

Factor Graph Fusion of Raw GNSS Sensing with IMU and Lidar for Precise Robot Localization without a Base Station

Jonas Beuchert^{1,2}, Marco Camurri², and Maurice Fallon²

Abstract—Accurate localization is a core component of a robot’s navigation system. To this end, global navigation satellite systems (GNSS) can provide absolute measurements outdoors and, therefore, eliminate long-term drift. However, fusing GNSS data with other sensor data is not trivial, especially when a robot moves between areas with and without sky view. We propose a robust approach that tightly fuses raw GNSS receiver data with inertial measurements and, optionally, lidar observations for precise and smooth mobile robot localization. A factor graph with two types of GNSS factors is proposed. First, factors based on pseudoranges, which allow for global localization on Earth. Second, factors based on carrier phases, which enable highly accurate relative localization, which is useful when other sensing modalities are challenged. Unlike traditional differential GNSS, this approach does not require a connection to a base station. On a public urban driving dataset, our approach achieves accuracy comparable to a state-of-the-art algorithm that fuses visual inertial odometry with GNSS data—despite our approach not using the camera, just inertial and GNSS data. We also demonstrate the robustness of our approach using data from a car and a quadruped robot moving in environments with little sky visibility, such as a forest. The accuracy in the global Earth frame is still 1–2 m, while the estimated trajectories are discontinuity-free and smooth. We also show how lidar measurements can be tightly integrated. We believe this is the first system that fuses raw GNSS observations (as opposed to fixes) with lidar.

I. INTRODUCTION

A key enabler for autonomous navigation is accurate localization using only a robot’s onboard sensors. Effective sensor fusion is essential to maximize information gain from each sensing modality. For example, proprioception based on inertial measurement units (IMUs) or encoders together with exteroception from cameras or lidars can be used to achieve a smooth, but slowly drifting, estimate of the motion of a mobile platform in a local environment. In contrast, satellite navigation can be used to estimate positions in a global Earth frame. These estimates are drift-free, but require sky visibility. Therefore, fusion of satellite navigation, proprioception, and exteroception is desirable for long-term autonomy.

The most popular way to fuse data from global navigation satellite systems (GNSS) with onboard perception sensing is two-stage fusion. First, an independent estimator running on the GNSS receiver computes global position fixes from raw




Fig. 1. Setups for experimental validation. **Left:** a car driving on public roads. **Center:** a quadruped robot moving in urban and forest environments. **Right:** a GNSS receiver carried handheld. Raw GNSS measurements were tightly fused with inertial sensing and optionally lidar using factor graphs.

GNSS observations. These fixes are then used as position priors within a second-stage pose estimator with proprioceptive and exteroceptive measurements [1], [2]. However, this two-stage fusion has several disadvantages:

- The GNSS-based state estimation is only loosely coupled with the fusion of the other measurements, which usually causes lower accuracy and higher uncertainty. In particular, when a GNSS receiver acquires only a few satellites, it can fail to calculate a position estimate by itself. As a result, down-stream processing would not be able to use any information from the GNSS at all.
- Maintaining a persistent network connection for differential GNSS (DGNSS) is impossible or inconvenient in many applications. Without DGNSS, receiver position uncertainty is in the order of several meters. This is magnitudes higher than the uncertainties of proprioception and exteroception and means that GNSS can only be used to anchor a robot’s trajectory in the global Earth frame or to correct long-term drift, but cannot aid accurate local navigation, e.g., if the robot’s exteroception fails.
- The position fixes are provided by a black box (the GNSS receiver), whose internal estimator might have unknown dynamics. This makes it difficult to rigorously fuse its output with other sensor data, especially when transitioning between open and covered areas.

We follow an alternative approach that addresses these disadvantages by instead fusing the raw observations of the individual satellites with proprioceptive and exteroceptive measurements in a single state estimator, which is more tightly coupled. This allows us to leverage information from even just a few satellites (e.g., less than four)—information which would be discarded in the two-stage approach.

Jonas Beuchert is supported by the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems. Maurice Fallon is supported by a Royal Society University Research Fellowship. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Accepted Manuscript version arising.

¹Dept. of Computer Science, Univ. of Oxford, UK

²Oxford Robotics Inst., Dept. of Eng. Science, Univ. of Oxford, UK
{beuchert, mcamurri, mfallon}@robots.ox.ac.uk

For each visible satellite and a given frequency band, a conventional GNSS receiver can provide three types of observables: pseudoranges, carrier phases, and Doppler shifts. These measurements are affected by a number of errors, rendering fusion particularly challenging and requiring robust optimization. In addition, the different observables have distinct properties: pseudoranges have a comparably high uncertainty while carrier phases are accurate, but challenging to use in real time without a permanent communication link to a base station. Our algorithm addresses both challenges.

