Architectures for Joint GPS/LEO Satellite Carrier Phase Receivers Designed for Rapid Robust Resolution of Carrier Cycle Ambiguities on Mobile Platforms

M. Rabinowitz, B.W. Parksinson, K. Gromov, *Stanford University*,
C.E.Cohen, *IntegriNautics Corporation*

September 20, 2000

Abstract

The emerging satellite telecommunication industry has seen substantial progress in the design and manufacture of low-cost, low-power transceivers which interface directly with Low Earth Orbit (LEO) Telecom Satellites. Herein, we describe the architecture of a joint GPS-LEO navigation receiver which tracks LEOS in conjunction with GPS in order to rapidly resolve cycle ambiguities on both the GPS and LEO signals. The objective is to achieve centimeter-level navigation performance using pre-existent LEO transceiver hardware, designed for low-cost data communication. The obstacles to be overcome include: synchronization of the LEO and GPS receivers, precise time-tagging of carrier-phase measurements without the use of hardware accumulators and latches, accommodating different frequency-dependent phase lags for each of the receiver paths, and the ability to scale the number of trackable LEO satellites and/or antennas in firmware. Different techniques for performing the integrated phase measurement, consistent with these objectives, are discussed. We describe the radio-frequency hardware, together with an efficient information smoother which we apply to the estimation of cycle ambiguities.

1 Introduction

Accuracy, robustness, power consumption and cost are some crucial performance criteria for precision navigation systems. In recent years, the makers of differential navigation receivers have been pushing the performance envelope by advancements in VLSI design, signal processing techniques, enhanced differential correction data, and precision radio-frequency circuitry. As background, we will consider some of the leading techniques for precision navigation.

Navcom Technologies *NCT-2000D* [12] is an integrated L-band INMARSAT/GPS receiver for high-accuracy differential GPS applications. The dual-band GPS receiver tracks C/A, P1 and P2 code with a purported carrier-measurement precision on the order of 1 centimeter. Navcom's WADGPS system achieves positioning accuracy on the order of 30 centimeters over the CONUS, using carrier-smoothed code [5] corrected for ionospheric effects, based on differential corrections from 8 reference stations, and broadcast by Geostationary satellites. Trimble's *BD750* is a dual frequency RTK receiver which decrypts P-code using the *Supertrak*^{TM} technology. Position accuracies of 2 centime-

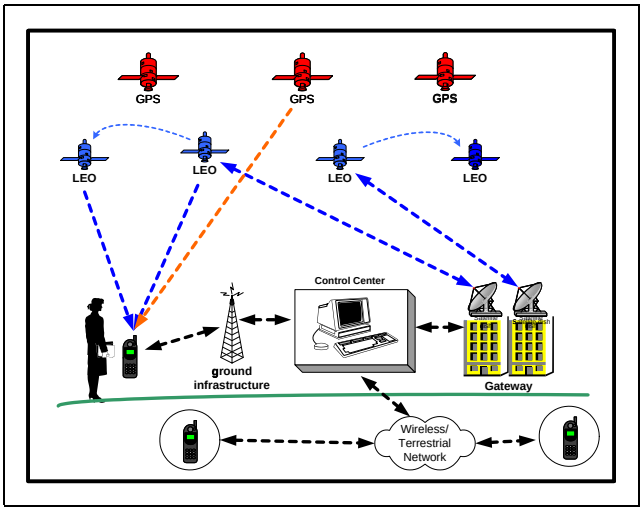

Figure 1: Architectural overview

ters horizontal and 3 centimeters vertical are achieved with latencies of less than 20 milliseconds, and baseline distances of up to 10km. These, and similar, high-end receivers are priced for retail at several thousand Dollars. In the low cost, low-power arena, TChip Semiconductor has issued an advance information data sheet describing the *TJ1004*, a superheterodyne GPS front end chipset, which integrates an on-chip oscillator, and generates a sign-magnitude output. The chip purportedly draws 6mA of current in the fully active state.

The performance of the receivers described by this paper hinges on the enhanced geometry achieved by combining GPS with LEO satellites. We describe algorithmic techniques which facilitate reduced cost and power consumption in these centimeter-level navigation receivers. The system we describe can achieve centimeter-level accuracy without any need for expensive dual-frequency GPS hardware. For example, a front end with similar performance criteria to those of TChip, described above, could be employed for the GPS segment.

2 System Overview

Fig. 1 presents a schematic overview of the system [9, 10, 3, 7, 8]. The central components are the Navstar GPS satellites, the LEO satellites, the user and reference receivers. The user and reference track the absolute carrier phase of the Navstar satellite signals together with the absolute carrier phase of multiple LEO satellite signals. The term "absolute" means that the phase measurement is accumulated over time and is not modulus 2π. The geometric diversity achieved primarily by the motion of the LEOS enables the user receiver to resolve the integer cycle ambiguities on the Navstar satellite signals as well as parameters related to the cycle ambiguities on the LEO signals, and consequently to position itself with cm-level accuracy relative to the reference receiver. The reference receivers are similar to the user receiver, and track only a single frequency downlink from each of the visible satellites. The reference receivers then convey their carrier-phase measurements to a control center where the differential corrections are computed. This differential data is conveyed to the user receiver via a satellite data channel, or via terrestrial infrastructure. Thus far, only a terrestrial data connection has been employed. By using reference phase prediction techniques, the data channel needs to be active in bursts every 5-10 seconds, where each burst conveys less than 1 Kbit of data.

