Integrating GLONASS with GPS for Drone Orientation Tracking

Mahanth Gowda[1(✉)], Justin Manweiler[2], Ashutosh Dhekne[1], Romit Roy Choudhury[1], and Justin D. Weisz[2]

[1] University of Illinois at Urbana-Champaign, Urbana, IL, USA
gowda2@illinois.edu
[2] IBM Research, Yorktown Heights, NY, USA

Abstract. In addition to position sensing, GPS receivers can be leveraged for orientation sensing too. We place multiple GPS receivers on drones and translate their relative positions into orientation. Such an orthogonal mode of orientation sensing provides failsafe under Inertial sensor failures – a primary cause of drone crashes today. This paper integrates GLONASS satellite measurements with GPS for enhancing the orientation accuracy.

Accurate estimate of orientation depends upon high precision relative positioning of the GPS receivers. While GPS carrier phases provide high precision ranging data, the phases are noisy and wrap after every wavelength which introduces ambiguity. Moreover, GPS signals experience poor SNR and loss of satellite locks under aggressive flights. This can severely limit both the accuracy and the amount of carrier phase data available. Fortunately, integrating the ubiquitously available Russian GLONASS satellites with GPS can double the amount of observations and substantially improve the robustness of orientation estimates. However, the fusion is non-trivial because of the operational difference between FDMA based GLONASS and CDMA based GPS. This paper proposes a temporal differencing scheme for fusion of GLONASS and GPS measurements, through a system called *SafetyNet*. Results from 11 sessions of 5–7 min flights report median orientation accuracies of 2° even under overcast weather conditions.

1 Introduction

Multiple GPS receivers can be used for tracking drone orientation. A rigid body is formed by placing 3 or more GPS receivers on the body of a drone. The orientation of the drone is a simple function of relative positions of the GPS receivers on the rigid body. While inertial sensors (IMU) like accelerometers, gyroscopes and compasses are conventionally used for orientation tracking, they can fail – leading to crashes [1,2,4,8,33]. Many of the failures are correlated from common sources of engine vibration and on-board electromagnetic interference [18,21,22,26,31]. Hence, it is not possible to cope with such failures by having redundant sensors. Since GPS provides a completely orthogonal modality of orientation sensing, it can serve as a failsafe for IMU failures.

N. Sastry and S. Chakraborty (Eds.): COMSNETS 2017, LNCS 10340, pp. 77–92, 2017.
DOI: 10.1007/978-3-319-67235-9_7

Achieving an orientation accuracy of 2° with GPS requires the relative positions between the GPS receivers to be known to an accuracy of 2 cm. Moreover, the accuracy needs to be upheld constantly over time, especially during periods of aggressive maneuvers, where the SNR is poor, locks with satellites break and other errors manifest. This is a challenging task since a basic GPS location is only 3 m accurate.

Even the best differential GPS technique only provides a relative positioning accuracy of 15 cm [14,15]. They leverage carrier phase data, which is a precise estimate of range from satellites, but polluted by ambiguity due to phase wrapping. Our recent Mobicom paper [10] shows how carrier phase data can be differenced over time and space and fused using an augmented particle filter framework. This results in accurate orientation estimate with a median accuracy of 2° and 95^{th} percentile accuracy of 8°.

In this paper, we expand the GPS/GLONASS integration component of SafetyNet, which is a critical factor in achieving a high accuracy. Under aggressive flights, the carrier phase data is not only noisy, but also sparse in number because of poor SNR and loss of satellite locks. Fortunately, integrating the ubiquitously available GLONASS (GLObal NAvigation Satellite System – [19]) satellite measurements with GPS doubles the number of measurements and enhances the robustness of orientation estimates. However, the fusion is non trivial since GLONASS operates on FDMA whereas GPS uses CDMA.

In the rest of the paper, we first provide a background of GPS, GLONASS and the basics of satellite positioning systems. We then revisit the math behind orientation sensing from [10]. We derive orientation as a function of various differences of carrier phases over, receivers, satellites and time. Finally, we integrate GLONASS measurements into our model. Evaluation results show substantial gains from the integration, especially in eliminating errors in the tail.

2 GPS Foundations

We present GPS foundations from first principles and end with discussions on modern techniques. As a result, this section is long. However, given that this paper builds over core GPS algorithms, the material is necessary. We also believe the material is easy to follow.

2.1 Global Navigation Satellite System (GNSS)

GNSS is the generic name given to satellite systems that provide localization services to receiver's on earth. The Global Positioning System (GPS) [24] is one example of a GNSS, developed by the US Government during 1970–1980s. GPS consists of a constellation of 31 satellites orbiting the earth at a height of 20,000 km. The satellites are continuously transmitting unique pseudo-random noise (PRN) sequences using CDMA at 2 different frequencies –1575.42 (L1) and 1227.60 MHz (L2). They also broadcast ephemeris data using which the (satellite) position and time of transmission can be calculated. A GPS receiver

on the ground localizes itself by trilateration, i.e., measuring and combining the time-of-flight (ToF) from different satellites. Velocity is computed from the doppler shifts from various satellites.