Specifically, the contributions of our work are:

- A novel factor graph design incorporating two types of raw GNSS observables (pseudoranges for drift-free localization in the global Earth frame and differential carrier phases for accurate and smooth local positioning) together with inertial measurements and optionally lidar—to our knowledge the first system to do so.
- A single real-time optimization phase, which implicitly handles GNSS initialization, normal operation, and GNSS drop-out. This eliminates the need to switch between different modes for the aforementioned phases and leads to fast convergence.
- An extensive evaluation on a car and a quadruped robot moving in challenging scenarios and comparison against the state-of-the-art on a public dataset.

II. RELATED WORK

The idea of fusing individual GNSS satellite observations (rather than pre-computed GNSS fixes) with proprioceptive and exteroceptive measurements in a single estimation framework has been pursued in previous work, including works using factor graph optimization. Most approaches used the pseudoranges provided by the GNSS receiver [3]–[9]. A pseudorange is the observed travel time of a radio signal from a satellite to the receiver multiplied by the speed of light. They can be seen as range-only observations of distant landmarks; although, pseudoranges usually have much larger meter-scale uncertainties than the centimeter-scale ones of visual landmarks [9]. For example, Gong et al., Wen et al., and Cao et al. fused pseudoranges with combinations of visual or inertial measurements to show that a tightly coupled approach can be robust in urban canyon scenarios with limited sky visibility [5], [9], [10]. They achieved mean positioning errors of a few meters. Due to their meter-scale measurement uncertainty, pseudoranges cannot be employed for centimeter-accurate localization—only for anchoring a trajectory in a global Earth frame and eliminating odometry drift. For accurate local navigation, these approaches rely on a combination of proprioception and exteroception.

To overcome this limitation, carrier-phase observations can also be used. The satellites transmit their data on carrier signals in the gigahertz range, i.e., sine waves with fixed frequencies. If the receiver could measure the number of sine-wave periods between the satellite and itself, this would serve as an additional observation that is proportional to the receiver-satellite distance. It would be more accurate than the pseudorange because signal wavelengths are in the range

18–26 cm. However, the receiver cannot count the absolute number of sine-wave periods. Instead, it can observe the phase of the carrier wave and count the change in the number of waves since the receiver first locked onto the signal. (The sum of these values is usually referred to as the *carrier phase*.) Thus, the range observation can only be inferred up to an unknown integer number of wavelengths, known as the *integer ambiguity*, which is the number of full waves when the signal was locked. Usually, real-time kinematic positioning (RTK) is used to resolve the integer ambiguity in real time. However, this requires a persistent connection between the moving GNSS receiver (the *rover*) and a nearby stationary second receiver at a known location (the *base*).

In the literature, approaches have been described that could make use of carrier phases in real time *without the need for a base station*. For example, Suzuki used carrier phases to create factors between states at different times to obtain relative distance measurements w.r.t. past epochs [11]. The author named this method *time-relative RTK* because of its similarity to RTK—with the current observations as rover observations and a set of previous observations as observations to a virtual base station. For each continuously tracked satellite, the approach estimated the integer ambiguity using the LAMBDA method [12], [13]. If the method could not resolve the integer ambiguity, then no factor was added to the graph. In this way, the integer ambiguity estimation was not tightly coupled with the factor-graph-based state estimation. Combining these carrier-phase factors with pseudorange and Doppler-shift factors in a single factor graph, Suzuki achieved mean positioning errors of 2–5 cm after UAV flights over 200 m or 100 s with good sky visibility and post-processing of data from a single-frequency receiver. There was no real-time evaluation of this method and it did not address fusion with non-GNSS measurements.

Lee et al. described a second method called *sequential-differential GNSS* in which they also create differential carrier-phase factors between different states in time [14]. This approach also canceled the integer ambiguities. It fused the carrier phases with pseudorange, Doppler shift, visual, and inertial measurements in a multi-state constraint Kalman filter (MSCKF). However, they also used time-relative factors for the pseudoranges. They anchored the local trajectory in the global Earth frame during an initialization phase and afterwards only used the pseudorange and carrier-phase observations for relative localization w.r.t. previously estimated states. This limits the usefulness of pseudorange observations for reducing *long-term* drift, especially, if sky visibility is lost intermittently or very limited. They demonstrated an RMSE of 0.32 m on a handheld dataset where sky visibility was never interrupted and differential GNSS factors could always be created between the current and previous state.

In summary, no method has been presented yet that tightly fuses pseudoranges and carrier phases with proprioception and/or exteroception for *long-term* autonomous localization *in real time* using a single GNSS receiver. Furthermore, to the best of our knowledge, tight fusion of raw GNSS data with lidar measurements has not been addressed in the literature.




Fig. 2. Reference frames: the Earth-centered-Earth-fixed (ECEF) frame E, the local frame W, and the frames on the platforms, including base B, which coincides with the GNSS antenna frame A, IMU frame I, and lidar frame L.

III. PROBLEM STATEMENT

We aim to estimate the position of a mobile platform that is equipped with a GNSS receiver, an IMU, and optionally a lidar—in real time and in a global Earth frame. The estimated trajectory ought to be smooth, i.e., discontinuity-free.

