Performance Analysis and Architectures for INS-Aided GPS Tracking Loops

Santiago Alban, Dennis M. Akos, Stephen M. Rock
Department of Aeronautics and Astronautics, Stanford University

Demoz Gebre-Egziabher
Department of Aerospace Engineering and Mechanics, University of Minnesota

BIOGRAPHIES

Santiago Alban is a Ph.D. candidate in the Aeronautics and Astronautics Department of Stanford University. He received his B.S. in Aerospace Engineering in 1996 from the University of Texas at Austin, and his M.S. from Stanford in 1998. His research involves the development of a low-cost GPS/INS attitude system for automobiles, and deep integration of INS systems with GPS tracking loops.

Dr. Demoz Gebre-Egziabher is an assistant professor of Aerospace Engineering and Mechanics at the University of Minnesota, Twin Cities. His research interests are in the areas of navigation, guidance and control with an emphasis on sensor fusion and system integration issues. Prior to joining the faculty at the University of Minnesota, he was an officer in the US Navy and subsequently a member of the GPS Laboratory at Stanford University where he received a Ph.D. in Aeronautics and Astronautics.

Dr. Dennis M. Akos completed the Ph.D. degree in Electrical Engineering at Ohio University within the Avionics Engineering Center. He has since served as a faculty member with Luleå Technical University, Sweden, and is currently a Research Associate with the GPS laboratory at Stanford University.

Dr. Stephen M. Rock is an Associate Professor of Aeronautics and Astronautics at Stanford University. He received his S.B. and S.M. in Mechanical Engineering from MIT in 1972 and Ph.D. in Applied Mechanics from Stanford University in 1978. Dr. Rock joined the Stanford faculty in 1988 where he teaches courses in dynamics and control and pursues research in developing and applying advanced control techniques for vehicle and robot applications. Prior to joining the Stanford faculty, Dr. Rock led the advanced controls group of Systems Control Technology.

ABSTRACT

GPS and inertial sensors have complementary characteristics, which have been exploited in the design of integrated GPS-inertial navigation and guidance systems. Traditionally, most hybrid GPS-inertial systems have been mechanized by combining the information from GPS and an Inertial Navigation System using either *loose integration* (i.e., integration at the position, velocity and/or attitude level) or *tight integration* (integration at the pseudorange, Doppler, or carrier phase level). Such integration schemes provide users with limited immunity against momentary GPS outages and also allow detection of certain classes of GPS signal failures. A third scheme of integration can be used, in which the inertial sensors are used to aid the GPS phase/frequency and code tracking loops directly. In this paper, this level of coupling is referred to as *ultra-tight integration*, and it offers potential improvements to GPS performance, such as higher phase-tracking bandwidth, and more resistance to radio frequency interference or multipath noise.

In this paper we propose a simple architecture for GPS-inertial systems with ultra-tight integration and present the results of some trade studies and simulations quantifying the performance of such systems. Performances of the ultra-tight GPS-inertial system are evaluated using a simulation tool developed specifically for this study. The metrics used for the evaluation are allowable reduction in the carrier tracking loop-filter bandwidth for improved signal-to-noise ratio, and robustness against carrier-phase cycle-slips. The sensitivity of these metrics to inertial sensor quality and GPS receiver clock noise is discussed and quantified. These studies show that an ultra-tightly coupled system using low cost/performance inertial sensors and a typical temperature-compensated crystal oscillator can function with a carrier tracking loop-filter bandwidth as low as 3Hz. This structure shows a 14dB improvement in phase-noise suppression when compared to a traditional 15Hz

loop filter, and comparable carrier-phase tracking bandwidth to that of the inertial sensors (>30Hz).

I. INTRODUCTION

The Global Positioning System (GPS) has become one of the most popular methods of navigation for users worldwide. The widespread use of GPS in terrestrial, marine, and airborne applications has been precipitated by its accuracy, global availability and the low cost of user equipment. Some applications that require enhanced performance or robustness from GPS navigation systems employ the use of inertial sensors (rate gyros and accelerometers) to boost the system bandwidth, provide better noise characteristics, or allow navigation through brief GPS outages.

Three are various architectures for fusing inertial sensors with GPS. Loose integration is the simplest method of coupling, and is depicted in Fig. 1. In this scheme, GPS and the inertial sensors generate navigation solutions independently (position, velocity and attitude). The two independent navigation solutions are subsequently combined to form a blended (or filtered) GPS-inertial navigation solution. One of the benefits of loose integration is that the blended navigation solution tends to have a higher bandwidth and better noise characteristics than the GPS solution alone. This configuration is best implemented with higher quality inertial sensors (navigation or tactical grade) if the GPS outages are long in duration. Lower quality inertial sensors can also provide some immunity against momentary GPS outages, especially if their outputs were calibrated using GPS prior to the outage. In this instance, the GPS-inertial loose integration is said to include feedback, whereby the difference between the GPS and inertial solutions is fed back to the inertial sensors to carry out the calibration. In general, lower quality inertial sensors (consumer or automotive grade) are suited for applications where GPS outages are infrequent and short in duration.

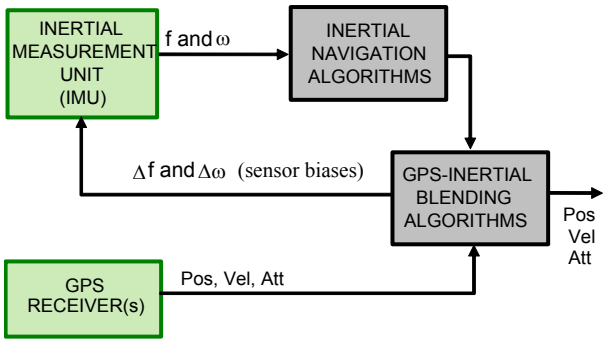

Figure 1. GPS/INS System with Loose Integration

A more complex level of coupling is tight integration, where GPS pseudoranges (PR), Doppler, or carrier phase (CP) measurements are blended with the navigation solution generated by the inertial sensors. In addition to the benefits of loose coupling, a tightly integrated system can have a more accurate navigation solution because the basic GPS observables used in the blending process (i.e. PR and Doppler measurements) are not as correlated as the position and velocity solutions used in loose integration [7]. Furthermore, tight integration provides a means for implementing a more sensitive fault detection and isolation scheme that can be used to verify the quality of PR and Doppler measurements [4,5]. Applications which use carrier-phase output (attitude determination and carrier-phase positioning) especially benefit from tight integration because integer ambiguities can be recovered and verified quickly from the navigation outputs, despite cycle slips and increased carrier-phase noise [1]. Figure 2 shows the general structure of a system with tight integration.