GLObal NAvigation Satellite System (GLONASS) is another GNSS system launched by Russia in the 1980's [19]. GLONASS has 24 operational satellites. Unlike GPS which uses CDMA for multiplexing, GLONASS uses a 15 channel FDMA with a center frequency of 1602 MHz and an inter-channel separation of 0.5625 MHz. Since, there are only 15 channels, but 24 operational satellites, those satellites that are antipodal (on opposite side of earth), reuse the same frequency, such that they are not simultaneously visible to a GPS receiver. Similarly, GALILEO [3] is an European GNSS system currently under development. Many GNSS receivers are capable of decoding signals from multiple GNSS systems, providing increased accuracy. This paper will also use such receivers and exploit the advantages of satellite diversity.

2.2 GPS Localization and Error Sources

A GNSS receiver on the ground can compute its 3D location, time, and velocity. The key idea is to measure various attributes of the arriving signal (e.g., time of flight, phase, etc.), and then apply statistical algorithms to estimate the errors and ambiguities in measurements. We give an overview of the techniques that underpin GPS; other GNSS systems are similar.

Pseudorange

When a satellite transmits a signal, it includes the starting time of the transmission (obtained from its atomic clock). The ground receiver records the time of reception also using its less accurate local clock. The time-of-flight (ToF) is the difference between these timestamps. When multiplied by the speed of light, the result gives the *rough* distance to the satellite, called *pseudorange*.

$$\text{Pseudorange} = \text{ToF} * (\text{speed of light})$$

Of course, the measured ToF is inaccurate because the clock of a typical GPS receiver is not synchronized to the GPS satellites. The resulting error can be up to 300 km. In addition, when the GPS signal enters the Earth's atmosphere, it can get delayed due to refractions in the Ionosphere and Troposphere. A signal also passes through a multipath channel, adding more errors. Assuming the true range between a satellite s and receiver i is ρ_i^s, the measured Pseudorange P_i^s, inclusive of all error sources, can be modeled as

$$P_i^s = \rho_i^s + ct_i - ct^s + A + M_i + \epsilon_i^s \tag{1}$$

Here, t_i and t^s are receiver and satellite clock biases, respectively, with respect to true time. A represents the range error due to refractions in the atmosphere. M_i^s denotes Multipath, ϵ_i^s is receiver's hardware noise and c is the speed of light. In today's systems, the satellite clock error t^s is small[1], and ct_i proves to be the

[1] Satellites estimate the errors themselves from mutually exchanged signals as well as from ground sources.

major source of error. Hence, ct_i is modeled as an unknown, and ρ_i^s is written as a function of the *unknown* 3D receiver location ρ_i and the *known* 3D satellite position ρ^s. This results in a total of 4 unknowns. Once we have a lock with 4 satellites, the time bias and the 3D locations can be jointly estimated resulting in a position fix. The ignored error sources, i.e., t^s, A, M, and ϵ, could contribute to an error of 1–4 m, depending upon the environmental conditions.

GPS receivers on phones and car dashboards use the above techniques. However, higher accuracy applications such as 3D orientation tracking require better performance. To this end, modern GPS research has leveraged the *phase* of the arriving signals, as detailed next.

Carrier Phase

Once a satellite lock is acquired, the phase of the arriving signal, ϕ_i^s, is constantly tracked by a *phase lock loop* (PLL). The true range between the satellite and the receiver, ρ_i^s, can be expressed in terms of wavelength λ.

$$\rho_i^s = \lambda N_i^s + \lambda \phi_i^s \tag{2}$$

N_i^s is an unknown integer, meaning that the PLL measurement of ϕ_i^s only captures the fractional part of the range. However, due to atmospheric effects, multipath, and clock issues, the above equation can be updated:

$$\lambda \phi_i^s = \rho_i^s + ct_i - ct^s + A + M_i + \epsilon_i - \lambda N_i^s \tag{3}$$

Estimating N_i^s is non-trivial and several algorithms exist [20,32].

One advantage of carrier phase is that *its changes over time* can be tracked reliably by utilizing the doppler shift in the signal [16]. Hence, $\phi_i^s(t_2)$ takes the same mathematical form as Eq. 3. Thus, if the initial value of N_i^s can somehow be estimated, the tracking thereafter can be good. We now explain how today's systems like Differential GPS (DGPS) with this integer ambiguity, N_i^s, and other error sources.

2.3 Computing Differentials

Environmental error sources in Eq. 3 are correlated over short time periods and within small geographical areas (200 km). Thus, two GPS measurements across time can be subtracted (or differenced) to eliminate some of these factors. Similarly, simultaneous measurements from multiple GPS receivers can also be differenced. Differential GPS (DGPS) [25] performs such operations on *pseudoranges*, while Real-Time Kinematic (RTK) [23] applies differentials to *carrier phase*. For our purpose of precise orientation tracking, the latter is more relevant. To this end, we outline 4 kinds of carrier phase differentials.