3 The Effects of Geometry

This section present a brief intuitive explanation for the geometric benefits of the LEO signal. Fig. 2 describes a hypothetical scenario where a user tracks a single LEO satellite, together with two GPS satellites. For the sake of simplicity, the reference station is ignored, the signals are assumed noise-free, the user receiver clock has no offset from GPS time, and the user as well as the GPS satellites are assumed to be stationary in EBEF co-ordinates over the course of tracking. Based on the phase measurement of the GPS signals, the user receiver can be placed at any of the vertices of the indicated lattice. The structure of this lattice is determined by the geometry of the GPS satellites, as well as the wavelength of the GPS L1 signals. Since the cycle ambiguity, or bias parameter, for the LEO is not initially known, all the user can directly measure is the change in phase as the LEO moves from one position to another. Each of these measurements place the user on a hyperboloid in space. After several such measurements, the user may be placed at the intersection of all of these hyperboloids. Notice that the user's position estimate is infinitely sensitive to disturbance on the axis perpendicular to the plane of the page. Nonetheless, the additional information from only a single LEO is sufficient to resolve on which vertex of the lattice the user resides.

4 Experimental Results

References [7, 8] describe in some detail the experimental results achieved by this technique, as applied to the

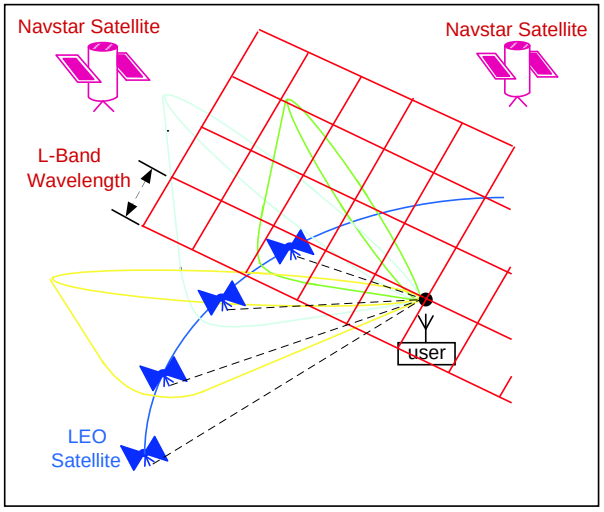

Figure 2: Geometric Interpretation of the Advantages of a LEO Satellite Signal

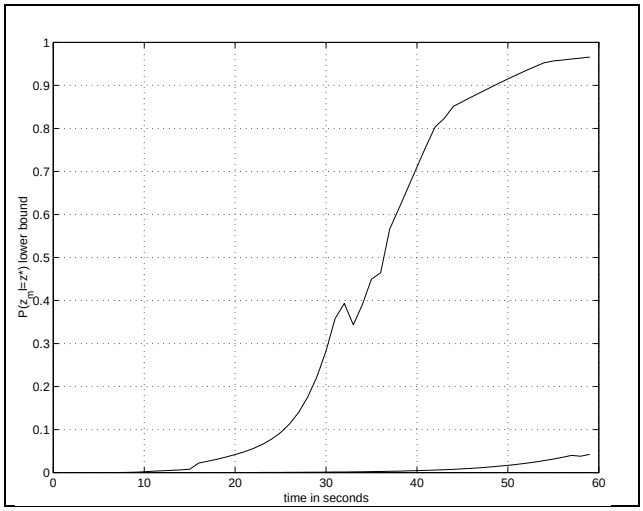

Figure 3: Evolution in the lower bound on the probability of selecting the correct set of integer cycle ambiguities. Probability is plotted using the GPS signals alone, as well as for the GPS signals combined with a single Orbcomm satellite signal.

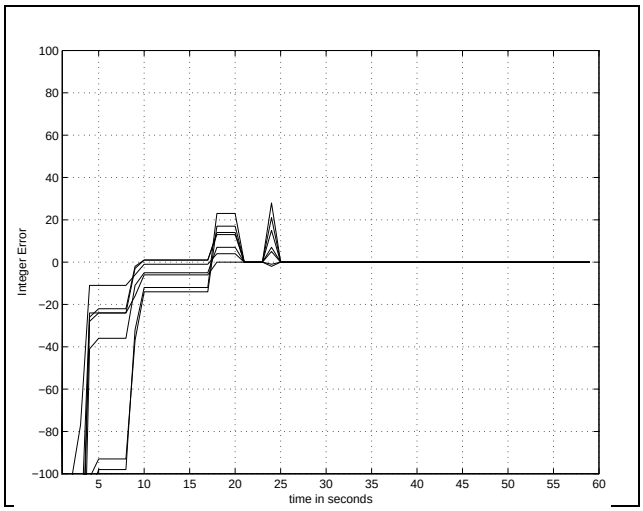

Figure 4: Maximum-likelihood integer estimates over 60 seconds of tracking GPS satellites together with a single Orbcomm satellite. Notice that the cycle ambiguities are resolved after roughly 25 seconds of tracking

Orbcomm satellites in combination with GPS. We will not discuss these results in detail here. Data from a single pass of one Orbcomm satellite will serve to indicate the level of performance for which the receiver was designed. Notice that the LEO passed almost directly overhead when this data was taken. This type of pass is, of course, atypical. Ideally, multiple LEO satellites would be tracked simultaneously. The evolution, in fig 3, of a lower bound on the probability of selecting the correct set of integers illustrates the performance improvement that is provided by the LEO satellite. Note from fig. 4 and fig. 5 respectively that cycle ambiguities are resolved in well under one minute, and that the resulting position errors fit roughly within the volume of a golf-ball.

5 Phase measurement firmware

In this section, we describe the technique by which phase is accumulated and measured for the LEO satellite downlinks. The two techniques which will be discussed are designed for the following objectives:

- The time-tagging of the phase measurements must be synchronized to GPS time with a precision of $100nS$ in order to achieve distance-equivalent errors less than $1mm$ [3, 7].
- The phase-tracking assembly should be implemented without the need for additional hardware, such as phase accumulators or registers which would need to be strobed by signals synchronized to the GPS hardware.

