2020b), tracking the GPS L1 C/A, GPS L2C (combined CL+CM codes), Galileo E1 (combined B+C codes), and L1 SBAS

(WAAS) signals on the reference and both rover antennas. Data bit prediction and wipeoff were performed on the GPS L1 C/A and SBAS signals. Reference and rover GNSS observables were produced at a rate of 5 Hz. PpEngine's baseline configuration parameters are given in Table 2.

TABLE 2
Baseline PpEngine Configuration Parameters

CDGNSS parameters		
Carrier-to-noise ratio threshold	C/N_0	≥ 40 dB-Hz
Phase lock statistic threshold	S_θ	≥ 0.8
Elevation mask	θ_{el}	$\geq 10^\circ$
Integer aperture (IA) validation test	Fixed-failure-rate difference test (Wang & Verhagen, 2015)	
IA fixed-failure rate	\bar{P}_f	$= 0.001$
Undifferenced zenith pseudorange std	σ_ρ	$= 1.5$ m
Undifferenced zenith phase std	σ_ϕ	$= 0.006$ m
Elevation-dependent weighting	w_s	$= (1 + 10 \exp(-\theta_{el} / 10^\circ))^{-1}$
Pseudorange outlier threshold std	γ	$= 1.5\sigma$
False-fix detection window length	l	$= 10$
False-fix detection threshold	$\bar{P}_{f,\Psi}$	$= 10^{-15}$
IMU parameters		
Accelerometer noise density	$\sqrt{S_a}$	$= 100,300 \mu g / \sqrt{\text{Hz}}^*$
Accelerometer bias steady-state std	σ_{b_a}	$= 0.5, 10 \text{ mg}^*$
Accelerometer bias time constant	τ_a	$= 100$ s
Gyroscope noise density	$\sqrt{S_g}$	$= 0.01, 0.05 (^\circ/\text{s}) / \sqrt{\text{Hz}}^*$
Gyroscope bias steady-state std	σ_{b_g}	$= 8, 30 ^\circ/\text{hr}^*$
Gyroscope bias time constant	τ_g	$= 100$ s
NHC parameters		
Lateral std	$\sigma_{nhc,y}$	$= 0.1$ m/s
Vertical std	$\sigma_{nhc,z}$	$= 0.2$ m/s
ZUPT parameters		
Longitudinal std	σ_{zx}	$= 0.05$ m/s
Lateral/vertical std	σ_{zy}, σ_{zz}	$= 0.01$ m/s
Accelerometer noise threshold	$\gamma_{\text{zupt},a}$	$= 0.8$ m/s ²
Gyroscope noise threshold	$\gamma_{\text{zupt},g}$	$= 0.006, 0.018$ rad/s*
ZUPT detection window	N_{zupt}	$= 10, 30$ *
Innovations test threshold	$\bar{P}_{f,\text{zupt}}$	$= 10^{-30}, 10^{-6}$ *

Note: Pairs marked with “*” indicate separate parameter values used with the industrial-grade and consumer-grade IMU, respectively. IMU noise parameters were increased from datasheet values for consistency with empirical observations. The elevation-dependent GNSS measurement weighting scheme is derived from Verhagen et al. (2010). Details of the NHC and ZUPT schemes are provided in this paper's supplemental information.

The phase center variation of the Sensorium antennas with respect to signal elevation angle was calibrated as in Humphreys et al. (2020b). The orientation, body-frame position, axis scale factors, and steady-state biases of both inertial sensors were calibrated offline using a short period of dynamic open-sky GNSS data at the beginning of the May 12 data set.

The position and attitude states of the tightly coupled navigation estimator were initialized with the first available batch of GNSS observables and inertial measurements. The position state was initialized with the standard (single-ended) pseudo-range position solution for the primary antenna. The attitude state was initialized by combining the gravity vector as determined by the IMU's accelerometers with a constrained-baseline snapshot CDGNSS solution for the baseline connecting the primary and secondary vehicle antennas, as determined with a brute-force attitude-domain search.