(1) Single Differentials Across Receivers (SD_{ij})

Consider the carrier phase equations for two GPS receivers i and j from the same satellite s. We ignore multipath and noise in the rest of the equations.

$$\lambda\phi_i^s = \rho_i^s + ct_i - ct^s + A^s - \lambda N_i^s \tag{4}$$
$$\lambda\phi_j^s = \rho_j^s + ct_j - ct^s + A^s - \lambda N_j^s \tag{5}$$

Differencing the above two equations yields the relative position between i and j with fewer error terms. Correlated error sources of atmospheric delays and satellite clock biases disappear.

$$\lambda\Delta\phi_{ij}^s = \Delta\rho_{ij}^s + c\Delta t_{ij} - \lambda\Delta N_{ij}^s \tag{6}$$

Figure 1 illustrates the scenario. Assuming that the satellite is far away, $\Delta\rho_{ij}^s$ can be approximated as $\rho_{ij}Cos\theta$, where ρ_{ij} is the true relative position between i and j (called **baseline** vector). Replacing $\rho_{ij}Cos\theta$ as a vectorial projection of ρ_{ij} on to the line of sight unit-vector $\hat{l_s}$ of satellite s, we have:

$$\lambda\Delta\phi_{ij}^s = \rho_{ij}.\hat{l_s} + c\Delta t_{ij} - \lambda\Delta N_{ij}^s \tag{7}$$

While some errors have disappeared[2], a function of the clock bias errors, $c\Delta t_{ij}$ still remains, motivating the need for double differentials.

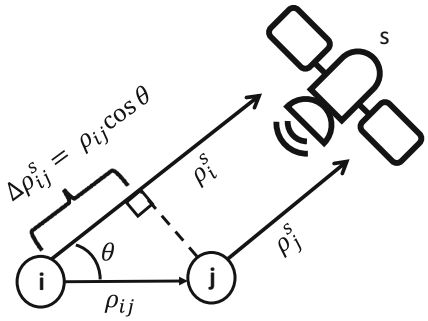

Fig. 1. $\rho_i^s - \rho_j^s = \Delta\rho_{ij}^s = \rho_{ij}.\hat{l_s}$

(2) Double Differentials Across Receivers and Satellites (DD_{ij}^{sk})

Given multiple satellites in range, the GPS ground receivers i and j can perform the same measurements with satellite k. Equation 7 can then be rewritten as:

$$\lambda\Delta\phi_{ij}^k = \rho_{ij}.\hat{l_k} + c\Delta t_{ij} - \lambda\Delta N_{ij}^k \tag{8}$$

[2] We are aware that differentials can amplify noise and multipath, however, since carrier phase noise is in the granularity of few mm [7], we ignore noise in the rest of the paper. We also assume multipath is not excessive, such as the drone flying low in Manhattan-like areas.

Subtracting the single differential Eqs. 7 and 8, we have a double differential (DD) as follows:

$$\lambda \nabla \Delta \phi_{ij}^{sk} = \rho_{ij}.(\hat{l_s} - \hat{l_k}) - \lambda \nabla \Delta N_{ij}^{sk} \tag{9}$$

The double differential (DD) eliminates the clock biases and the residue is only the integer ambiguity terms. More details on integer ambiguity resolution is done in [10], but assuming that the ambiguities are magically fixed, this provides us a reasonably precise estimate of relative positions (called baselines) between receiver pairs. Ignored factors, including multipath, noise, and antenna phase center errors, add up to a few centimeters of error. Thus, if one of the receiver's absolute position is known in the granularity of millimeters, the absolute position of the other receiver can be estimated precisely as well. Real Time Kinematics (RTK) technology operates exactly as above – *it uses the accurately known location of the reference receiver to calculate the location of the other.* Figure 2 shows the relative distance between two receivers placed roughly 45 cm apart. Differencing techniques convincingly outperform naive subtraction of 3D GPS positions.

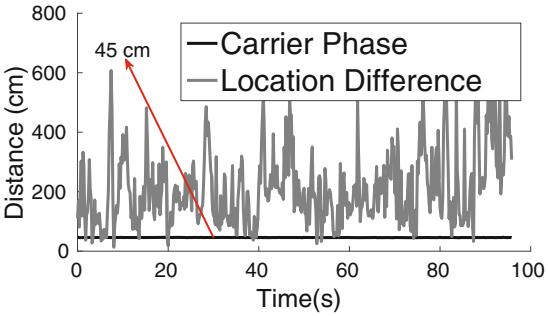

Fig. 2. Relative positioning using carrier phase and double differentials is an order of magnitude more accurate than naïve location differencing. Here, a 45 cm baseline is correct within a few cm.

(3) Double Differentials Across Receivers and Time ($DD_{ij}^{t_{12}}$)

Our final double differential combines receivers and time. We compute the single differential between receivers ij from Eq. 7 but write them for consecutive time points t_1 and t_2.

$$\lambda \Delta \phi_{ij}^{s}(t_1) = \rho_{ij}(t_1).\hat{l_s} + c\Delta t_{ij}(t_1) + \lambda \Delta N_{ij}^{s}(t_1) \tag{10}$$

$$\lambda \Delta \phi_{ij}^{s}(t_2) = \rho_{ij}(t_2).\hat{l_s} + c\Delta t_{ij}(t_2) + \lambda \Delta N_{ij}^{s}(t_2) \tag{11}$$

Assuming no cycle slips (we will relax this assumption later), $\Delta N_{ij}^{s}(t_2) = \Delta N_{ij}^{s}(t_1)$, hence subtracting the above two equations eliminates integer ambiguity:

$$\lambda \nabla \Delta \phi_{ij}^{s}(t_{12}) = (\rho_{ij}(t_1) - \rho_{ij}(t_2)).\hat{l_s} + c.\Delta t_{ij}(t_{12}) \tag{12}$$

One may physically interpret this equation as the subtraction of two vectors, where the first vector is a drone baseline at time t_1 and the second vector is the same baseline at t_2. In other words, this captures the relative motion of the drone across time.