The latter requirement is motivated by the desire that centimeter-level navigation should be achievable on prebuilt receiver hardware (such as portable data devices). While changes to the microprocessor firmware are relatively cost-efficient and easily implemented, changes to device hardware can require considerable additional investment. In addition, as more processing is performed in firmware, rather than hardware, the size, cost and power consumption of receivers goes down.

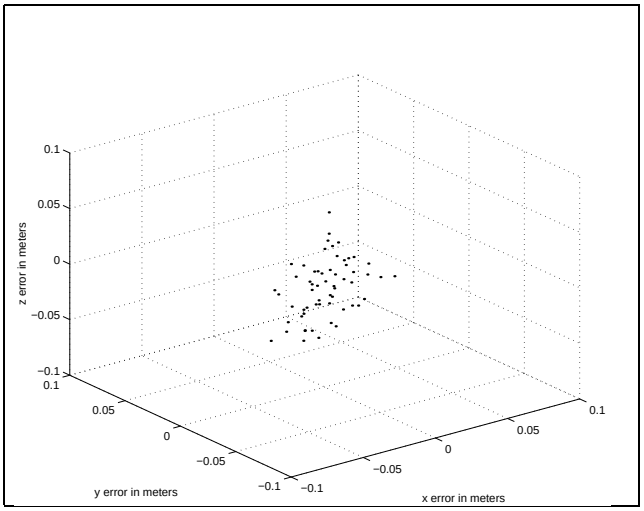

Figure 5: Position errors after resolving the cycle ambiguities on GPS and a parameter related to the cycle ambiguity on an Orbcomm satellite. Notice that 95% errors fit roughly within the size of a golf-ball.

5.1 Implementing an Open Loop NCO

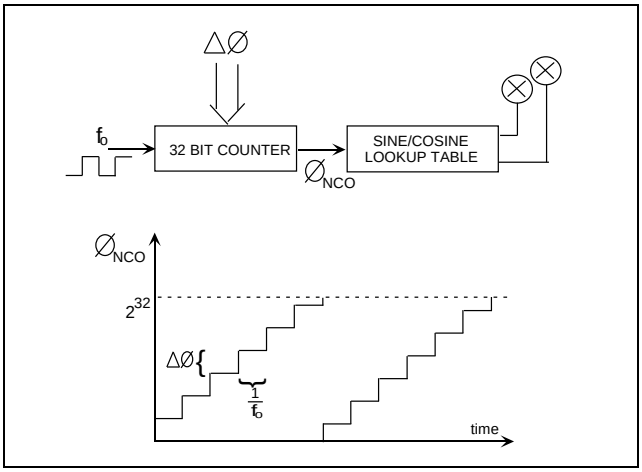

Figure 6: Block structure of a generic hardware NCO

The first phase-measurement technique involves implementing a firmware model of the NCO, which is driven by the same commands that are uploaded to the hardware NCO. Consider the operation of the NCO illustrated in fig. 6. The phase of the NCO is incremented by $\Delta\phi$ every $\frac{1}{f_o}$ seconds, where f_o is the oversampling frequency. The phase output by the NCO is modulus 2^{32}, consequently 2^{32} corresponds to 2π in the sine/cosine lookup table.

The microprocessor is driven by 4 signals:

- Symbol Ready (sym-rdy) - From the LEO correlator hardware to indicate that an integrated symbol is ready for downloading to the microprocessor.
- GPS Marker (gps-mrk) - This is a digital clock signal, of 50% duty cycle, the leading edge of which marks one GPS epoch, synchronized to GPS time.
- Slow Synchronization Signal - This signal is also synchronous to GPS time. It is a pulsed signal, with period substantially slower than gps-mrk. It is used to synchronize the initial sampling times of the LEO and GPS receivers.
- Clock Signal (clk) - This signal is tied to the oscillator of the GPS receiver. The timing standard used by the LEO microprocessor, as well as f_o must be synchronized to this signal. The simplest approach is to drive all the LEO hardware with clk signals coherent to f_{clk}.

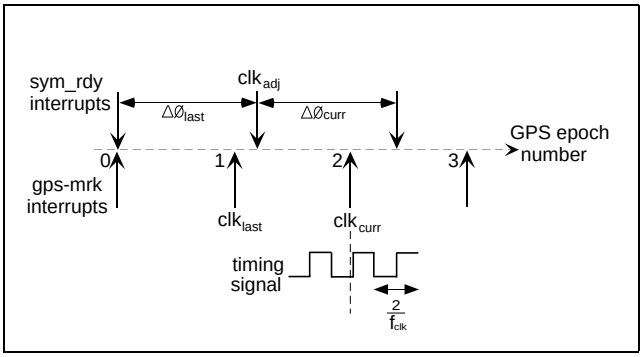

Figure 7: Interrupts which drive the open loop model of the NCO implemented in firmware

In the following description, illustrated in fig. 7 we focus attention on the sym-rdy and the gps-mrk signals, which are fundamental to the phase measurement routine. The sym-rdy marker is a strobe generated by the LEO hardware at the symbol rate of the LEO signal, *symbol-rate*. Sym-rdy causes a hardware interrupt which initiates an iteration of the phase-locked-loop control law, at the end of which a new $\Delta\phi$ is uploaded to the 32-bit counter. The time at which this upload occurs, measured according to the *clk* signal, is clk_{adj}.

