Plug and Play Sensor Fusion for Lane-Level Positioning of Connected Cars in GNSS-Challenged Environments

Andrey Soloviev, Michael Veth, Chun Yang, QuNav

ABSTRACT

For automotive safety applications, connected cars exchange navigation data between vehicles and apply this information for collision avoidance. One of the key challenges is to enable the required navigation accuracy for any driving environments at any time.

GNSS signal availability and navigation data quality decrease rapidly once operational environments shift from open sky to degraded signal scenarios such as urban and tree-covered roads. Multi-sensor augmentations of GNSS can maintain the required localization capabilities. However, current multi-sensor implementations are rather ad hoc and sensors-specific.

This paper presents a truly plug-and-play navigation solution that automatically reconfigures itself as sensors are connected to (disconnected from) the system, without the need to redesign the system architecture or its specific components.

For experimental demonstrations, test data were collected in urban canyons of downtown San Francisco, CA in January 2016. The paper provides experimental results for various sensor configurations including carrier phase GNSS, consumer-grade IMU, video camera, and the use of vehicle motion constraints. Consistent positioning in urban canyons is demonstrated in support of automotive safety applications.

INTRODUCTION

Many existing and perspective applications of navigation systems would benefit notably from the ability to navigate accurately and reliably in difficult environments. Examples of difficult navigation scenarios include urban canyons, indoor applications, radio-frequency (RF) interference and jamming environments. In addition, different segments of a mission path can impose significantly different requirements on the navigation sensing technology and data processing algorithms. To exemplify, Figure 1 shows a mission scenario of an autonomous aerial vehicle (UAV).

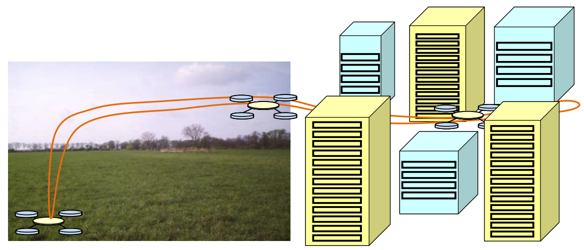

Figure 1. UAV mission example

For this example, the UAV is deployed in an open field; next, the vehicle enters an urban canyon to perform tasks such as surveying and inspection; and, finally, it returns to the deployment point. To enable operation of the UAV at any point on the flight path, a precision navigation, attitude, and time capability on-board the vehicle is required.

Currently, the majority of navigation products rely on Global Navigation Satellite Systems (GNSS) as a primary navigation aid. However, GNSS performance decreases rapidly once operational environments shift from open sky to degraded signal scenarios. Multi-sensor augmentations of GNSS can maintain desired navigation capabilities [1]. However, existing multi-sensor implementations are sensors-specific. Short-term gains in implementation efficiency are soon offset by non-recurring engineering costs of initial development and long-term higher integration costs whenever changes or upgrades are required.

To address limitations of sensor fusion technologies in GNSS-degraded applications, we have been developing plug-and-play (PnP) sensor fusion mechanizations [2]. As shown in Figure 2, the PnP navigation solution automatically reconfigures itself as sensors are connected to (disconnected from) the system, without the need to redesign the system architecture or its specific components. PnP sensor fusion is supported by a reconfigurable integration filtering engine (RIFE). The navigation filter mechanization is abstracted into object-oriented multi-sensor estimation. Various sensors are represented by generic classes in the RIFE library. Each class is designed for a generic type of sensors (rather than for a specific sensor) wherein sensor types are defined by the type of measurement. When a sensor is connected to the system, the RIFE is reconfigured by identifying the sensor's measurement type and activating a sensor object using a corresponding class from the RIFE library.

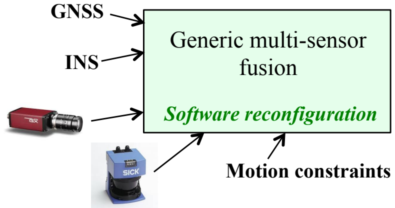

Figure 2. Plug and play solution: Generic sensor fusion automatically reconfigures itself for a chosen set of sensor

RIFE utilizes a self-contained inertial navigation system (INS) as its core sensor. INS does not rely on any type of external information and can thus operate in any environment. However, INS solution drifts over time. To mitigate inertial drift, this core sensor is augmented with reference navigation data sources.

It is difficult, if not impossible, to create an exhaustive list of all potential aiding measurements. Yet, it is possible to categorize aiding measurements into generalized types. To illustrate, Figure 3 exemplifies a generic class of relative position. This class is defined as a projection of position change vector on a specified axis or axes of navigation-frame or body-frame. As shown in Figure 3, three aiding measurements (odometer, 2D lidar, and 3D lidar) can be represented by this generic formulation.