2.4 The Bigger Picture

We take-away 2 key points from the discussion on the differentials:

- The double differentials across receivers and satellites (DD_{ij}^{sk}) yields the drone's baseline vectors at any given time point (Eq. 9). However, this estimate is still polluted by integer ambiguity.
- The double differential across receivers and time ($DD_{ij}^{t_{12}}$) yields the drone's *relative change* in the baseline vectors during flight (Eq. 12). Importantly, this relative estimate is free of the integer ambiguities.

Thus, we now have *two* separate estimates of the drone's 3D baseline vectors, each with different error properties. SafetyNet recognizes the opportunity of combining these two noisy estimates to precisely track the drone baselines, ultimately tracking orientation.

3 System Model

SafetyNet will model orientation tracking as a state estimation problem, where the state is defined as 3D orientation. We formally define "3D orientation" first and then design the model.

Figure 3 pretends 4 GPS receivers have been placed on a drone – their locations denoted as ρ_1, ρ_2, ρ_3 and ρ_4. The baseline vectors joining one of the receivers (say ρ_4) to the others can be defined as $\rho_{ij} = \rho_i - \rho_j$. When the "baselines" are aligned with the North-East reference axes, the baseline matrix B_o can be written as:

$$B_o = [\rho_{41}\ \rho_{42}\ \rho_{43}] \tag{13}$$

Assuming that the magnitude of the baseline vectors are d_1 and d_2, we can expand B_o as:

$$B_o = \begin{bmatrix} d_1 & 0 & d_1 \\ 0 & d_2 & d_2 \\ 0 & 0 & 0 \end{bmatrix} \tag{14}$$

Figure 3 also shows the rotation conventions of the 3 Euler angles – *pitch, roll and yaw*. Applying these rotations on B_o will obviously yield the new baseline matrix, B. For all our results, we will express rotations in terms of the Euler angles (i.e., degrees), which are intuitive to understand. However, for the purpose of mathematical efficiency, we will use *quaternion mathematics*, an alternative representation to Euler angles. Briefly, the baselines at an arbitrary orientation,

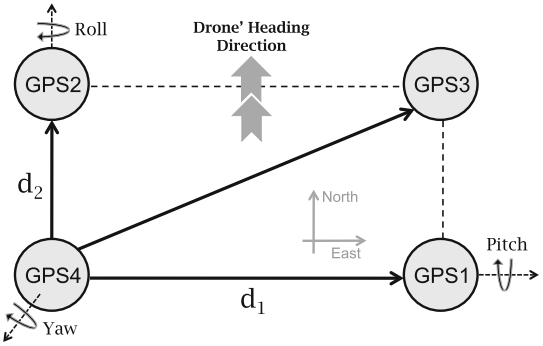

Fig. 3. The drone baseline vectors ρ_{41} and ρ_{42} aligned with Earth's reference frame.

called quaternion q, can be expressed in terms of the initial orientation q_o (aligned with reference axes), as below:

$$\rho_{ij}(q) = A(q)'\rho_{ij}(q_o) \tag{15}$$

Here $A(q)$ is the *rotation matrix* associated with the orientation quaternion q. Similarly, extending the effect of rotations to the entire baseline matrix, we have:

$$B(q) = A(q)'B_o \tag{16}$$

Effectively, orientation is about estimating the rotation matrix $A(q)$ using projected measurements of $B(q)$ on various satellite directions. Details on conversion between quaternions, Euler angles and rotation matrices can be found in [9]. Regardless of these mathematical conversions, *the core conceptual question still pertains to estimating the matrix B at any given time.*

4 System Design: Phase 1

We adopt a Bayesian filtering approach for tracking drone orientation. Figure 4 shows the model: (1) A state transition function, derived from the incremental changes in orientation ($DD_{ij}^{t_{12}}$), models the next state of the drone. Recall that these changes are affected by the hardware noise and multipath errors. (2) A measurement function, (DD_{ij}^{sk}), reflects the absolute orientation of the drone at any given time. Of course, this measurement is polluted by integer ambiguity. We adopt a Kalman Filter to combine the uncertainties from the transition and measurement functions, and track the most likely state of the system through time. We describe this basic design first. Then we focus on resolving the error sources (such as integer ambiguity, cycle slips, and missing data), and redesign the framework to accommodate these optimizations. Our final design is an "adjusted" particle filter algorithm that tracks drone orientation with consistent accuracy.

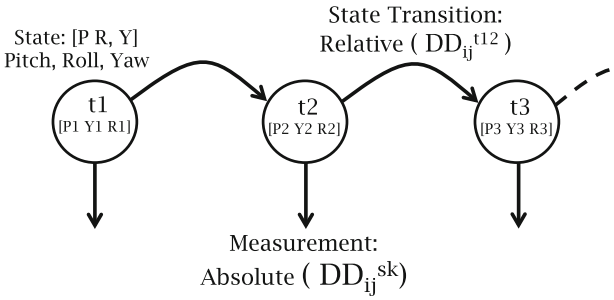

Fig. 4. Bayesian filtering approach to tracking orientation state over time.