A. Frames

Ultimately, we are interested in position estimates in geographic coordinates (latitude, longitude). However, for computational reasons, we internally use the Cartesian Earth-centered-Earth-fixed (ECEF) frame E as global frame. Its z -axis is the Earth's rotation axis, and its x -axis points to the prime meridian, cf. Fig. 2.

If the system has exteroceptive sensors, then we are not only interested in a global position estimate, but also in maintaining a smooth local trajectory with no discontinuities during initial GNSS convergence. Therefore, we establish a local frame W that we align with the platform's pose at the start of the experiment. We estimate the transformation $\mathbf{T}_{EW} \in \text{SO}(3) \times \mathbb{R}^3$ jointly with the platform position in frame W such that \mathbf{T}_{EW} absorbs any large changes/discontinuities due to GNSS convergence while the position estimate in W remains unaffected.

If the sensing system has only GNSS and inertial sensing, then we set W to be identical to E. The frames rigidly attached to the robot are base B (coincident with the GNSS antenna frame A), IMU I, and lidar L, see Fig. 2.

B. State

The state of the system at time t_i is

$$\mathbf{x}_i \triangleq [\mathbf{R}_i, \mathbf{p}_i, \mathbf{v}_i, \mathbf{b}_i^g, \mathbf{b}_i^a, \delta \mathbf{t}_i^r] \in \text{SO}(3) \times \mathbb{R}^{12+n}, \quad (1)$$

where $\mathbf{R}_i \in \text{SO}(3)$ is the orientation, $\mathbf{p}_i \in \mathbb{R}^3$ is the position, and $\mathbf{v}_i \in \mathbb{R}^3$ is the linear velocity of the base frame B in the local frame W. The IMU biases are expressed as $\mathbf{b}_i^g \in \mathbb{R}^3$ and $\mathbf{b}_i^a \in \mathbb{R}^3$ and $\delta \mathbf{t}_i^r \in \mathbb{R}^n$ is a vector that contains the clock offsets between each satellite system and the receiver clock. Our GNSS receiver accesses four satellite systems (i.e., $n = 4$): GPS, BeiDou, GLONASS, and Galileo. We need to model the offset of the receiver clock w.r.t. the time frames of each satellite system since the receiver clock is too imprecise to accurately measure the travel time of satellite signals that travel at the speed of light.

If the setup includes a lidar, we estimate the transformation \mathbf{T}_{EW} in addition to the system states \mathbf{x}_i . The history of all unknowns is then

$$\mathcal{X}_k \triangleq \{\mathbf{x}_i, \mathbf{T}_{EW}\}_{i \in \mathcal{T}_k}, \quad (2)$$

where \mathcal{T}_k is the set of all state indices in a certain time window.

C. Measurements

The measurements include acceleration and angular velocity in the IMU frame I and lidar point clouds \mathcal{L}_i . Inertial measurements are received as a set between times t_i and t_j as \mathcal{I}_{ij} . The GNSS receiver observes pseudoranges \mathcal{P}_i and double-differenced carrier phases \mathcal{C}_{ij} . (The latter are explained in Sec. IV-D.) Thus, all measurements in a time window \mathcal{T}_k are

$$\mathcal{Z}_k \triangleq \{\mathcal{I}_{ij}, \mathcal{L}_i, \mathcal{P}_i, \mathcal{C}_{ij}\}_{i,j \in \mathcal{T}_k}. \quad (3)$$

We create a state \mathbf{x}_i whenever there is a GNSS observation at time t_i and motion correct a potentially received lidar point cloud at that timestamp [15]. If there is no GNSS observation for a certain amount of time, e.g., due to no sky visibility, we create a state with lidar and inertial measurements only.

D. Maximum-a-Posteriori Estimation

We maximize the likelihood of the measurements \mathcal{Z}_k , given the history of states \mathcal{X}_k

$$\mathcal{X}_k^* = \arg \max_{\mathcal{X}_k} p(\mathcal{X}_k | \mathcal{Z}_k) \propto p(\mathcal{X}_0) p(\mathcal{Z}_k | \mathcal{X}_k), \quad (4)$$

where measurements are assumed to be conditionally independent and corrupted by white Gaussian noise. Therefore, we can express Eq. (4) as a least-squares minimization [16]

$$\begin{aligned} \mathcal{X}_k^* = \arg \min_{\mathcal{X}_k} & \|\mathbf{r}_0\|_{\Sigma_0}^2 \\ & + \sum_{i,j \in \mathcal{T}_k} \left(\|\mathbf{r}_{\mathcal{I}_{ij}}\|_{\Sigma_{\mathcal{I}_{ij}}}^2 + \|\mathbf{r}_{\mathcal{L}_i}\|_{\Sigma_{\mathcal{L}_i}}^2 \right. \\ & \left. + \sum_{\rho_i^k \in \mathcal{P}_i} \|r_{\rho_i^k}\|_{\sigma_{\rho_i^k}}^2 + \sum_{\Delta \Delta \phi_{ij}^{kl} \in \mathcal{C}_{ij}} \|r_{\Delta \Delta \phi_{ij}^{kl}}\|_{\sigma_{\Delta \Delta \phi_{ij}^{kl}}}^2 \right), \end{aligned} \quad (5)$$

where \mathcal{T}_k is the set of all state indices in the sliding smoothing window up to time t_k . Each term is the residual associated with a factor type, weighted by the inverse of its covariance matrix. Residuals include a prior, IMU factors, relative odometry factors from lidar, and two types of GNSS factors, which are detailed in the following section.