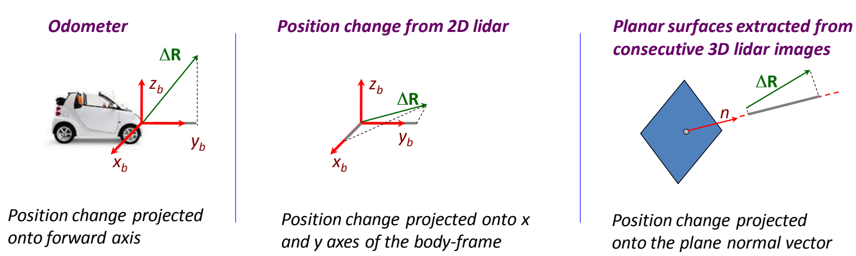

Figure 3. Examples of Relative Position Observables

RIFE was initially introduced in reference [2]. This previous work verified RIFE functionality with various experimental data sets. The verification was primarily focused on the use of higher-grade inertial sensors (tactical and navigation grade). The current paper extends previously demonstrated RIFE capabilities by applying its PnP software to consumer-grade sensors in GNSS-degraded environments.

One of the key aspects of fusion with consumer-grade sensors in obstructed signal environments (such as urban canyons) is the use of carrier phase. Pseudoranges are generally too noisy to be efficiently integrated with consumer-grade IMU. Therefore, carrier phase measurements were included into the RIFE architecture.

The remainder of the paper is organized as follows. First, we describe the incorporation of GNSS carrier phase into RIFE. Next, we introduce a specific use case where the application of RIFE is particularly beneficial from the perspectives of navigation accuracy and system re-configurability. As such we consider connected cars for automotive safety where lane-level accurate positioning is required. Finally, experimental results are shown for various sensor configurations and representative test scenarios in downtown San Francisco, CA. Test results demonstrate that RIFE enables consistent and reliable positioning in dense urban areas while using consumer-grade sensors suitable for automotive applications.

ADDITION OF GNSS CARRIER PHASE TO ROCONFIGIRABLE FILTER DESIGN

Relative position observables that are shown in Figure 3 can be also applied to utilize carrier phase measurements in GNSS-degraded environments. When GNSS is partially available, the use of carrier phase is extremely beneficial, especially, for integration with consumer-grade inertial units. The reason is that carrier phase-based relative ranging is sub-cm accurate, which provides at least two orders of magnitude noise reduction as compared to meter-accurate pseudoranges.

In degraded GNSS environments, the total number of satellites is generally insufficient to enable phase ambiguity resolution. Therefore, integer ambiguities are eliminated by differencing carrier phase over time and applying temporal phase differences as filter measurements. This approach is referred to as the dynamic state INS calibration [3]. It observes projections of position changes (instead of absolute position) on platform-to-satellite line-of-sight (LOS) and estimates the rest of inertial error states (including velocity errors, attitude errors and sensor biases) from these observations.

To incorporate GNSS carrier phase into the RIFE measurement structure, temporal differences (TD) of carrier phase are first be pre-processed into the relative position format and then directly utilized by RIFE. Equation (1) formulates the carrier phase TD equation [3]:

$$\Delta\tilde{\phi}_j = \tilde{\phi}_j(t_k) - \tilde{\phi}_j(t_{k-1}) = \Delta r_j + \Delta\delta t_{rcvr} + \Delta errors_j + \Delta\eta_j \quad (1)$$

In (1):

$\Delta\tilde{\phi}_j$ is the carrier phase TD for satellite j;

$\tilde{\phi}_j$ is the carrier phase measurement for satellite j;

$t_k = t_0 + k \cdot \Delta t_{GNSS}$ is the discrete time and Δt_{GPS} is the GNSS measurement update interval;

$\Delta r_j = r_j(t_k) - r_j(t_{k-1})$ is the TD of the range r_j between the GNSS antenna and satellite j;

$\Delta\delta t_{rcvr}$ is the TD of the receiver clock bias, or, equivalently, the receiver clock drift accumulated over the Δt_{GNSS} interval;

the $\Delta errors_j$ term represents changes in deterministic error components of stand-alone GNSS measurements and includes changes in ionospheric and tropospheric delays, changes in the satellite clock bias, and drift components of relativistic corrections; and,

$\Delta\eta_j$ is the joint noise and multipath term, which includes carrier noise and multipath.

The TD in satellite/receiver range is expressed as follows:

$$\Delta r_j = \text{SV Doppler}_j - \Delta geometry_j - \left(\mathbf{e}_j(t_k), \Delta\mathbf{R}\right) \quad (2)$$

where:

SV Doppler_j is a change in the range due to the satellite motion along the line-of-sight (LOS);

$\Delta geometry_j$ accounts for changes in the relative satellite/receiver geometry;

$\mathbf{e}_j$ is the unit vector pointed from the receiver to the satellite, this vector is generally referred to as the LOS unit vector;

$\Delta\mathbf{R}$ is the receiver position change vector for the interval $[t_{k-1}, t_k]$; and,

(,) is the vector dot product.