4.1 State Transition Model

The relative baseline changes over time are directly obtained from Eq. 12, copied for convenience:

$$\lambda \nabla \Delta \phi_{ij}^{s}(t_{12}) = (\rho_{ij}(t_1) - \rho_{ij}(t_2)).\hat{l_s} + c.\Delta t_{ij}(t_{12}) \tag{17}$$

Omitting details, we rewrite with quaternions:

$$\lambda \nabla \Delta \phi_{ij}^{s}(t_{12}) = \rho_{ij}(q_o)\lfloor A(q_1).\hat{l_s}_{\times} \rfloor \delta\theta + c.\Delta t_{ij}(t_{12}) \tag{18}$$

where, q_1 is the orientation quaternion at time t_1, $\delta\theta$ is the rotation vector [5] associated with quaternion δq. $\lfloor_{\times}\rfloor$ is the vector cross operator [6]. We now solve Eq. 18 for various satellites s and GPS receiver-pairs ij using least-squares estimation. The result yields an estimate of the rotation vector $\delta\theta$ (hence δq) between two time points t_1 and t_2. We can thus estimate the new orientation quaternion q_2 as:

$$q_2 = \delta q \otimes q_1 \tag{19}$$

Here, $\otimes$ is the quaternion multiplication operator. When translated back to Euler angles, the result is the relative orientation change – pitch, yaw, and roll – from one state to the next. As mentioned earlier, this estimate is polluted by hardware noise and multipath.

4.2 Absolute Orientation Measurement

Equation 9 showed that double differentials across receivers and satellites (DD_{ij}^{sk}) are estimates of the *absolute* baseline vectors of the drone. While they cancel out clock bias errors, they leave the integer ambiguities as follows:

$$\lambda \nabla \Delta \phi_{ij}^{sk} = \rho_{ij}.(\hat{l_s} - \hat{l_k}) + \lambda \nabla \Delta N_{ij}^{sk} \tag{20}$$

Translating to quaternions again, using very similar conventions as described earlier, we have:

$$\begin{aligned} \lambda \nabla \Delta \phi_{ij}^{sk} - \rho_{ij}(q_o).A(q_n).(\hat{l_s} - \hat{l_k}) = \\ \rho_{i,j}(q_o).\lfloor A(q_n).(\hat{l_s} - \hat{l_k})_{\times} \rfloor \delta\theta + \lambda.\nabla \Delta N_{ij}^{sk} \end{aligned} \tag{21}$$

For now, let's assume the integer ambiguities are resolved (details in [10]) – then, we are left with a set of linear equations over different satellite pairs sk and baselines ij. By solving them using standard Least Squares Estimation (LSE), the rotation vector $\delta\theta$ and associated quaternion δq can be obtained. Hence the orientation q is.

$$q = \delta q \otimes q_n \tag{22}$$

Here q_n is an initial orientation estimate for the purposes of linearization of Eq. 21 (usually comes from the transition model). Note that this estimate q is absolute in the earth's reference frame, since the satellite locations sk are both known in that reference frame.

4.3 Filtered Orientation Estimation

We combine the state transition and absolute measurement models from the previous two subsections using an Augmented Particle Filter (APF) framework. Details are elaborated in [10].

4.4 Boosting Transition Function Reliability: GLONASS

Unlike GPS satellites, GLONASS uses FDMA and the satellites transmit in different frequencies. Hence, the *Double Differentials across Receivers and Satellites* (DD_{ij}^{sk}) equations of GLONASS have integer ambiguities that are more complex to resolve compared to GPS. Let us rewrite Eqs. 7 and 8 for two GLONASS satellites s and k

$$\lambda^s \Delta \phi_{ij}^s = \Delta \rho_{ij}.\hat{l_s} + c\Delta t_{ij} + \lambda^s \Delta N_{ij}^s \tag{23}$$

$$\lambda^k \Delta \phi_{ij}^k = \Delta \rho_{ij}.\hat{l_k} + c\Delta t_{ij} + \lambda^k \Delta N_{ij}^k \tag{24}$$

Here, λ^s and λ^k are wavelengths of two GLONASS satellites. Differencing the two equations above similar to the derivation of Eq. 9 results in

$$\begin{aligned} \lambda^s \Delta \phi_{ij}^s - \lambda^k \Delta \phi_{ij}^k = \\ \Delta \rho_{ij}.(\hat{l_s} - \hat{l_k}) + \lambda^s \Delta N_{ij}^s - \lambda^k \Delta N_{ij}^k \end{aligned} \tag{25}$$

The ambiguity terms do not group themselves into a single integer. This makes it difficult to use the GLONASS DD_{ij}^{sk} equation above towards derivation of absolute drone orientation, as outlined in Sect. 4.2.

As discussed in Sect. 6, it might be possible to use finer integer ambiguity resolution effectively in both integers, but it would be expensive. Alternatively, it might be possible to use two satellites with adjacent channels and assume a common frequency, but it would be inaccurate for our purposes. Hence, we take a completely alternative approach and use temporal differencing instead. Equation 12 in Sect. 2 has only one term with a wavelength factor. It gives relative changes in the geometry over time, which can be easily exploited for computing the transition function in Sect. 4.1. This will not only add more confidence to the transition function, but also helps in quickly converging on the set of outliers caused due to cycle slips.