IV. FACTOR GRAPH FORMULATION

Figure 3 shows the factor graph structure. Each measurement factor is associated with a residual, which is the difference between a model-based prediction given the connected variables and the observation. We will summarize the IMU and lidar residuals in Sec. IV-A and IV-B before detailing the pseudorange and carrier-phase residuals in Sec. IV-C and IV-D.




Fig. 3. Factor graph structure with variable nodes (large circles) for states x_i and transformation T_{EW} between local frame and global Earth frame and factor nodes (small, colored) for priors, and proprioceptive (IMU), exteroceptive (lidar), and GNSS measurements (pseudorange, carrier phase).

A. Pre-integrated Inertial Measurements

We follow the standard method of IMU measurement pre-integration to constrain the pose, velocity, and biases between two consecutive nodes of the graph, providing high-frequency state updates between nodes. The residual has the form

$$\mathbf{r}_{\mathcal{I}_{ij}} = \begin{bmatrix} \mathbf{r}_{\Delta \mathbf{R}_{ij}}^T, \mathbf{r}_{\Delta \mathbf{v}_{ij}}^T, \mathbf{r}_{\Delta \mathbf{p}_{ij}}^T, \mathbf{r}_{\mathbf{b}_{ij}^a}, \mathbf{r}_{\mathbf{b}_{ij}^g} \end{bmatrix}. \quad (6)$$

For a description of pre-integration, see Forster et al. [17].

B. ICP Registration

We register the lidar measurements to a local submap [15] using iterative closest point (ICP) odometry [18] and add the registration output as relative pose factors between times t_i and t_m with the residual

$$\mathbf{r}_{\mathcal{L}_i} = \Phi \left(\tilde{\mathbf{T}}_i^{-1} \tilde{\mathbf{T}}_m^{-1} \mathbf{T}_i^{-1} \mathbf{T}_m \right), \quad (7)$$

where $\mathbf{T}_i = [\mathbf{R}_i, \mathbf{p}_i]$ is the estimated pose, $\tilde{\mathbf{T}}_i \in \text{SE}(3)$ is the estimate from the ICP module, and Φ is the lifting operator defined by Forster et al. [17].

C. Pseudoranges

For each acquired satellite signal k at time t_i , a GNSS receiver reports a pseudorange $\rho_i^k \in \mathcal{P}_i$, which is the observed signal travel time multiplied by the speed of light. It is close to the spatial satellite-receiver distance, but affected by additional terms, in particular signal delays in different layers of the atmosphere and time offsets of the clocks that are used to measure the signal travel time [19]. The pseudorange residual for a single received satellite signal k is approximately

$$\begin{aligned} r_{\rho_i^k} = & \|\mathbf{s}_i^k - \mathbf{T}_{EW} \mathbf{p}_i\| \\ & + \delta \rho_T(\mathbf{s}_i^k, \mathbf{T}_{EW} \mathbf{p}_i) + \delta \rho_I(\mathbf{s}_i^k, \mathbf{T}_{EW} \mathbf{p}_i) \\ & + c \cdot \delta t_{i,g}^r - (\rho_i^k + c \cdot (\delta t^k + \nu^k)), \end{aligned} \quad (8)$$

where $\mathbf{s}_i^k \in \mathbb{R}^3$ is the satellite position in frame E at the signal transmit time that corresponds to the observation time t_i . It is corrected for the Earth rotation. The spatial signal delay in the troposphere is $\delta \rho_T: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}_0^+$, the delay in the ionosphere is $\delta \rho_I: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}_0^+$, the speed of light is c , and the offset of the receiver clock from the clock of the GNSS

with index $g \in \{0 \dots 3\}$ is $\delta t_{i,g}^r \in \mathbb{R}$. The pseudorange ρ_i^k is adjusted to correct for satellite clock bias $\delta t^k \in \mathbb{R}$ and relativity $\nu^k \in \mathbb{R}$.

We model satellite position, satellite clock offset, and both atmospheric delays based on data broadcasted by the satellites. Due to these effects, modeling errors can be up to a few meters. The antenna position and the receiver clock offset are the unknowns that we need to estimate.

If we use a multi-band receiver, then there can be multiple residuals for a single satellite, one for each band in which the receiver acquired a signal.

