Chapter 9 Stand-Alone Positioning

9.0 Introduction

Stand-alone positioning refers to the ability to determine one's location using only the signals from the core navigation satellites. No additional information is used to enhance the quality of that position. Stand-alone positioning is the most basic service offered by any constellation of navigation satellites. This chapter will outline how measurements from multiple satellites are used to form an estimate of the user's position. Later chapters describe how that service may be enhanced or augmented by information that is separate from the core constellation broadcast signals and data.

9.1 Code and Carrier-Phase Measurement Model

Each GNSS signal has two components that are tracked by the receiver: the code and the carrier. The code is modulated on top of the carrier wave and provides a measure of the length of time it took the signal to travel from the satellite to the receiver. The carrier is a sinusoidal waveform that provides a low noise but ambiguous measure of the relative distance between the satellite and receiver since tracking began.

9.1.1 Code

As described in Chapter 7, the delay-locked tracking-loop aligns the receiver's internal replica of the satellite's pseudorandom code with the incoming signal. Each GPS code is unique and is used to identify the satellite as well as to determine distance. The tracking-loop outputs the time of reception of the start of the code sequence according to the receiver's clock. The patterns of ones and zeros in the code allows the receiver to distinguish one rising edge from another and therefore the absolute position of each (with a 1 millisecond ambiguity due to the length of the C/A code). It took the signal some time to travel from the satellite to the receiver. The signal left the satellite at an earlier time in an absolute GNSS timeframe. The satellite has its own estimate of the time of transmission according to its own clock. The pseudorange measurement made by the receiver is the difference between the time of transmission according to the satellite and the time of reception according to the receiver, multiplied by the speed of light. This measurement contains the true range to the satellite, but it also contains any synchronization error between the clocks as well as the effects of other error terms that will be described below.

Each GPS C/A code has 1023 bits, or "chips" and it is broadcast at 1.023 M chips per second. The entire code sequence takes 1 millisecond to broadcast and is therefore repeated every millisecond. Each chip is slightly less than one microsecond (or about 300 m) in duration. The code-tracking loop can generally resolve the timing of the chip transitions to within a fraction of a meter.

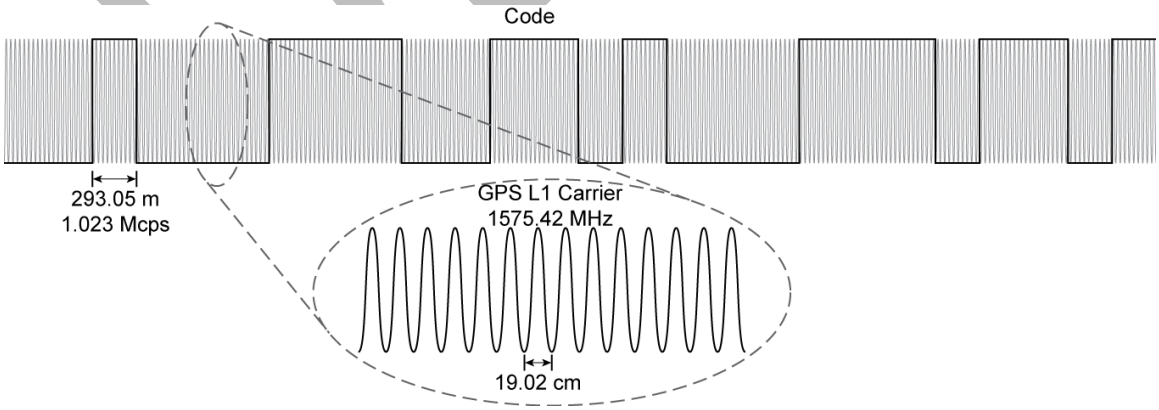

Figure 9-1 The GPS L1 C/A code and carrier signals.

9.1.2 Carrier

In addition to the code measurement the receiver implements a tracking loop that measures the phase of the underlying carrier signal. The output of this carrier phase tracking loop is a very precise measurement of the difference in phase of the signal from the time of transmission until the time of reception. The GPS L1 C/A carrier wavelength is approximately 19 cm in length. The carrier-tracking loop can generally resolve these to within a fraction of a centimeter. However, there is an ambiguity as to how many complete carrier wavelengths have passed since the time of transmission. Thus, this is a very precise measurement, but not an absolute one. Later we will describe how these two measurements can be combined to improve the overall measurement accuracy.

9.1.3 Multiple Frequencies

The GPS L1 C/A signal at 1,575.42 MHz is currently the only GPS signal used operationally by aircraft. The L2 frequency at 1,276.60 MHz has signals that are encrypted and reserved for military use. However, these signals can still be tracked, although imperfectly, and are used by ground receivers to improve GPS performance as we will see in Chapter 14. There is also a civil L2 signal on newer GPS satellites that can be fully tracked by non-military receivers. However, the new signal that draws the most aviation interest is L5 at 1,176.45 MHz. This signal, like the L1 signals, is in a frequency band designated for aeronautical radio-navigation services (ARNS). This designation is essential to authorize use of the frequency for use in air navigation.

When these L5 signals become fully operational, a receiver may have code and carrier measurements not only for the L1 C/A signal, but separate code and carrier measurements for the L5 signals as well. Later in this chapter we will discuss ways to combine these different measurements to improve overall performance.

9.1.4 Multiple Constellations

Nearly all aircraft in operation today make use of the GPS L1 signals. There is also limited use of GLONASS. Galileo and Beidou are still in their implementation phase and are not yet used by aircraft. It is expected that they too will be used for aviation as they become fully operational. The Russian Federation has issued a mandate on the use of GLONASS by Russian flagged air carriers. This is likely to increase its use, at least in Russia. It is possible that other mandates will be issued for the other constellations as well. The signals on these other constellations also provide similar code and carrier measurements. However, each constellation uses its own time reference and its own coordinate frame. As will be shown later in the chapter, when combining signals from different constellations these time and reference frame differences must be taken into account.

GLONASS distinguishes it satellites by their center frequencies rather than by their codes as is done on the other constellations. Its openly available signals are near L1 and range from 1,598.0625 to 1,605.375 MHz. This range is far enough from GPS L1 that a separate receiver front end is required to collect the signal for processing. Having different frequencies for each signal also introduces additional biases that need to be taken into account. GLONASS currently utilizes the PZ-90.11 reference system, which is aligned within centimeters to ITRF. GLONASS time is linked to Russia's implementation of UTC, UTC(SU), but unlike the other constellations, it incorporates leap seconds and is centered on the Moscow time zone. It is therefore three hours and 16 seconds ahead of GPS time. In addition, there is a clock steering error that is allowed to be as large as 1 millisecond, but is typically several hundred nanoseconds.

Galileo broadcasts signals that overlap the GPS L1/C/A and L5 signals. These signals have very similar behavior as the GPS signals. They have the same center frequencies and therefore are compatible with equipment that can receive the GPS signals. Operational service is scheduled to begin in 2020. Galileo uses the Galileo terrestrial reference frame (GTRF), which is aligned with ITRF to within 3 cm (2-sigma). Galileo system time is linked to UTC and is specified to be within 50 ns 95% of the time.