5 Evaluation

We conduct all experimentation using a 3DR X-8 octocopter (8-rotor), pictured in Fig. 5. "X" implies 4 arms with unequal spacing – yielding greater stability of roll angle versus pitch (recall that pitch is the dominant motion when an airplane takes off or lands, roll is dominant when the plane makes a left/right turn). Two rotors attach to each arm: one above, one below. To the X-8, we mounted 4 u-Blox NEO-M8T multi-GNSS receivers, adjacent to the top motor of each arm. We also mounted a Raspberry Pi 2 near the drone's center of mass – the GPS receivers transfer the data to the Pi via USB at 5 Hz. IMU data at 10 Hz is recovered as binary logs from the X-8's USB interface.[3]

A GoPro camera, affixed to the underbelly of the airframe, points vertically downwards. The GoPro captures video at 30 fps with a fast shutter (to minimize motion blur and rolling shutter effects). We use `ffmpeg` to sample the video into stills at 7.5 fps (every fourth frame). We post-process the images using *Pix4D*:

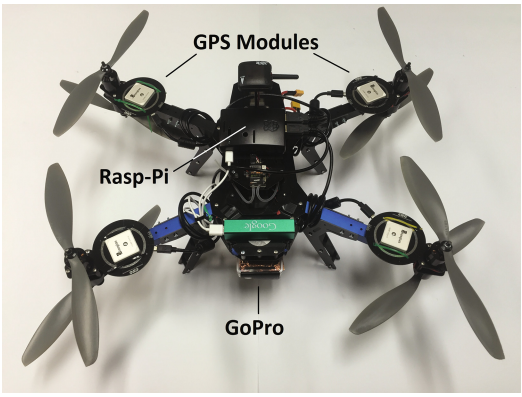

Fig. 5. Description of our drone platform

[3] Experiments were conducted as per FAA regulations. Operator trained by a general aviation pilot, over relatively vacant fields, with flight height not exceeding 45 m.

advanced commercial photogrammetry software that uses *structure from motion* (SfM) to perform 3D registration of each frame. From Pix4D's outputs, we can recover a high-precision estimate of drone attitude (accuracy ≈0.05° [13,28]). Note: structure from motion would not be practical as a realtime IMU replacement – each 5–7 min flight requires several hours of processing on a server of 16 CPU cores, 32 GB RAM, and CUDA on a high-end NVIDIA GRID K2 GPU.

We refer to *pitch/roll* angles together as *tilt* angles and the *yaw* angle as the *heading* angle. Figures 6 and 7 show the distribution of tilt and heading orientation errors before and after incorporation of GLONASS satellite measurements. Median errors before and after GLONASS integration is 2.26° and 1.79° respectively. The corresponding 95th percentile errors are 26° and 7° respectively. Evidently, GLONASS measurements increase the accuracy substantially, and more importantly towards the tail of the distribution. Cutting down the tail errors will enhance the reliability of SafetyNet.

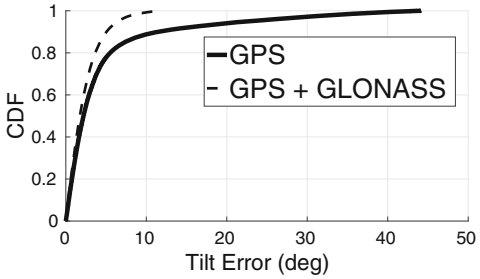

Fig. 6. Tilt error CDF before/after GLONASS integration

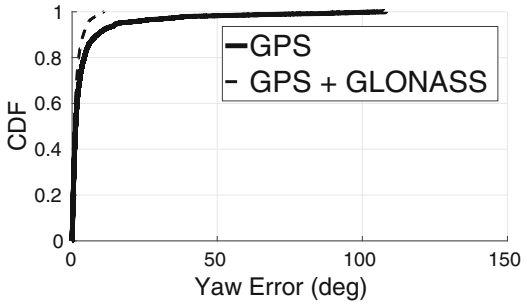

Fig. 7. Heading error CDF before/after GLONASS integration

Figure 8(a), (b) gives a break up of the 95^{th} percentile and median errors of tilt orientation angles over several flights conducted under varying weather conditions. Flights 7–11 were conducted in fog. As evident, GLONASS integration contributes to an error reduction of $3.5x$ at the 95^{th} percentile and 27% at the

median. Similarly, Fig. 9(a), (b) gives a break of the 95^{th} and median errors for the yaw angle. GLONASS integration is contributing to a gain of $11x$ at the 95^{th} percentile and 37% at the median. Substantial boost in orientation estimation accuracy by GLONASS would increase the reliability of *SafetyNet* to serve as IMU failsafe.