5.3 | Baseline Performance

5.3.1 | Ambiguity Resolution and Positioning

The achieved integer-fix availability was 97.52% and 96.62% when tightly coupled with the industrial-grade and consumer-grade IMUs, respectively, across both days of the data set. When tightly coupled with the industrial-grade IMU, the 95th-percentile horizontal positioning error was 8.4 cm when fixed and 10.1 cm overall (fixed and float). Using the consumer-grade IMU, the 95th-percentile horizontal error was 9.2 cm when fixed and 12.0 cm overall. The empirical CDF of the 3D positioning errors using both grades of inertial sensor is shown in Figure 8. Detailed statistics of the estimator's positioning performance in the baseline configuration are given in Table 3, and statistics of its attitude performance in Table 4. Figures 9 and 10 show the position and attitude error, respectively, over time for the May 9 portion of TEX-CUP. In Figure 9, the longest two float periods seen with the industrial-grade IMU lasted 47 and 36 seconds, and exhibited position errors of a few meters. Positioning error spikes are visible during these float epochs as, under this paper's single-epoch ambiguity resolution scheme, only pseudorange




FIGURE 8 CDFs of 3D positioning error of baseline estimator configuration across both days of the TEX-CUP urban positioning data set

TABLE 3

Baseline Estimator Ambiguity Resolution and Positioning Performance on Each Day of the TEX-CUP Data Set

		Overall positioning performance (fix & float epochs)							
		Ambiguity Resolution		3D		Horizontal		Vertical	
Data set		P_V (%)	P_f (%)	d_{95} (cm)	RMSE (cm)	d_{95h} (cm)	RMSE (cm)	d_{95v} (cm)	RMSE (cm)
May 9	Unaided	77.31%	0.33%	742.1	1799.6	471.4	1713.3	209.3	550.5
	Consumer-grade IMU	94.56%	0.19%	20.2	34.4	16.4	28.6	7.3	19.2
	Industrial-grade IMU	97.28%	0.21%	13.8	20.6	10.6	18.9	4.0	8.1
May 12	Unaided	76.94%	0.49%	860.5	571.9	525.5	406.1	467.9	402.7
	Consumer-grade IMU	98.63%	0.56%	13.7	7.7	10.6	6.0	5.7	4.8
	Industrial-grade IMU	97.74%	0.43%	12.3	13.8	9.7	12.5	4.9	5.8

Note: Position errors are reported relative to the TEX-CUP ground truth trajectory, which was generated with a survey-grade mobile mapping system as described in Section 5.1, with a self-reported accuracy varying between 2 and 15 cm (1σ) along the route. *Unaided* indicates the use of the motion model of Humphreys et al. (2020b) in lieu of inertial tight coupling, as described in Section 5.4. Quoted 95th percentile and RMS error quantities are over the entire data set (i.e., for both float and fixed epochs). P_V denotes the availability of an aperture-test-validated fixed solution for conditioning in the primary filter. P_f denotes the false-fix rate as determined by an excursion of the primary filter beyond 30 cm from the ground truth when conditioned on fixed ambiguities.

TABLE 4

Baseline Estimator Attitude Performance on Each Day of the TEX-CUP Data Set

Dataset	Roll (°)		Pitch (°)		Yaw (°)	
	p_{95}	RMS	p_{95}	RMS	p_{95}	RMS
May 9						
Consumer-grade IMU	0.99	0.59	0.99	0.60	0.99	0.26
Industrial-grade IMU	0.28	0.21	0.28	0.18	0.28	0.22
May 12						
Consumer-grade IMU	0.84	0.43	0.84	0.41	0.84	0.27
Industrial-grade IMU	0.29	0.21	0.29	0.17	0.29	0.18

Note: The quoted 95th percentile and RMS error quantities are over the entire data set (i.e., for both float and fixed epochs).

measurements that have meter-level accuracy have any effect during float epochs. Some smaller error excursions are visible during fixed epochs due to periods of poor satellite visibility and increased multipath error.

One may note that Table 3 appears to show that the consumer-grade IMU slightly outperformed the industrial-grade IMU by some metrics on the May 12 data set. While the consumer-grade IMU did achieve a slightly higher P_V (and correspondingly lower RMS position error) than the industrial-grade, it should be noted that the P_f value was also slightly higher (i.e., worse), so a direct comparison cannot be drawn between the two P_V values. However, the consumer-grade IMU allowed a few more large positioning error excursions than the industrial-grade, resulting in both P_f and d_{95} larger than those of the industrial-grade IMU.