The current Beidou signals are offset from GPS L1 and do not overlay GPS L5. The open B1 signal is centered at 1,561.098 MHz. It is close enough to L1 that some aviation receiver front ends can capture the

signal. BeiDou-3 is currently being fielded and this phase adds signals centered at L1 and L5. Phase III is expected to be complete by 2020. Beidou uses the Compass geodetic system (CGS), which is aligned to within a cm of ITRF. Beidou system time is linked to UTC(NTSC) and is specified to be within 100 ns 95% of the time. Beidou time is offset from UTC by only two leap seconds and is therefore 14 seconds ahead of GPS time.

9.2 Errors Affecting Pseudorange Measurements

Chapter 5 described various error sources that affect the broadcast signal. In the next subsections we will discuss how these errors affect positioning accuracy.

9.2.1 Satellite Clock and Orbit

The position and time offset estimates for the satellite at the time of transmission are necessary to determine the user position. This information is supplied to the user through the broadcast navigation data described in Chapter 3. This data contains a description of the satellite position as a function of time and the satellite clock offset from system time also as a function of time. The parameters that describe these functions were estimated some time earlier by the control segment. Therefore they will contain some amount of error that will affect the ability of the user to form an accurate position fix.

9.2.2 Ionosphere

As the signals propagate from the satellite to the terrestrial user, they pass through the Earth's atmosphere. The constituents of the atmosphere affect the propagation of radio signals in several ways. The ionosphere is a region of the upper atmosphere where solar radiation has stripped electrons away from their atoms. This creates a region of free electrons against a backdrop of positively charged ions. Because of their electrical charge and their low mass, the electrons strongly interact with the electromagnetic fields of the GNSS signals and alter their propagation. The code and carrier components of the signal are affected differently. The code is delayed while the carrier is advanced by nearly equal amounts. The magnitude of the effect is inversely proportional to square of the frequency of the carrier signal. Higher frequency signals suffer less delay. More details about these effects can be found in Chapter 5.

The use of the more precise carrier to smooth the noisier code is limited due to the differences in the effects of the ionosphere. Further, the unknown size of the delay can create sizable errors for single frequency positioning. In the future, two frequencies will be available at the aircraft. Both measurements together may be used to directly estimate and remove the effects of the ionosphere, dramatically limiting its impact on positioning.

9.2.3 Troposphere

As the signals draw nearer to the Earth, the neutral components of the atmosphere become sufficiently dense to also affect signal propagation. Primarily it is molecular nitrogen and oxygen, together with water vapor, that cause the majority of the delay. Unlike the ionosphere, the troposphere affects all signal components and frequencies identically. The dry gases cause a relatively predictable amount of signal delay. The water vapor causes less overall delay, but the amount of water vapor is much more variable and harder to predict.

9.2.4 Multipath

As the signal reaches the users antenna, it often arrives by more than one path. Nearby objects reflect copies of the signal toward the antenna. These generally weaker copies are delayed compared to the direct path, but nonetheless cause the receiver some confusion about the precise time of arrival. This leads to a highly variable error that varies as the satellites and/or the user move in relationship to the reflectors. The code signals are much more significantly affected with errors that can be several meters. The carrier signal is also affected but only at the centimeter or millimeter level.

9.2.5 Receiver Noise

The pseudorange measurement is also affected by noise from multiple sources, including received RF radiation through the antenna and thermal noise from the receiver electronics. This noise creates a small amount of quickly varying error that also affects the code measurements much more significantly than the carrier measurements.

9.2.6 Measurement Models

The code and carrier measurements can be expressed as a sum of the true range plus the contributing error terms. The code phase measurements can be written as:

$$\rho_{j,s}^{i} = r_{j,s}^{i} + \Delta \boldsymbol{r}_{s}^{i} \cdot \mathbf{1}_{j}^{i} + \Delta B_{s}^{i} - \Delta b_{j,s} + I_{j,s}^{i} + T_{j}^{i} + M_{j,s}^{i} + \nu_{j,s}^{i} \quad (9.1)$$

where:
i is the satellite index,
j is the station/user index,
s is the signal index for specific signals on specific frequencies,
$\rho_{j,s}^{i}$ is the pseudorange measurement from user j to satellite i using signal s,
$r_{j,s}^{i}$ is the distance between the user's antenna and the reported satellite antenna phase center for signal s,
$\Delta \boldsymbol{r}_{s}^{i}$ is the vector between the true and the true phase center for satellite i on signal s,
$\mathbf{1}_{j}^{i}$ is the line of sight vector between satellite i and user j,
ΔB_{s}^{i} is the difference between the true and the reported clock for satellite i on signal s,
$\Delta b_{j,s}$ is the difference between the user clock and the system time on signal s,
$I_{j,s}^{i}$ is the ionospheric delay on the signal between user j and satellite i for signal s,
T_{j}^{i} is the tropospheric delay on the signal between user j and satellite i for (common for all signals),
$M_{j,s}^{i}$ is the multipath error at user j on satellite i for signal s, and
$\nu_{j,s}^{i}$ is the local receiver noise error at user j on satellite i for signal s.

The corresponding phase measurement can be written as

$$\phi_{j,s}^{i} = r_{j,s}^{i} + \Delta \boldsymbol{r}_{s}^{i} \cdot \mathbf{1}_{j}^{i} + \Delta B_{s}^{i} - \Delta b_{j,s} + I_{j,s}^{i} + T_{j}^{i} + N_{j,s}^{i}\lambda_{s} + M_{j,s}^{i} + \nu_{j,s}^{i} \quad (9.2)$$

where:
$N_{j,s}^{i}$ is an unknown number of integer wavelengths between the user j and satellite i for signal s, and
λ_{s} is the wavelength for signal s.

For GPS, the signals of interest are the C/A code on L1, the C/A carrier on L1, the code on L5 and the phase on L5. These will be denoted respectively as C1, L1, C5, and L5. Let us look in more detail at some of these components and their differences among the signals. Although both the user's and the satellite's antenna phase centers vary depending on the signal, this variation is usually on the order of 10 - 50 cm and is commonly ignored for aviation. We will approximate that $r_{j,s}^{i}$ and $\Delta \boldsymbol{r}_{\boldsymbol{s}}^{\boldsymbol{i}}$ are independent of the underlying signal and account for the error when we quantify uncertainty. Similarly, the timing of the different signals are not perfectly aligned with each other. The broadcast data includes inter-signal bias terms to correct for these offsets, but small errors remain. For aviation use, we can approximate that after applying these corrections ΔB_{s}^{i} is also independent of the underlying signal.

The receiver introduces signal dependent timing delays due to differential delays through its antenna and front-end filter. However, these delays primarily create a common delay across all satellites for each signal type. The principle effect is on the time solution rather than on the position estimate. This timing difference is typically measured in tens to hundreds of nanoseconds. For aviation applications, such common mode timing errors are negligible.

