GPS Precise Absolute Positioning via Kalman Filtering

Brian W. Tolman
Applied Research Laboratories,
The University of Texas at Austin

BIOGRAPHY

Brian W. Tolman is a research scientist at Applied Research Laboratories, The University of Texas at Austin, with over 20 years experience in GPS-related research, data analysis and software development. He holds a Ph.D. in theoretical physics from the University of Texas at Austin.

ABSTRACT

Applied Research Laboratories, The University of Texas at Austin (ARL:UT), under sponsorship of the National Geospatial-Intelligence Agency (NGA), has developed precise positioning software that uses dual frequency GPS carrier phase and pseudorange data to obtain 2-3 centimeter level positioning accuracy for a single static receiver. The paper describes the design of the algorithm, its implementation in platform-independent software, and the results of a large test effort, including positioning accuracy and convergence behavior.

Dual frequency pseudorange and carrier phase data are combined in preprocessing using a new smoothing algorithm that cleanly separates the pseudorange noise, carrier phase biases and first order ionospheric delay from the geometric range plus non-dispersive errors. This algorithm produces a best estimate of the pseudorange noise and removes it without the use of any filtering. It also produces optimal bias estimates for the carrier phases without otherwise altering them, specifically without changing their noise characteristics.

The positioning estimator is a full implementation of the extended Kalman filter algorithm. The filter processes one ionosphere-free phase measurement per satellite; the filter state includes receiver position, clock bias, a residual zenith wet tropospheric delay, and a phase bias for each satellite. Adaptation of the filter to realtime operation and single frequency data processing will be straightforward.

Several standard corrections must be applied to the data before filtering to ensure the accuracy of the result. These include phase windup, antenna phase center offsets, three relativity corrections, and the site displacement effects of solid Earth tides, ocean loading and atmospheric loading. The software implementation is built on The GPS Toolkit (GPSTk), an extensive open source platform-independent C++ library of GNSS-related algorithms. The positioning program is a single stand-alone executable that runs on Linux, Windows, Solaris, Mac and UNIX platform.

The results of extensive testing of the ARL:UT software reiterates the well known result that the accuracy of the position solution is dominated by the accuracy of the satellite ephemeris. Filter convergence is very fast, typically less than 2-3 hours. Averaging over a sufficient timespan (about 1 day) can reduce the residual error to below 2 cm in each position horizontal component. These results have been verified with over 360 days of data at 12 different sites worldwide.

INTRODUCTION

The pioneering work of Zumberge et.al. (1997) presented the original precise point positioning (PPP) algorithm, which processed ionosphere-free combinations of dual frequency pseudorange and carrier phase data with precise satellite ephemerides in the extended Kalman filter algorithm. That paper was also the first to state the important conclusion that the accuracy of the PPP position solution is dominated by the accuracy of the satellite ephemeris. A major difficulty with the algorithm was the relatively slow convergence of the filter due the large amount of noise in the pseudorange data.

Since that important work, several groups have developed new algorithms, and precise positioning results have improved significantly (Bisnath and Langley 2003, Gao and Shen 2002, Kouba and Heroux 2001, and references therein). Of course much of the improvement has resulted directly from improvement in the accuracy of satellite ephemeris products. The focus seems to have been on adding and improving preprocessing algorithms that remove much of the pseudorange noise before the estimation process, including use of a Hatch filter for this purpose, as well as new estimation algorithms that are modifications or simplifications of the Kalman filter algorithm. It seems that none of this work has retained the full Kalman algorithm. Also, new data types have been studied, including the code-phase average (Gao and Shen 2002) which allows for simple extension of the PPP algorithms to single frequency data. The major difficulties in PPP work still seem to be convergence times and removing the pseudorange noise.

ARL:UT has developed PPP software that introduces a new, optimal algorithm to remove pseudorange noise in preprocessing, and that includes a full implementation of the Kalman filter algorithm. The smoothing algorithm is able to cleanly separate the geometric information from the pseudorange noise and the ionospheric delay, while maintaining the low noise characteristics of the phase. It helps enlighten the relation among the standard data types used in other PPP work, showing that they differ only in the placement and handling of the pseudorange noise. Both the smoothing algorithm and the Kalman filter contribute to the algorithm's superior convergence behavior.

This paper begins with a presentation of the smoothing algorithm, including a demonstration of how the standard data types are related, in the first section. Then the Kalman filter design is presented, and specifics of the tuning of the filter and its implementation in software are given. The data corrections, models and site displacement effects are briefly discussed in the next section, with emphasis on what is specific to the ARL:UT development. A large dataset has been used to test the software and generate results; this is discussed in the next section, including repeatability at the sub-5-cm level, absolute accuracy at the same level, and a direct comparison of the NGA and International GNSS Service (IGS) ephemeris products. Finally, conclusions are summarized in the final section.

SMOOTHING ALGORITHM

This section presents a simple algorithm that obtains an optimal estimate of the biases on the phase. These biases are then used to construct both "debiased phases" and "smoothed pseudoranges." The algorithm properly handles the ionospheric delay, which is frequency-dependent and affects pseudorange and phase differently, and cleanly separates the ionospheric delay and pseudorange noise from the geometric range information. It introduces only a small constant bias error, without changing the (phase) noise characteristics of the results.

Fundamental equations. The four fundamental GPS observables, pseudorange and carrier phase on each of two carriers with frequencies f_1 and f_2, for a single receiver and satellite, are modeled as

$$\begin{aligned} P_1 &= R+T+cdt_r-cdt^s+I+\epsilon_{P_1} \\ P_2 &= R+T+cdt_r-cdt^s+\left(\frac{f_1^2}{f_2^2}\right)I+\epsilon_{P_2} \\ L_1 &= R+T+cdt_r-cdt^s-I+B_1+\epsilon_{L_1} \\ L_2 &= R+T+cdt_r-cdt^s-\left(\frac{f_1^2}{f_2^2}\right)I+B_2+\epsilon_{L_2} \end{aligned}$$

where P is the pseudorange, L is the carrier phase with length units, R is the geometric range, dt is the clock bias at the receiver (sub r) and the satellite (super s), T is the tropospheric delay, I is the ionospheric delay at the f_1 frequency, and B is the phase bias (not an integer number of cycles). The noises ϵ are assumed to include receiver noise and multipath.