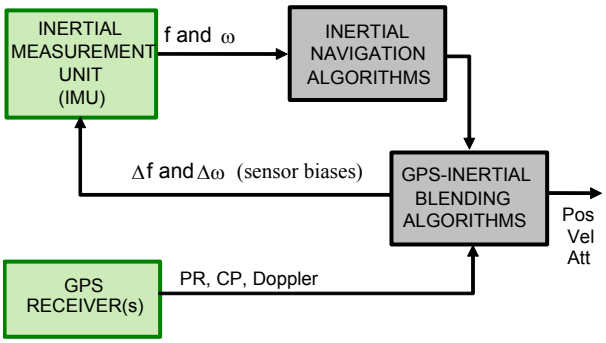

Figure 2. GPS/INS System with Tight Integration

A more complex and potentially most beneficial level of GPS-inertial integration occurs at the GPS tracking-loops level, as shown in Fig. 3. In this paper, this level of coupling is referred to as ultra-tight integration. This configuration is more complex than the other architectures discussed because it changes the structure of the traditional GPS tracking loops. In terms of performance, ultra-tight integration also offers the most benefits in terms of accuracy and robustness improvements to the GPS receiver and overall system. Ultra-tight integration can improve acquisition time [14] as well as the tracking performance of the phase-lock loop in terms of bandwidth and noise rejection, thus producing more accurate Doppler and phase measurements. As will be discussed later in the paper, the use of inertial sensors in ultra-tight integration allows the reduction of the carrier-phase tracking-loop bandwidth by eliminating the need to track the vehicle platform dynamics. This integration scheme results in cleaner carrier-phase measurements and faster tracking of the carrier phase. Furthermore, estimation of the drift-rate of the GPS receiver clock permits uninterrupted tracking of a channel

despite a brief line-of-sight blockage. This capability can prevent carrier-phase cycle slips, and therefore has potential applications in high performance navigation systems where robustness to cycle slips is of paramount importance [12,15].

It should be noted that the term "tight integration" has been used in a broader context to encompass both the tight and ultra-tight architectures introduced here, as done in [16].

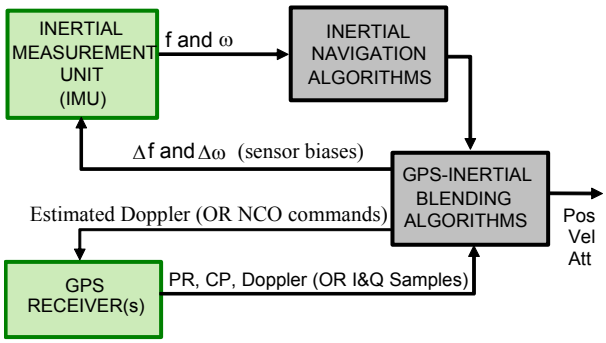

Figure 3. GPS/INS System with Ultra-Tight Integration

The topic of ultra-tight integration is relatively new, and is being developed by a few organizations, including Interstate Electronics Corporation [3], The Aerospace Corporation, and Draper Laboratories [11]. The approach taken in their implementations involves the use of a single large filter, or smaller multiple filters that use in-phase and quadrature samples from the GPS receiver channels as measurements for updating the filter states. Control of the replicated carrier and code generators also comes from navigation filter outputs, which are propagated with IMU measurements to achieve high bandwidth carrier-phase tracking and anti-jam capability. This more complex architecture variant of ultra-tight integration is also termed "deep integration", and is characterized by implementing the closed-loop signal tracking for all channels through the navigation filter itself, thus precluding the need to maintain separate code and carrier tracking loops [16].

The integration structure presented in this paper shows a simpler approach to ultra-tight integration that utilizes classical control theory. The concepts that will be presented can be used to study the theory and performance gains of a GPS/INS navigation system with ultra-tight integration, without the added complexity of optimal estimation. The straightforward architecture of the system is not limited to be a simple learning tool, but can also potentially be implemented to realize a basic ultra-tightly coupled system. The remaining sections of this paper will focus on the development of this simple architecture and its performance. The GPS receivers of the system feature a modified phase-lock loop that utilizes external Doppler estimates, and a rate-aided delay-lock loop for maintaining code tracking during brief signal outages. Although much of the discussion involves details of the GPS tracking loops, some attention is also given to external Doppler-frequency and clock-error frequency estimation. Other important design issues such as GPS receiver clock drift and inertial sensor quality are also covered, and evaluated versus system performance.

This analysis was performed with the use of a GPS software receiver, which was developed specifically to study the details of ultra-tight integration. The software receiver tracks GPS sampled signals (real or simulated) stored in data files, and gives standard GPS outputs such as position, velocity, and attitude (if two signals are used).

II. GPS TRACKING LOOPS

This section reviews some of the basic concepts of traditional GPS phase and code tracking loops, including modeling and implementation. The ideas that will follow in the development of tracking loops that incorporate inertial aiding will build upon these traditional loop structures to realize the structure of a GPS receiver with ultra-tight integration. A more detailed treatment of these concepts can be found in [6].

A. The Phase-Lock Loop (PLL)

After the GPS signal passes through the RF front end and is sampled, the discrete signal is centered at an intermediate frequency (f_{IF}), typically around 1-20 MHz. The GPS signal consists of the superposition of all GPS signals in view of the antenna, and is processed by parallel channels in the receiver, where each channel tracks (or searches for) the signal from one satellite. After a satellite has been acquired, the PLL for the channel tracks the phase and frequency of the carrier signal. A typical model of a second-order PLL is shown in Fig. 4.

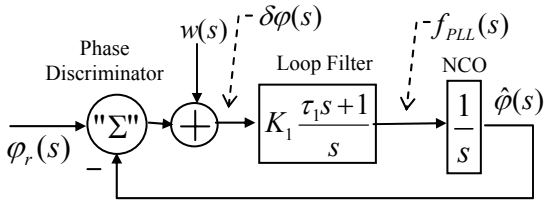

$\varphi_r(s)$ = reference phase
$w(s)$ = external phase noise
$\delta\varphi(s)$ = measured phase error
$f_{PLL}(s)$ = frequency tracked by PLL
$\hat{\varphi}(s)$ = tracked phase

Figure 4. Model of Second-Order PLL

In this model, the plant in the PLL model is a numerically controlled oscillator (NCO), which is controlled by f_{PLL}. The loop filter is a compensator designed to track the reference phase with a typical bandwidth of 15-30Hz. It is important to note that the loop filter design is a compromise between bandwidth and noise suppression, because the reference phase and the phase noise are treated by the same transfer function, shown below:

$$\hat{\varphi}(s) = H_1(s)\big(\varphi_r(s) + w(s)\big) \tag{1}$$

$$H_1(s) = \frac{K_1(\tau_1 s + 1)}{s^2 + K_1\tau_1 s + K_1} \tag{2}$$

By maintaining lock on the phase of the carrier signal, the PLL also tracks any deviations of the carrier frequency from the intermediate frequency. In Fig. 4, f_{PLL} represents the frequency deviation from f_{IF}. The components of f_{PLL} include the Doppler frequency of the signal (f_{dopp}), due to relative motion between the receiver and the satellite, as well as frequency errors due to the local reference oscillator (f_{clk}). The clock frequency errors are introduced through the down-conversion and sampling process in the RF front end. Mathematically, f_{PLL} can be expressed as:

$$f_{\mathrm{PLL}} = f_{dopp} + f_{clk} + f_{noise} \tag{3}$$

The frequency content of f_{dopp} can be significantly higher than that of f_{clk}, depending on the user platform dynamics. As will be discussed later in the development of an ultra-tight integration structure, the knowledge of the satellite motion and user platform dynamics can be used to estimate the Doppler frequency externally. Providing this estimate to the PLL will then preclude the need to track the fast dynamics of f_{dopp} within the loop, and will allow tightening the loop-filter bandwidth to improve noise performance. This principle is the crux of the ultra-tight integration concept.