The ionospheric delay is different for each signal. The effects on the code and carrier are nearly equal and opposite at each frequency. Further, the delays are larger for lower frequencies. The effects can be summarized as follows:

$$I^i_{j,L1} \cong -I^i_{j,C1},\ I^i_{j,C5} \cong \frac{f_1^2}{f_5^2} I^i_{j,C1},\ I^i_{j,L5} \cong -I^i_{j,C5} \tag{9.3}$$

where f_1 and f_5 are the GPS L1 and L5 center frequencies respectively. As previously mentioned, the tropospheric delay is common to all signals. The multipath and noise errors are much larger for the code signals than they are for the carrier signals. That is:

$$M^i_{j,C1} \gg M^i_{j,L1},\ M^i_{j,C5} \gg M^i_{j,L5},\ \nu^i_{j,C1} \gg \nu^i_{j,L1},\ \nu^i_{j,C5} \gg \nu^i_{j,L5} \tag{9.4}$$

Finally, the integer ambiguity $N^i_{j,s}$ on the carrier measurements is a constant value as long as the phase tracking loop maintains lock on the signal. Effectively, the tracking loop is integrating the measured frequency and this term is the constant of integration. As long as it stays constant, the relative change in the range from one time to another is being measured very precisely. As we will see in a later section, this property is very useful for reducing the effects of code noise and multipath.

9.2.7 Uncertainty Models

The variance on the stand-alone single-frequency measurement of a particular signal, *s*, can be written as

$$(\sigma^i_{j,s})^2 = (\sigma^i_{URA})^2 + (\sigma^i_{iono,j,s})^2 + (\sigma^i_{tropo,j})^2 + (\sigma^i_{cnmp,j,s})^2 \tag{9.5}$$

where:
$(\sigma^i_{URA})^2$ is the variance on the clock and ephemeris errors on satellite i,
$(\sigma^i_{iono,j,s})^2$ is the variance on the ionospheric delay between user *j* and satellite *i* for signal *s*,
$(\sigma^i_{tropo,j})^2$ is the variance on the tropospheric delay between user *j* and satellite *i*, and
$(\sigma^i_{cnmp,j,s})^2$ is the variance on the noise and multipath error at user *j* on satellite *i* for signal *s*.

σ^i_{URA} is included in the broadcast of the GPS satellite and describes the uncertainty of the clock and ephemeris errors after using the broadcast navigation data. $\sigma^i_{iono,j,C1}$ is defined for aviation use in Appendix J of the SBAS MOPS. This term describes the residual ionospheric errors after applying the single frequency ionospheric delay model that is also broadcast from the GPS satellites. $\sigma^i_{tropo,j}$ is described in Appendix A of the SBAS MOPS as a function of elevation angle. It describes the residual uncertainty after applying the standard tropospheric delay model described in the same document. For an airborne user $\sigma^i_{cnmp,j,C1}$ can be set equal to σ_{air} as defined in Appendix J of the SBAS MOPS.

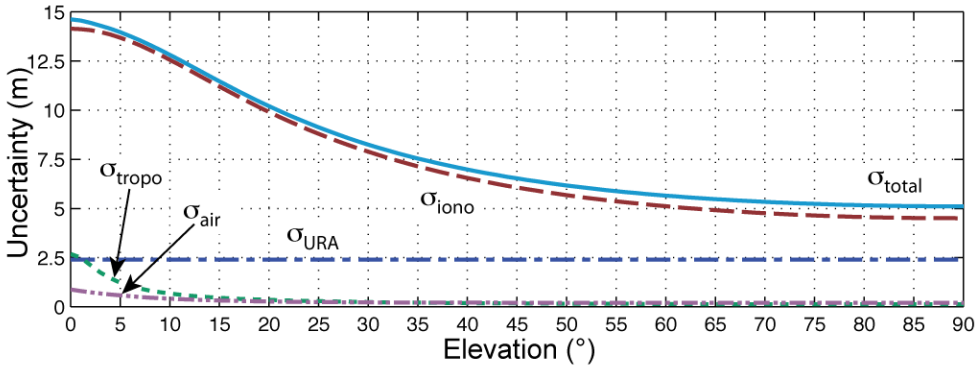

Figure 9-2 The GPS L1 C/A stand-alone uncertainty components.

Figure 9-2 depicts these uncertainty components as a function of elevation angle to the satellite. The σ_{URA} value is most commonly set to 2.4 m currently and does not depend on elevation angle. The ionopspheric uncertainty is by far the largest term, while the tropospheric, noise and multipath terms are almost negligible. The values plotted here are very conservative and are intended to describe uncertainty under near worst-case conditions.

9.3 Carrier-Smoothing of Code Measurements

The concept of carrier smoothing dates back to the early 1980's. The idea is a very clever one: use the noisy code measurements to estimate the bias on the comparatively quiet carrier measurements. Thus, we take the best of each: the absolute measurement of the code and the low noise of the carrier. This works exceptionally well because the code and carrier are nearly identically affected by many of the same error sources. As we will see, carrier smoothing is an extremely effective way to reduce the effects of multipath and receiver termal noise on the position solution.

9.3.1 Single Frequency Carrier-Smoothing

On any individual frequency, the code and carrier are not affected identically by the ionosphere. The code is delayed by an amount proportional to the Total Electron Content (TEC) along the line of sight from the satellite to the user. The carrier, however, is advanced by nearly the same amount. In mid-latitudes (approximately between 20 and 60 degrees), the ionosphere normally is slowly varying. Therefore, for short smoothing times, this difference can be incorporated into the estimate of the carrier bias and the resulting error is negligible. However, there are times when the ionosphere may change by a large amount in a comparatively short time. These events are rare in mid-latitudes, but are more common in other parts of the world, particularly in equatorial regions. Rapid ionospheric changes can introduce significant error to the smoothing process. For very large variations, the error can be significantly larger than the multipath errors the smoothing is intended to mitigate.

To simplify notation, we will omit the satellite, user, and signal indices of the prior section and concentrate onln on a single line of sight. We will denote the instantaneous pseudorange measurement at time t_k by ρ_k and the corresponding phase measurement by ϕ_k. The smoothed pseudorange at time t_k is designated $\bar{\rho}_k$ and is given by

$$\bar{\rho}_k = \frac{\rho_k + (n-1)\cdot(\bar{\rho}_{k-1} + \phi_k - \phi_{k-1})}{n} \tag{9.6}$$

Here we can see that the smoothing estimate from the prior epoch, $\bar{\rho}_{k-1}$, is updated, using the difference of the carrier measurements, and averaged with the current pseudorange measurement. The propagated term is given more weight; here n represents the amount of smoothing. Larger values indicate a greater amount of smoothing with a longer associated time constant.

Often at the beginning of the smoothing process n is set equal to k. The first value cannot be smoothed and is set to ρ_0. The second value is the average of $\rho_0 + \phi_1 - \phi_0$ and ρ_1. Adding the term $\phi_1 - \phi_0$ to ρ_0 accounts for any changes to the signal between the times t_1 and t_0. This process is repeated for time 2 through n. At some point, the much earlier measurements are no longer representative of the current state. Small differences between the code and carrier errors accumulate and can become significant. For aviation, n is limited to 100 (assuming a one second interval between measurements) as after 100 seconds the difference of the ionospheric effect can be significant.