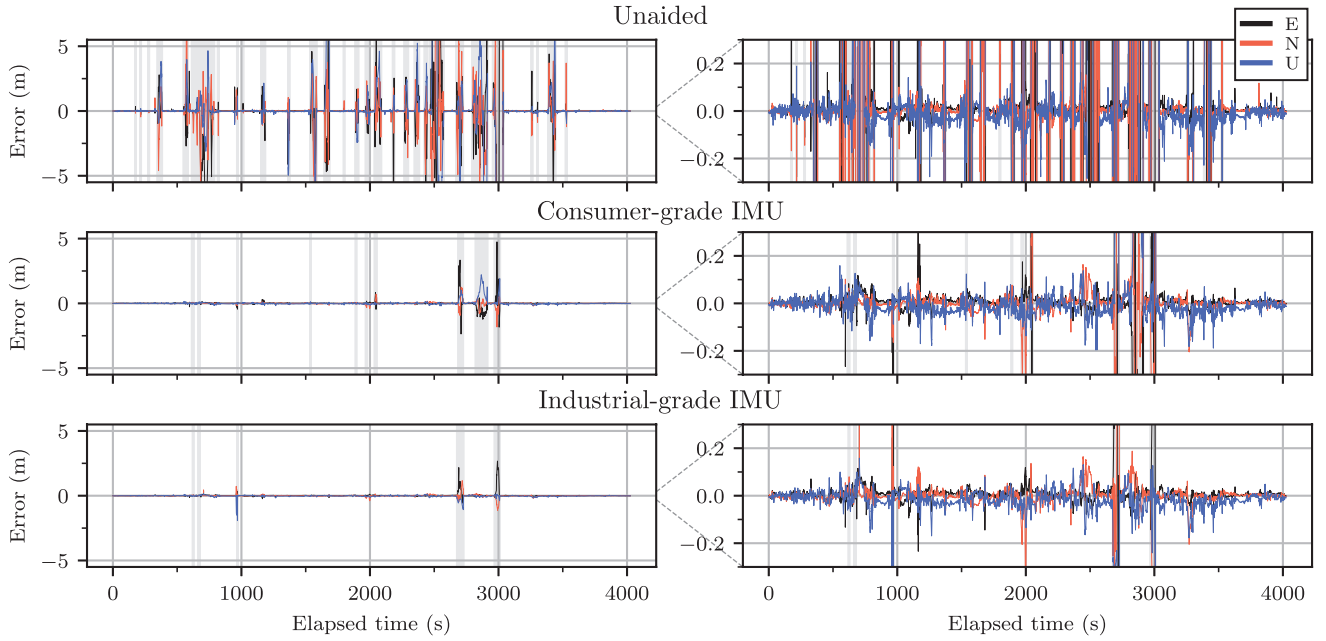


FIGURE 9 Baseline estimator positioning error over time in the east, north, and up directions for the May 9 TEX-CUP data set: Gray shading indicates float epochs (periods when the estimator accepted the float CDGNSS solution). *Elapsed time* indicates time since the data set start at GPS TOW 411,003 s.

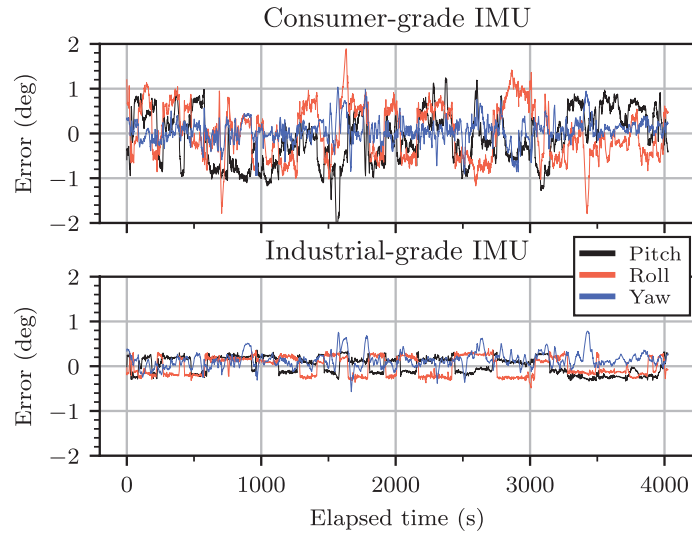


FIGURE 10 Baseline estimator attitude estimation error over time for the May 9 TEX-CUP data set; heading-dependent pitch and roll errors are evident when tightly coupled with the industrial-grade IMU; these are likely due to residual calibration errors.

Of the two data sets, the May 9 data set had more difficult GNSS satellite geometry, yielding a lower P_V for both IMU grades vs. May 12. The tighter motion constraints provided by the industrial-grade IMU were able to overcome the degraded geometry to a greater degree than those provided by the consumer-grade IMU.