Defining the constant

$$\alpha \equiv \frac{f_1^2}{f_2^2}-1$$

and calling the geometric range plus all the non-dispersive errors simply "the range,"

$$\rho = R+T+cdt_r-cdt^s$$

the four fundamental equations simplify to

$$\begin{aligned} P_1 &= \rho+I+\epsilon_{P1} \\ P_2 &= \rho+(\alpha+1)I+\epsilon_{P2} \\ L_1 &= \rho-I+B_1+\epsilon_{L1} \\ L_2 &= \rho-(\alpha+1)I+B_2+\epsilon_{L2} \end{aligned}$$

There are 4 fundamental quantities here:

1. the range ρ,
2. the ionospheric delay I on f_1; on f_2 it is $(\alpha+1)I$,
3. phase biases B on the two frequencies,
4. pseudorange errors on the two frequencies.

The range is the most important, because it contains the geometric information needed for positioning and timing. The ionospheric delay is useful because it yields information about the upper atmosphere, but most often it is eliminated before further processing. The errors on the pseudoranges contain information on multipath, receiver noise and other effects, but also would be removed before processing. The noise on the carrier phase is neglected compared to that of the pseudorange.

Many techniques have been found to model, mitigate, or isolate these four quantities. For example the ionospheric delay is eliminated in the *ionosphere-free pseudorange* and *ionosphere-free phase* linear combinations, defined by

$$P_{IF} \equiv \frac{\alpha+1}{\alpha}P_1-\frac{1}{\alpha}P_2 = \rho+\epsilon_{P_1+P_2}$$

$$\begin{aligned} L_{IF} &\equiv \frac{\alpha+1}{\alpha}L_1-\frac{1}{\alpha}L_2 \\ &= \rho+\frac{\alpha+1}{\alpha}B_1-\frac{1}{\alpha}B_2+\epsilon_{L_1+L_2} \end{aligned}$$

The ionosphere-free phase carries only the phase noise, but is biased. The ionosphere-free pseudorange is unbiased but dominated by the large noise of the pseudoranges. Another approach to removing the ionospheric delay (Gao and Shen 2002) is to form a *code-phase average* on each frequency (i=1,2)

$$\tilde{R}_i = \frac{1}{2}(P_i + L_i) = \rho + \frac{1}{2}\lambda_i N_i + \epsilon_{P_i+L_i}.$$

This quantity is also ionosphere-free, but it retains both a bias and the large pseudorange noise.

The smoothing algorithm presented here uses the fundamental equations to derive optimal estimates of the phase biases B, and then remove them; it is based on Blewitt (1998). The fundamental equations, written in matrix form, are

$$\begin{bmatrix} P_1 \\ P_2 \\ L_1 \\ L_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & \alpha+1 & 0 & 0 \\ 1 & -1 & 1 & 0 \\ 1 & -\alpha-1 & 0 & 1 \end{bmatrix} \begin{bmatrix} \rho \\ I \\ B_1 \\ B_2 \end{bmatrix} + noise$$

This matrix equation can be inverted to give

$$\begin{bmatrix} \rho \\ I \\ B_1 \\ B_2 \end{bmatrix} = \frac{1}{\alpha} \begin{bmatrix} \alpha+1 & -1 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ -\alpha-2 & 2 & \alpha & 0 \\ -2\alpha-2 & \alpha+2 & 0 & \alpha \end{bmatrix} \begin{bmatrix} P_1 \\ P_2 \\ L_1 \\ L_2 \end{bmatrix} + noise.$$

The noise on this vector is dominated by the pseudorange noise, and equal to (assuming the pseudorange noises to be of equal magnitude, and neglecting the phase noise)

$$noise \begin{bmatrix} \rho \\ I \\ B_1 \\ B_2 \end{bmatrix} = \begin{bmatrix} 2.98 \\ 2.91 \\ 5.13 \\ 6.53 \end{bmatrix} \epsilon_P.$$

This approach is not a good one for processing, due to the large noise, however it does isolate the phase bias terms from the ionospheric delay and the range, which allows them to be estimated.

Bias solution. Define the bias vector

$$\vec{B} \equiv \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}$$

and consider the matrix operator

$$D = \frac{1}{\alpha} \begin{bmatrix} \alpha+2 & -2 \\ 2\alpha+2 & -\alpha-2 \end{bmatrix}$$

where vector notation denotes a simple two-vector of elements for frequencies L1 and L2. Then the lower two equations in the matrix solution above are simply

$$\vec{B} = \vec{L} - D\vec{P} \quad .$$

Since the biases B are constants, an estimate of their value is formed simply by averaging the data quantity on the right hand side of this equation. This estimate is optimal if the largest possible segment of unbroken phase is used to compute it, ideally an entire satellite pass from horizon to horizon.

The noise on this data is just the pseudorange noise; it follows that the standard error of the bias estimates is

$$\epsilon \approx \frac{1}{\sqrt{N}} \sigma_P$$

where σ_P is the magnitude of the pseudorange noise and N is the number of data points in the satellite pass. This bias error is larger than the phase noise, but much smaller than the pseudorange noise. Figure 1 presents the computed bias error, as a function of the size of the satellite pass, for a typical dataset with 10-second rate.

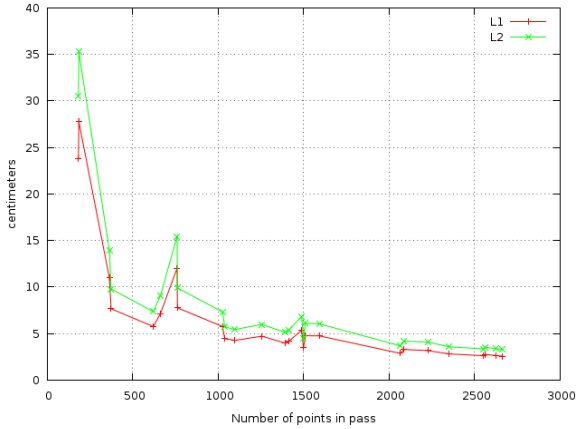

Figure 1. Smoothing bias error as a function of satellite pass size, for both GPS carrier frequencies, one dataset.

Debiased phases. The estimated biases are now removed from all the phases in the dataset to form the debiased phases

$$\vec{L}_{db} = \vec{L} - \vec{B}.$$

It is important to note that debiasing the phases leaves a small residual bias error, but does not change the noise characteristics of the phase.