Figure 9-3 provides an example of the smoothing process

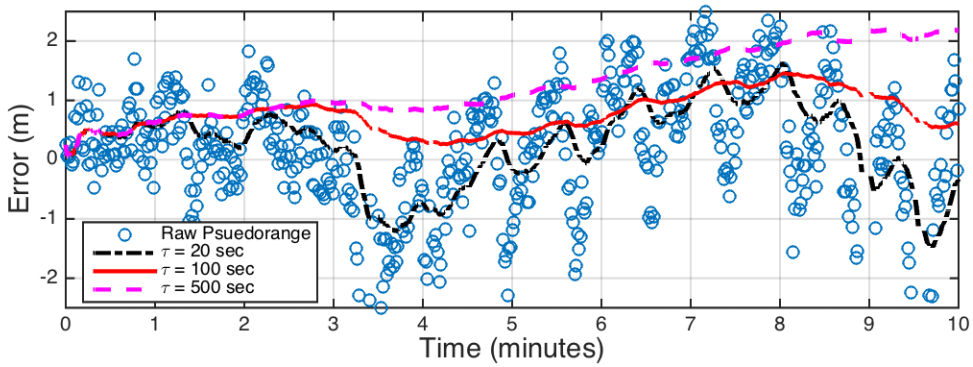

Figure 9-3 Smoothed and unsmoothed GPS L1 C/A pseudorange measurements.

We can investigate the effects of carrier smoothing in more detail by studying the most relevent error components of the code and carrier measurements, the ionosphere and the code multipath

$$\begin{aligned} \rho_k &= r_k + I_k + M_k \\ \phi_k &= r_k - I_k + bias \end{aligned} \tag{9.7}$$

Here r_k represents common terms such as the range to the satellite, clock offsets, and tropospheric delay, I_k is the ionospheric delay/advance at L1, M_k is the code multipath and *bias* is the carrier bias. Substituting these values into (9.6) and subtracting $r_k + I_k$ from both sides yields the error in the smoothed value ($\varepsilon_k = \bar{\rho}_k - r_k - I_k$) as a function of time

$$\varepsilon_k = \frac{(n-1)\cdot(\varepsilon_{k-1}) - 2(n-1)\cdot(I_k - I_{k-1}) + M_k}{n} \tag{9.8}$$

With a little rearranging, the above equation can be approximated by the continuous expression

$$\frac{\partial \varepsilon}{\partial t} = -\frac{\varepsilon}{\tau} - 2\frac{\partial I}{\partial t} + \frac{M}{\tau} \tag{9.9}$$

Where $\tau = (n - 1)(\Delta t)$ is the smoothing time constant. This is a first order linear differential equation whose solution is given by

$$\varepsilon = e^{-\frac{t}{\tau}}\left\{\varepsilon_0 - \int_{t'=0}^{t} e^{\frac{t'}{\tau}} \cdot \left(2\frac{\partial I}{\partial t'} - \frac{M}{\tau}\right) dt'\right\} \tag{9.10}$$

This is a non-iterative expression for the error in the smoothing filter as a function of the ionospheric delay and multipath.

To see the effect of multipath on the filter, let us examine the case with no ionospheric change. We set the ionospheric divergence term to zero and ignore the initial value to obtain

$$\varepsilon = e^{-\frac{t}{\tau}}\left\{\int_{t'=0}^{t} e^{\frac{t'}{\tau}} \cdot \frac{M(t')}{\tau} dt'\right\} \tag{9.11}$$

It is instructive to examine two extremes: a constant bias and white noise. If the multipath has a very long period compared to τ, then it is effectively a constant term in the integral. The resulting error term equals the

constant multipath. Thus, the smoothing filter is completely ineffective against long-period multipath. If the multipath is modeled as uncorrelated, zero-mean, white noise, then the expected value of the error will be zero while the expected variance of the error will the variance of the multipath divided by $\tau = 2(n - 1)$.

Actual multipath typically falls somewhere between these two models. There is some time correlation, which for airborne users should be less than the 100-second time constant. Thus, in absence of ionospheric variations, longer time constants lead to smaller errors. However, as we will see in the next sections, rapid rates of change in the ionosphere limit the performance of these filters.

As we can see from (9.10), the error in the smoothing filter depends very strongly on the rate of ionospheric variation. If we model the ionosphere as having a constant slope, and we neglect the multipath error, we can evaluate the integral explicitly

$$\begin{aligned} I &= a \cdot t \\ \frac{\partial I}{\partial t} &= a \\ \varepsilon &= e^{-\frac{t}{\tau}} \left\{ \varepsilon_0 - 2 \int_{t=0}^{t} e^{\frac{t}{\tau}} \cdot a dt \right\} \\ \varepsilon &= e^{-\frac{t}{\tau}} \left\{ \varepsilon_0 - 2a \cdot \boldsymbol{\tau} \cdot e^{\frac{t}{\tau}} \Big]_{t=0}^{t} \right\} \\ \varepsilon &= \varepsilon_0 \cdot e^{-\frac{t}{\tau}} - 2a \cdot \boldsymbol{\tau} \cdot \left(1 - e^{-\frac{t}{\tau}} \right) \end{aligned} \tag{9.12}$$

This shows that any initial offset decays to zero while the filter error asymptotically approaches a constant offset equal to negative two times the rate of change multiplied by the time constant. For example, a rate of change of +10 mm/second would approach a –2 m error in a filter with a 100-second time constant. Of course, for the error to come close to the limit, the rate of change would have to be sustained for a period lasting several time constants.

9.3.2 Dual Frequency Carrier-Smoothing

With the availability of two frequencies the risk ionospheric divergence can be overcome. Instead of using the pseudorange measurement from just one frequency, the information from two frequencies can be combined to eliminate the effects of the ionosphere. Because the ionospheric delay is proportional to $1/f^2$, the effect of the ionosphere on the L1 code and carrier can be directly estimated

$$\begin{aligned} I_\rho &= \left(\rho_{L5} - \rho_{L1} \right) \frac{f_{L5}^2}{f_{L1}^2 - f_{L5}^2} \\ I_\phi &= -\left(\phi_{L5} - \phi_{L1} \right) \frac{f_{L5}^2}{f_{L1}^2 - f_{L5}^2} \end{aligned} \tag{9.13}$$

These estimates can be subtracted from the measurements to create ionosphere-free measurements

$$\begin{aligned} \rho_{iono\_free} &= \rho_{L1} - I_\rho = \frac{f_{L1}^2 \rho_{L1} - f_{L5}^2 \rho_{L5}}{f_{L1}^2 - f_{L5}^2} \\ \phi_{iono\_free} &= \phi_{L1} + I_\phi = \frac{f_{L1}^2 \phi_{L1} - f_{L5}^2 \phi_{L5}}{f_{L1}^2 - f_{L5}^2} \end{aligned} \tag{9.14}$$

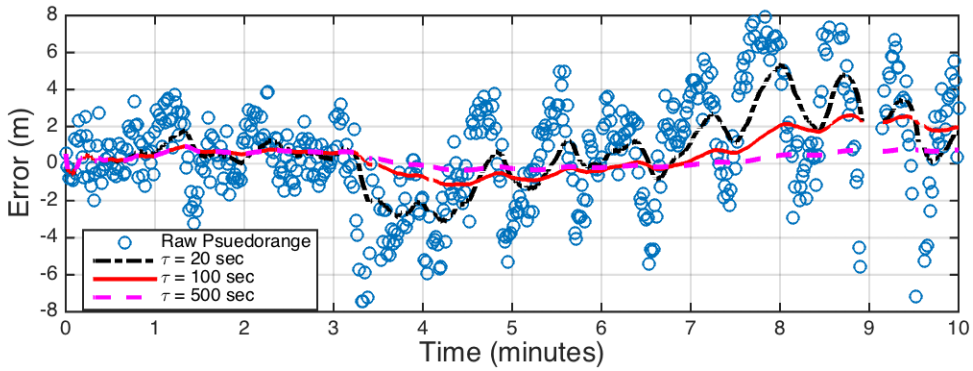

Figure 9-4 Smoothed and unsmoothed GPS dual-frequency, ionosphere-free pseudorange measurements.