Figure 5 shows the actual structure of a PLL, including components like mixers and "accumulate and dump" (A&D) operations which perform averaging over one or more code periods.

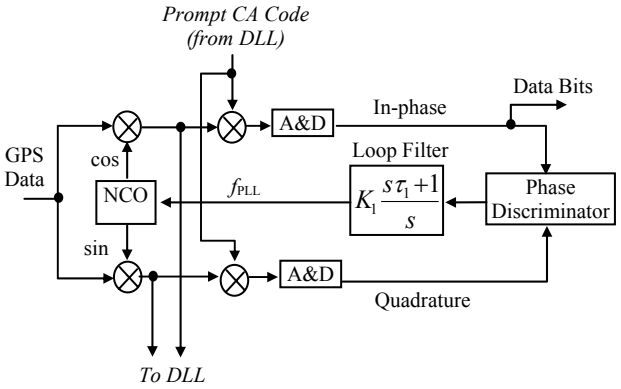

Figure 5. Structure of PLL

The oscillator generates in-phase and quadrature replicas of the carrier signal for the channel, with a frequency equal to the sum of f_{PLL} and f_{IF}. The signal replicas are used to mix the incoming signal to baseband for extracting the navigation bits and measuring the phase error. As shown in Fig. 5, the PLL also utilizes the prompt CA code provided by the delay-lock loop, which in turn, depends on the in-phase and quadrature baseband signals provided by the PLL.

B. The Delay-Lock Loop (DLL)

The function of the DLL is to track the CA code component of the GPS signal. In addition to generating the prompt code needed for tracking the carrier in the PLL, the code phase from the code generator is also used to determine the pseudoranges for position determination. Although the structure of the DLL will not change with the implementation of ultra-tight integration, some explanation is necessary as to why a rate-aided DLL should be used to complement the robustness enhancements in the PLL through ultra-tight integration.

Figure 6 shows the model of a noncoherent DLL. The code generator (CG) produces a replica of the CA code at a frequency imposed by the loop compensator, which is usually a simple gain. The error signal in this model ($\delta\tau(s)$) represents the difference between the phase of the reference CA code ($\tau_r(s)$) and that of the replicated CA code ($\tau_r(s)$).

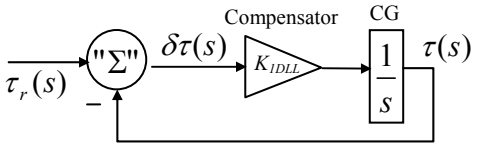

Figure 6. Model of DLL

A rate-aided DLL uses the less noisy frequency estimate determined by the PLL to aid the code-tracking loop, as shown in Fig. 7.

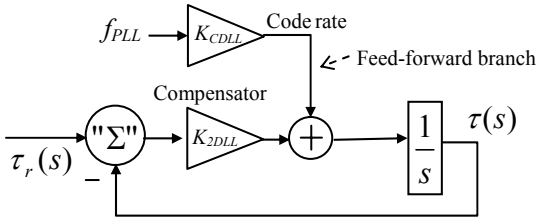

Figure 7. Model of Rate-Aided DLL

Figure 8 shows the structure of a rate-aided DLL. As shown, the error signal is the difference of the early and late powers, which are computed by mixing the

replicated carrier (in-phase and quadrature) with early and late code replicas.

The gain that converts the PLL frequency to the code rate is denoted K_{CDLL}. With this configuration, the gain of the compensator can be reduced to improve noise performance, since the necessary bandwidth needed to maintain code lock is provided by the feed-forward branch. However, the closed-loop configuration is still needed to track the slow and unpredictable drifts caused by code-carrier divergence [6]. In addition to improved noise performance, a rate-aided DLL can also provide a synchronized code despite brief outages in the GPS signal, as long as the external frequency aiding is uninterrupted and the outages are not long enough to allow significant code-carrier divergence. This advantage will be exploited when uninterrupted aiding is available with ultra-tight integration.

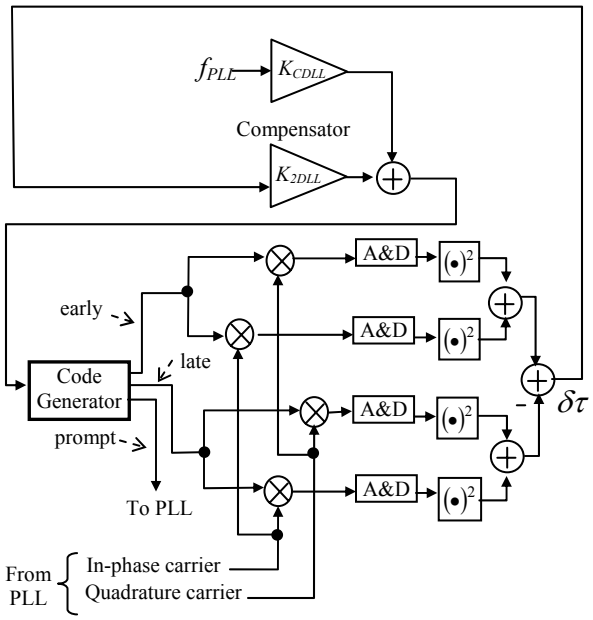

Figure 8. Structure of Rate-Aided DLL

III. DOPPLER AND CLOCK-ERROR FREQUENCY ESTIMATION

As explained in section IIA, a traditional PLL tracks the Doppler and clock error components of the carrier signal. If these components can be estimated externally, they can be used to aid the PLL to track with more efficiency and robustness. The Doppler frequency of the carrier can be estimated from the outputs of an Inertial Navigation System (INS), an attitude reference, and satellite ephemeris data. As will be discussed later, the accuracy of these Doppler estimates depends on the quality of inertial sensors used.

A. Doppler Frequency Estimation

The Doppler frequency of the carrier signal can be expressed simply as the velocity of the receiver relative to the transmitter, projected onto the line-of-sight (LOS) vector. The relationship is expressed in Eq. 4:

$$f_{dopp} = \frac{1}{\lambda}(\vec{V}^{RX} - \vec{V}^{S}) \cdot \vec{1}^{S} \tag{4}$$

where

λ = Wavelength of carrier at L1 frequency
$\vec{V}^{RX}$ = Velocity of receiver antenna
$\vec{V}^{S}$ = Velocity of satellite
$\vec{1}^{S}$ = Unit LOS vector from receiver to satellite

The components of Eq. 4 can be computed in any Earth-fixed reference frame, such as Earth-Centered-Earth-Fixed (ECEF) or local East-North-Up (ENU). The transformation matrix between these two reference frames is easily computed as a function of latitude and longitude.