P_f , however, was worse for all estimators on May 12 than on May 9, despite generally better GNSS geometry. This appears to be because the May 12 data set happened to contain more trouble spots: portions of the data set with very heavy multipath that yielded frequent incorrect fixes, typically when the vehicle slowed down or came to a complete stop in the deepest urban canyon portions of the route.

5.3.2 | *Attitude*

The attitude performance of the baseline estimator is excellent when tightly coupled with either the industrial-grade or consumer-grade inertial sensor, achieving single-degree-level precision in all three axes with the consumer-grade sensor and sub-degree precision with the inertial sensor. Better performance with the industrial-grade sensor is as expected due to its significantly better gyroscope noise properties.

5.4 | *Effect of Inertial Tight Coupling*

To evaluate the benefit of inertial tight coupling, PpEngine was run in an unaided mode, using the nearly-constant-velocity motion model described in Humphreys et al. (2020b) for propagation in place of an inertial sensor. Attitude dynamics were modeled as a simple integrated white noise process, with noise intensity of $0.3 \text{ }^\circ/\sqrt{\text{s}}$ in the vehicle pitch and roll axes, and $5.7 \text{ }^\circ/\sqrt{\text{s}}$ in the vehicle yaw axis. Because vehicle pitch (rotation about the axis connecting the primary and secondary Sensorium antennas) is not strongly observable, a weak pseudo-measurement of zero pitch was added with standard deviation of 10° at a rate of 5 Hz, which was found to provide good estimation performance.

The positioning and ambiguity resolution performance of the unaided estimator is shown in Table 3. Tight coupling with even a consumer-grade IMU has a clearly beneficial effect on ambiguity resolution, greatly increasing the fraction of fixed-integer epochs for comparable P_f , and reducing 95th-percentile and RMS positioning errors from meters to centimeters.

5.5 | *Performance in Alternate Configurations*

Next, the estimator was run in a collection of alternate configurations (with various features disabled) in order to study its performance's sensitivity to the presence of various algorithmic components. Results are given in Table 5.

5.5.1 | *Performance in Single-Antenna Mode*

In Configuration 2, the estimator was run using only a single vehicle-mounted antenna. CDGNSS observables were formed for only a single baseline, between the reference antenna and the primary Sensorium GNSS antenna. Despite having fewer integer ambiguities (only one baseline's worth) to fix at each measurement epoch, the integer fix availability was lower than for the multi-antenna case in this configuration because the estimator was no longer able to exploit the measurement noise cross-covariance between baselines due to the shared antenna. Positioning performance is slightly worse in the single-antenna case by all metrics, as would be expected due to the loss of half of all GNSS measurements.

5.5.2 | *Performance Without Vehicle Motion Constraints*

The non-holonomic constraints were disabled in Configuration 3, and zero-velocity updates were disabled in Configuration 4. It can be seen that the

TABLE 5

Estimator Ambiguity Resolution and Positioning Performance on the Combined TEX-CUP May 9 and May 12 Data Sets With Various Estimator Features Disabled

Configuration	Industrial-grade IMU				Consumer-grade IMU			
	P_v (%)	P_f (%)	d_{95} (cm)	RMSE (cm)	P_v (%)	P_f (%)	d_{95} (cm)	RMSE (cm)
1 Baseline (all features enabled)	97.52	0.33	13.0	17.5	96.62	0.38	16.2	24.8
2 Single vehicle antenna	96.89	0.55	16.2	25.3	95.73	0.53	19.3	30.6
3 Sans non-holonomic constraints	97.48	0.33	13.2	18.3	94.86	0.60	26.1	116.0
4 Sans zero-velocity updates	97.39	0.34	13.8	38.2	96.39	0.40	16.9	27.1
5 Sans pseudorange outlier exclusion (§4.3)	93.09	0.48	19.1	45.5	87.73	1.25	241.8	133.6
6 Sans re-seed (§4.4.5)	97.00	0.91	15.0	36.8	96.58	0.46	16.5	34.2
7 Sans false-fix detection & recovery (§4.4)	99.00	24.28	590.8	194.7	97.27	10.03	121.6	46.2
8 Single frequency (L1 only)	97.65	0.96	14.7	24.2	90.35	0.38	138.9	87.1
9 Sans SBAS	95.73	0.55	17.1	25.6	88.86	0.86	191.2	80.3
10 EKF CDGNSS Update (linearizeEkf)	97.52	0.33	13.0	17.5	96.62	0.38	16.2	24.8

non-holonomic constraints have only a small effect on performance when tightly coupling with the industrial-grade IMU. With the consumer-grade IMU, however, the effect is much greater. Clearly, non-holonomic constraints provide a major performance boost in the consumer-grade IMU case, cutting the number of float GNSS measurement epochs in half (fix availability increased from 94.86% to 96.62%).