The debiased phases are not equivalent to the "smoothed" pseudoranges, due to the different dependence on the ionospheric delay. To find such a smoothed pseudorange, reconsider the fundamental equations with the raw phases replaced by the debiased phases, and make the pseudorange error terms explicit, calling them δ . Then the bias terms are moved to the left side and the fundamental equations become

$$\begin{bmatrix} P_1 \\ P_2 \\ L_{1\mathrm{db}} \\ L_{2\mathrm{db}} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & \alpha+1 & 0 & 1 \\ 1 & -1 & 0 & 0 \\ 1 & -\alpha-1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \rho \\ I \\ \delta_1 \\ \delta_2 \end{bmatrix} + noise$$

This matrix equation has the solution

$$\begin{bmatrix} \rho \\ I \\ \delta_1 \\ \delta_2 \end{bmatrix} = \frac{1}{\alpha} \begin{bmatrix} 0 & 0 & \alpha+1 & -1 \\ 0 & 0 & 1 & -1 \\ \alpha & 0 & -\alpha-2 & 2 \\ 0 & \alpha & -2\alpha-2 & \alpha+2 \end{bmatrix} \begin{bmatrix} P_1 \\ P_2 \\ L_{1\mathrm{db}} \\ L_{2\mathrm{db}} \end{bmatrix} + noise.$$

Smoothed pseudoranges. The second pair of these four equations is simply

$$\vec{\delta} \;=\; \vec{P} - D\,\vec{L}_{db} \;=\; \vec{P}_{raw} - \vec{P}_{sm}$$

where this equation defines the "smoothed" pseudoranges as

$$\vec{P}_{sm} \;=\; D\,\vec{L}_{db}.$$

Smoothed range and ionosphere. Using the debiased phases and the smoothed pseudoranges, the range (true range plus non-dispersive errors) is simply

$$\rho \;=\; \frac{\alpha+1}{\alpha} L_{1\mathrm{db}} - \frac{1}{\alpha} L_{2\mathrm{db}} \;=\; \vec{W} \cdot \vec{L}_{db}$$

where

$$\vec{W} = \begin{bmatrix} \dfrac{\alpha+1}{\alpha} & -\dfrac{1}{\alpha} \end{bmatrix},$$

and the debiased ionospheric delay is

$$I \;=\; \frac{1}{\alpha}(L_{1\mathrm{db}} - L_{2\mathrm{db}}).$$

The range and ionospheric delay here do not depend on the raw pseudorange at all, and so have only the noise of the phase. The noise vector for this solution is (assuming noises on the two carriers are identical)

$$noise \begin{bmatrix} \rho \\ I \\ \delta_1 \\ \delta_2 \end{bmatrix} \;=\; \begin{bmatrix} 2.98\,\epsilon_L \\ 2.91\,\epsilon_L \\ \epsilon_P \\ \epsilon_P \end{bmatrix}.$$

Of course the range error has exactly the noise of the pseudorange, but the other quantities have only about three times the noise of the phase. Note however, that because all these quantities depend on the debiased phases, there will be an error contribution to each of them involving the small bias errors that come from the smoothing algorithm, but, again, this is a bias, not a noise.

Standard derived quantities. Finally, it is of interest reconsider the derived range quantities mentioned earlier in this section. If we form these quantities using the debiased phases and the smoothed pseudoranges, the results is, for the ionosphere-free phase

$$L_{IF} \;=\; \frac{\alpha+1}{\alpha} L_{1\mathrm{db}} - \frac{1}{\alpha} L_{2\mathrm{db}} \;=\; \vec{W} \cdot \vec{L}_{db} \;=\; \rho$$

and for the ionosphere-free pseudorange (noting the identity $\vec{W} \cdot D \;=\; \vec{W}$)

$$\begin{aligned} P_{IF} \;&=\; \frac{\alpha+1}{\alpha} P_{1\mathrm{sm}} - \frac{1}{\alpha} P_{2\mathrm{sm}} \\ &=\; \vec{W} \cdot \vec{P}_{sm} \\ &=\; \vec{W} \cdot D\,\vec{L}_{db} \\ &=\; \vec{W} \cdot \vec{L}_{db} \;=\; L_{IF} \;=\; \rho \end{aligned}$$

and finally for the two code-phase averages

$$\begin{bmatrix} \tilde{R}_1 \\ \tilde{R}_2 \end{bmatrix} = \begin{bmatrix} \left(\dfrac{\alpha+2}{2\alpha}\right) + \dfrac{1}{2} & -\dfrac{1}{\alpha} \\ \dfrac{\alpha+1}{\alpha} & -\left(\dfrac{\alpha+2}{2\alpha}\right) + \dfrac{1}{2} \end{bmatrix} \cdot \begin{bmatrix} L_{1\mathrm{db}} \\ L_{2\mathrm{db}} \end{bmatrix}$$

$$= \frac{1}{\alpha} \begin{bmatrix} \alpha+1 & -1 \\ \alpha+1 & -1 \end{bmatrix} \cdot \begin{bmatrix} L_{1\mathrm{db}} \\ L_{2\mathrm{db}} \end{bmatrix} \;=\; \begin{bmatrix} \rho \\ \rho \end{bmatrix}.$$

Thus when smoothing is implemented properly, all of the conventional derived range quantities are identical (even in their dependence on the bias error), and equal to the smoothed range quantity.

KALMAN FILTER DESIGN AND TUNING

The design of the Kalman filter attempts to maximize the rate of convergence while maintaining simplicity and numerical stability.

Filter state. The filter state $\boldsymbol{X}$ consists of three Earth-centered Earth-fixed Cartesian components of position, a clock bias, a residual zenith wet tropospheric delay, and a single phase bias for each visible satellite:

$$\vec{X} = \begin{bmatrix} \vec{x} & t & T & B^0 & \ldots & B^j \end{bmatrix}$$

The position and clock states are modeled as random walks. The position states are allowed to vary in order to allow for the effect of ephemeris error in the solution (cf. Results section below). The clock is expected to vary much more than the position states. The residual tropospheric delay is modeled as an exponentially time-correlated (first order Markov) process (Bierman 1977).

The phase biases model several constant effects, including the residual bias error arising from the smoothing, the initial (non-integral) phase biases at the receiver and the satellite, and any residual bias in the phase wind-up correction. These biases should be pure constants with no

variation whatever, however they are modeled as random walks as well (with very small process noise), to avoid numerical problems when the filter is run for long times.