The computation of the ECEF satellite velocity from ephemeris data is a well-known procedure, described in [13]. Another very simple and relatively accurate method for computing satellite velocity is to estimate the time derivative of its position. The computation of the ECEF satellite position from ephemeris data is described in [19]. Using this method to compute the satellite position as a function of GPS time-of-week, the velocity of the satellite can be approximated as the discrete time-derivative of its position, with a small time-step ($\Delta t \approx 1$ms):

$$\underline{V}_E^S(t) \approx \frac{\underline{R}_E^S(t+\Delta t) - \underline{R}_E^S(t)}{\Delta t} \tag{5}$$

where

$\underline{V}_E^S(t)$ = ECEF velocity of satellite at GPS time of week t
$\underline{R}_E^S(t)$ = ECEF position of satellite at GPS time of week t

The ENU velocity of the antenna is one of the common measurements from GPS receivers, but it does not have sufficient bandwidth to aid the PLL in a feedback fashion. However, an integrated GPS/INS system (as shown in Fig. 3) can provide velocity measurements with sufficient bandwidth and time stability to aid the PLL through an intermediate Doppler calculation. Since the inertial sensors provide accelerations in a body-fixed frame, an attitude reference is needed to transform them to an Earth-fixed frame. The attitude system can be implemented in a number of ways. The options include an unaided inertial navigator (provided it has relatively stable rate gyro and accelerometers), a loosely or tightly coupled GPS-inertial system, or an Attitude Heading Reference System (AHRS) which uses lower quality rate gyros aided by a

multi-antenna GPS attitude determination system [1,9]. The AHRS option is more economical even though it can be more complex to implement than the others.

The GPS/INS velocity filter must also account for the fact that the GPS antennas and INS sensors are generally not collocated. Thus, the innovations of the filtering process may also include terms originating from the cross product of the angular velocity of the vehicle and the vector between the INS sensors and the GPS antenna. More detail on the structure of a GPS/INS velocity filter can be found in [7,18].

Earlier in the paper, it was noted that the crux of ultra-tight integration is the notion that the high-frequency content of f_{PLL} due to platform motion, f_{dopp}, can be estimated with the use of inertial sensors. It is also important to determine how well the Doppler frequency can be estimated with inertial sensors. The accuracy of the Doppler estimates will obviously depend on the quality of the inertial sensors used.

In mathematical terms, we are interested in determining the 1-σ standard deviation in the Doppler estimate, σ_{dopp}, and how it is affected by inertial sensor quality. Assuming that most of the error in the external Doppler estimation comes from the inertial sensor estimate of $\vec{V}^{RX}$, a perturbation of Eq. 4 results in the following expression for σ_{dopp} :

$$\sigma_{dopp}^2 = \frac{\vec{1}^{S} \cdot E\left(\delta\vec{V}^{RX}(\delta\vec{V}^{RX})^{T}\right) \cdot \vec{1}^{S}}{\lambda^2} \qquad (6)$$

The error equations for $\delta\vec{V}^{RX}$ are well known [2,7,10] and will not be discussed here, except to point out that the *psi-angle* error equations discussed in [10] are used in the subsequent analysis.

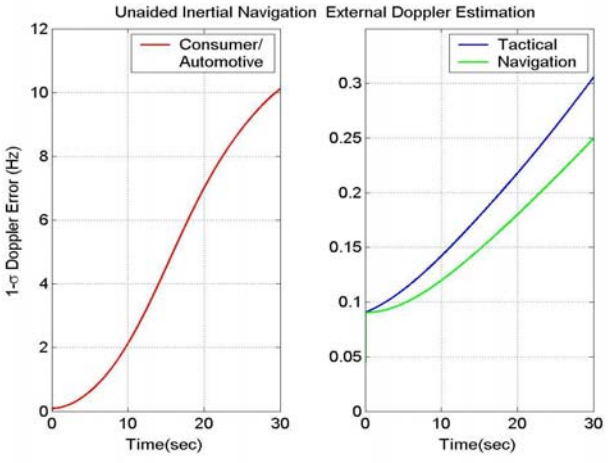

Figure 9. Errors in Doppler Estimation with Unaided Inertial Navigation

The plots shown in Fig. 9 are generated with the use of these standard inertial navigation-system error models in conjunction with sensor-error models developed in [8], for navigation, tactical, and automotive/consumer grade inertial navigation sensors. These plots show the 1-σ of the Doppler estimation error as a function of time for an unaided inertial navigation system. The large error accrued with consumer or automotive-grade inertial sensors implies that lower quality sensors require constant calibration from the GPS outputs.

B. Clock-Error Frequency Measurement

The clock-error frequency component of f_{PLL} can be measured concurrently with the GPS velocity. Since it is a common error term on all the channels, four or more satellites are necessary to solve for three components of the user velocity (Vx, Vy, Vz) and the clock error, as shown in Eq. 7:

$$\begin{bmatrix} -f_{PLL1}\lambda + \vec{V}^{S1}\cdot\vec{1}^{S1} \\ -f_{PLL2}\lambda + \vec{V}^{S2}\cdot\vec{1}^{S2} \\ \vdots \\ -f_{PLLN}\lambda + \vec{V}^{SN}\cdot\vec{1}^{SN} \end{bmatrix} = \begin{bmatrix} \vec{1}^{S1T} & 1 \\ \vec{1}^{S2T} & 1 \\ \vdots & \vdots \\ \vec{1}^{SNT} & 1 \end{bmatrix} \begin{bmatrix} Vx \\ Vy \\ Vz \\ V_{clk} \end{bmatrix} \qquad (7)$$

where

$$V_{clk} = \lambda f_{clk} \qquad (8)$$

The behavior of the clock-error frequency (f_{clk}) varies significantly with the type of reference oscillator in the GPS receiver. Figure 10 shows measurements of f_{clk} for a common temperature-compensated crystal oscillator (TCXO) used in most consumer GPS receivers, and for a higher quality oven-controlled crystal oscillator (OCXO).

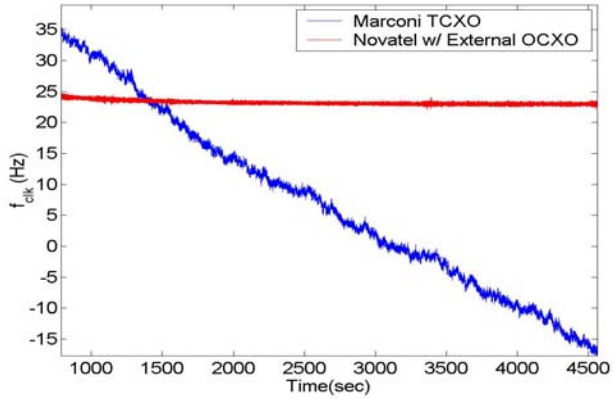

Figure 10. Clock-Error Frequency Measurements

In these measurements, the OCXO has very little drift with time after an initial transient, and f_{clk} is composed primarily of a constant bias and broadband

noise. In contrast, the TCXO frequency error appears to be made up of three components: broadband noise, exponentially correlated noise, and a slow drift with time. The clock behavior is important in the design of phase-lock loops for ultra-tight integration, and this topic will be revisited in section IV.