D. Carrier Phases

The residual of a carrier phase $\phi_i^k \in \mathbb{R}^+$ (in units of lengths) could be written similarly to the pseudorange residual [19]

$$\begin{aligned} r_{\phi_i^k} = & \|\mathbf{s}_i^k - \mathbf{T}_{EW} \mathbf{p}_i\| \\ & + \delta \rho_T(\mathbf{s}_i^k, \mathbf{T}_{EW} \mathbf{p}_i) - \delta \rho_I(\mathbf{s}_i^k, \mathbf{T}_{EW} \mathbf{p}_i) + \lambda^k \omega^k \\ & + c \cdot \delta t_{i,g}^r - (\phi_i^k + \lambda^k N_i^k + c \cdot (\delta t^k + \nu^k)). \end{aligned} \quad (9)$$

Most of the terms are the same as in Eq. (8). However, the delay in the ionosphere $\delta \rho_I$ has the opposite sign. There is also a wind-up effect $\omega^k \in \mathbb{R}$ resulting from the interplay between the changing satellite orientation and the circularly polarized carrier wave (with wavelength $\lambda^k \in \mathbb{R}^+$), which has a magnitude ranging from centimeters to a few decimeters [20]. The most significant difference is an unknown offset in the number $N_i^k \in \mathbb{N}$ of wavelengths.

The advantage of carrier-phase observations over pseudorange ones is that their measurement noise is only ~ 5 mm. However, modeling them precisely in real time is challenging because satellite positions and atmospheric delays are not necessarily known with submeter accuracy. Furthermore, the integer ambiguity N_i^k is an additional unknown that needs to be estimated. Therefore, we apply double-differencing, i.e., combine multiple observations to cancel out terms that are unknown or imprecisely known. Specifically, we combine two observations of two different satellites of the same GNSS and signal band at the current time t_i with two older observations of the same satellites at a time $t_j < t_i$ [11]

$$\Delta \Delta \phi_{ij}^{kl} = \left(\tilde{\phi}_j^l - \tilde{\phi}_j^k \right) - \left(\tilde{\phi}_i^l - \tilde{\phi}_i^k \right), \quad (10)$$

where $\tilde{\phi}_i^k = \phi_i^k + c \cdot (\delta t^k + \nu^k)$ is the carrier phase ϕ_i^k corrected for satellite clock bias and relativity. If the change of the atmospheric delays of the satellite signals k and l between times t_i and t_j is negligible and the signals are still locked (i.e., $N_i^k = N_j^k$ and $N_i^l = N_j^l$), then the residual is

$$\begin{aligned} r_{\Delta \Delta \phi_{ij}^{kl}} = & (\|\mathbf{s}_j^l - \mathbf{T}_{EW} \mathbf{p}_j\| - \|\mathbf{s}_j^k - \mathbf{T}_{EW} \mathbf{p}_j\|) \\ & - (\|\mathbf{s}_i^l - \mathbf{T}_{EW} \mathbf{p}_i\| - \|\mathbf{s}_i^k - \mathbf{T}_{EW} \mathbf{p}_i\|) \\ & - \Delta \Delta \phi_{ij}^{kl}. \end{aligned} \quad (11)$$

For each pair of band and satellite system (GPS, GLONASS, Galileo, BeiDou), we select the satellite that has been continuously visible for the longest time for the first satellite signal k . For the second satellite signal l , we iterate over all remaining satellite observations in the same signal band. We create a residual for each satellite pair obtained this way.

TABLE I

MEAN (TOP) AND MEDIAN (BOTTOM) HORIZONTAL LOCALIZATION ERRORS [M] IN THE GLOBAL EARTH FRAME WITH RTK AS GROUND TRUTH.

dataset	platform	duration [s]	baseline methods			GVINS	our proposed method	
			GNSS-fix	IMU, GNSS-fix	IMU, ICP, GNSS-fix		IMU, raw-GNSS	IMU, ICP, raw-GNSS
Hong Kong	car	2454				1.75 1.63	1.54 1.44	
Jericho	car	923	4.01	4.04	failure		2.22	2.03
			2.67	2.82			1.88	1.71
Park Town	car	264	4.57	4.46			2.79	
			2.56	2.54			2.00	
Bagley	quadruped	1120	3.40	3.46	4.78		2.34	2.07
			2.84	3.16	5.50		1.97	2.02
Thom	quadruped	440	12.31	10.93	4.24		1.66	1.33
			7.42	9.32	2.92		0.86	1.00

V. IMPLEMENTATION

We integrated the GNSS factors from Sec. IV-C and IV-D into the VILENS state estimator, which includes IMU and lidar registration factors already [15]. We created lidar registration factors at half the rate of the GNSS factors and solved the factor graph with a fixed lag smoother based on the incremental optimizer iSAM2 [21] in the GTSAM library [22]. The transformation \mathbf{T}_{EW} between the local and global frame was continuously estimated, while the states outside the optimization window \mathcal{T} were marginalized.

Raw GNSS observations are prone to outliers; for example, the multi-path effect occurs when satellite signals are reflected by surrounding buildings or vegetation before they are received. This induces a longer travel distance than the direct line of sight. To mitigate this, we first applied a threshold-based outlier detector and then used the Huber loss function to reduce the effect of remaining outliers [4].