The phase measurement in the GPS receiver is associated with the leading edge of the gps-mrk signal. In the LEO receiver, the gps-mrk leading edge initiates an interrupt, which records the current time, clk_{cur}, and measures the change in oscillator phase since the last epoch. Consider the time-line of interrupts as shown in fig. 7, where we have assumed that $\frac{1}{symbol_rate} > 1$ GPS epoch. The labels clk_{last}, clk_{curr}, $\Delta\phi_{last}$ and $\Delta\phi_{curr}$ are assigned assuming the phase increment is being calculated at the epoch marker 2.

At each gps-mrk interrupt, the calculation of phase for the open-loop NCO proceeds as follows:

$$
\begin{aligned}
&If\ clk_{adj} > clk_{last} \\
&\quad \phi_{curr} = \phi_{last} + (clk_{adj} - clk_{last})\Delta\phi_{last} \\
&\qquad +(clk_{curr} - clk_{adj})\Delta\phi_{curr} \\
&else \\
&\quad \phi_{curr} = \phi_{last} + (clk_{curr} - clk_{last})\Delta\phi_{curr} \qquad (1)
\end{aligned}
$$

A discussion of the accuracy of this open loop NCO is in the appendix.

5.2 Implementing the Phase Tracking Assembly in Firmware

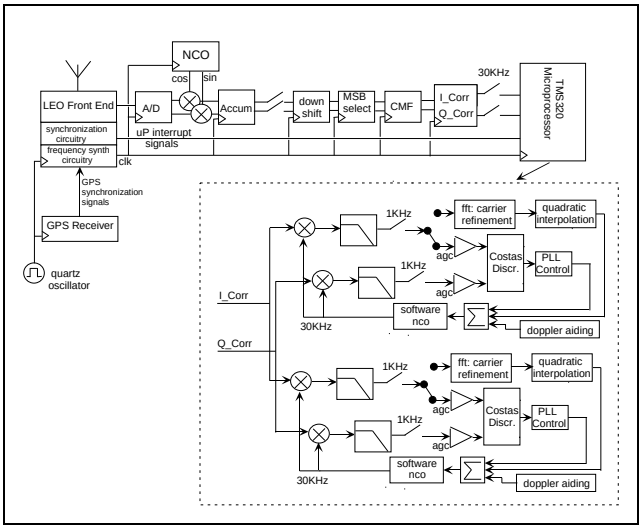

Figure 8: Design of a software PLL for tracking multiple Orbcomm satellites

While the previous approach is computationally cheap, the system is not scalable i.e. a dedicated hardware NCO is necessary for each satellite being tracked. In order to surmount this problem, we have developed firmware to implement the entire Phase Locked Loop (PLL) for multiple parallel channels. The block diagram of fig. 8 describes the operation of the software PLL. The technique described here was implemented on a TMS320 Microprocessor operating at a clock speed $f_{clk} \approx 20MHz$, and with the assumption of a maximum signal bandwidth (including Doppler) of 15kHz [1]. It will be clear from the description how the relevant parameters would change for different hardware and signal configurations. The first digital mixer multiplies the incoming signal with a constant tone $1000Hz$ below the lowest expected spectral component of the signal. The signal is down-shifted to avoid saturation of the correlators, and again filtered to enhance SNR out of the correlator outputs. [2] The correlator output is sampled at 30KHz, just above the Nyquist rate [11]. The three segments of each channel's PLL which are clocked at 30kHz, and for which code needs to be very efficient, are the software NCO, the mixing, and the filtering of the mixed signal.

The NCO is implemented using a cosine lookup table, the filter is implemented using a direct-form II [11] implementation of a $200Hz$ Butterworth lowpass filter. The rate of phase accumulation by the NCO is determined by 3 inputs: the Doppler-aiding parameters, the frequency offset estimated by the carrier refinement algorithm, and the output of the pll control law. The Doppler-aiding parameters are the coefficients of a third-order polynomial model of the Doppler-shifted frequency as a function of time. This model is derived from simple ephemeris prediction software using Keplarian TLE's, and a second order Earth Gravity Model [14]. The Doppler-aiding parameters determine the initial frequency of the NCO, which is accurate to within $100Hz$ for Orbcomm. The bandwidth of the digital filter must be large enough to accommodate the initial NCO frequency error, but small enough to filter out the signals of other satellites. The firmware implementing the mixers, filters and NCO for the two channels executes each iteration in 250 cycles of f_{clk}. Five LEO channels can co-exist without exceeding the processing capabilities of the TMS320 microprocessor. The output of the filters is sampled at a rate of $1KHz$. Since the bandwidth of the pll is roughly $10Hz$, the NCO frequency must be refined before the pll control law can be closed. This refinement is implemented by means of a 128-point FFT of the filtered, downsampled signal, followed by a quadratic interpolation to identify the residual carrier offset. Once the carrier is refined, the phase-offset is sensed with a Costas Discriminator. A digital AGC is used to achieve unity magnitude for the signal entering the discriminator. The gain variation of the AGC is filtered to a bandwidth of roughly 7Hz. The high-speed segment of the PLL is driven by the sym-rdy interrupts, the slower segment is driven by the gps-mrk interrupts. In order to avoid memory conflicts, the phase computation is performed only in the sym-rdy interrupt routine. The gps-mrk interrupt routine records a timestamp, and the phase associated with that timestamp is found by interpolation.

[1] Based on the UHF downlink of the Orbcomm satellites, operating in the UHF band assigned to *Little LEOS*

[2] In the case of Orbcomm, a non-CDMA signal, the correlator code is set to all 1's to implement a low-pass filter.

5.3 Construction of Receiver Hardware

Fig. 9 displays the *Mark I* hardware which was implemented for tracking the Transit satellites of the (now retired) Navy Navigation Satellites System.