These iono-free measurements may then be substituted for the single frequency values in (9.6). They will have no ionospheric discrepancy between the code and carrier. The smoothing time constant can be made much longer. Figure 9-4 shows the performance of dual frequency smoothing. Notice that longer smoothing times now lead to better accuracy and are not susceptable to being pulled off due to ionospheric rates of change. However, in a moving platform such as an airplane, there may be other discrepencies, such as undetected cycle slips and carrier phase wind-up, that discourage the use of very long smoothing times. The ideal airborne dual-frequncy smoothing time is still under evaluation, but is likely to be between 100 and 1000 seconds.

Also notice in Figure 9-4 that the iono-free combination increases the magnitude of the multipath and thermal error. The factors multiplying the pseudorange measurements in (9.14) are approximately 2.26 for L1 and 1.26 for L5 leading to a significant increase in contribution from these two error sources. However, the elimination of the ionosphere error is worth this trade in many situations. Further the increased errors can be mitigated by increasing the smoothing time.

9.4 Least-Squares PVT Computation

GNSS positioning is accomplished by finding the position that best matches the measured pseudoranges. This process is comparable to plotting spheres around the known satellite locations and finding the point of intersection, a technique known as trilateration. However, because the user's clock is not synchronized to GPS, and it is part of the pseudorange measurements, the radii of the spheres vary with the solution as well. Fortunately, the user clock error is common to all measurements and is then included as one of the unknowns to be determined. There are four of these unknowns: the user's true three dimensional position coordinates which we will call x, y, and z (where $\mathbf{x} = [x\ y\ z]^{\mathrm{T}}$), and the user clock offset from GPS time, $c\delta t_u$, which we will denote as b.

If $\mathbf{x}^{(i)} = [x^{(i)}\ y^{(i)}\ z^{(i)}]^{\mathrm{T}}$ is the true position of satellite i, and $B^{(i)}$ is the true satellite clock offset from GPS time, the true range to the satellite would be given by

$$r^{(i)} = \sqrt{\left(x^{(i)} - x\right)^2 + \left(y^{(i)} - y\right)^2 + \left(z^{(i)} - z\right)^2} = \left\|\mathbf{x}^{(i)} - \mathbf{x}\right\| \tag{9.15}$$

and the error-free pseudorange would be

$$\rho^{(i)} = r^{(i)} + b - B^{(i)} = \left\|\mathbf{x}^{(i)} - \mathbf{x}\right\| + b - B^{(i)} \tag{9.16}$$

The pseudoranges have a non-linear dependence on the four unknowns. The easiest method to solve these equations is to linearize them about an initial guess at the user location and clock, $\mathbf{x_0}$ and b_0. The true position and clock can be expressed as offsets from this initial state: $\mathbf{x} = \mathbf{x_0} + \boldsymbol{\delta}\mathbf{x}$, and $b = b_0 + \delta b$. The pseudorange can then be written as

$$\rho^{(i)} \approx r_0^{(i)} + b_0 - B^{(i)} - \frac{\mathbf{x}^{(i)} - \mathbf{x_0}}{\left\| \mathbf{x}^{(i)} - \mathbf{x_0} \right\|} \cdot \boldsymbol{\delta}\mathbf{x} + \delta b = r_0^{(i)} + b_0 - B^{(i)} - \mathbf{1}^{(i)} \cdot \boldsymbol{\delta}\mathbf{x} + \delta b \tag{9.17}$$

where the derivation of this equation can be found in [ref]. The difference between the observed pseudorange and the expected value given the assumed position and clock states is given by

$$\delta\rho^{(i)} = \rho^{(i)} - r_0^{(i)} - b_0 + B^{(i)} \tag{9.18}$$

where $r_0^{(i)}$ and b_0 correspond to the starting estimate for the user position and clock. The *N* pseudorange measurements can be stacked and written in matrix notation as

$$\boldsymbol{\delta\rho} = \begin{bmatrix} \delta\rho^{(1)} \\ \delta\rho^{(1)} \\ \vdots \\ \delta\rho^{(N)} \end{bmatrix} = \begin{bmatrix} \left(-\mathbf{1}^{(1)}\right)^{\mathrm{T}} & 1 \\ \left(-\mathbf{1}^{(2)}\right)^{\mathrm{T}} & 1 \\ \vdots & \vdots \\ \left(-\mathbf{1}^{(N)}\right)^{\mathrm{T}} & 1 \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{\delta}\mathbf{x} \\ \delta b \end{bmatrix} = \mathbf{G} \cdot \begin{bmatrix} \boldsymbol{\delta}\mathbf{x} \\ \delta b \end{bmatrix} + \boldsymbol{\varepsilon} \tag{9.19}$$

where **G** is called the geometry matrix and it describes how changes in the position and clock estimates, from the initial estimates, affect the expected pseudoranges. Note that we have now added an error term, $\boldsymbol{\varepsilon}$, to (9.19) to reflect that there will be some unmodeled measurement errors affecting the pseudorange. We wish to find the inverse relationship, so that given the differences in the measured pseudoranges from the expected values we can find the position and clock offsets. This inverse operation can be found through

$$\begin{bmatrix} \boldsymbol{\delta}\hat{\mathbf{x}} \\ \delta\hat{b} \end{bmatrix} = \left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1} \cdot \mathbf{G}^{\mathrm{T}} \cdot \boldsymbol{\delta\rho} \tag{9.20}$$

These estimates can be used to update the initial position and clock estimates and the process can be iterated until the solution converges. This process is very well behaved and should converge reliably with only a few iterations. For a cold start the starting time and position can be all zeros (center of the Earth with the clock assumed to be fully synchronized to system time). If there has been a recent position solution then it can serve as the starting point for the next time update. In this latter case, no iteration is typically needed as the previous estimates should be sufficiently close that the linear approximation is sufficiently accurate.

The prior expression for the inverse assumes that all measurements are equally accurate. However, as we have already seen, lower elevation satellites typically have more error. It is better to include knowledge of the expected accuracy of the pseudorange measurements when balancing the relative contribution from one measurement versus the other ones. A weighting matrix can be formed from the confidences previously specified

$$\mathbf{W^{-1}} = \begin{bmatrix} \left(\sigma^{(1)}\right)^2 & & & \cdots & 0 \\ & \left(\sigma^{(2)}\right)^2 & & \cdots & 0 \\ \vdots & \vdots & \ddots & & \vdots \\ 0 & 0 & \cdots & & \left(\sigma^{(N)}\right)^2 \end{bmatrix} \tag{9.21}$$

The optimum position change estimate reflecting this weighting is given by

$$\begin{bmatrix}\boldsymbol{\delta\hat{\mathbf{x}}}\\ \delta\hat{b}\end{bmatrix} = \left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\boldsymbol{\delta\rho} \tag{9.22}$$

Note that the weighting only comes into play if there are more than four pseudorange measurements. If there are only four, then the equations have an exact solution and the relative weighting has no effect. Also notice that the absolute weighting is not important. The solution is the same if the weighting matrix is multiplied by any arbitrary non-zero scalar. It is only the relative weights between the different measurements that affect the outcome.