Carrier phase TDs are adjusted for the satellite motion terms, geometry terms, and delta error terms prior to their exploitation as Kalman filter measurement observables. For the TD adjustment, satellite motion and geometry terms are computed as follows [4]:

$$\text{SV Doppler}_j = \left(\mathbf{e}_j(t_k), \mathbf{R}_{SVj}(t_k)\right) - \left(\mathbf{e}_j(t_{M-1}), \mathbf{R}_{SVj}(t_{k-1})\right) \quad (3)$$

$$\Delta geometry_j = \left(\mathbf{e}_j(t_k) - \mathbf{e}_j(t_{k-1}), \mathbf{R}(t_{k-1})\right) \quad (4)$$

In (3) and (4):

$\mathbf{R}_{SVj}$ is the satellite position vector; and,

$\mathbf{R}$ is the receiver position vector.

For geometry compensation, the receiver position vector $\mathbf{R}$ at the previous update (t_{M-1}) is estimated based on GNSS pseudorange measurements. For those cases where not enough pseudorange measurements are available, the position estimate is propagated using inertial data. Note that a sub-hundred-meter level accurate position estimate is generally sufficient to support mm-level accuracy in the carrier phase TDs [4]. The satellite position $\mathbf{R}_{SVj}$ vector is computed from ephemeris data, and the LOS unit vector $\mathbf{e}_j$ is computed based on ephemeris data and the pseudorange-based receiver position estimate. Tropospheric drift terms are compensated based on tropo models [4]. Iono delta errors are normally compensated using dual frequency measurements [4]. However, generally, iono drift terms stay at a mm/s level or less unless ionospheric scintillations are present [4]. Thus, for most operational scenarios, uncompensated iono drift does not significantly influence the accuracy of carrier phase TDs. For this reason, iono corrections were not implemented.

From equation (1) and (2), carrier phase TDs that are adjusted for the satellite motion, geometry changes, and delta error terms are expressed as follows:

$$\Delta\tilde{\varphi}_j^{adj}(t_k) = -\left(\mathbf{e}_j(t_k), \Delta\mathbf{R}\right) + \Delta\delta t_{rcvr} + \Delta\eta_j \quad (5)$$

Adjusted carrier phase observation is directly supported by the generic relative position observable of RIFE, which is formulated as a projection of position change on specified axes that are resolved in the navigation or body frame. This observation also contains the receiver drift term that is added to the RIFE state vector for GNSS carrier phase processing.

APPLICATION CASE STUDY

As a specific case study, the plug and play navigation technology is applied for automotive vehicle-to-vehicle (V2V) safety systems, whose adoption by the automotive mass market is expected to start within next three to five years [5]. Figure 4 shows a generalized architecture of a V2V safety system.

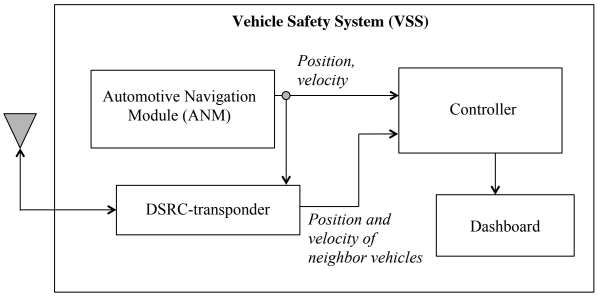

Figure 4. Generalized architecture of the V2V safety system

An automotive navigation module (ANM) performs real-time estimation of vehicle position and velocity states. Inter-vehicle exchange of navigation data is enabled by a transponder, which uses dedicated short-range communication (DSRC) radios to support the data exchange between neighboring vehicles (generally, within a 200-250 meter range). A controller compares the vehicle's own trajectory with trajectories of surrounding vehicles in order to predict potential collisions. Collision prediction results are used to generate audio, visual, and/or haptic warning signals to the driver.

Operating in an advisory mode, the V2V safety system does not perform automatic collision avoidance per se. Rather the final decision is still made by the driver. Yet, the system significantly increases the driver's situational awareness, which is instrumental in reducing the accident rate. Indeed, recent studies indicate that the advisory-mode cooperative safety principle would prevent 81% of accidents [6]. Examples of such accidents are illustrated in Figure 5.