We implemented the GNSS processing components using the GPSTk library [23]. To avoid a cold start of at least 30 s where satellite positions and clocks are unknown, we preloaded publicly available satellite navigation data before the start of an experiment. No online data was used thereafter.

VI. EXPERIMENTAL RESULTS

To evaluate our proposed algorithm, we conducted three experiments using both, public and self-recorded datasets.

Experiment 1 (Sec. VI-A - raw GNSS and IMU fusion): We compared against the state-of-the-art using the *urban driving* sequence of the public GVINS dataset [10]. The setup for this dataset included a 200 Hz consumer-grade ANALOG DEVICES ADIS16448 IMU, two 20 Hz APTINA MT9V034 cameras, and a 10 Hz entry-level multi-band U-BLOX C099-F9P GNSS receiver. Because there was no lidar data, we evaluated only the fusion of IMU and raw GNSS with our algorithm. The sequence has a duration of 41 min and includes sections with little sky visibility (urban canyons) and brief complete GNSS drop-outs (underpasses). We also collected four sequences with limited sky visibility of our own by mounting a 5 Hz C099-F9P GNSS receiver with a multi-band GNSS antenna and a consumer-grade 200 Hz BOSCH BMI085 IMU on an electric vehicle and a BOSTON DYNAMICS SPOT quadruped robot, see Fig. 1. The sequences evaluated were:

- *Jericho*: a 4.4 km driving loop through narrow streets in Oxford’s city center at speeds up to 11 m/s. This

includes several sections with GNSS drop-out.

- *Park Town*: a 1.2 km drive along tree-lined avenues.
- *Bagley*: the quadruped moving through a dense commercial forest. This is known to be challenging for satellite navigation due to the very limited sky visibility, many outliers in the GNSS measurements because of signal reflections by surrounding vegetation (multi-path effect), and signal degradation caused by the electromagnetic interference of the robot with the GNSS signals.
- *Thom*: an urban scenario, where the quadruped walked around a high-rise building, passed through a tunnel and ended in a yard between tall buildings, see Fig. 4-B.

For all sequences, we used RTK to obtain ground truth.

Experiment 2 (Sec. VI-B - raw GNSS, IMU, and lidar fusion): We evaluated fusion with data from a 10 Hz HESAI XT32 lidar that was part of the setup for the sequences *Jericho*, *Bagley*, and *Thom*.

Experiment 3 (Sec. VI-C - carrier-phase-only fusion): Finally, we tested the carrier-phase factor’s usefulness for accurate local navigation. For this, we hand carried the GNSS receiver around a 12 m tall tree for 105 m and 2 min, cf. Fig. 4-E, and created the sequence *Tree*.

A. Experiment 1: Raw GNSS and IMU fusion

The column **GVINS** of Tab. I shows the localization errors of the open-source GVINS algorithm, which fuses pseudoranges and Doppler shifts from the GNSS receiver with inertial measurements and vision constraints from a camera. The column **IMU, raw-GNSS** presents results for our algorithm fusing pseudoranges, carrier phases, and inertial measurements only. Our algorithm performs slightly better than GVINS despite the fact that we do not use a camera. This demonstrates that carrier-phase observations can replace exteroception for accurate local navigation when the latter is unavailable, but some sky visibility exists.

There is a practical difference between the algorithms, too: GVINS requires a dedicated initialization phase with sufficient sky visibility that lasts for 15.3 s for the *urban driving* sequence. In contrast, our algorithm does not require such a separate phase and converges to the correct pose of B in the Earth frame E as soon as the first slight motion occurs, after less than 4 s. We assume that this is partially due to the utilization of carrier phases: a small motion is sufficient to estimate the orientation of B in E from the precise differential