IV. A CONTROLS APPROACH TO ULTRA-TIGHT INTEGRATION

This section introduces a simple feed-forward method of using the external Doppler and clock-error frequency estimates to aid the PLL. The concept of feed-forward is useful when a component or form of the reference signal can be measured, and processed in some way to aid the controller. A textbook example of this idea is when the DC bias of the reference signal can be measured; the bias is then multiplied by the inverse of the plant model and added to the output of the controller. This scheme allows the use of a controller that requires less power, since its output would no longer include the DC bias required to track the reference signal.

A. PLL Model for Ultra-Tight Integration

The feed-forward implementation of a PLL for ultra-tight integration utilizes the fact that the sum of external Doppler and clock-error frequencies is a measurement of the *rate of change* of the reference carrier phase. In fact, the sum of these two frequencies is precisely the quantity that the controller in Figs. 4 and 5 generates, and uses to control the NCO to maintain phase lock. For modeling purposes, the PLL for ultra-tight integration depicted in Fig. 11 shows the feed-forward branch originating from the reference signal. It should be noted that most of this branch only represents the relationship between the reference phase and the external frequency estimates, and is not part of a realization of the loop.

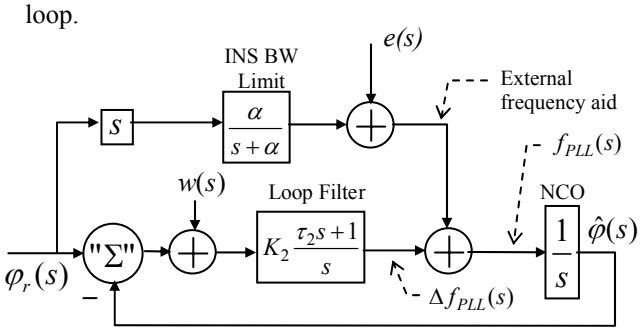

Figure 11. Model of PLL with Ultra-Tight Integration

The feed-forward branch includes a single-pole low-pass filter and an additive error term, *e(s)*. The low pass filter accounts for the finite bandwidth of the inertial sensors, and the error term represents errors in the external frequency estimates. These errors can be due to several factors, which may include imperfect sensor calibration, errors in the platform attitude measurements, and clock-error frequency.

The PLL configuration in Fig. 11 offers advantages over the traditional structure (Fig. 4) in terms of phase-noise suppression and tracking bandwidth. These benefits are possible because the reference input and phase noise are no longer treated by the same transfer function; the phase-noise input is still related to the output by a similar transfer function as Eq. 2 (with different poles and zeros), but the relationship between the reference input and output is now largely dependent on the inertial sensors. Equations 9-11 show these relationships, assuming zero input from *e(s)*.

$$\hat{\varphi}(s) = H_3(s)\varphi_r(s) + H_2(s)w(s) \tag{9}$$

where

$$H_3(s) = \frac{\frac{\alpha}{s+\alpha} + \frac{K_2(\tau_2 s+1)}{s^2}}{1 + \frac{K_2(\tau_2 s+1)}{s^2}} \tag{10}$$

$$H_2(s) = \frac{K_2(\tau_2 s+1)}{s^2 + K_2\tau_2 s + K_2} \tag{11}$$

Equation 10 is expressed in such a way to illustrate that as the INS bandwidth gets large ($\alpha \to \infty$), H_3(s) approaches unity. This result suggests that the phase dynamics originating from user motion can be tracked with the bandwidth provided by the inertial sensors. The design of the loop filter no longer needs to account for high-bandwidth user dynamics, but it must still be able to track any errors in the external frequency estimates (*e(s)*). The term Δf_{PLL} in Fig. 11 represents the signal that tracks *e(s)*:

$$\Delta f_{PLL}(s) = -H_2(s)e(s) \tag{11}$$

Thus, the tracking bandwidth versus noise-rejection tradeoff still takes place, but with the important difference that the required tracking bandwidth is lower, so noise rejection improves. The lower limit on the bandwidth of the loop filter is now determined by the dynamics of *e(s)* that must be tracked.

To illustrate the benefits of ultra-tight integration in terms of phase-noise rejection and phase-tracking bandwidth, Fig. 12 shows the magnitude of the closed-loop transfer function $H_1(s)$ used in a traditional 15Hz PLL versus $H_3(s)$ and $H_2(s)$ used in an ultra-tightly integrated system with 100Hz bandwidth inertial sensors and 3Hz bandwidth for error tracking.

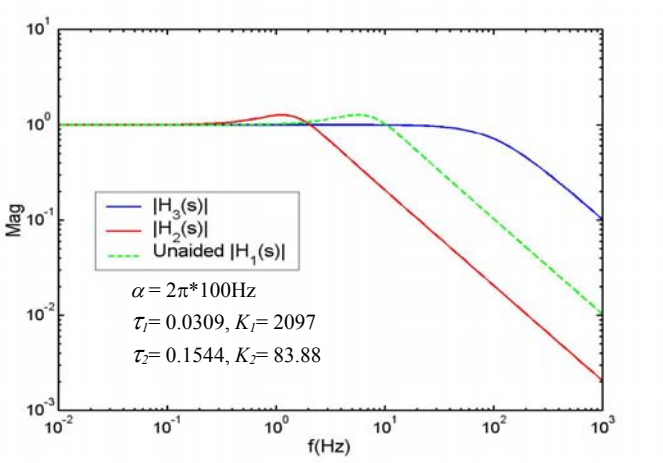

Figure 12. PLL Transfer Functions

B. PLL Filter Design For Ultra-Tight Integration

As mentioned in the last section, the bandwidth of the PLL loop filter for ultra-tight integration will be determined by the highest frequency components of f_{PLL} that cannot be estimated externally. These components may include low frequency noise in the clock-error, or errors in the external Doppler estimates. Since most consumer GPS receivers are equipped with an internal TCXO reference oscillator, the clock errors may well be the driving factor in the design of the PLL loop filter.

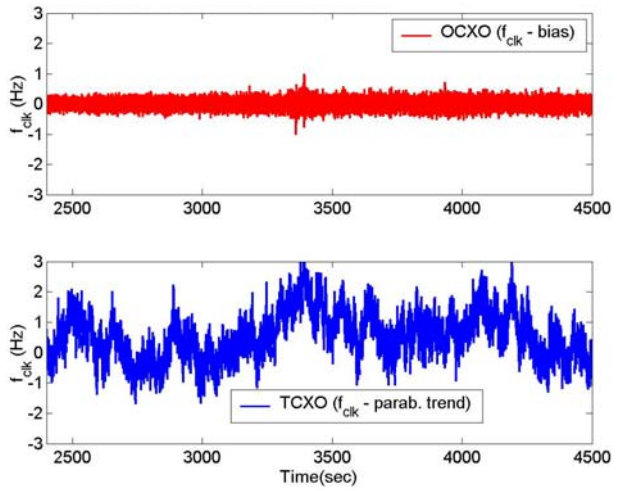

Figure 13. Noise Components of Clock-Error Frequencies

Figure 13 shows the same clock noise data as Fig. 10, but with the bias (for the OCXO) or parabolic trend (for the TCXO) removed to emphasize the noise component in each case. The OCXO noise appears to be mostly white, while the TCXO noise contains a salient low frequency component (exponentially correlated noise) in addition to white noise. Since the low frequency trends of the clock noise must be tracked by the PLL, it is evident that a GPS system with a TCXO will require a higher bandwidth PLL than that with an OCXO.