The weighted least squares approach minimizes the error in the pseudorange residuals as described by

$$\chi^2 = \hat{\boldsymbol{\varepsilon}}^{\mathrm{T}}\cdot\mathbf{W}\cdot\hat{\boldsymbol{\varepsilon}} \tag{9.23}$$

where $\hat{\boldsymbol{\varepsilon}}$ is obtained from

$$\hat{\boldsymbol{\varepsilon}} = \boldsymbol{\delta\rho} - \mathbf{G}\cdot\begin{bmatrix}\boldsymbol{\delta\hat{\mathbf{x}}}\\ \delta\hat{b}\end{bmatrix} \tag{9.x}$$

The covariance of this error was already described by (9.21). We also wish to find the expected covariance of the position solution, **P**, which is given by

$$\begin{aligned}\mathbf{P} &= \langle\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\boldsymbol{\varepsilon}\cdot\left[\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\boldsymbol{\varepsilon}\right]^{\mathrm{T}}\rangle\\ &= \langle\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\boldsymbol{\varepsilon}\cdot\boldsymbol{\varepsilon}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\cdot\left[\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\right]^{\mathrm{T}}\rangle\\ &= \left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\langle\boldsymbol{\varepsilon}\cdot\boldsymbol{\varepsilon}^{\mathrm{T}}\rangle\cdot\mathbf{W}\cdot\mathbf{G}\cdot\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\\ &= \left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{W}^{-1}\cdot\mathbf{W}\cdot\mathbf{G}\cdot\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\\ &= \left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\end{aligned} \tag{9.x}$$

where the brackets $< x >$ indicate the expected value of x. The final result is already calculated as part of position solution estimate in (9.22). We will return to this estimate later

9.4.1 Single-Frequency Solution

The equations in the previous section did not explicitly describe the effects of the error sources. For a single-frequency user, models to describe the ionosphere and the troposphere should be applied to reduce the effects of these error sources. GPS broadcasts parameters for a simple single frequency ionospheric model [ref] that, on average, is able to remove about half of the overall ionospheric delay. The aviation community has also adopted a standard tropospheric model that is able to eliminate approximately 95% of the overall tropospheric delay [ref]. For a single-frequency user, (9.18) is modified to

$$\delta\rho^{(i)} = \bar{\rho}_{L1}^{(i)} - r_0^{(i)} - b_0 + B_{L1}^{(i)} - \hat{I}^{(i)} - \hat{T}^{(i)} + \varepsilon^{(i)} \tag{9.24}$$

where $\bar{\rho}_{L1}^{(i)}$ is the smoothed pseudorange measurement, and $\hat{I}^{(i)}$ and $\hat{T}^{(i)}$ are the model estimates for the ionospheric and tropospheric delays respectively. If the initial position and clock estimates are too far off, it may be necessary to recalculate, $B_{L1}^{(i)}$, $\hat{I}^{(i)}$ and $\hat{T}^{(i)}$ as part of the iteration. Usually, the last calculated position as the starting point for the next update is sufficiently accurate provided that updates are calculated frequently and the user velocity is not too large. The satellite clock term should include all of the elements described in the GPS interface specification [ref] Sections 20.3.3.3.3.1 and 20.3.3.3.3.2, multiplied by the speed of light to convert the units from seconds to meters.

The satellite variance values to be used for the weighting matrix (9.21) are given in Section 9.2.7 by (9.5). $(\sigma_{URA}^{i})^2$ is broadcast as part of the satellites navigation message. In aviation, these variances are made

extremely large to overbound the worst-case situations. Safety-of-life is more important than accuracy. GPS broadcasts an user range accuracy (URA) index that can be used to find a range of possible values. In aviation, the upper bound is most commonly used for this variance value. The upper limit on the ionospheric variance, $(\sigma^i_{iono})^2$, after applying the single frequency (Klobuchar) model is specified in Appendix J of the SBAS MOPS [ref]. Similarly, the upper limit on the ionospheric variance, $(\sigma^i_{tropo})^2$, is specified in Appendix A of the SBAS MOPS. Finally, the airborne, post-smoothing code noise and multipath (CNMP) variance, $(\sigma^i_{air})^2$, is specified in Appendix J of the SBAS MOPS. Together, these four variances are added to find the weighting matrix.

Equation (9.22) is then used to refine the user position and clock estimates. Equations (9.15), (9.24), and (9.22) can be recalculated until $\boldsymbol{\delta\hat{x}}$ and $\delta\hat{b}$ are sufficiently small and the solution has converged. Figure 9-5 shows the position solution results for an hour's worth of data. The dominant error source is the ionosphere that can remain relatively constant over many tens of minutes to hours. The next largest error source is the satellite clock error that also remains relatively constant over this time period. The multipath and thermal noise are comparatively smaller and have been effectively mitigated by the smoothing. The sudden jumps are due to changes in visible satellites as satellites rise and set above a minimum 5° threshold.

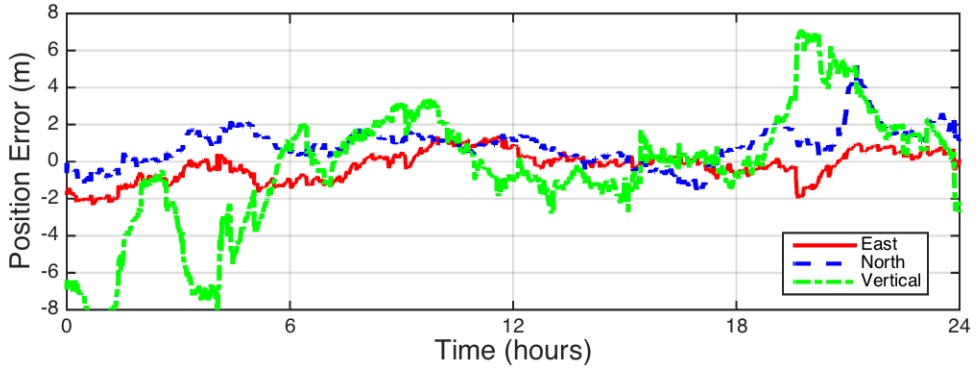

Figure 9-5 Single-frequency position estimation errors

9.4.2 Dual-Frequency Solution

The ionospheric errors and uncertainty variance can be greatly minimized by using the ionospheric-free combination of pseudoranges. Using the pseudorange and carrier combinations of (9.14), in place of the GPS L1 C/A single frequency measurements, virtually eliminates the ionospheric effects. Equation (9.24) can be replaced by

$$\delta\rho^{(i)} = \rho^{(i)}_{iono\_free} - r^{(i)}_0 - b_0 + B^{(i)}_{L5L1} - \hat{T}^{(i)} + \varepsilon^{(i)} \tag{9.25}$$

Note that the satellite clock term, $B^{(i)}$, is slightly different for the different frequency combinations. In the single-frequency case, it includes the broadcast group delay term: T_{GD}. The iono-free combination replaces $c\ T_{GD}$ with

$$c\left(\frac{ISC^{(i)}_{\mathrm{L5Q5}} - \gamma_{15} ISC^{(i)}_{\mathrm{L1CA}}}{1-\gamma_{15}} - T_{GD}\right) \tag{9.26}$$

As described in Section 20.3.3.3.1.2.2 of L5 GPS interface specification [ref], where $\gamma_{15} = (f_{L1}/f_{L5})^2 = (1575.42/1176.45)^2 \approx 1.8$.