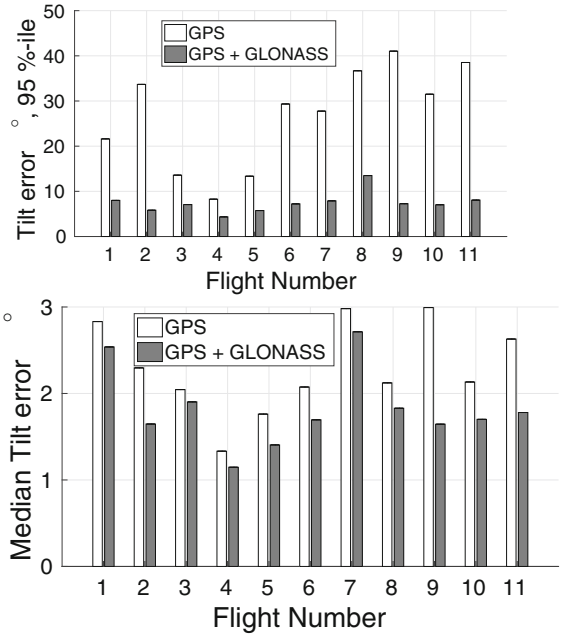

Fig. 8. GLONASS integration dramatically cuts down the tail tilt errors (a) 95^{th} percentile tilt errors with/without GLONASS (b) Median tilt errors with/without GLONASS

6 Related Work

Resolving Integer ambiguities in GLONASS is challenging because the satellites operate at different frequencies. We summarize the main techniques in literature which attempt to solve this problem.

Common Clock: The second term in single differencing Eq. 7 represents the clock bias difference between the two receivers whose phases are being differenced. If we connect the two receivers with a common clock, this bias term is completely eliminated. Works in [11,17] exploit this opportunity for integrating GLONASS with GPS. Since there is no differencing performed across satellites, the wavelength difference will not matter. Such an approach however requires tight synchronization between receivers with a common clock. This increases the hardware complexity, whereas SafetyNet doesn't have this requirement.

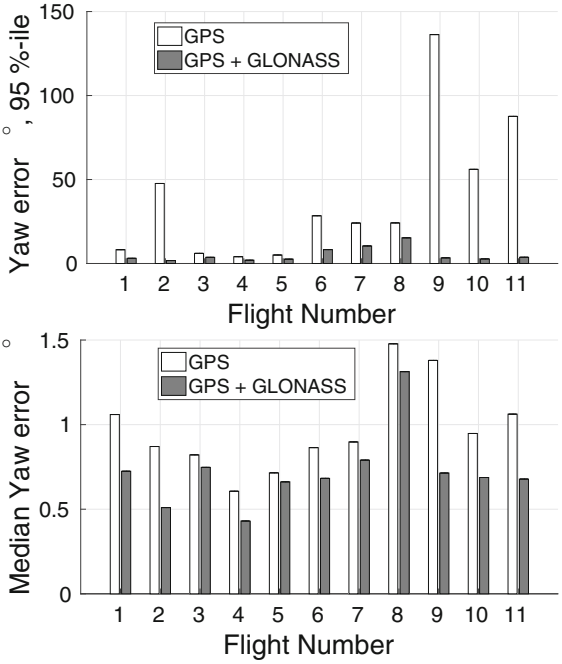

Fig. 9. GLONASS integration dramatically cuts down the tail heading errors (a) 95^{th} percentile heading errors with/without GLONASS (b) Median heading errors with/without GLONASS

Search: Work in [29] follows an extensive searching and optimization techniques for resolving single differenced (SD_{ij}) and double differenced integer ambiguities (DD_{ij}^{sk}). Such an approach increases the complexity especially under heavy cycle slip rate during aggressive flights. Work in [30] also uses an extensive search procedure with a resolution of 64 um which becomes impractical under low SNR and heavy cycle slip rate.

Pseudorange: Works in [12,27] use pseudoranges for resolving integer ambiguities. However, pseudorange data is low in precision (30 cm), hence not suitable for integer ambiguity resolution for drone orientation estimate which needs an accuracy of around 2 cm.

7 Discussion

Energy Consumption: GPS receivers are known to be power hungry. While this is a concern for smart-phone like applications, a flying drone consumes 1000 times higher power than GPS. Therefore supporting GPS receivers constantly is feasible.

Cost: Commercial drones already use multiple GPS receivers for redundancy. Moreover, the need for safety is higher for bigger drones carrying heavier pay-

loads. Given the cost of GPS receivers is negligible relative to the cost of such drones, we believe the additional cost of GPS receivers is justifiable.

8 Conclusion

GPS sensors offer substantial promise in orientation sensing, a desirable feature for serving as failsafe for IMU failures. Augmenting GPS with GLONASS measurements can offer enormous gain in the orientation accuracy, particularly in cutting down the tail of the error distribution and thereby enhancing the robustness. While integrating FDMA based GLONASS with CDMA based GPS is non-trivial, we show that incorporating time differenced GLONASS carrier phase data with GPS provides effective fusion.