The power spectral densities (PSD) of this data are shown in Fig. 14, and are a revealing indicator of the bandwidth needed to track the colored components of the clock noise. The transition between colored noise and white noise specifies the required bandwidth to track the signal [6], and occurs at the frequency where the PSD stops decreasing with increasing frequency. In Fig. 14, the piecewise linear fits through the data show that this transition takes place at about 3Hz for the TCXO, and at about 0.1 Hz for the OCXO.

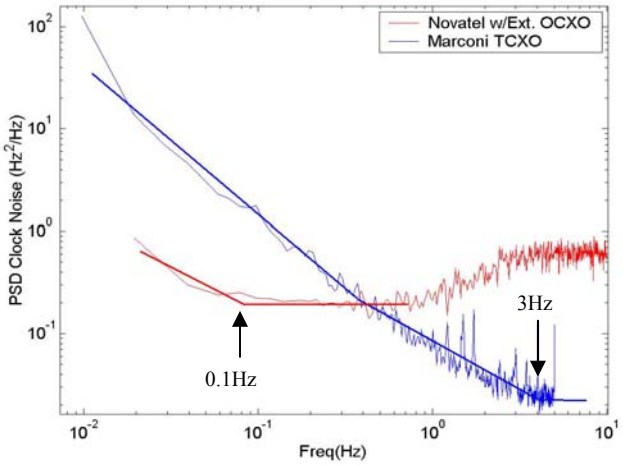

Figure 14. Power Spectral Densities of Clock-Error Frequencies

Other components of the system, especially the inertial sensors and attitude system, must also be considered when specifying the bandwidth of the loop filter. Specifically, any errors introduced by these external components into the loop through the estimated Doppler, must be tracked by the PLL. For example, a system equipped with high-grade inertial sensors with very steady and predictable biases would allow for a lower bandwidth PLL than that with automotive or consumer-grade inertial sensors, whose biases may also have significant colored noise components. If these biases were not calibrated precisely at a particular point in time, the error hence introduced would have to be corrected by the PLL.

C. Implementation of PLL for Ultra-Tight Integration

In many low-cost applications of ultra-tight integration, it may be necessary to have a GPS system in operation before the external Doppler frequency estimates are available. The most common example of this

limitation is using inexpensive inertial sensors with unknown biases, and using normal GPS attitude and velocity outputs to calibrate the biases. After the inertial sensors are properly calibrated, the attitude and velocity of the vehicle can be measured with the bandwidth needed for ultra-tight integration.

The transition between using a traditional loop filter and ultra-tight integration can take place as shown in Fig. 15. When the external Doppler and clock-error frequency estimates are available, the switch that represents the input to the NCO can simply toggle from the traditional filter output to the branch containing the combined external frequency estimates and output of the new loop filter. The external frequency estimates in Fig. 15 are those discussed in section III ($\hat{f}_{dopp}$, $\hat{f}_{clk}$), or can be filtered versions of these measurements.

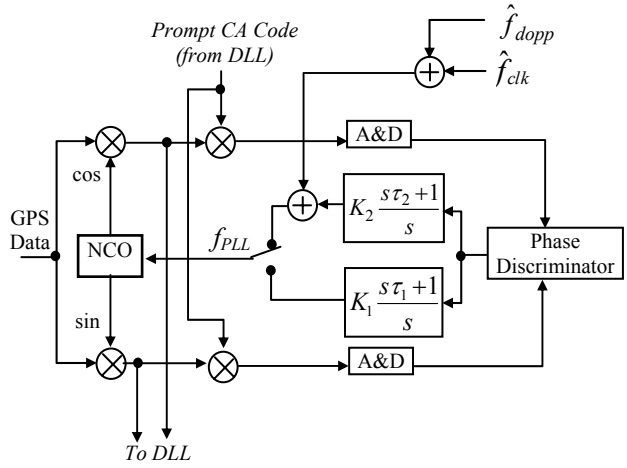

Figure 15. Implementation of a PLL for Ultra-Tight Integration

D. Performance of PLL with Ultra-Tight Integration

The performance results presented here were obtained with the use of the GPS software receiver mentioned briefly in section I, and with simulated GPS data for the trajectory shown in Fig. 16. The simulation of the GPS data included clock noise similar to that of the TCXO shown in Figs. 12 and 13. In addition, the simulations included inertial system outputs with noise characteristics of Bosch DRS-MM1 automotive grade sensors, and GPS/INS blending filters for velocity and attitude. The phase-lock loop implemented in the software receiver has the structure shown in Fig. 15, and the bandwidths shown in Fig. 12 for the two loop filters. The transition to ultra-tight integration occurs when the attitude and velocity filters have converged, shortly after the first PVT measurement.

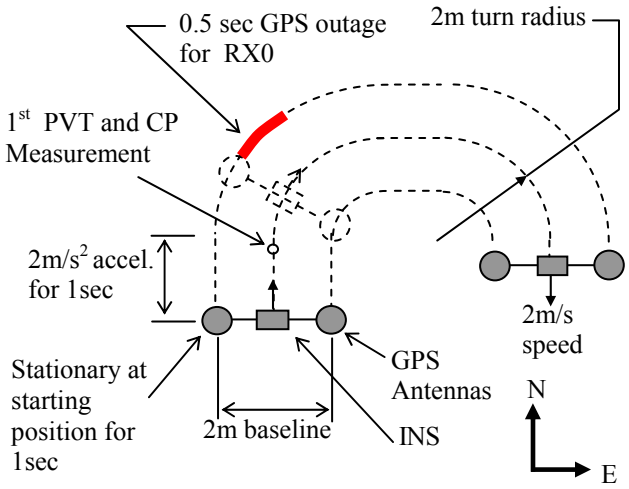

Figure 16. Simulated Test Trajectory

Figure 17 shows f_{PLL} for one channel and for both antennas as tracked by the phase-lock loop before and after ultra-tight integration is invoked. As expected, the noise in the tracked frequency is reduced by about 14dB with ultra-tight integration, yet the loop is able to track the rapid changes in Doppler frequency that occur as turns begin and end.