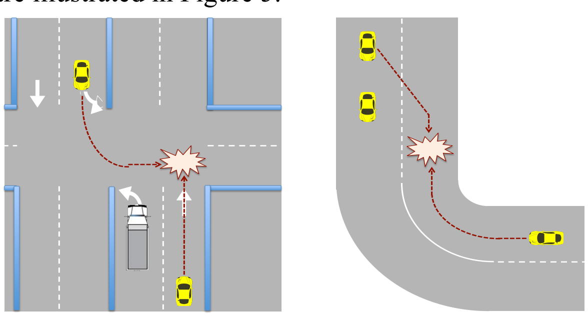

Left turn at an unregulated intersection when the visibility of the oncoming traffic is (partially) obstructed, which leads to a possibility of the side collision

Passing on a two-lane road with a limited visibility of the oncoming traffic, which creates a possibility of the head-on collision

Figure 5. Examples of accidents prevented by the V2V

For these examples, the V2V cooperative technology completely "sees" the oncoming traffic based on the inter-vehicle exchange of navigation data. The system predicts a possibility of collision and recommends a corresponding prevention action: engaging brakes for the first example and merging back to the lane for the second example, respectively.

Currently, one of the key challenges for such advanced V2V systems is the lack of navigation solutions that satisfy V2V accuracy requirements for all driving scenarios (i.e. open-sky, tree-covered roads, benign and dense urban) and meet the cost limitations of the automotive market. An automotive navigation module (ANM) that supports accurate navigation capabilities in real time is required to enable reliable prediction and prevention of traffic accidents. The currently anticipated accuracy requirement is for horizontal positioning in the 1.5-meter range at 95% confidence level. GNSS technology can satisfy this requirement in open-sky areas. However, in many signal-challenged environments, GNSS performance degrades rapidly and cannot support lane-level accuracy for reliable prediction and avoidance of automotive accidents.

To address this limitation, our PnP sensor fusion solution has been applied to low-cost sensors suitable for automotive market. The application specifically focuses on the use of sensors that are already installed in cars for other purposes such as odometer and video-camera. As better sensors become available and/or the sensor mix shifts due to technological advancements, the plug and play approach is able to rapidly incorporate these advancements to lower unit cost and improve performance. In addition, the same software package can be utilized for different vehicle models (from basic to luxury) from different automakers to optimally accommodate available sensors. Finally, the solution can be used for development and testing purposes to rapidly assess the influence of specific sensors on the localization accuracy (e.g., when new sensors become available).

EXPERIMENTAL DEMONSTRATION

For experimental demonstrations of PnP navigation capabilities, a ground vehicle test setup was developed and implemented. This setup is designed to include a variety of sensors and support development and demonstration efforts for various use cases and accuracy requirements from centimeter-level urban positioning for surveying and mapping to meter-level localization for automotive safety.

Current setup uses of-the-shelf sensors for initial demonstration. The sensor quality is similar to automotive sensors. Off-the-shelf sensors were utilized to accelerate the interface development. Future efforts will address PnP demonstrations with automotive sensors that are installed in the vehicle by manufacture.

Figure 6 shows the ground vehicle test setup that includes GPS receivers (survey-grade Novatel and consumer-grade NVS receivers with pinwheel and patch antennas); a higher-grade MEMS IMU (~10 deg/hr, STIM-300 manufactured by Sensonor); a PX4 autopilot with consumer-grade IMU (~100 deg/hr unit manufactured by ST Micro), magnetometer and baro-altimeter; Prosilica video cameras, SICK LMS-200 scanning lidar, Microsemi chip-scale atomic clock (CSAC); data synchronization and data collection units. The physical layout of the data collection system has been designed to support the real-time demonstration objective. Figure 7 shows the annotated sensor board. The data collection and processing system has been palletized and mounted in a transportable 19"-6U equipment case as shown in Figure 8. This test setup was utilized to evaluate various low-cost solutions for automotive safety applications.

Figure 6. Ground vehicle test setup for demonstration of plug-and-play navigation capabilities

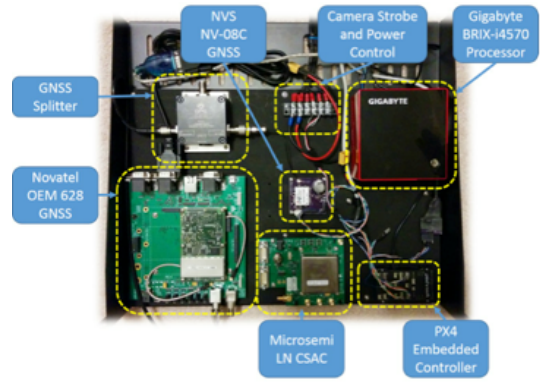

Figure 7. Annotated sensor board

Figure 8. Portable data collection system: the system is designed to mount in a standard 19" rack.

For experimental demonstrations, experimental data were collected in urban canyons of downtown San Francisco, CA in January 2016. Figure 9 illustrates typical test environments.