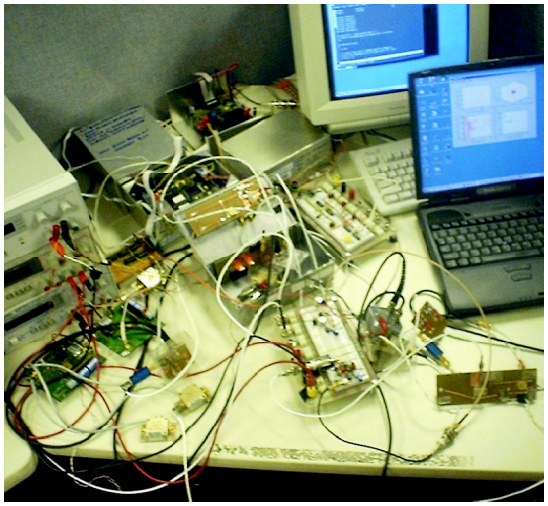

Figure 9: Receiver hardware, *Mark I*, designed for Transit Satellite Downlink. Both user and reference receivers are shown.

The hardware for Orbcomm was re-designed to fit onto a single pc-board, so that the joint GPS and LEO receivers could be co-located in a single RF-shielded $20cm \times 12cm \times 10cm$ box. Fig. 10 shows the *Mark II* hardware for tracking the UHF downlink of the Orbcomm satellites, and other *Little LEO* constellations operating at UHF.

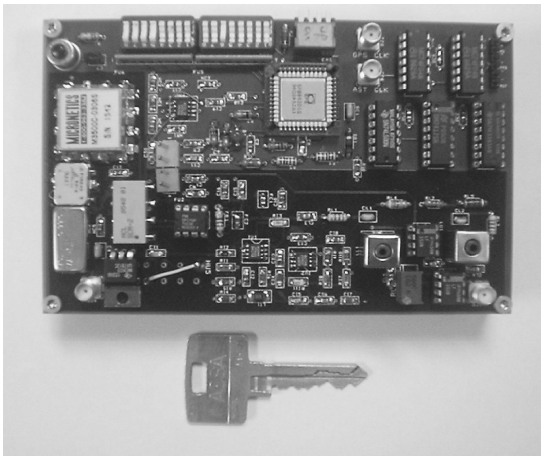

Figure 10: Mark II: Receiver hardware for tracking the Orbcomm Satellite downlink synchronously with GPS.

The board consists of 4 layers; all RF lines requiring shielding are on layers two and three. These lines are enclosed in a Faraday cage constructed by creating copper pour zones on the top and bottom layer, and connecting them by vias. Care was taken to spatially separate the RF and digital signals on the board, as well as to maintain a separate analog and digital ground. The RF (analogue) ground is floating, and is connected to the digital ground by a 10 Ohm chip resistor.

6 Receiver Software

6.1 Estimation Strategy

In this section, we describe an iterative smoothing technique which estimates integer cycle ambiguities together with the relative phase-lags between different paths of the receiver front end. The technique is very efficient computationally, so that integers can be resolved in near real-time, and phase can be sampled at a rate of $10Hz$ or more. For an in depth discussion of the data reduction technique, refer to [3, 7]. In this explanation, we will skip the details of how the phase measurements are post-processed and how the clock and cycle ambiguity parameters are redefined. We will simply assume that we can decompose the set of parameter updates into a set of time-dependent updates, namely position and clock estimates $\Delta\theta(t)$, and time-independent updates, $\Delta\mathbf{n}$:

$$\Delta\mathbf{x}(t) = \begin{bmatrix} \Delta\theta(t) \\ \Delta\mathbf{n} \end{bmatrix} \tag{2}$$

The elements of $\Delta\mathbf{n} = [\Delta N_2 \dots \Delta N_S]^T$ are treated as real numbers by the smoothing routine, however all save one are integer valued [7]. We assume, without loss of generality, that the real-valued parameter, which accommodates the difference in phase lag between the LEO and GPS channels, is ΔN_S. We model the evolution of parameters between phase measurements as a Markov process

$$\Delta\mathbf{x}(t_k) = \Delta\mathbf{x}(t_{k-1}) + \mathbf{w}(t_k) \tag{3}$$

where

$$\begin{aligned} E\left\{\mathbf{w}(t_k)\mathbf{w}(t_k)^T\right\} &= \mathbf{Q}(t_k) \\ &= \lim_{\sigma_w \to \infty} \begin{bmatrix} \mathbf{I}_4{\sigma_w}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \end{aligned} \tag{4}$$

This process [2] is based upon the idea that the parameters $\mathbf{n}$ are constant for all phase measurements and that no assumption is made about the user's motion, nor about the clock drift of the receiver, between measurements. After post- processing the integrated phase measurements, we may model the phase prediction error at time t_k with a linear approximation:

$$\Delta\mathbf{y}(t_k) = \mathbf{H}(t_k)\Delta\mathbf{x}(t_k) + \mathbf{v}(t_k) \tag{5}$$

where $\mathbf{v}_{t_k}$ is a disturbance with distribution $\tilde{N}(\mathbf{0}, \mathbf{R}(t_k))$. $\mathbf{H}(t_k)$ is the *estimation matrix* [3, 7] and $\Delta\mathbf{x}(t_k)$ is the set of updates to the parameter estimates. Denoting the covariance of our post-measurement parameter estimates $\Delta\hat{\mathbf{x}}(t_k)$ as $\mathbf{P}(t_k)$, the Kalman filtering equations for this system are [6]:

$$\begin{aligned}\mathbf{P}(t_k)^{-1} &= (\mathbf{P}(t_{k-1}) + \mathbf{Q}(t_{k-1}))^{-1} \\ &\quad +\mathbf{H}(t_k)^T\mathbf{R}(t_k)^{-1}\mathbf{H}(t_k)\end{aligned}$$

$$\begin{aligned}\Delta\hat{\mathbf{x}}(t_k) &= \Delta\hat{\mathbf{x}}(t_{k-1}) + \mathbf{P}(t_k)\mathbf{H}(t_k)^T\mathbf{R}(t_k)^{-1} \\ &\quad (\Delta\mathbf{y}(t_k) - \mathbf{H}(t_k)\Delta\hat{\mathbf{x}}(t_{k-1}))\end{aligned} \tag{6}$$