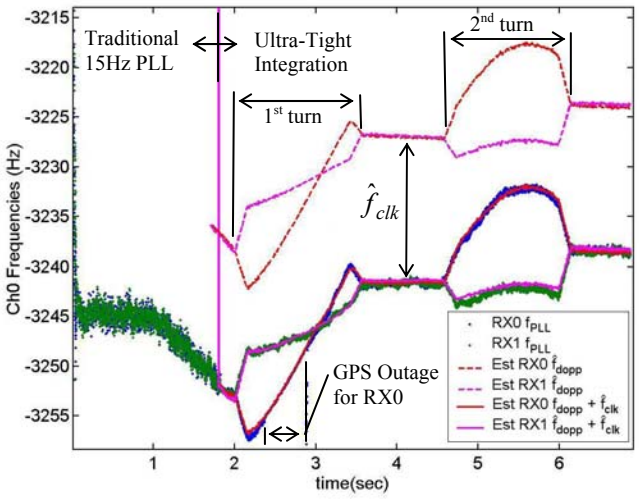

Figure 17. PLL Noise Performance with Ultra-Tight Integration

Figure 17 also illustrates how f_{PLL} is offset from the external Doppler frequency estimates ($\hat{f}_{dopp}$) by $\hat{f}_{clk}$. The $\hat{f}_{clk}$ quantity is estimated online by low-pass filtering the measurements obtained with Eqs. 7-8. Furthermore, it can be seen that the quantity $\hat{f}_{dopp} + \hat{f}_{clk}$ (feed-forward component) approximates f_{PLL} very well, such that the output of the loop filter (feedback

component) has mean of virtually zero when ultra-tight integration is enabled. The feedback component corrects errors in the feed-forward component in the long term, but in the short term it may not be needed to predict the value of f_{PLL} required to maintain the phase of the replicated carrier aligned. This concept can be used during a short GPS outage to maintain the replicated carrier within a fraction of a half-cycle of the true carrier of the GPS signal, and the integrated carrier phase measurement will have no cycle slip when the GPS signal is recovered. Robustness against cycle slips is of particular importance to any application that uses integrated carrier phase, such as attitude determination [1,9] or carrier-phase positioning [12,15]. Such applications require integer ambiguity searches each time a cycle slip occurs.

The trajectory depicted in Fig. 16 includes a half-second GPS outage for one of the receivers during the first turn; the outage is also labeled in Fig. 17, during which the feedback component is switched off and f_{PLL} consists only of the feed-forward component. Figure 18 shows the differential carrier phase between the two antennas for one of six channels being tracked. The enlarged detail of the plot shows a small transient that occurs when the GPS signal for the channel is reacquired, as the phase of the replicated carrier is realigned. No cycle slip occurred in this case, however, since the feed-forward component was able to maintain the phase of the replicated carrier within half a cycle of the true carrier.

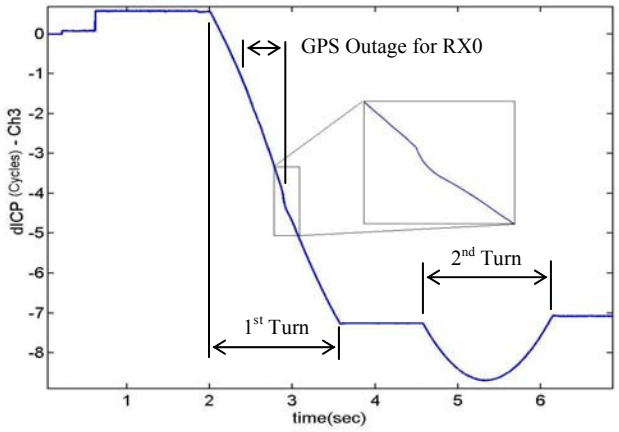

Figure 18. Differential Carrier Phase Measurement with Ultra-Tight Integration

The transient in f_{PLL} after the GPS outage is due to divergence in phase between the true and replicated carrier signals while the feedback component is disabled. Divergence of the replicated carrier phase can be caused by the drifting clock-error frequency or inertial sensor biases, since estimates for these quantities cannot be updated during total GPS outages. Therefore, the duration of a GPS outage that can be tolerated by the system without a cycle slip depends on the quality of the GPS reference oscillator and inertial sensors. That is, a system using an OCXO in conjunction with navigation or tactical grade inertial sensors will be able to withstand a longer outage than that with a TCXO and consumer or automotive grade inertial sensors.

The noise characteristics of the reference oscillator and the inertial sensors can also be used to assess the system's ability to avert cycle slips through GPS signal outages. The more stable the devices' output, the more likely that their behavior can be predicted, and the less phase error that will be accrued by the replicated carrier during the outage.

The Allan variance plot of the devices' output noise component is indicative of their long-term stability [8]. Figure 19 shows the Allan variance for the TCXO data in Fig. 13, and for the sampled output from a stationary Bosch DRS-MM1 rate gyro.

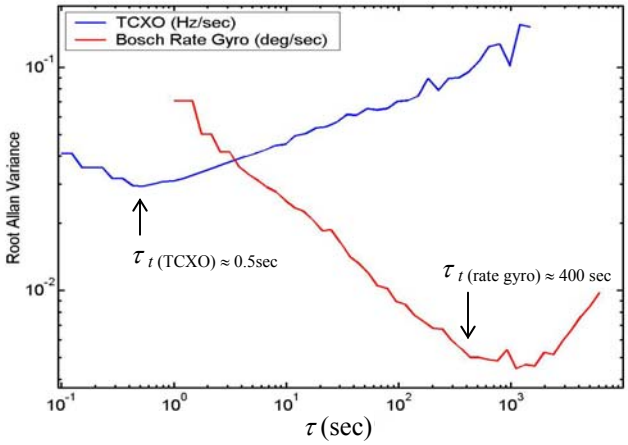

Figure 19. Allan Variance Plots for TCXO Clock-Error Frequency and Bosch DRS-MM1 Rate Gyro

The Allan variance plot shows the variance of averaged sections of the data (on the vertical axis) as a function of the length of the averaging, τ (on the horizontal axis). For signals containing a combination of broadband noise and exponentially correlated noise (Gauss-Markov process), the Allan variance plot is characterized by a section of decreasing variance with increasing τ, followed by a transition region, and finally a section of increasing variance with increasing τ. The latter section of the plot is due to the correlated component of the noise, and the transition region indicates the point at which longer averaging will not reduce the Allan variance. The stability of the device is characterized by the value of τ where the transition region begins (τ_t), and more stable devices will have higher values of τ_t.

Since the system's cycle slip robustness during a GPS outage will be limited by the least stable component, a system with the TCXO and gyros examined in Fig. 19 could handle an outage on the order of 0.5 seconds (τ_t for the TCXO) without a cycle slip. During a longer outage, the correlated noise in the clock-error frequency is more likely to have drifted enough to cause a cycle slip. This inexpensive set of components may not provide sufficient cycle-slip robustness for multiple second GPS outages, but it would be adequate in applications where most cycle slips are caused by momentary LOS blockage which may only last fractions of a second, such as in GPS applications for automobiles in urban areas [1].