The incorporation of zero-velocity updates improves all of the presented statistics when tightly coupling with either the industrial-grade or consumer-grade IMU. But for the consumer-grade IMU, the benefit of ZUPTs is minimal. This is likely because the poor noise properties of the consumer-grade IMU required such strict thresholds to limit false detections that many opportunities to apply a ZUPT were missed. However, ZUPTs have a clearly positive benefit on overall RMS position error, as they help to constrain against float-solution position error when the vehicle is stopped, which is when urban code multipath errors are largest.

5.5.3 | Performance Without Pseudorange Outlier Exclusion

The pseudorange innovations-based outlier exclusion mechanism (Section 4.3) was disabled in Configuration 5. Without this mechanism, the availability of validated integer-fixed solutions decreased drastically due to the presence of outlier measurements caused by multipath. Interestingly, the false-fix rate, P_f , was

elevated, but not to extreme levels. This was likely because false fixes due to multipath-induced outliers were reverted by the false-fix recovery mechanism.

5.5.4 | Performance Without False-Fix Detection and Recovery

In Configuration 6, the re-seed mechanism described in Section 4.4.5 was disabled and, in Configuration 7, the entire false-fix detection and recovery mechanism (Section 4.4) was disabled. Integer fix performance without the re-seed mechanism is appreciably reduced, as the estimation performance of the float-only estimator suffers without the ability to re-seed from especially trustworthy integer fixes. Disabling the false-fix detection and recovery mechanism has a catastrophic effect on the false-fixing rate, P_f , for the reasons given in Section 4.4.

5.5.5 | Performance on Subsets of GNSS Signals

In Configuration 8, the estimator was run using only L1 GNSS signals (i.e., GPS L2C was disabled), and in Configuration 9, SBAS L1 signals were disabled. Ambiguity resolution and positioning performance on these GNSS signal subsets was fairly close to the baseline case when using the industrial-grade IMU, but with the consumer-grade IMU, a substantial loss of integer-fix availability occurred (down from 96.62% to 90.35% and 88.86% in each of these configurations, respectively). The weaker motion constraints provided by the consumer-grade IMU cause the estimator to require more GNSS signals for acceptable integer fix availability.

5.5.6 | Performance With EKF-Based Linearization

The estimator was run using `linearizeEKF` in place of `linearizeUKF` in Configuration 10. Due to the excellent attitude performance of the estimator with either inertial sensor, no significant difference in performance arises on the TEX-CUP data set by using the UKF update. This can be explained by referring to the results of Figure 4, which show that the benefit of the UKF linearization is significant for the Sensorium's inter-antenna distance only when attitude uncertainty exceeds approximately 8° on a single axis, whereas the attitude error on the TEX-CUP data set never exceeded 2° .

6 | CONCLUSION

A vehicular pose estimation technique has been presented and evaluated that tightly couples multi-antenna CDGNSS, a low-cost MEMS IMU, and vehicle dynamics constraints (non-holonomic constraints and zero-velocity updates). The unscented transform was used to linearize the multi-antenna CDGNSS update, allowing for the use of a linear integer least-squares solver for ambiguity resolution while exploiting between-baseline correlations and respecting the constraints provided by known vehicle antenna geometry, even under large attitude uncertainties. Robust estimation techniques were developed to mitigate the effects of urban multipath and signal blockage, and to recover from false integer fixes. The estimator was evaluated using the publicly available TEX-CUP urban positioning data set,

yielding a 96.6% and 97.5% integer fix availability, and a 12.0-cm and 10.1-cm overall (fix and float) 95th-percentile horizontal positioning error with a consumer-grade and industrial-grade inertial sensor, respectively, over more than two hours of driving in the urban core of Austin, Texas. A performance sensitivity analysis showed that the false-fix detection and recovery scheme is key to achieving an acceptably low false integer fixing rate of 0.3% and 0.4%, respectively. Having a second vehicle-mounted GNSS antenna significantly increased integer-fix availability, decreased false-fix rate, and improved both root-mean-square and 95th-percentile positioning performance as compared to a single-baseline CDGNSS configuration.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Transportation under Grant 69A3552047138 for the CARMEN University Transportation Center, and by the Army Research Office under Cooperative Agreement W911NF-19-2-0333. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. Map data © OpenStreetMap contributors (<https://www.openstreetmap.org/copyright>).