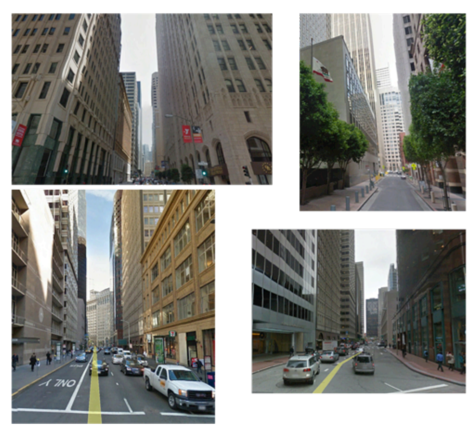

Figure 9. Example test environments; Google Earth and Street View mode were used to obtain the photographs

First, we evaluated GNSS-only performance. As expected, GNSS positioning capabilities were found to be extremely limited. Figure 10 shows GNSS (GPS+GLONASS) position solution displayed in Google Earth. Sparse and unreliable position fixes are obtained, which is clearly unsatisfactory for automotive safety applications.

Figure 10. Typical performance of GNSS position solution in downtown environments

Next, PnP navigation software was automatically reconfigured to evaluate GNSS/INS integration with consumer-grade inertial. Carrier phase measurements were processed using temporal phase differences (to eliminate integer ambiguities) and relative position observables of RIFE. Figures 11 and 12 show example test results. Integration of GNSS with inertial significantly improves solution availability and quality as compared to the GNSS-only option. Yet, in dense urban canyons it does not allow for lane-level positioning. In fact, position errors can grow to the level of 20 meters when the vehicle remains in urban canyon over a significant period of time (such as 5 minutes) and the GNSS signal availability is limited to one or two visible satellites.

Finally, integration of carrier phase GNSS, monocular video-camera and vehicular motion model (velocity constraints) was evaluated. The PnP software was automatically reconfigured for this sensor configuration using generic relative position observables (for GNSS carrier phase); velocity observables (for motion model); and, relative bearing observations for monocular video.

Figure 11. Position solution of carrier phase GNSS integrated with consumer-grade inertial for example test scenario 1: continuous trajectory reconstruction is obtained; however, in dense urban canyons significant deviations can be present (as indicated in the zoomed image in the right-hand side)

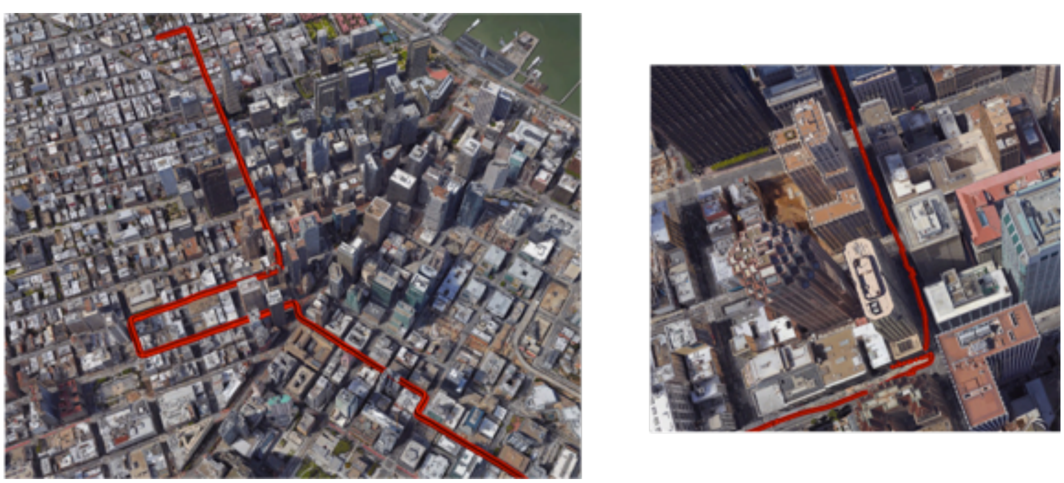

Figure 12. Position solution of carrier phase GNSS integrated with consumer-grade inertial for example test scenario 2

Figures 14 and 15 show example test results. Consistent and reliable positioning in urban canyons is demonstrated. It is noted that visual inspection of the position solution was performed. First, trajectory was represented in Google Earth. Next, we manually verified that it corresponds to the correct lane driven. Future efforts will compare consumer-grade position solution against reference trajectory obtained with higher-grade sensors in order to quantify the navigation performance.