Equation (9.25) does not use the single frequency ionospheric model as the ionospheric-free combination already eliminates these effects. The advantage of this approach is greates when the single frequency model is least accurate. Figure 9-6 shows the position accuracy for the same dataset as Figure 9-6 except that it uses the iono-free combination in place of L1 only measurements. The ionospheric model errors are largest at at the beginning and end of the 24 hour period. Notice that the dual-frequency processing has better accuracy in those regions. However, towards the middle of the data set, the ionopsheric model accuracy is good and the dual frequency processing suffers in comparison as it magnifies the effects of multipath and noise. The real advantage is that the ionospheric combination eliminates the uncertainty of whether or not the ionospheric model is good. Although it may be accurate the majority of the time, the few times it is very inaccurate can lead to unacceptably large position errors. The ionospheric-free combination limits the large errors even though it may make the nominal errors slightly worse.

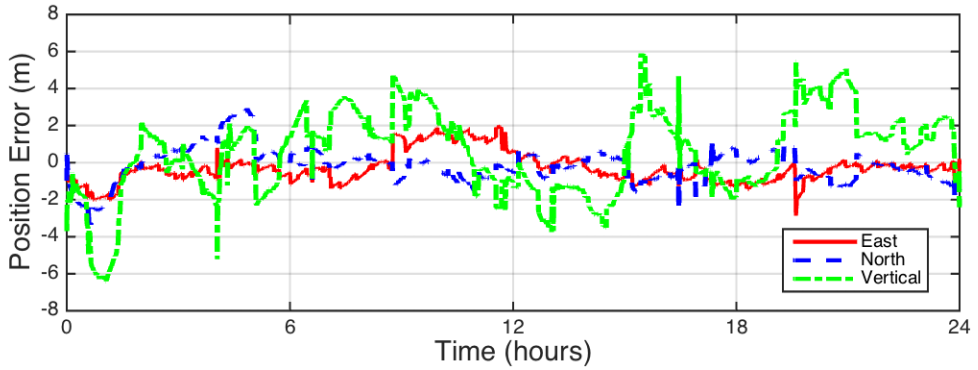

Figure 9-6 Dual-frequency, ionospheric-free position estimation errors

9.4.3 Multiple Constellation Solution

Pseudoranges from more than one constellation can be combined to create a better position solution estimate. However, one must be very careful as different systems may use different coordinate reference frames and time-bases. The coordinate frames are more easily handled as the major constellation service providers are all using systems that are very similar to the international terrestrial reference frame ITRF-93. For aviation purposes these frames may be considered to be identical. For more precise positioning or if using a system that is not essentially identical to ITRF-93, then the satellite positions from one system may need to be converted into the other systems reference frame. For many reference systems there are well-defined conversions down to some level of precision. The user would need to find such a conversion that meets their needs. The official reference frame for GPS is WGS-84 [ref], for GLONASS it is PZ-90.11 [ref], for Galileo it is the Galileo terrestrial reference frame (GTRF) [ref], and for BeiDou it is the BeiDou terrestrial reference frame (BTRF) [ref]. Fortunately, all of these reference frames are aligned to within centimeters of each other [ref] and for aviation purposes, their differences can be neglected.

The conversion between time references is more difficult. These differences need to be updated continuously. They can be monitored and predicted hours ahead as are the orbits and clock of the individual constellations. All of the constellations provide the expected offsets of their systems clock from their states implementation of UTC. However, the resulting offset from one system to another may be on the order of 100 nsec (30 m). Galileo will broadcast the offset of Galileo time from GPS time, but it has not yet been established how accurate this broadcast value will be. Further, due to differences in the signal structures and in receiver filtering, the offset from one constellation to another may be dependent on the specific receiver implementation. The best approach to combining pseudorange measurements from different constellations is to include an unknown clock state for each constellation. Thus, a solution using only one constellation has four total unknowns, where one using two constellations should have five unknowns. The geometry matrix for the multi-constellation case is given by

$$\mathbf{G} = \begin{bmatrix} \left(-\mathbf{1}^{(1,1)}\right)^{\mathrm{T}} & 1 & 0 & \dots & 0 \\ \left(-\mathbf{1}^{(1,2)}\right)^{\mathrm{T}} & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \left(-\mathbf{1}^{(1,K_1)}\right)^{\mathrm{T}} & 1 & 0 & \dots & 0 \\ \left(-\mathbf{1}^{(2,1)}\right)^{\mathrm{T}} & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \left(-\mathbf{1}^{(J,N_J)}\right)^{\mathrm{T}} & 0 & \dots & 0 & 1 \end{bmatrix} \tag{9.27}$$

and there are *J* unknown clock offsets, one for each constellation supplying a valid measurement. It is also possible to model the extra states as clock differences relative to one of the constellations. These two approaches are mathematically identical.

As of this writing only GPS and GLONASS have full constellaitons. Figure 9-7 Shows the position solution for the same location and times as was used in Figures 9-5 and 9-6. Unfortuntely, GLONASS has fewer healthy satellites than does GPS and its ranging accuracy is worse. As a result, combining GPS and GLONASS together only results in a modest improvement compared to using GPS by itself in Figure 9-6. Using a second constellation whose performance matches GPS would result in a noticibly larger improvement. One would expect the position errors to be smaller by a factor of 1/sqrt(2).

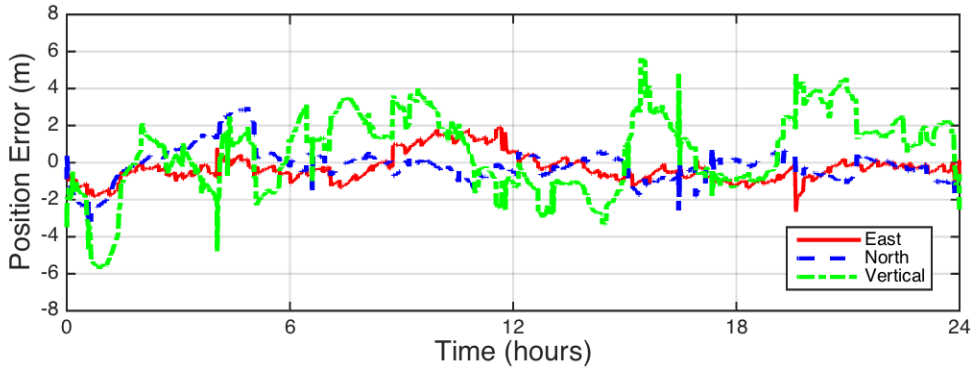

Figure 9-7 Dual-constellation, dual-frequency, ionospheric-free position estimation errors

9.5 Kalman Filter PVT Computation

The least squares solution of the previous section uses a set of measurements that were all taken at the same time. This is commonly referred to as a snapshot solution. That is, it only uses measurements from one particular time step to estimate the state at that time step. An alternative approach is to include past information in the estimation of the current state. A Kalman filter updates state estimates over time, combining past state estimates with new measurements. The Kalman filter starts with an initial estimate of the state and an initial confidence in that state (in the form of a covariance matrix). This state is updated to the current time, as is the covariance matrix. The current set of measurements is compared against the predicted measurements given the expected state. The difference between the two is called the innovation and is used to update the state (and covariance) to the best current estimate. At the next time step, the process repeats itself: the previous state is predicted forward in time and then updated with measurements. The Kalman filter is a careful blend between the known dynamics of the system and the accuracy of the measurements. If both are well characterized, the filter outputs an optimum estimate of the user position. If the system dynamics are limited, the Kalman filter can substantially reduce noise in the estimate.