If phase-lock is attained on a new satellite s during parameter estimation, the initial covariance of the parameter estimate $\Delta\hat{N}_s$ is very large; this can cause computational difficulties. Consequently, we use instead an information form of the Kalman filter ([2], [13], [4]) where we define an information matrix

$$\begin{aligned}\mathbf{S}(t_k) = \mathbf{P}(t_k)^{-1} &= \begin{bmatrix} \mathbf{P}_\theta(t_k) & \mathbf{P}_{\theta n}(t_k) \\ \mathbf{P}_{\theta n}(t_k)^T & \mathbf{P}_n(t_k) \end{bmatrix}^{-1} \\ &= \begin{bmatrix} \mathbf{S}_\theta(t_k) & \mathbf{S}_{\theta n}(t_k) \\ \mathbf{S}_{\theta n}(t_k)^T & \mathbf{S}_n(t_k) \end{bmatrix}\end{aligned} \tag{7}$$

and an information vector $\mathbf{v}(t_k) = \mathbf{S}(t_k)\Delta\hat{\mathbf{x}}(t_k)$. These definitions, together with the disturbance model of equ. (5) allow us to considerably reduce the smoother computationally. By taking the limit of $\sigma_w \to \infty$, the updates (6) can be implemented using block-matrix operations:

$$\begin{aligned}\mathbf{S}(t_k) &= \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{P}_n(t_k)^{-1} \end{bmatrix} + \mathbf{H}(t_k)^T\mathbf{R}(t_k)\mathbf{H}(t_k) \\ &= \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_n(t_{k-1}) - \mathbf{S}_{\theta n}(t_{k-1})^T\mathbf{S}_\theta(t_{k-1})^{-1}\mathbf{S}_{\theta n}(t_{k-1}) \end{bmatrix} \\ &\quad +\mathbf{H}(t_k)^T\mathbf{R}(t_k)\mathbf{H}(t_k),\ \mathbf{S}(t_0) = \mathbf{0}\end{aligned} \tag{8}$$

$$\begin{aligned}\mathbf{v}(t_k) &= \mathbf{S}(t_k)\Delta\hat{\mathbf{x}}(t_{k-1}) + \mathbf{H}(t_k)\mathbf{R}(t_k)^{-1} \\ &\quad (\Delta\mathbf{y}(t_k) - \mathbf{H}(t_k)\Delta\hat{\mathbf{x}}(t_{k-1})) \\ &= \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_n(t_{k-1}) - \mathbf{S}_{\theta n}(t_{k-1})^T\mathbf{S}_\theta(t_{k-1})^{-1}\mathbf{S}_{\theta n}(t_{k-1}) \end{bmatrix} \\ &\quad \begin{bmatrix} \Delta\hat{\theta}(t_k) \\ \Delta\hat{\mathbf{n}}(t_k) \end{bmatrix} + \mathbf{H}(t_k)^T\mathbf{R}(t_k)^{-1}\Delta\mathbf{y}(t_k) \\ &= \begin{bmatrix} \mathbf{0} \\ \hat{\mathbf{v}}_n(t_{k-1}) - \mathbf{S}_{\theta n}(t_{k-1})^T\mathbf{S}_\theta(t_{k-1})^{-1}\mathbf{v}_\theta(t_{k-1}) \end{bmatrix} \\ &\quad +\mathbf{H}(t_k)^T\mathbf{R}(t_k)^{-1}\Delta\mathbf{y}(t_k), \mathbf{v}(t_0) = \mathbf{0}\end{aligned} \tag{9}$$

In order to emulate the batch solution, the filter is passed forward and backward over the data. For the backward pass, simply interchange the t_{k-1} and t_k in equ. (8), and start with initial conditions $\mathbf{S}(t_N) = \mathbf{0}$ and $\mathbf{v}(t_N) = \mathbf{0}$. The $\mathbf{S}(t_k)$ and $\mathbf{v}(t_k)$ from forward and backward passes are linearly combined

$$\mathbf{S}(t_k) = \mathbf{S}(t_k)^F + \mathbf{S}(t_k)^B \tag{10}$$

$$\mathbf{v}(t_k) = \mathbf{v}(t_k)^F + \mathbf{v}(t_k)^B \tag{11}$$

and the parameter updates are then found according to $\Delta\hat{\mathbf{x}}(t_k) = \mathbf{S}(t_k)^{-1}\mathbf{v}(t_k)$.