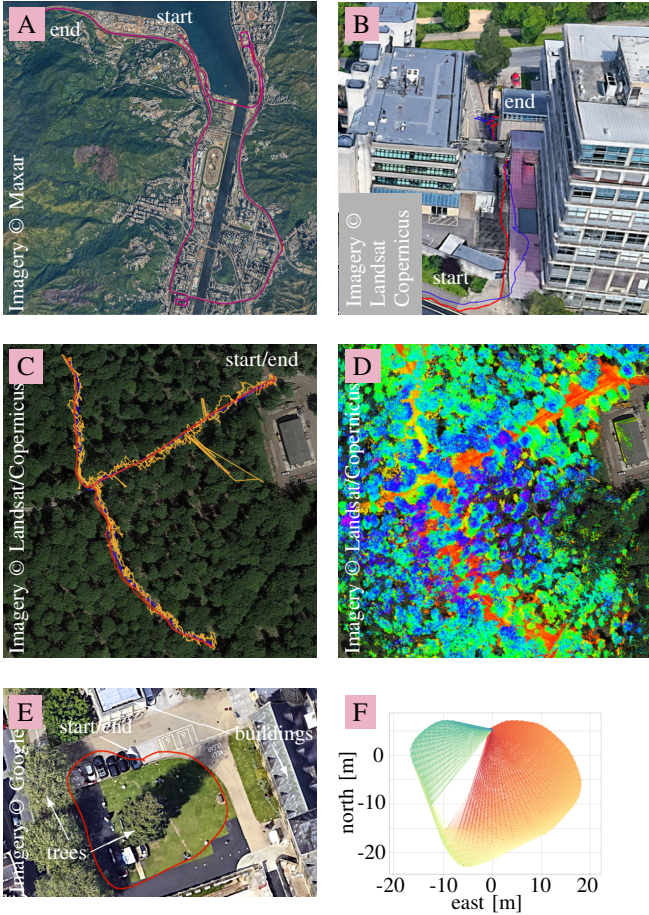


Fig. 4. Experimental results. (A)¹: Trajectory of a car in Hong Kong estimated with our algorithm fusing inertial and GNSS measurements (red) compared to RTK (ground truth, blue) [10]. (B): Trajectory of a quadraped traversing between two buildings estimated using IMU, ICP, and GNSS (red) compared to RTK (blue), part of sequence *Thom*. The RTK trajectory drifted into the building when there was little GNSS data while our estimate did not. (C): Trajectory of a quadraped in the Bagley Wood estimated using IMU, ICP, and GNSS (red) in comparison to RTK (blue) and single GNSS fixes (orange). (D): Same trajectory with lidar scans overlaid. (E): Trajectory of a handheld device estimated using time-differential carrier-phase factors only and given the starting position. (F): The corresponding factors between different states in time, represented by lines and crosses respectively. Few factors were created to states near the trees on the left. Still, the horizontal localization error at the end was less than 10 cm.

carrier phases, while more motion is required to achieve the same with the less accurate pseudoranges and Doppler shifts.

In Tab. I, we compare the accuracy of separately computed non-differential GNSS fixes (column **GNSS-fix**), fusion of these fixes with inertial measurements (**IMU, GNSS-fix**), and our own algorithm fusing raw GNSS observations with inertial measurements (**IMU, raw-GNSS**) for our sequences. The median global accuracy is 1–2 m, which is sufficient for many autonomous vehicle applications, e.g., to initialize a fine-grained lidar localization system or to reject incorrect place proposals from such a system. The accuracy is also close to that achieved on the *Hong Kong* sequence despite the GNSS receiver operating at half the rate and our sequences being shorter and having no sections with as good sky visibility. Both issues hinder global convergence. Furthermore, the trajectory estimates are smooth after initial convergence, as

can be seen in the supplementary material¹. In particular, the 30–60 % smaller error on *Bagley*¹ demonstrates the advantage of using individual GNSS observations and inertial measurements in a single optimization framework as opposed to separating the computation of GNSS fixes and the fusion. The gains mainly come from improved robustness in this scenario, which has fewer and more noisy GNSS observations.

To summarize, these results show that the pseudorange factors enable robust localization in the global Earth frame. In addition, the carrier-phase factors help to create a locally smooth and accurate trajectory when exteroception is absent.

B. Experiment 2: Raw GNSS, IMU, and Lidar Fusion

For the sequences *Jericho*, *Bagley*¹, and *Thom*, we also compare the accuracy of fusing separately computed GNSS fixes with IMU measurements and ICP (**IMU, ICP, GNSS-fix**) versus our own algorithm when fusing raw GNSS observations with inertial measurements and ICP (**IMU, ICP, raw-GNSS**) in Tab. I. The global accuracy is similar to *Exp. 1*, while we also obtain a georeferenced map of the local environment, see Fig. 4-D. Furthermore, the results for *Thom* show that our algorithm, with its single optimization stage, can implicitly handle GNSS drop-out and smoothly switch between GNSS-aided and non-GNSS navigation, cf. Fig. 4-B.

C. Experiment 3: Carrier-Phase-Only Fusion

Lastly, we tested our algorithm¹ with only double-differential carrier-phase factors and without proprioception or exteroception on the *Tree* sequence, see Fig. 4-E and 4-F. Here the double-differenced carrier phase measurements provided an input roughly equivalent to relative local odometry. The horizontal positioning error in the end was less than 10 cm, indicating performance close to Suzuki’s method [11], despite the fact that we only used one type of GNSS factor and did not have clear sky visibility.

VII. CONCLUSIONS

We have presented a novel system to localize a mobile robot in an unknown environment by tightly fusing GNSS observations, proprioceptive measurements, and optionally exteroceptive measurements into a single factor graph optimization. The approach employs raw GNSS observations from individual satellites instead of pre-computed position fixes to maximize the information taken from the GNSS receiver. This includes not only pseudorange observations for absolute positioning on Earth, but also accurate carrier-phase observations for precise local localization to support the situation where exteroceptive measurements are either unavailable or degraded. We showed that the proposed approach improves accuracy by up to several meters in comparison to the two-stage algorithms where a pre-computed fix is used in the factor graph instead of raw data—especially

¹We provide open-source code to reproduce the results in Fig. 4-E and 4-F on <https://github.com/JonasBchrt/raw-gnss-fusion> together with the dataset sequences *Bagley* and *Tree* and interactive maps with the estimated trajectories corresponding to some of the results in Tab. I.

if the view of the sky is limited. Typical median localization errors in the Earth frame are still just 1–2 m.