REFERENCES

- Abd Rabbou, M., & El-Rabbany, A. (2015). Tightly coupled integration of GPS precise point positioning and MEMS-based inertial systems. *GPS Solutions*, 19(4), 601–609. <http://doi.org/10.1007/s10291-014-0415-3>
- Bar-Shalom, Y., Li, X. -R., & Kirubarajan, T. (2001). *Estimation with applications to tracking and navigation*. John Wiley and Sons. <http://doi.org/10.1002/0471221279>
- Elmezaen, A., & El-Rabbany, A. (2021). Ultra-low-cost tightly coupled triple-constellation GNSS PPP/MEMS-based INS integration for land vehicular applications. *Geomatics*, 1(2), 258–286. <http://doi.org/10.3390/geomatics1020015>
- Fan, P., Li, W., Cui, X., & Lu, M. (2019). Precise and robust RTK-GNSS positioning in urban environments with dual-antenna configuration. *Sensors*, 19(16), 3586. <http://doi.org/10.3390/s19163586>
- Gao, Z., Zhang, H., Ge, M., Niu, X., Shen, W., Wickert, J., & Schuh, H. (2017). Tightly coupled integration of multi-GNSS PPP and MEMS inertial measurement unit data. *GPS Solutions*, 21(2), 377–391. <https://doi.org/10.1007/s10291-016-0527-z>
- Giorgi, G., & Teunissen, P. J. G. (2010). Carrier phase GNSS attitude determination with the multivariate constrained LAMBDA method. *2010 IEEE Aerospace Conference*, Big Sky, MT. <https://doi.org/10.1109/AERO.2010.5446910>
- Green, G. N., & Humphreys, T. E. (2018). Data-driven generalized integer aperture bootstrapping for high-integrity positioning. *IEEE Transactions on Aerospace and Electronic Systems*, 55(2), 757–768. <http://doi.org/10.1109/TAES.2018.2864770>
- Henkel, P., & Günther, C. (2012). Reliable integer ambiguity resolution: multi-frequency code carrier linear combinations and statistical a priori knowledge of attitude. *NAVIGATION*, 59(1), 61–75. <http://doi.org/10.1002/navi.6>
- Henkel, P., Sperl, A., Mittmann, U., Fritzel, T., Strauss, R., & Steiner, H. (2020). Precise 6D RTK positioning system for UAV-based near-field antenna measurements. *2020 14th European Conference on Antennas and Propagation (EuCAP)*, Copenhagen, Denmark. <http://doi.org/10.23919/EuCAP48036.2020.9135771>
- Hirokawa, R., & Ebinuma, T. (2009). A low-cost tightly coupled GPS/INS for small UAVs augmented with multiple GPS antennas. *NAVIGATION*, 56(1), 35–44. <https://www.ion.org/publications/abstract.cfm?articleID=102486>
- Hong, S., Lee, M. H., Chun, H.-H., Kwon, S.-H., & Speyer, J. L. (2005). Observability of error states in GPS/INS integration. *IEEE Transactions on Vehicular Technology*, 54(2), 731–743. <http://doi.org/10.1109/TVT.2004.841540>
- Humphreys, T. E., Kor, R. X. T., Iannucci, P. A., & Yoder, J. E. (2020a). Open-world virtual reality headset tracking. *Proc. of the 33rd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2020)*, 2931–2947. <http://doi.org/10.33012/2020.17635>
- Humphreys, T. E., Murrian, M. J., & Narula, L. (2020b). Deep-urban unaided precise global navigation satellite system vehicle positioning. *IEEE Intelligent Transportation Systems Magazine*, 12(3), 109–122. <http://doi.org/10.1109/imits.2020.2994121>