Measurement model. The measurement equation includes a single range measurement per satellite in the form of the ionosphere-free phase, built from the debiased phases (see above). Specifically, the linearized measurement equation is

$$\begin{bmatrix} \vec{\beta}^0 & 1 & w^0 & 1 & 0 & \dots & 0 \\ \vec{\beta}^1 & 1 & w^1 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & & \vdots \\ \vec{\beta}^i & 1 & w^i & 0 & 0 & \dots & \delta^i_j \end{bmatrix} \cdot \begin{bmatrix} \vec{x} \\ t \\ T \\ B^0 \\ \vdots \\ B^j \end{bmatrix} = \begin{bmatrix} L^0_{IF} - \tilde{R}^0_E \\ L^1_{IF} - \tilde{R}^1_E \\ \vdots \\ L^i_{IF} - \tilde{R}^i_E \end{bmatrix}$$

where the superscript identifies satellites. On the left side of this equation '1' and 'w' are scalars, w being the wet tropospheric mapping function evaluated at the satellite elevation, and $\vec{\beta}$ is a row vector of direction cosines for satellite 'i',

$$\vec{\beta}^i = \left[\alpha^i_x, \quad \alpha^i_y, \quad \alpha^i_z \right].$$

On the right side, L_{IF} is the debiased ionosphere-free phase (i.e. the range ρ), and $\tilde{R}_E$ is the filter's prediction of the range from the previous timestep.

The measurement covariance matrix is diagonal with a single element equal to

$$\left(\frac{\alpha+1}{\alpha} \right)^2 \sigma^2_{L1} + \frac{1}{\alpha^2} \sigma^2_{L2}$$

for each satellite. The measurement covariance is weighted using either the cosecant of the satellite elevation angle, or the cosecant squared.

State propagation. The filter's state propagation equation is simply

$$\Phi = diagonal\left[\vec{1}, \quad 1, \quad m, \quad 1^0, \quad \dots, \quad 1^j \right]$$

(recall that the states are 3-vector of position, clock, residual tropospheric delay, one bias per satellite). The residual zenith tropospheric delay is modeled as a constant with exponentially correlated noise (first order Markov process); m is the exponential factor for the colored noise (Bierman 1977)

$$m = e^{-dt/\tau}$$

where dt is the filter time step and τ is the time constant of the correlation.

Process noises. The process noise matrix for this state has only diagonal elements equal to the magnitude of the process noise for that element. The exception is the residual tropospheric delay state, which includes a term expressing the exponential correlation of the noise; it is

$$\left(1 - e^{-2dt/\tau}\right) \sigma^2_{rztd}$$

which multiplies the amplitude of the noise by a factor involving the timestep dt of the filter and the time constant of the correlation.

Tuning the filter consists of choosing values for these process noises and the measurement noises that optimize the filter performance. (For random walk states, the noise rate is scaled by the filter timestep.) A nominal set of tuning parameters are presented in Table 1.

Position	2.0e-5 meter/sqrt(sec)
Clock	6.0 meter/sqrt(sec)
Residual Troposphere	0.1 meter with timescale 1 hour
Phase biases	2.0e-6 meter/sqrt(sec)
Measurements	1.0e-2 meter

Table 1. A nominal set of tuning parameters for the ARL:UT Kalman filter.

Software implementation. The ARL:UT precise positioning software has been developed using The GPS Toolkit (GPSTk 2008), an extensive open source C++ library of GNSS-related algorithms. The software, like the GPSTk, is very platform independent, and runs on Linux (probably any form of UNIX), Windows, Solaris and Mac.

A serious problem arises in Kalman filters that process GNSS data in that each bias state is determined solely by the data in a single satellite pass, yet usually that portion of the data covers only a small part of the time span of the full dataset. Thus at any given time, many of the bias states are “unobservable” by the filter because the satellite data is not yet available (and thus the filter is singular), and then after a satellite pass has been processed, the corresponding bias state is thereafter “dead,” meaning it never changes because there is no more data for it.

This problem can be solved by an implementation of the Kalman algorithm which allows states to be dynamically added and deleted. The ARL:UT Kalman filter is implemented using the square root information filter or SRIF (Bierman 1977) as implemented in the GPSTk. The SRIF does allow modifications of the state vector. Thus any particular bias is not included in the filter state until the corresponding satellite data is available, and then after the end of the satellite pass, the dead bias state is removed. This allows the filter to always be non-singular, and improves the efficiency and numerical stability of the algorithm.

DATA CORRECTIONS AND SITE DISPLACEMENTS

There are many sources of error and perturbing effects that must be accounted for in the algorithm in order to produce centimeter-level results. These corrections are applied to

the data in preprocessing. All of these effects have been discussed in detail elsewhere (Kouba and Heroux 2001, McCarthy 1996); only brief comments will be added here.

Corrections include *phase windup* and *relativity* effects (including the gravitational redshift, second order Doppler shift and the Sagnac effect (Ashby and Spilker 1996)). The remaining effects are listed below, along with comments specific to this work.

Antenna phase center offset (PCO) and variation (PCV). The software reads ANTEX format files (ANTEX 2008), for both satellite and receiver antennas. The IGS satellite PCOs are estimated values, based on large amounts of live-sky data; these values should be used with the IGS ephemeris products. The NGA ephemeris, on the other hand, uses the satellite manufacturers' value for the antenna PCOs; these values should be used with the NGA ephemeris products. These two sets of antenna PCOs are significantly different. Inclusion of the satellite PCOs allows the ARL:UT processor to make use of either the center-of-mass based satellite ephemerides or the antenna-phase-center based ones.

Cycle slip detection and correction. Cycle slips must be removed from the phase data before smoothing and filtering, as the filter assumes that a single constant bias applies to the phase. This is accomplished in the ARL:UT software using the discontinuity corrector which is part of the GPSTk.

Multipath is not addressed specifically by the algorithm. Pseudorange multipath is essentially removed by the smoothing algorithm presented above. Phase multipath is small but potentially a source of position error, although averaging over long times should be sufficient to mitigate it strongly.

Tropospheric delay. The nominal delay is removed using any of several standard tropospheric models, as provided in the GPSTk. The software takes input from RINEX format meteorological files for weather information, or, in their absence, uses default values, which also may be provided by the user. The tropospheric delay state in the Kalman filter accounts for any residual delay. The wet mapping function in the Kalman filter measurement update may also be chosen to be any of several standard models, as provided by the GPSTk tropospheric models.