REFERENCES

- [1] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, “LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping,” in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2020, pp. 5135–5142.
- [2] Z. Gong, P. Liu, F. Wen, R. Ying, X. Ji, R. Miao, and W. Xue, “Graph-based adaptive fusion of GNSS and VIO under intermittent GNSS-degraded environment,” *IEEE Trans. on Instrumentation and Measurement*, vol. 70, pp. 1–16, 2020.
- [3] R. M. Watson and J. N. Gross, “Evaluation of kinematic precise point positioning convergence with an incremental graph optimizer,” in *IEEE/ION Position, Location and Navigation Symp.*, 2018, pp. 589–596.
- [4] R. M. Watson, “Enabling robust state estimation through covariance adaptation,” Ph.D. dissertation, West Virginia Univ., Morgantown, 2019.
- [5] W. Wen, X. Bai, Y. C. Kan, and L.-T. Hsu, “Tightly coupled GNSS/INS integration via factor graph and aided by fish-eye camera,” *IEEE Trans. on Vehicular Technology*, vol. 68, no. 11, pp. 10 651–10 662, 2019.
- [6] W. Wen, T. Pfeifer, X. Bai, and L.-T. Hsu, “It is time for factor graph optimization for GNSS/INS integration: Comparison between FGO and EKF,” *arXiv: 2004.10572*, 2020.
- [7] N. Sünderhauf and P. Protzel, “Towards robust graphical models for GNSS-based localization in urban environments,” in *IEEE Intl. Multi-Conf. on Systems, Signals & Devices*, 2012.
- [8] N. Sünderhauf, M. Obst, S. Lange, G. Wanielik, and P. Protzel, “Switchable constraints and incremental smoothing for online mitigation of non-line-of-sight and multipath effects,” in *IEEE Intelligent Vehicles Symp.*, 2013, pp. 262–268.
- [9] Z. Gong, R. Ying, F. Wen, J. Qian, and P. Liu, “Tightly coupled integration of GNSS and vision SLAM using 10-DoF optimization on manifold,” *IEEE Sensors Journal*, vol. 19, no. 24, pp. 12 105–12 117, 2019.
- [10] S. Cao, X. Lu, and S. Shen, “GVINS: Tightly coupled GNSS–visual–inertial fusion for smooth and consistent state estimation,” *IEEE Trans. on Robotics*, pp. 1–18, 2022.
- [11] T. Suzuki, “Time-relative RTK-GNSS: GNSS loop closure in pose graph optimization,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4735–4742, 2020.
- [12] P. J. G. Teunissen, P. J. De Jonge, and C. C. J. M. Tiberius, “The LAMBDA method for fast GPS surveying,” in *Intl. Symp. GPS Technology Applications*, 1995.
- [13] T. Takasu and A. Yasuda, “Development of the low-cost RTK-GPS receiver with an open source program package RTKLIB,” in *Intl. Symp. on GPS/GNSS*, 2009.
- [14] W. Lee, P. Geneva, Y. Yang, and G. Huang, “Tightly-coupled GNSS-aided visual-inertial localization,” in *International Conference on Robotics and Automation (ICRA)*, 2022.
- [15] D. Wisth, M. Camurri, and M. Fallon, “VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots,” *IEEE Trans. on Robotics*, pp. 1–18, Aug. 2022.
- [16] F. Dellaert and M. Kaess, “Factor graphs for robot perception,” *Foundations and Trends in Robotics*, vol. 6, pp. 1–139, 2017.
- [17] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, “On-manifold preintegration for real-time visual-inertial odometry,” *IEEE Trans. on Robotics*, vol. 33, no. 1, pp. 1–21, 2017.
- [18] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, “Comparing ICP variants on real-world data sets,” *Autonomous Robots*, vol. 34, no. 3, pp. 133–148, 2013.
- [19] E. Kaplan and C. Hegarty, *Understanding GPS: principles and applications*. Artech House, 2005.
- [20] J.-T. Wu, S. C. Wu, G. A. Hajj, W. I. Bertiger, and S. M. Lichten, “Effects of antenna orientation on GPS carrier phase,” *Astrodynamics*, pp. 1647–1660, 1992.
- [21] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, “iSAM2: Incremental smoothing and mapping using the Bayes tree,” *Intl. Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [22] F. Dellaert, M. Kaess, *et al.*, “Factor graphs for robot perception,” *Foundations and Trends in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [23] R. B. Harris and R. G. Mach, “The GPSTk: an open source GPS toolkit,” *GPS Solutions*, vol. 11, no. 2, pp. 145–150, 2007.