- Jackson, J., Davis, B., & Gebre-Egziabher, D. (2018). An assessment of low-cost RTK GNSS receivers. *2018 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Monterey, CA. <http://doi.org/10.1109/PLANS.2018.8373438>
- Julier, S. J., & Uhlmann, J. K. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3), 401–422. <http://doi.org/10.1109/JPROC.2003.823141>
- Kennedy, S., Hamilton, J., & Martell, H. (2006). Architecture and system performance of SPAN—NovAtel's GPS/INS solution. *Proc. of IEEE/ION Position, Location, and Navigation Symposium*, San Diego, CA, 266–274. <https://www.ion.org/publications/abstract.cfm?articleID=6656>
- Li, T., Zhang, H., Gao, Z., Chen, Q., & Niu, X. (2018). High-accuracy positioning in urban environments using single-frequency multi-GNSS RTK/MEMS-IMU integration. *Remote Sensing*, 10(2), 205. <http://doi.org/10.3390/rs10020205>
- Medina, D., Vilà-Valls, J., Hesselbarth, A., Ziebold, R., & García, J. (2020). On the recursive joint position and attitude determination in multi-antenna GNSS platforms. *Remote Sensing*, 12(12), 1955. <http://doi.org/10.3390/rs12121955>
- Mohiuddin, S., & Psiaki, M. L. (2007). High-altitude satellite relative navigation using carrier-phase differential Global Positioning System techniques. *Journal of Guidance, Control, and Dynamics*, 30(5), 1628–1639. <https://doi.org/10.2514/1.27827>
- Murrian, M. J., Gonzalez, C. W., Humphreys, T. E., & Novlan, T. D. (2016). A dense reference network for mass-market centimeter-accurate positioning. *2016 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Savannah, GA. <http://doi.org/10.1109/PLANS.2016.7479708>
- Nagai, K., Spenko, M., Henderson, R., & Pervan, B. (2021). Evaluating INS/GNSS availability for self-driving cars in urban environments. *Proc. of the 2021 International Technical Meeting of the Institute of Navigation*, 243–253. <http://doi.org/10.33012/2021.17830>
- Narula, L., Iannucci, P. A., & Humphreys, T. E. (2022). All-weather sub-50-cm radar-inertial positioning. *Field Robotics*, 2, 525–556. <http://doi.org/10.55417/fr.2022019>
- Narula, L., LaChapelle, D. M., Murrian, M. J., Wooten, J. M., Humphreys, T. E., de Toldi, E., Morvant, G., & Lacambre, J.-B. (2020). TEX-CUP: The University of Texas Challenge for Urban Positioning. *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Portland, OR. <http://doi.org/10.1109/PLANS46316.2020.9109873>
- Odijk, D. (2002). *Fast precise GPS positioning in the presence of ionospheric delays* (No. 52). NCG, Nederlandse Commissie voor Geodesie.
- Ong, R. B., Petovello, M. G., & Lachapelle, G. (2009). Assessment of GPS/GLONASS RTK under various operational conditions. *Proc. of the 22nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS 2009)*, Savannah, GA, 3297–3308. <https://www.ion.org/publications/abstract.cfm?articleID=8743>
- Petovello, M. G., Cannon, M. E., & Lachapelle, G. (2004). Benefits of using a tactical-grade IMU for high-accuracy positioning. *NAVIGATION*, 51(1), 1–12. <http://doi.org/10.1002/j.2161-4296.2004.tb00337.x>
- Psiaki, M. L. (2010). Kalman filtering and smoothing to estimate real-valued states and integer constants. *Journal of Guidance, Control, and Dynamics*, 33(5), 1404–1417. <https://doi.org/10.2514/1.48567>
- Psiaki, M. L., & Mohiuddin, S. (2005). Relative navigation of high-altitude spacecraft using dual-frequency civilian CDGPS. *Proc. of the 18th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS 2005)*, Long Beach, CA, 1191–1207. <https://www.ion.org/publications/abstract.cfm?articleID=6316>
- Psiaki, M. L., & Mohiuddin, S. (2007). Modeling, analysis, and simulation of GPS carrier phase for spacecraft relative navigation. *Journal of Guidance, Control, and Dynamics*, 30(6). <https://doi.org/10.2514/1.29534>
- Psiaki, M. L., O'Hanlon, B. W., Powell, S. P., Bhatti, J. A., Wesson, K. D., Humphreys, T. E., & Schofield, A. (2014). GNSS spoofing detection using two-antenna differential carrier phase. *Proc. of the 27th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2014)*, Tampa, FL, 2776–2800. <https://www.ion.org/publications/abstract.cfm?articleID=12530>
- Reid, T. G., Houts, S. E., Cammarata, R., Mills, G., Agarwal, S., Vora, A., & Pandey, G. (2019). Localization requirements for autonomous vehicles. *SAE International Journal*, 2(3), 173–190. <https://doi.org/10.4271/12-02-03-0012>
- Scherzinger, B. M. (2006). Precise robust positioning with inertially aided RTK. *NAVIGATION*, 53(2), 73–83. <http://doi.org/10.1002/j.2161-4296.2006.tb00374.x>
- Solà, J., Deray, J., & Atchuthan, D. (2018). A micro Lie theory for state estimation in robotics [Technical Report IRI-TR-18-01]. Institut de Robòtica i Informàtica Industrial. <http://www.iri.upc.edu/files/scidoc/2089-A-micro-Lie-theory-for-state-estimation-in-robotics.pdf>
- Teunissen, P. J. G. (1995). The least-squares ambiguity decorrelation adjustment: a method for fast GPS integer ambiguity estimation. *Journal of Geodesy*, 70, 65–82. <https://doi.org/10.1007/BF00863419>
- Teunissen, P. J. G. (2003). Integer aperture GNSS ambiguity resolution. *Artificial Satellites*, 38(3), 79–88. <http://doi.org/10.11937/27352>