Ionospheric delay. The delay due to the ionosphere is removed to first order in the formation of the ionosphere-free phase. Higher order ionospheric delays may be as large as 10's of centimeters, and therefore may be a source of error, although averaging will strongly mitigate their effect. They have not been incorporated into the ARL:UT filter, although there are plans to do so.

The C1-P1 bias. This is a systematic bias between the C/A code and P code pseudorange on L1. Since precise ephemeris products are usually derived from P code pseudorange data, if C/A code data is used with such an ephemeris to estimate precise positions, the bias contributes directly to the solution error. In this work, however, the pseudoranges are only used to debias the phase (cf. above), therefore any residual bias will contribute only a small error to the initial phase bias, and of course this will be accounted for in the corresponding Kalman filter state. A single-frequency filter would have to account for this effect, however.

Solid Earth tides. The model of McCarthy (1996) is implemented in the software. This is consistent with the model used in the production of the NGA ephemerides.

Ocean loading. This effect is modeled with the equations in McCarthy (1996) using the site-specific coefficients obtained at http://www.oso.chalmers.se/~loading.

Atmospheric loading. This is almost certainly too small to be of significant effect; however it is quite simple, and the model of McCarthy (1996) has been implemented in the software.

Solar and Lunar ephemerides are necessary for the implementation of several of these effects, including phase windup, solid earth tides and ocean loading. In this work, the JPL DE403 solar system ephemeris was implemented and tested using the GPSTk. The DE403 ephemeris is consistent with the models described in McCarthy (1996).

TESTING AND RESULTS

The ARL:UT precise positioning software has been extensively tested in order to study the convergence behavior of the Kalman filter, and to generate positioning results with regard to both repeatability and accuracy. Comparison of the residual errors for a large dataset processed with the National Geospatial-Intelligence Agency (NGA) ephemeris, and then with the International GNSS Service (IGS) ephemeris, demonstrates that ephemeris error does indeed dominate the final solution accuracy, and provides a direct comparison of those two ephemeris products.

The test data set. A world-wide dataset, consisting of 30 days (days 70-99 of 2007) at each of 9 NGA monitor stations, plus 30 days (days 90-119 of 2008) from a site in New Mexico, were used to test the ARL:UT precise positioning software. The NGA (P-code) data was collected by an Ashtech ZY12 with cesium frequency standard at a 30-second rate; the New Mexico data came from a Trimble NetRS receiver running on internal clock and collecting 10 second C/A L1 code and codeless L2 data. All these stations run continuously on very stable mounts. Meteorological data is available for the NGA sites, but not for the New Mexico site.

Convergence. The testing shows that the Kalman filter is

very stable and relatively insensitive to tuning parameters. The convergence characteristics of the filter are excellent, with convergence being achieved after processing less than about two hours of data. Figure 2 shows elements of the square root of the covariance for a typical run.

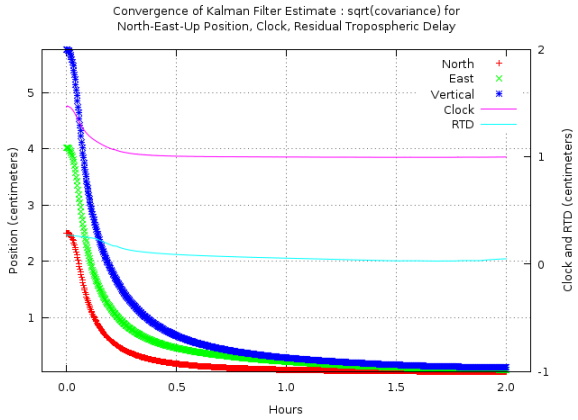

Figure 2. Convergence of the Kalman filter; well under 2 hours for a typical dataset. Sqrt(covariance) is plotted for each filter state.

Residual error sequence. In addition to the covariance information, the output of the Kalman filter will include timeseries of the position, clock and tropospheric delay states. One expects the residual error in the position states to be the combined result of all unmodeled or mismodeled effects; normally for a Kalman filter this is very like a white noise sequence.

In this case the residual error sequence will include errors from the ephemeris and satellite clock input, combined together using the relative satellite-receiver geometry. In fact this error should dominate, since the ephemeris is the dominant source of the position accuracy. The NGA and IGS ephemeris products are nominally expected to be accurate at the 5-10cm level. This is expected to produce a systematic, quasi-periodic sequence with 5-10cm amplitude and a time scale determined by the timescales of the ephemeris and of the changing satellite geometry as seen at the receiver.

This timeseries of residual error will produce position errors at the same level, about 5-10cm. Averaging over time may improve the final position estimate, with the precision of the results improving as the averaging interval is increased.

Processing very long datasets using the ARL:UT positioning software shows just the expected behavior, as presented in Figure 3, which shows the filter behavior for a 4-day dataset at the New Mexico site using the NGA ephemeris. Note the convergence of the filter, from the initial position error at meter level, during the first few hours. After convergence the position states contain noise at the few-centimeter level, presumably due to ephemeris and clock error.

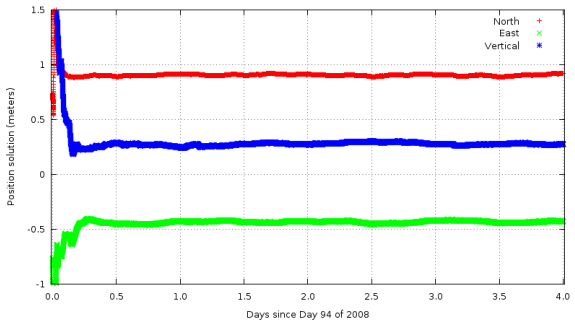

Figure 3. Kalman filter position states for North (red), East (green) and Vertical (blue) components (including an arbitrary offset), for a 4-day dataset, using the NGA ephemeris. The receiver is a Trimble NetRS, a commercial survey-grade receiver; the site is in New Mexico. Note the convergence behavior in the first few hours.

To examine the assumption that ephemeris and clock error dominates the residual error sequence, the dataset of Figure 3 was reprocessed using exactly the same input except using the IGS (final) ephemeris+clock product. The components of the resulting error sequences are presented in Figures 4a,b,c.

Note that there is strong correlation in the vertical component (Fig. 4c). This correlation is expected, however, and is due to the influence of the clock and atmospheric errors, which alias strongly into the vertical component of any position solution.