Assuming the noise assumptions are accurate, it makes no difference which $\Delta\mathbf{x}(t_k)$ in the sequence $k = 1 \ldots N$ is used to update the constant parameters $\hat{\mathbf{n}}$. This is born out in practice. It is simplest to take the parameter updates from the last measurement epoch t_N:

$$\hat{\mathbf{n}} \leftarrow \hat{\mathbf{n}} + \Delta\mathbf{n}(t_N) \tag{12}$$

The smoothing and updating process is repeated until the elements of $\Delta\hat{\mathbf{x}}(t_k)$ become negligibly small. When that condition occurs, we check that we may select integers with acceptable integrity by lower-bounding the probability of correct integer detection, as described in [7, 8, 1]. If the integrity is sufficient, one final pass of the smoother is initiated. This final pass sets up the parameter estimates and covariances for an integer search routine. For those parameters that are integer-valued,

we seek to preserve their integer nature in calculating the phase prediction error, so that this information can be exploited in the integer-search operation. Consequently, on the final smoothing pass, before calculating the phase prediction error of equ. (5), the relevant parameters are rounded off:

$$\hat{n}_i \leftarrow \lceil \hat{n}_i \rfloor, \; i = 1 \ldots S-2 \tag{13}$$

After the final smoother pass, we find the covariance of the constant parameters

$$\mathbf{P}_n(t_N) = \left(\mathbf{S}_n(t_N) - \mathbf{S}_{\theta n}(t_N)^T \mathbf{S}_\theta(t_N)^{-1} \mathbf{S}_{\theta n}(t_N)\right)^{-1}$$

The covariance matrix for the integer-valued parameters can now be found from the relevant rows and columns of $\mathbf{P}_n$: $\mathbf{P}_z \leftarrow \mathbf{P}_n(t_N)_{1:S-2,1:S-2}$. Similarly, the preliminary estimates of the integer parameters, as discussed above, are obtained from the estimated bias updates of the final smoother pass $\hat{\mathbf{z}} = [\Delta n_1(t_N) \ldots \Delta n_{S-2}(t_N)]$. We are now ready to search for the integer solution as described in [8].

7 Conclusions: Highlights of this Receiver Architecture

The joint GPS-LEO receiver is able to resolve cycle ambiguities with greater integrity, in less time, than is achieved by stand-alone GPS receivers since it makes use of the geometric diversity provided by the rapid angular motion of the LEOS. In order to minimize the cost and power consumption of the joint receiver, the precisely timetagged integrated phase measurements can be performed in firmware. This facilitates the use of pre-existed hardware designed for data communications, and enables one to scale the receiver to track multiple LEO satellites and/or multiple antennas without the addition of hardware. The information smoothing technique employed to resolve cycle ambiguities on a mobile platform is designed to be computationally efficient. This data reduction procedure typically determines the bottleneck for the rate at which phase can be sampled. The technique for estimating frequency-dependent phase lags of the dual frequency system enables one to achieve centimeter-level accuracy without the use of expensive, pre-calibrated rf components in the receiver front end.

8 Appendix

8.0.1 Accuracy of the Open Loop NCO

It is not intuitive that an open-loop NCO can maintain a measure of the integrated phase which remains accurate enough for cm-level navigation. However, it has been experimentally verified that this is achievable. f_{clk} Must be sufficiently high that time-tagging can be performed to a precision of roughly $0.1\mu S$. Consider the case, illustrated in fig.7 at gps-mrk 2, where the leading edge of the timing signal occurs just after clk_{curr} is recorded. Since this leading edge will be included in the phase measurement at gps-mrk 3, the fractional cycle missed in measuring the current epoch will be included in the subsequent measurements. Hence errors due to time-tagging precision are not cumulative. The accumulated timing pulses counted over several gps-mrk interrupts will be accurate to within 2 cycles of f_{clk}. Since the $\Delta\phi$ are continually changing, inaccuracy will arise if each $\Delta\phi$ is not applied for the correct period of time in equ.(1). Since the phase acceleration of the LEO downlink doesn't exceed $100/Hz/sec$, for the pathological case, this phase error would be roughly $100Hz/sec \times 10^{-7}sec = 10^{-5}cycles/sec$ or $10^{-3}cycles$ after integrating phase for 1 minute 40 seconds. This error is below the system noise floor.

8.0.2 Common-mode Timing Errors for the Open Loop NCO

The TMS320 microprocessor registers one interrupt if the relevant pin is held low for one or two falling edges of an internal clock, which operates at a frequency of $\frac{f_{clk}}{2}$. If the pin is held low for more than two falling edges, multiple interrupts may be recorded, with catastrophic effects on phase tracking. In order to generate the correct pulse at each rising edge of the gps-mrk signal, the hardware of fig. 11 is used. Notice that the delay from the instant of the gps-mrk leading edge to the recording of clk_{cur} in the microprocessor is composed of the digital delay in generating the interrupt pulse, as well as the firmware delay in initializing the interrupt service routine. So long as this delay remains constant over the course of tracking, it is incorporated into the frequency-dependent phase-lag estimated for the LEO front end - see section (6.1). The only variable delay is that between the leading edge of the gps-mrk and the leading edge of the interrupt pulse, indicated as Δclk in the figure. The hardware in fig. 11 maintains this delay below $100nS$.

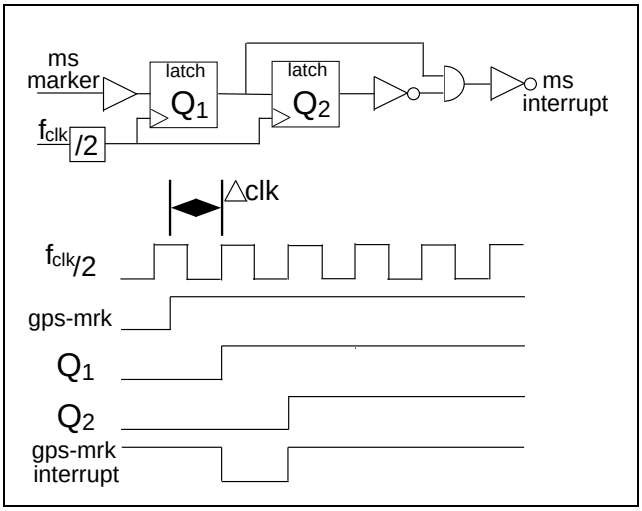

Figure 11: Segment of synchronization hardware for generating the millisecond interrupt signal

8.0.3 Numerical Precision of the Open Loop NCO

The primary error affecting accuracy of the open-loop NCO is quantization noise. The numerical precision with which the calculations of equ. (1) are performed can result in considerable quantization errors for 32-bit floating point processors. This numerical precision issue is resolved using extra 32-bit words, and purely integer multiplication in assembly firmware as illustrated in fig.12.