CONSLUSION

In this paper, we applied the reconfigurable integration filtering engine (RIFE) for accurate positioning in GNSS-degraded environments. Particularly, experimental test results demonstrate that RIFE can be automatically reconfigured for low-cost sensor configurations in order to enable accurate positioning in urban canyons thus supporting accuracy requirements of connected cars for automotive safety. Ongoing and future efforts include (i) RIFE demonstration with automotive sensors, (ii) configuration of RIFE using higher-quality sensors for high-precision applications (centimeter to decimeter accuracy

Figure 13. Position solution of carrier phase GNSS integrated with consumer grade inertial, video-camera and motion constrains for test scenario 1; reliable trajectory reconstruction is achieved for the entire duration of the test

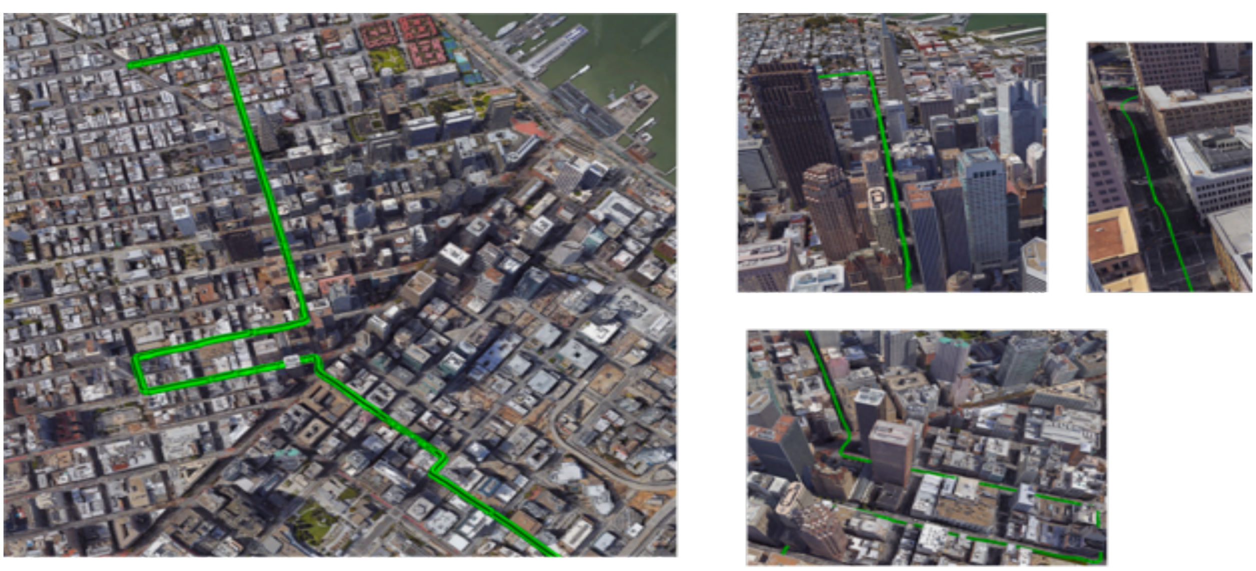

Figure 14. Position solution of carrier phase GNSS integrated with consumer grade inertial, video-camera and motion constrains for test scenario 2; reliable trajectory reconstruction is achieved for the entire duration of the test

level) such as surveying, mapping and inspection in urban environments and forestry areas; (iii) application for unmanned aerial vehicles; and, (iv) conversion of post-processing software modules into the real-time operation mode.

REFERENCES

[1] Groves, P.D., Wang, L., Walter, D., Martin, H., Voutsis, K., Jiang, Z., "The Four Key Challenges of Advanced Multisensor Navigation and Positioning," Proceedings of IEEE/ION PLANS 2014, Monterey, CA, May 2014, pp. 773-792.

[2] Soloviev, A., Yang, C., "Reconfigurable Integration Filter Engine (RIFE) for Plug-and-Play Navigation," Proceedings of the 26th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2013), Nashville, TN, September 2013, pp. 2075-2083.

[3] J. L. Farrell, "GPS/INS-Streamlined," NAVIGATION, Journal of the Institute of Navigation, Vol. 49, No. 2, 2002.

[4] F. van Graas, A. Soloviev, "Precise Velocity Estimation Using a Stand-Alone GPS Receiver," NAVIGATION, Journal of the Institute of Navigation, Vol. 51 No. 4, 2004.

[5] "U.S. Department of Transportation Announces Decision to Move Forward with Vehicle-to-Vehicle Communication Technology for Light Vehicles," DOT press release, February 2014.

[6] U.S. Department of Transportation, 2010: "Frequency of Target Crashes for IntelliDrive Safety Systems".

LOW-COST GNSS/INS INTEGRATION
Conquers Harsh Environments

A software-driven navigation engine makes consistent, reliable navigation possible in tunnels, garages and urban canyons.

In difficult GNSS signal environments for driving—and here urban canyons, tunnels and parking structures are the standouts—GNSS performance may be severely degraded or completely denied. Inertial aiding has become the standard for ground vehicle navigation. Requirements for autonomous navigation in these circumstances will be rigorous, but those for map-matching, telematics and fleet vehicle tracking are much less so. Simply a consistent and reliable positioning on the meter level, capable of withstanding 5 to 10-minute outages, will answer a growing need. Low cost, ease of installation and self-containment, with no need to connect to the vehicle's CAN bus, are the key drivers in this market.