However in the horizontal components there is a lack of correlation, which can be attributed to the different satellite ephemeris products used in the two runs of the filter. Thus these results suggest that the ephemeris indeed dominates the residual error sequence, and that the Kalman filter is well designed and has converged.

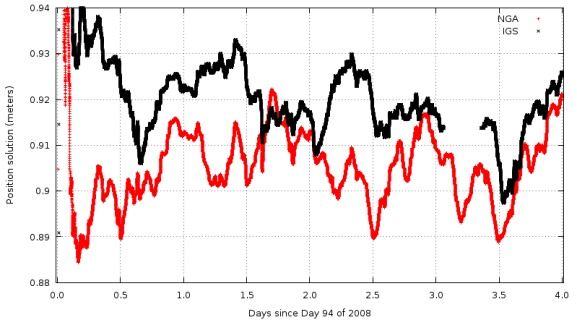

Figure 4a. Residual error sequence for the North position component, from the Kalman filter state after convergence, for a 4-day dataset from New Mexico (as in Fig. 3), using both the NGA ephemeris (colored) and the final IGS ephemeris (black). (The gap in the IGS results are due to an unknown problem in the ephemeris file igs14740.sp3.) Note the lack of correlation among the two results.

Note that two timescales are apparent in the data, at about ¼ day and also at about 1 day. This probably arises from the satellite orbit period, the period over which the satellite geometry repeats (~1 day) and the timescale of significant changes in the satellite geometry as seen at the receiver (~6 hours).

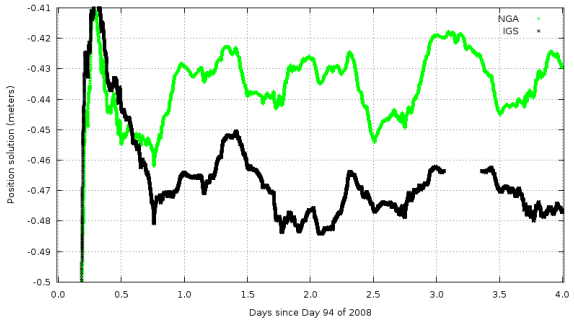

Figure 4b. Same as Figure 4a, except on a different vertical scale, for the East position component.

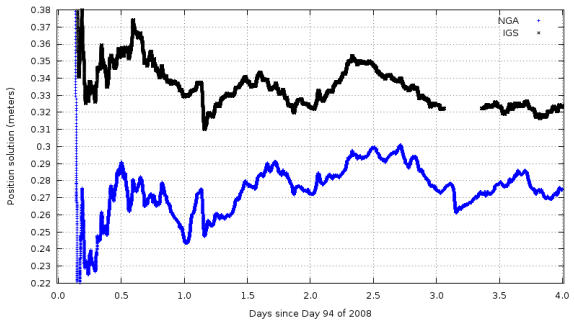

Figure 4c. Same as Figure 4a, except on a different vertical scale, for the Vertical position component. Note the correlation here, which is most likely due to clock and atmospheric effects.

Also note that these final solutions are offset at the 5-10cm level in each component; this is consistent with the statement that the NGA and IGS ephemeris are believed to agree at the level of 4.5-5cm radial, 14cm along track, and 10cm cross track (NGA web site, 2008).

Repeatability. The ARL:UT positioning software was extensively tested using the 30-day test dataset in order to generate statistics on the repeatability of the final solution. The data was processed for each day in the dataset, using different dataset lengths, specifically in 4-hour increments, then again in 4.8, 6, 8, 12, 18, 24 and 48 hour intervals. Statistics were then computed for each component of all the solutions at each site, separately at each dataset length. The standard deviations of these component solution sets give an indication of the repeatability of the solution produced by the software.

Figures 5-8 present the repeatability results. They show the standard deviation of the position solutions as a function of dataset length for four different static sites, located in South Korea, New Mexico, Bahrain and England. These results are typical of those at the other 6 sites used in the testing (not shown here). (In these Figures, it is important to point out that averaging was not done over the entire dataset, but only the part after filter convergence. For simplicity and uniformity in this test, all runs allowed for a 3-hour convergence time; thus a dataset length of 12 hours consisted of 12 – 3 = 9 hours of averaging.)

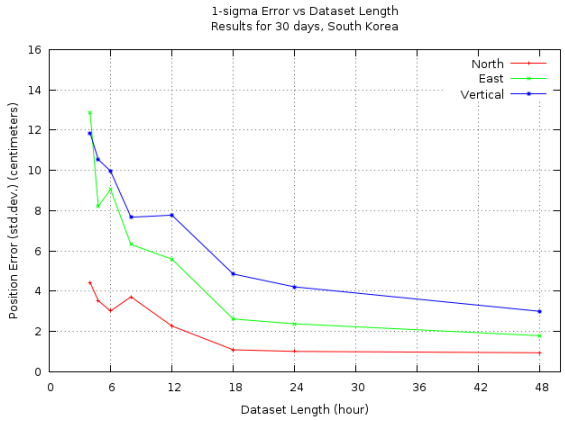

Figure 5. Position error, expressed as the standard deviation of all position solutions (North, East and Vertical components) derived from 30 days of data, versus the length of the dataset. This site is the NGA Monitor Station in South Korea.

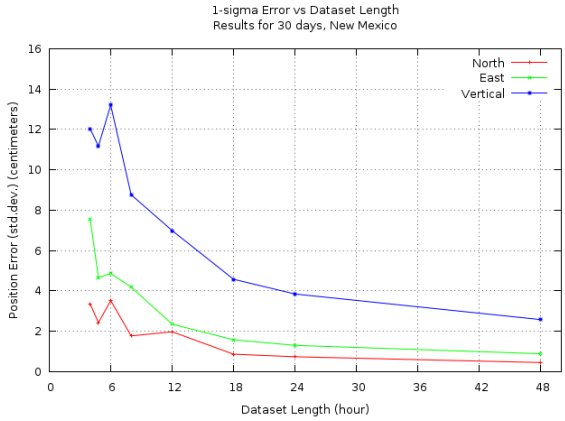

Figure 6. Position error, expressed as the standard deviation of all position solutions (North, East and Vertical components) derived from 30 days of data, versus the length of the dataset. This site is in New Mexico.