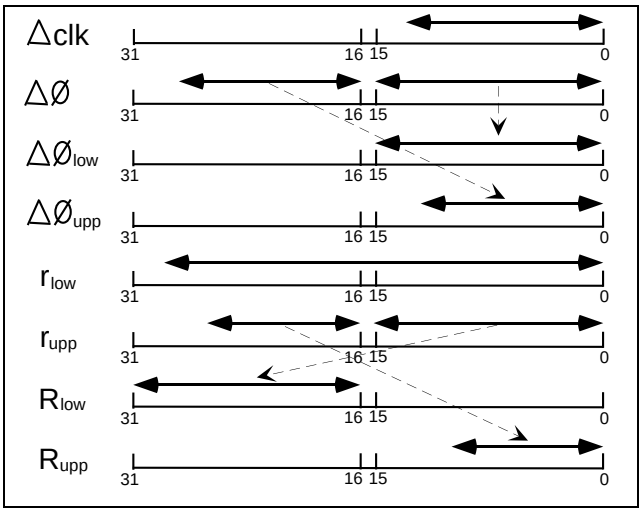

Figure 12: Sequence of steps used to avoid quantization noise for the open loop NCO

The register Δclk stores the number of cycles between ms-mrk interrupts, which requires 14 bits. The register $\Delta\phi$ stores the increment made to the counter upon each leading edge of the oversampling clock, f_o. $\Delta\phi$ requires 29 bits. We initialize a register $\Delta\phi_{low}$ which stores the lower 16 bits of $\Delta\phi$, and $\Delta\phi_{upp}$ which contains the upper 13 bits of $\Delta\phi$, down-shifted. We then perform the integer multiplications:

$$\begin{aligned} r_{low} &= \Delta\phi_{low}\Delta clk \\ r_{upp} &= \Delta\phi_{upp}\Delta clk \end{aligned} \tag{14}$$

The upper 16 bits of r_{upp} are moved to the lower 16 bits of R_{upp}, the lower 16 bits of r_{upp} are moved to bits 16-31 of R_{low}. The accumulation of the phase is then performed

$$\phi_{low} = \phi_{low} + R_{low} + r_{low} \tag{15}$$

and the any carry is stores in c. Then,

$$\phi_{upp} = \phi_{upp} + R_{upp} + c \tag{16}$$

ϕ_{low} Now stores the fractional component of the phase (2^{32} corresponds to 2π), ϕ_{upp} stores the number of complete cycles. Note that these phase measurements need to be scaled by $\frac{f_{clk}}{2f_o}$, the number of cycles of the oversampling clock for each cycle of the timing signal, in order to obtain the true phase measurement.

References

[1] A.Hassibi. Integer paramter estimation in linear models with application to GPS. Technical Report ECS-9222391, Information Systems Laboratory, Stanford University, 1996.

[2] B.S.Pervan C.E.Cohen and B.W. Parkinson. Integrity in cycle ambiguity resolution for GPS-based precision landing. *Third International Conference on Differential Satellite Navigation Systems (DSNS-94), London, UK*, 1994.

[3] M.Rabinowitz et al. A system using leo telecommunication satellites for rapid acquisition of integer cycle ambiguities. *IEEE PLANS 98*, April 1998.

[4] A. Gelb and Technical Staff. *Applied Optimal Estimation*, volume 1. Analytic Sciences Corporation, 1974.

[5] J.J. Spilker Jr. Fundamentals of signal tracking theory. *Global Positioning System: Theory and Applications*, pages 245–327, 1994.

[6] R.E. Kalman and R.S. Bucy. New results in linear filtering and prediction. *ASME Transactions, Journal of Basic Engineering*, 83:95–108, 1961.

[7] M.Rabinowitz. *PhD Thesis: Precision Navigation Using Low Earth Orbit Telecommunication Satellites.* PhD thesis, Department of Electrical Engineering, Stanford University, 2000.

[8] B.W.Parkinson M.Rabinowitz, C.E.Cohen and J.J.Spilker. Performance of a differential carrier-phase gps-leo navigation system. *Proceedings of the Institute of Navigation ION GPS-2000*, September 2000.

[9] C.E.Cohen M.Rabinowitz and B.W.Parkinson. The application of leos to cycle anbiguity resolution on navstar transmissions for kinematic carrier-phase positioning. *Institute of Navigation, ION97*, 1(1), September 1997.

[10] C.E.Cohen M.Rabinowitz and B.W.Parkinson. Patent application: Resolving integer cycle ambiguities with LEO satellites. Technical Report 60/041,184, Application: US Patent and Trademark Office, March 1998.

[11] A.V. Oppenheim and R.W.Schafer. *Discrete-Time Signal Processing.* Prentice-Hall, Inc., 1989.

[12] F.Nelson R.R.Hatch, R.T.Sharpe. John deere's starfire wadgps. *Proceedings of the Institute of Navigation ION GPS-2000*, September 2000.

[13] R.F. Stengel. *Optimal Control and Estimation*, volume 1. Dover Publications, Inc., 1994.

[14] J.R. Wertz. *Spacecraft attitude determination and control.* D.Reidel Publishing Company, 1985.