One such solution is the GNSS/Inertial Vehicular Engine (GIVE) for automotive navigation. Using consumer-grade, relatively low-cost GNSS chipsets with microelectromechanical systems (MEMS) inertial sensors of cell-phone quality, it enables robust navigation capabilities. Its strength is in the sensor-agnostic software, which is customizable to the implementation platform. Tight coupling of carrier-phase GNSS with inertial measurements and a vehicle motion model and software-based multipath mitigation yield consistent position, velocity and attitude data in challenging environments.

The navigation engine fuses GNSS carrier-phase measurements and pseudoranges with IMU measurements and a vehicular motion model featuring non-holonomic constraints and zero-velocity updates. Non-holonomic means it is implicitly dependent upon parameters such as, in this case, constraints in the velocity domain of not leaving the Earth's surface or moving suddenly to the side, as is consistent with ground vehicle behavior. The zero-velocity updates detect when the car stops, and use that for updating the inertial position, thus halting inertial drift.

Altogether, the system can coast for a relatively long time—5 to 10 minutes or longer—without GNSS data, and without expensive hardware.

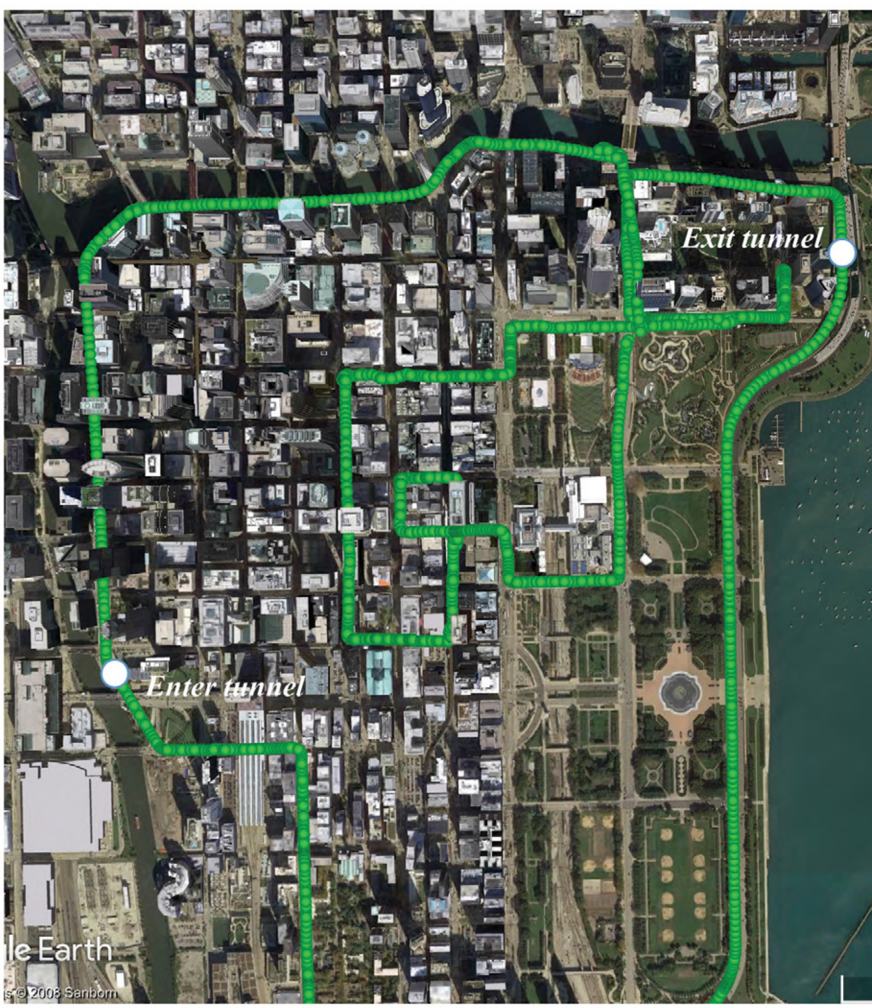

FIGURE 1 Consistent navigation performance in long tunnel, Lower Wacker Drive, Chicago.

Multipath Mitigation

When GNSS and inertial technologies are combined for challenging environments, identification of GNSS measurement outliers, mainly caused by multipath, and excluding them from the data fusion becomes key to the process. GIVE includes algorithms to detect and exclude bad measurements, thus negating multipath effects. Multipath mitigation in the engine has two levels of protection:

- INS-based statistical gating: residual screening of the tightly coupled Kalman filter
- Probabilistic data association filtering (PDAF): adaptive weighting that incorporates probability of missed-detection

Measurement quality control can be accomplished most efficiently by predicting GNSS measurements values based on the inertial solution; comparing predicted and actual measurements; and, then discarding measurements with large discrepancies. GIVE applies INS-based statistical gating to eliminate those GNSS measurements that do not agree with their values predicted based on inertial outputs.