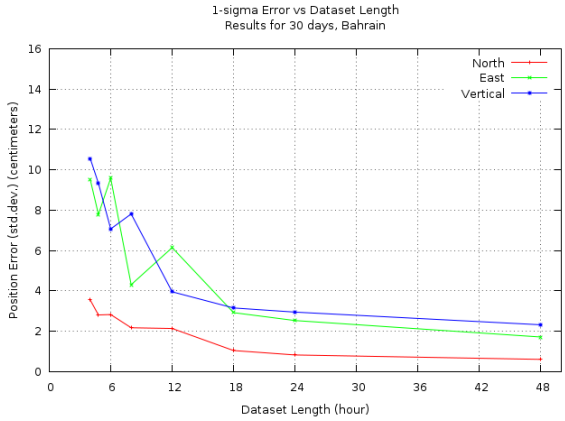

Figure 7. Position error, expressed as the standard deviation of all position solutions (North, East and Vertical components) derived from 30 days of data, versus the length of the dataset. This site is the NGA Monitor Station in Bahrain.

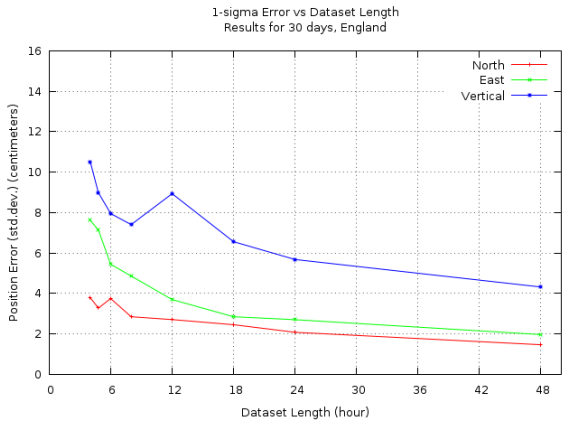

Figure 8. Position error, expressed as the standard deviation of all position solutions (North, East and Vertical components) derived from 30 days of data, versus the length of the dataset. This site is the NGA Monitor Station in England.

Repeatability clearly improves as the dataset length increases. This follows from the fact that averaging reduces the effect of the residual errors in the final solution.

Also, many of the results tend to show a drop in standard deviation beginning around 6 hours, but then rising again at 12 hours. This can be understood as the result of averaging over the shortest characteristic timescales in the residual error sequence. As the averaging time reaches about 6 hours, the standard deviation falls due to the effect of averaging over the short timescale variations; however as the dataset length approaches 12 hours it includes the variations over 12 and 24 hours, and so the standard deviation increases sharply. After 24 hours, the averaging interval is longer than any timescale of variations in the data, and so the standard deviation continues to drop with increasing dataset length.

Repeatability results for the ARL:UT precise positioning software may be summarized as follows. With less than 12 hours of data, the position solution reflects the ephemeris error, which is about 5 centimeters each component, with 10-15 cm in the vertical. Averaging will quickly improve these results, so that results at 12 hours of data improve significantly. Further improvement is possible at 24 hours; however the results at 48 hours probably reflect a practical limit on the repeatability or accuracy that can be obtained. Table 2 summarizes these results.

Dataset length	Horizontal	Vertical
< 12 hours	5cm	10-15cm
12 hours	4cm	6cm
24 hours	2cm	4cm
48 hours	1cm	2cm

Table 2. Results for the repeatability of the final position solution.

Accuracy. The results of the ARL:UT precise positioning software for the NGA monitor stations have been compared with the known positions of those sites. The differences between the solutions and the known positions are consistent with the repeatability results, meaning that the differences are of order the 1-sigma results presented above. Figures 9 and 10 show the difference of the final position solution (North, East and Vertical components) using the NGA ephemeris, with the known positions of the NGA Monitor stations, for the stations in England and Australia.

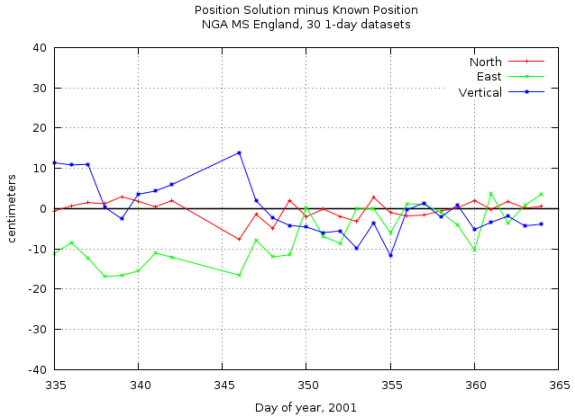

Figure 9. Position solution, differenced with the known NGA position, for 30 days of data at the NGA Monitor Station in England.

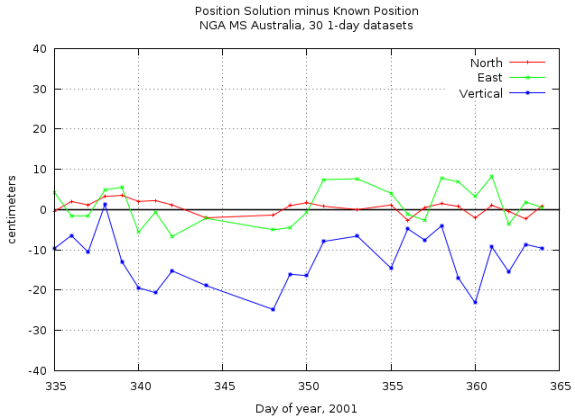

Figure 10. Position solution, differenced with the known NGA position, for 30 days of data at the NGA Monitor Station in Australia.

As a further test of accuracy, two IGS sites were chosen at random for comparison of the ARL:UT position solution, using the IGS final ephemeris product, with the known position of the site as published by IGS. Figures 11 and 12 show the difference of the final position solution (North, East and Vertical components) with the known IGS station position (epoch 2008.6189, SOPAC 2008) for the sites at Richards Bay, South Africa (rbay) and Riga, Latvia (riga). In both cases, 32 days of data (day 168-199 of 2008) were processed.

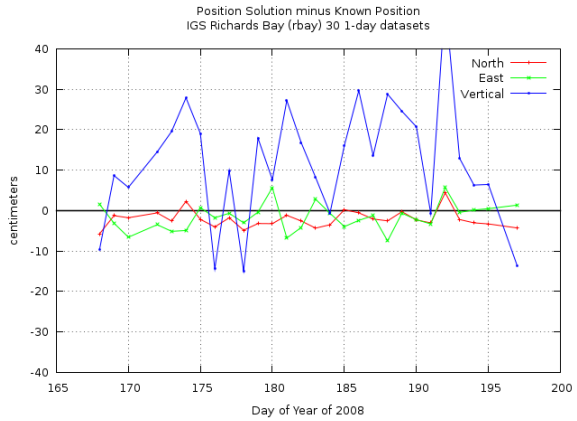

Figure 11. Position solution, differenced with the published IGS position, for 30 days of IGS data at Richards Bay (rbay).