E. CA Code Generation During GPS Outages

As mentioned in section II, the PLL and DLL are coupled tracking loops; the DLL provides the prompt code used in the PLL, and the PLL provides in-phase and quadrature carrier replicas to the DLL for code tracking. To successfully recover the GPS signal after an outage, the replicated code must also be maintained to within half of the correlator width of the true code, such that the feedback component of the loop can reacquire lock when the GPS signal is recovered. This task is considerably easier than maintaining the replicated carrier within half a cycle of the true carrier, since the typical width of a traditional correlator is half of a chip (1 chip=293m), much longer than a carrier cycle (1 cycle = 19cm).

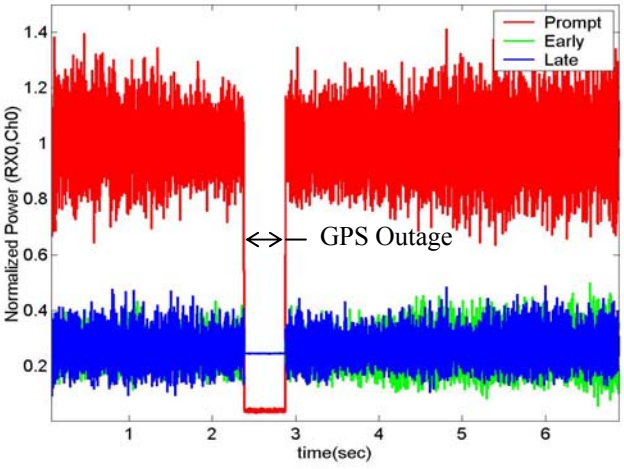

Figure 20. Code Tracking with Rate-Aided DLL During GPS Outage

The use of a rate-aided DLL in an ultra-tightly integrated system can provide sufficient means of maintaining the code phase aligned during an outage. As with the PLL, the feedback component of the DLL can be disabled during the outage, and the code generator can be driven only from the feed-forward branch shown in Fig. 7. Figure 20 shows the prompt, early, and late code powers for the same channel studied in Fig. 17, and for the receiver with the GPS outage. There is no transient when the GPS signal is recovered, indicating that the replicated CA code was nearly aligned with the true CA code during the outage, with only the feed-forward branch of the rate-aided DLL activated.

CONCLUSIONS

This paper has presented the development of a simple architecture that can be used to achieve ultra-tight integration of GPS and inertial sensors, and some of the important issues that influence the system design. The results of comprehensive system-level simulations have shown that this architecture is flexible enough to function with most types of inertial sensors and GPS reference oscillators. While higher quality components improve performance, the results of the case study investigated show that the benefits of ultra-tight integration can be achieved even with automotive-grade inertial sensors and common TCXO oscillators. The simulated ultra-tightly coupled system with these inexpensive components demonstrated meaningful advantages over a system with traditional tracking loops, including a 14dB improvement in phase-noise suppression, higher phase-tracking bandwidth (comparable to bandwidth of inertial sensors), and robustness against carrier-phase cycle slips.

Continuing efforts in this research will focus on the verification of the system performance with real GPS data, and a comparison with the performance attained with more complex, optimal estimation techniques of ultra-tight coupling.

ACKNOWLEGEMENTS

The authors gratefully acknowledge the contributions of friends and colleagues at Stanford University that have contributed to this work. We also give special appreciation to Bosch RTC for their continued interest and funding of this research.

REFERENCES

[1] Alban, Santiago, "An Inexpensive and Robust GPS/INS System for Automobiles", Proceedings of the ION-GPS 2002, Portland, OR, September 2002.

[2] Benson, Donald O., "A Comparison of Two Approaches to Pure-Inertial and Doppler-Inertial Error Analysis," IEEE Transactions on Aerospace Electronic Systems, Vol AES-11, No. 4, July 1975

[3] Besser, J. et al., "Trunav: A Low-Cost Guidance/Navigation Unit Integrating a SAASM-Based GPS and MEMS IMU in a Deeply Coupled Mechanization", Proceedings of the ION-GPS 2002, Portland, OR, September 2002.

[4] Brenner, Mats, "Integrated GPS/Inertial Fault Detection Availability", *Navigation*, Vol 43, No. 2, 1996. pp. 339-358.

[5] Diesel, J. W. and S. Lu, "GPS/IRS AIME: Calculation of Thresholds and Protection Radius Using Chi-Square Methods", Proceedings of ION-NTM, Palm Springs, CA. Sep. 1995. pp 1959-1964.

[6] Enge, Per and Pratap Misra, "Global Positioning System: Signals, Measurements, and Performance", Ganga-Jamuna Press, 2001

[7] Farrell, Jay and Matthew Barth, *The Global Positioning System and Inertial Navigation,* McGraw-Hill, New York, NY, 1999

[8] Gebre-Egziabher, Demoz, "Design and Performance Analysis of a Low-Cost Dead-Reckoning Navigator for General Aviation Applications", Ph.D. Thesis, Stanford University, 2001

[9] Gebre-Egziabher, Demoz, Roger Hayward and J. David Powell,, "A Low-Cost GPS/Inertial Attitude Heading Reference System (AHRS) for General Aviation Aircraft", Proceedings of the IEEE Position, Location and Navigation Symposium - PLANS '98, Palm Springs, California, April 1998

[10] Goshen-Meskin, Drora, Itzhack Bar-Itzhack, "Unified Approach to Inertial Navigation System Error Modeling," AIAA Journal of Guidance, Control and Dynamics, Vol 15, No. 3, May-June 1992

[11] Gustafson, Donald et al., "A High Anti-Jam GPS Based Navigator", Proceedings of the ION-NTM 2000, Anaheim, CA, January 2000.

[12] Heo, M.B, Boris S. Pervan, et. al. "Airborne Autonomous Fault Detection for Shipboard Landing Navigation Using Carrier Phase DGPS," Proceedings of the ION-NTM 2003, Anaheim, CA, January 2003.

[13] Hoffman-Wellenhof, Bernhard, H. Lichtenegger and J. Collins, *GPS: Theory and Practice*, 5th Edition, Springer-Verlag Wien, New York, NY, 2001, pp.40-45.

[14] Kraye, Christian et al., "Performance Analysis and Development of a Tightly Coupled GNSS/INS System", Institute of Geodesy and Navigation, University of the Federal Armed Forces, Munich, 2002

[15] Pervan, Boris S, Fang-Cheng Chan, et. al, "System Concepts for Cycle Ambiguity Resolution and Verification for Aircraft Carrier Landing," Proceedings of ION-GPS 2001, Salt Lake City, UT. September 2001.

[16] Phillips, Richard E. and George T. Schmidt, "GPS INS Integration", AGARD Lecture Series MSP LS 207 on System Implications and Innovative Application of Satellite Navigation, NATO, Paris, July 1996.

[17] Scherzinger, B., "Precise Robust Positioning with Inertial/GPS RTK", Proceedings of the ION-GPS-2000, Salt Lake City, UT, September 2000

[18] Siours, George M., *Aerospace Avionics Systems: A Modern Synthesis*, Academic Press Inc., San Diego, CA 1993. pp 324-333.

[19] Anon., *NAVSTAR GPS Space Segment/Navigation User Interfaces*, Interface Control Document No. ICD-GPS-200C, Department of Defense, June 2000