THE SYSTEM CAN COAST FOR A RELATIVELY LONG TIME—5 TO 10 MINUTES OR LONGER—WITHOUT GNSS DATA AND WITHOUT EXPENSIVE HARDWARE.

PDAF further mitigates the influence of outliers that pass through the residual check. PDAF modifies the Kalman Filter estimation step by directly incorporating a probability of outlier missed-detection into the measurement/prediction weighting scheme. The use of PDAF proves particularly beneficial for "transitioning" cases, from GNSS-denied to GNSS-challenged, such as exiting from a tunnel into an urban canyon. For such cases, residual sigma bounds are increased due to unmitigated inertial drift, and multipath errors can leak through leading to detrimental effects on the navigation performance. De-weighting of potentially corrupted measurements based on their respective missed-detection probabilities significantly improves the robustness of the multipath mitigation performance.

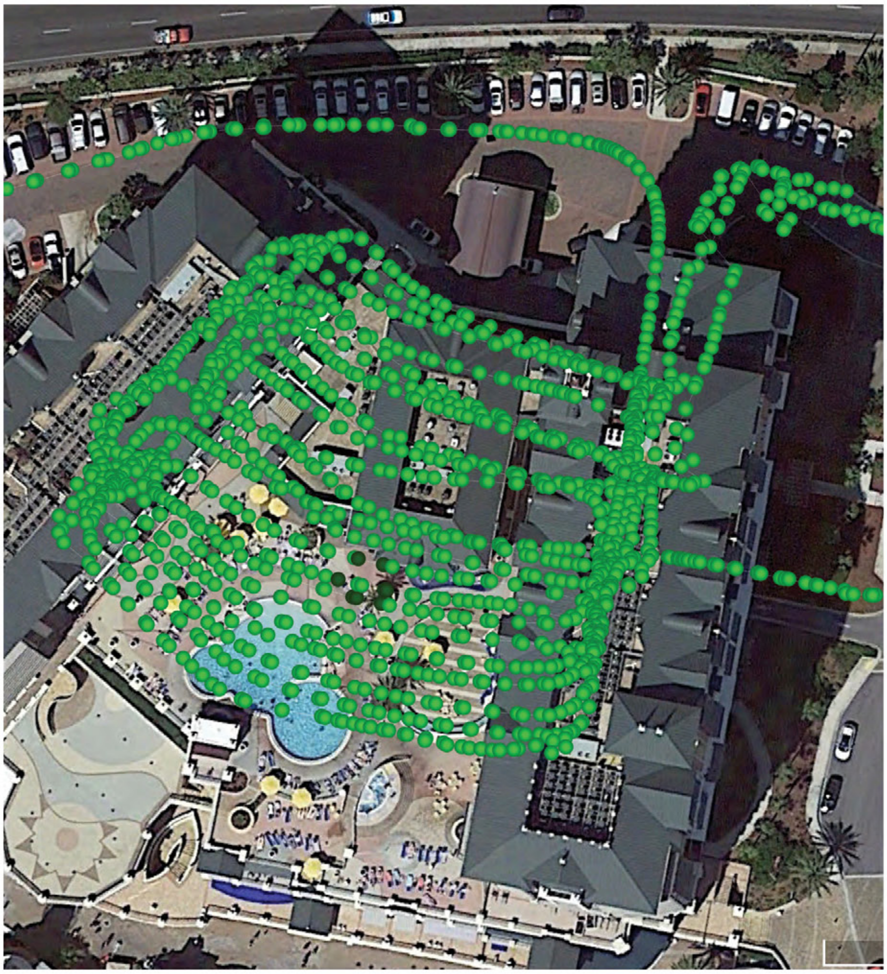

FIGURE 2 Loops inside a completely enclosed parking garage; GNSS outage for 5 minutes of each loop.

Carrier-Phase Processing

The engine also takes advantage of temporal carrier-phase differences over time. Temporal differencing eliminates integer ambiguities while capturing the underlying motion dynamic for the INS error-state estimation. The use of temporal differences is particularly beneficial for the identification and removal of close-range non-line-of-sight (NLOS) multipath errors. It may be challenging to identify multipath signals reflected by buildings within a close range in the range domain. Yet, significant Doppler differences are generally present, due to the difference in LOS between direct and multipath signals, which makes multipath errors readily distinguishable in the temporal phase domain.

Extensive test results in difficult urban environments in Chicago, San Francisco, Atlanta metropolitan centers and elsewhere have demonstrated consistent navigation performance in harsh GNSS environments. **Figures 1 and 2** show some of the results of those tests.

GIVE can be installed as an after-market device on almost any vehicle; it can also be ported onto the user's smartphone.

GIVE is a product of QuNav, founded by Andrey Soloviev and based in Fort Walton Beach, Florida. The technology was showcased at ION GNSS+ 2019.