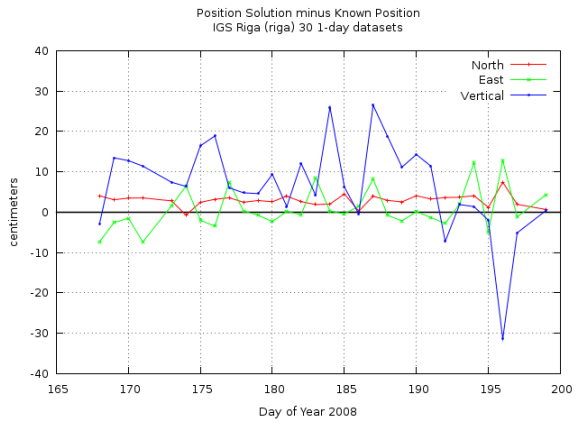

Figure 12. Position solution, differenced with the published IGS position, for 30 days of IGS data at Riga, Latvia (riga).

The full 24 hour data set was processed for each day. No attempt to edit the data was made, except that two days' data were rejected at Riga, one at Richards Bay, and two each at England and Australia, because they did not successfully process; investigation into the source of these data problems is not yet complete.

Statistics on these results show that the position solution is essentially consistent, to within 1-sigma, with the known positions, as follows.

	North	East	Vertical
England	-0.3 +/- 2.3cm	-6.7 +/- 6.5cm	-0.2 +/- 6.3cm
Australia	0.5 +/- 1.6cm	1.1 +/- 4.7cm	-12.4 +/- 6.4cm
R.Bay	-2.1 +/- 2.1cm	-1.6 +/- 3.3cm	12.1 +/- 15.1cm
Riga	3.1 +/- 1.9cm	3.4 +/- 10.2cm	4.8 +/- 14.4cm

Table 3. Statistical results (1-sigma) on the accuracy of the final position solution for four 30-day datasets (one solution per day).

Although there are some small biases in the differences, most often in the vertical component (particularly Australia), these are expected to be reduced or removed after a more careful analysis that includes data quality assurance.

CONCLUSION

The smoothing algorithm presented here is an optimal solution to the problem of removing most of the noise in the raw pseudorange data before filtering it, without changing the (phase) noise characteristics of the data. It improves the convergence of the filter dramatically and eliminates the need to apply a Hatch filter or other non-optimal preprocessing to the data.

The ARL:UT estimation algorithm is a full Kalman filter, without simplifications or modifications that reduce its optimality, yet it is very simple and robust, with excellent convergence properties. The Kalman filter is controlled by tuning, and allows the user to perform a smoother or backward Kalman filter on the data. The implementation of this filter uses the square root information (SRIF) factorization, which makes it possible to add and delete the bias states as the appropriate data becomes available, keeping the filter non-singular, efficient and numerically stable.

The accuracy and repeatability of the final position solution is dominated by the accuracy of the satellite ephemerides – a well known result demonstrated here in detail. Averaging over the residual errors in the solution yields repeatability at about 2cm in the horizontal, and 5cm in the vertical position components. Absolute accuracy is consistent with the repeatability results, for both the the NGA and IGS ephemeris products.

This work in on-going, and several improvements are anticipated in the near future. Other GNSS data will be incorporated into the processing. Also, a single frequency filter design using the code-phase average will be a relatively straightforward option to be added to the software. Adaptation of this algorithm to realtime operation should be relatively straightforward. Finally, many improvements to the algorithm and the robustness of the software are anticipated or currently in process.

ACKNOWLEDGMENTS

This work was supported by the National Geospatial-Intelligence Agency, for which the author is grateful. The author also wishes to thank Everett Swift of the Naval Surface Warfare Center for helpful discussions.

REFERENCES

ANTEX (2008), The Antenna Exchange Format, cf. ftp://igscb.jpl.nasa.gov/igscb/station/general/ or other IGS web sites for current ANTEX files; see file Antex13.txt for the format description.

Ashby, N. and J. J. Spilker Jr. (1996), "Introduction to Relativistic Effects on the Global Positioning System," in Global Positioning System: Theory and Applications, Progress in Astronautics and Aeronautics, Vol. 1, Chapter

18, eds. B. W. Parkinson, J. J. Spilker Jr, Volume 164, American Institute of Aeronautics and Astronautics, Inc., V, Washington, D. C.

Bierman, G. (1977), "Factorization Methods for Discrete Sequential Estimation," Academic Press, New York, 1977.

Bisnath, S. and R. Langley (2003), “High-Precision Single-Frequency GPS Point Positioning,” Proceedings of the Institute of Navigation ION-GPS/GNSS 2003, Portland.

Blewitt, G. (1998), GPS for Geodesy, 2nd edition, Ch. 6, edited by Teunissen, PJG, and A. Kleusberg, Springer-Verlag Berlin, 1998.

Gao, Y. and X. Shen (2002). “ A New Method for Carrier Phase Based Precise Point Positioning,” Navigation, Journal of the Institute of Navigation, Vol. 49, No. 2, 2002.

GPSTk (2008), The GPS Toolkit, http://www.gpstk.org.

Kouba, J. and P. Heroux (2001), “GPS Precise Point Positioning Using IGS Orbit Products,” GPS Solutions, Vol. 5, No. 2, 2001.

McCarthy, D.D. editor (1996), IERS Conventions (1996), IERS Technical Note 21.

NGA web site (2008), National Geospatial-Intelligence Agency, http://earth-info.nga.mil/GandG/sathtml/.

SOPAC (2008), Scripps Orbit and Permanent Array Center, http://sopac.ucsd.edu/cgi-bin/dbShowArraySitesMap.cgi?array=IGS

Zumberge, J. F., M.B. Heflin, D. C. Jefferson, M. M. Watkins, and F. H. Webb (1997), “Precise Point Positioning for the Efficient and Robust Analysis of GPS Data from Large Networks,” Journal of Geophysical Research, Vol 102, No. B3, pp 5005-5017.