- Teunissen, P. J. G. (2006). The LAMBDA method for the GNSS compass. *Artificial Satellites*, 41(3), 89–103. <https://doi.org/10.2478/v10018-007-0009-1>
- Teunissen, P. J. G. (2020). Best integer equivariant estimation for elliptically contoured distributions. *Journal of Geodesy*, 94(9), 82. <http://doi.org/10.1007/s00190-020-01407-2>
- Teunissen, P. J. G., & Giorgi, G. (2009). To what extent can standard GNSS ambiguity resolution methods be used for single-frequency epoch-by-epoch attitude determination? *Proc. of the 22nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS 2009)*, Savannah, GA, 235–242. <https://www.ion.org/publications/abstract.cfm?articleID=8426>
- Teunissen, P. J. G., & Montenbruck, O. (Eds.). (2017). *Springer handbook of global navigation satellite systems*. Springer. <https://doi.org/10.1007/978-3-319-42928-1>
- Teunissen, P. J. G., & Verhagen, S. (2009). The GNSS ambiguity ratio-test revisited: a better way of using it. *Survey Review*, 41(312), 138–151. <https://doi.org/10.1179/003962609X390058>
- Vana, S. (2021). Low-cost, triple-frequency multi-GNSS PPP and MEMS IMU integration for continuous navigation in urban environments. *Proc. of the 34th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2021)*, St. Louis, MO, 3234–3249. <http://doi.org/10.33012/2021.18093>
- Verhagen, S., Odijk, D., Teunissen, P. J. G., & Huisman, L. (2010). Performance improvement with low-cost multi-GNSS receivers. *2010 5th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*, Noordwijk, Netherlands. <https://doi.org/10.1109/NAVITEC.2010.5708015>
- Wang, L., & Verhagen, S. (2015). A new ambiguity acceptance test threshold determination method with controllable failure rate. *Journal of Geodesy*, 89(4), 361–375. <http://doi.org/10.1007/s00190-014-0780-2>
- Wu, S., Zhao, X., Pang, C., Zhang, L., Xu, Z., & Zou, K. (2020). Improving ambiguity resolution success rate in the joint solution of GNSS-based attitude determination and relative positioning with multivariate constraints. *GPS Solutions*, 24(1), 31. <http://doi.org/10.1007/s10291-019-0943-y>
- Yang, Z., Li, Z., Liu, Z., Wang, C., Sun, Y., & Shao, K. (2021). Improved robust and adaptive filter based on non-holonomic constraints for RTK/INS integrated navigation. *Measurement Science and Technology*, 32(10). <https://doi.org/10.1088/1361-6501/ac0370>
- Yoder, J. E. (2021). *Low-cost inertial aiding for deep-urban tightly-coupled multi-antenna precise GNSS* [Unpublished master's thesis]. The University of Texas at Austin. <https://radionavlab.ae.utexas.edu/wp-content/uploads/2022/02/tight-coupling-journal.pdf>
- Yoder, J. E., Iannucci, P. A., Narula, L., & Humphreys, T. E. (2020). Multi-antenna vision-and-inertial-aided CDGNSS for micro aerial vehicle pose estimation. *Proc. of the 33rd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2020)*, 2281–2298. <http://doi.org/10.33012/2020.17588>
- Zhang, H. T. (2006). *Performance comparison on kinematic GPS integrated with different tactical-grade IMUs* [Unpublished master's thesis]. The University of Calgary. <https://doi.org/10.11575/PRISM/206>

How to cite this article: Yoder, J. E., & Humphreys, T. E. (2023). Low-cost inertial aiding for deep-urban tightly coupled multi-antenna precise GNSS. *NAVIGATION*, 70(1). <https://doi.org/10.33012/navi.561>