References

1. Diydrones. http://diydrones.com/forum/topics/problem-quad-copter-spins-around-itself
2. Drone crash survey. https://docs.google.com/forms/d/1R4MaX8iZWRoyzKJ6rKr0U-p3x-6j0ZVm5a0p83hHty0/viewanalytics
3. Galileo. https://en.wikipedia.org/wiki/Galileo_(satellite_navigation)
4. Google forum. https://groups.google.com/forum/#!msg/drones-discuss/fnWYM48pGys/kwPI_q_qv0QJ
5. Rotation vectors. http://farside.ph.utexas.edu/teaching/301/lectures/node100.html
6. Vector cross product. http://soe.rutgers.edu/~meer/GRAD561/ADD/antisymm.pdf
7. Blewitt, G.: Basics of the GPS technique: observation equations. In: Johnson. B. (ed.) Geodetic Applications of GPS, pp. 10–54. Nordic Geodetic Commission, Gothenburg (1997)
8. De Pasquale, G., Somà, A.: Reliability testing procedure for MEMS IMUs applied to vibrating environments. Sensors **10**(1), 456–474 (2010)
9. Diebel, J.: Representing attitude: Euler angles, unit quaternions, and rotation vectors. Matrix **58**(15–16), 1–35 (2006)
10. Gowda, M., Manweiler, J., Dhekne, A., Choudhury, R.R., Weisz, J.D.: Tracking drone orientation with multiple GPS receivers. In: Mobicom. ACM, New York (2016)
11. Han, K.J., Gerard, L.: Determining Heading and Pitch Using a Single Difference GPS/GLONASS Approach. University of Calgary, Canada (1999)
12. Han, S., Dai, L., Rizos, C.: A new data processing strategy for combined GPS/GLONASS carrier phase-based positioning. In: Proceedings of the ION GPS 1999, pp. 1619–1627 (1999)
13. Harwin, S., Lucieer, A.: Assessing the accuracy of georeferenced point clouds produced via multi-view stereopsis from Unmanned Aerial Vehicle (UAV) imagery. Remote Sens. **4**(6), 1573–1599 (2012)
14. Hedgecock, W., et al.: Regtrack: a differential relative GPS tracking solution. In: Proceedings of the 11th Annual International Conference on Mobile Systems, Applications, and Services - MOBISYS 2013 (2013)

15. Hedgecock, W., Maroti, M., Ledeczi, A., Volgyesi, P., Banalagay, R.: Accurate real-time relative localization using single-frequency GPS. In: Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems - SENSYS 2014, pp. 206–220. ACM (2014)
16. Kaplan, E., Hegarty, C.: Understanding GPS: Principles and Applications. Artech House, Boston (2005)
17. Keong, J.: GPS/GLONASS attitude determination with a common clock using a single difference approach. In: ION GPS 1999 Conference, Nashville, pp. 14–17 (1999)
18. Lai, Y.-C., Jan, S.-S.: Attitude estimation based on fusion of gyroscopes and single antenna GPS for small UAVs under the influence of vibration. GPS Solut. **15**(1), 67–77 (2011)
19. Langley, R.B.: GLONASS: review and update. GPS World **8**(7), 46–51 (1997)
20. Laurichesse, D., Mercier, F., Berthias, J.-P., Broca, P., Cerri, L.: Integer ambiguity resolution on undifferenced GPS phase measurements and its application to PPP and satellite precise orbit determination. Navigation **56**(2), 135–149 (2009)
21. Liu, X., Randall, R.: Blind source separation of internal combustion engine piston slap from other measured vibration signals. Mech. Syst. Signal Process. **19**(6), 1196–1208 (2005)
22. Malyavej, V., Torteeka, P., Wongkharn, S., Wiangtong, T.: Pose estimation of unmanned ground vehicle based on dead-reckoning/GPS sensor fusion by unscented Kalman filter. In: 6th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology, ECTI-CON 2009, vol. 1, pp. 395–398. IEEE (2009)
23. Mekik, C., Arslanoglu, M.: Investigation on accuracies of real time kinematic GPS for GIS applications. Remote Sens. **1**(1), 22–35 (2009)
24. Misra, P., Enge, P.: Global Positioning System: Signals, Measurements and Performance, 2nd edn. Ganga-Jamuna Press, Lincoln (2006)
25. Parkinson, B.W., Enge, P.K.: Differential GPS. Glob. Position. Syst. Theor. Appl. **2**, 3–50 (1996)
26. Suh, Y.S.: Attitude estimation by multiple-mode Kalman filters. IEEE Trans. Industr. Electron. **53**(4), 1386–1389 (2006)
27. Tsujii, T., Harigae, M., Inagaki, T., Kanai, T.: Flight tests of GPS/GLONASS precise positioning versus dual frequency KGPS profile. Earth, Planets Space **52**(10), 825–829 (2000)
28. Vallet, J., Panissod, F., Strecha, C., Tracol, M.: Photogrammetric performance of an ultra light weight swinglet UAV. In: UAV-g, no. EPFL-CONF-169252 (2011)
29. Wang, J.: An approach to GLONASS ambiguity resolution. J. Geodesy **74**(5), 421–430 (2000)
30. Wang, J., Rizos, C., Stewart, M.P., Leick, A.: GPS and GLONASS integration: modeling and ambiguity resolution issues. GPS Solut. **5**(1), 55–64 (2001)
31. Wells, M.: Attenuating magnetic interference in a UAV system. Ph.D. thesis, Carleton University, Ottawa (2008)
32. Zhang, B., Teunissen, P.J., Odijk, D.: A novel un-differenced PPP-RTK concept. J. Navig. **64**(S1), S180–S191 (2011)
33. Zhang, Y., Chamseddine, A., Rabbath, C., Gordon, B., Su, C.-Y., Rakheja, S., Fulford, C., Apkarian, J., Gosselin, P.: Development of advanced FDD and FTC techniques with application to an unmanned quadrotor helicopter testbed. J. Franklin Inst. **350**(9), 2396–2422 (2013)