Further, it can update the state with fewer than four measurements. However, if the dynamic model is incorrect or the state covariance is too small, the filter can ignore the measurements and provide misleading and overly confident estimates in the users location.

The Kalman filter starts with an initial estimated state $\mathbf{x_0}$ and covariance $\mathbf{P_0}$, usually estimated from the snapshot solution (9.22). A state transition matrix, $\mathbf{F}$, is used to predict the states value at the next time step. For example, in a simple kinematic model, the state would include position and velocity. The updated state would assume that the velocity does not change and the position is the previous position plus the previous velocity multiplied by the time difference between steps. For a two dimensional kinematic example the state estimate, state covariance matrix, and state transition matrix are given by:

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ \dot{x} \\ \dot{y} \end{bmatrix}, \quad \mathbf{P} = \begin{bmatrix} \sigma_x^2 & \sigma_{x,y}^2 & \sigma_{x,\dot{x}}^2 & \sigma_{x,\dot{y}}^2 \\ \sigma_{x,y}^2 & \sigma_y^2 & \sigma_{y,\dot{x}}^2 & \sigma_{y,\dot{y}}^2 \\ \sigma_{x,\dot{x}}^2 & \sigma_{y,\dot{x}}^2 & \sigma_{\dot{x}}^2 & \sigma_{\dot{x},\dot{y}}^2 \\ \sigma_{x,\dot{y}}^2 & \sigma_{y,\dot{y}}^2 & \sigma_{\dot{x},\dot{y}}^2 & \sigma_{\dot{y}}^2 \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{9.x}$$

In this example the state transition matrix is only a function of the time between steps. For more complicated dynamic models it could be a function of time as well. The prediction step that updates the prior step to the current one is then given by

$$\mathbf{x}_{k|k-1} = \mathbf{F}_k \cdot \mathbf{x}_{k-1|k-1} \tag{9.x}$$

where the subscript $k|j$ indicates the estimate at time step k using measurement up to and including time step j. The covariance matrix estimate also needs to be predicted

$$\mathbf{P}_{k|k-1} = \mathbf{F}_k \cdot \mathbf{P}_{k-1|k-1} \cdot \mathbf{F}_k^{\mathrm{T}} + \mathbf{Q}_k \tag{9.x}$$

where $\mathbf{Q}_k$ is a covariance matrix describing the process noise of the dynamic model. In our simple kinematic model we have neglected acceleration, mapping this uncertainty into the state estimate covariance matrix yields

$$\mathbf{Q} = \sigma_{acc}^2 \begin{bmatrix} \frac{1}{4}\Delta t^4 & 0 & 0 & 0 \\ 0 & \frac{1}{4}\Delta t^4 & 0 & 0 \\ 0 & 0 & \Delta t^2 & 0 \\ 0 & 0 & 0 & \Delta t^2 \end{bmatrix} \tag{9.x}$$

After the prediction step in the Kalman filter the state is then updated with the measurements at time step k. The expected measurement values given our predicted state are given by $\mathbf{G}_k \cdot \mathbf{x}_{k|k-1}$. The innovation or difference between actual and expected measurements is then given by

$$\mathbf{y}_k = \boldsymbol{\rho}_k - \mathbf{G}_k \cdot \mathbf{x}_{k|k-1} \tag{9.x}$$

If the predicted state were perfectly accurate then the innovation should only contain the error in the pseudorange measurements. However, if the state prediction is inaccurate, the innovation also contains information about that error. The measurement error covariance matrix as given previously is $\mathbf{W^{-1}}$. The covariance matrix of the innovation is therefore given by

$$\mathbf{C}_k = \mathbf{G}_k \cdot \mathbf{P}_{k|k-1} \cdot \mathbf{G}_k^{\mathrm{T}} + \mathbf{W}_k^{-1} \tag{9.x}$$

The state estimate is updated by the measurements according to

$$\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + \mathbf{K}_k \cdot \mathbf{y}_k \tag{9.x}$$

and the state covariance matrix is updated by

$$\mathbf{P}_{k|k} = \left(\mathbf{I} - \mathbf{K}_k \cdot \mathbf{G}_k\right)\mathbf{P}_{k|k-1} \tag{9.x}$$

where $\mathbf{I}$ is the identity matrix, and $\mathbf{K}_k$ is the Kalman gain matrix. It is optimally given by

$$\mathbf{K}_k = \mathbf{P}_{k|k-1} \cdot \mathbf{G}_k^{\mathrm{T}} \cdot \mathbf{C}_k^{-1} \tag{9.x}$$

At the next time step the process of predicting and then updating with measurements is repeated. These steps are illustrated in figure 9-x.

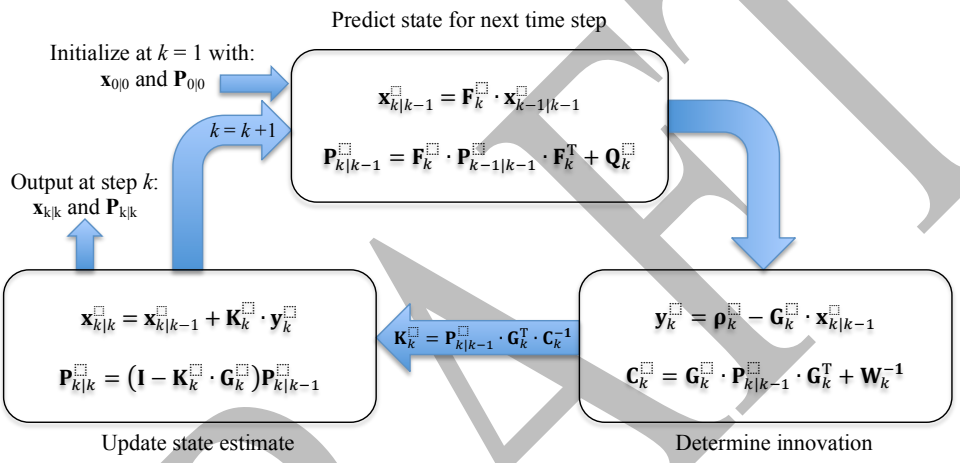

Figure 9-x Kalman filter processing cycle

Figure 9-x shows the results of implementing the Kalman filter on the dual frequency data used for Figure 9-6. We used a simple kinematic model with eight states: (x, y, z, b, $\dot{x}$, $\dot{y}$, $\dot{z}$, $\dot{b}$). The first four elements of the state were initialized with the initial snapshot solution using (9.22). The velocity states were all set to zero. The covariance matrix elements corresponding to the position and clock states were initialized using the covariance matrix from the snapshot solution. The remaining elements were set to zero, except for the last four diagonal elements, which were set to (10 m/s)^2 for the position velocity variance and (100 m/s)^2 for the clock velocity variance. The 1-σ acceleration uncertainty for the $\mathbf{Q}$ matrix was set to 5 mm/s^2 for the position and 5 cm/s^2 for the clock. These values were chosen as they are representative of a static user. For a moving platform, larger values are likely more appropriate. The Kalman filter estimates are smoother than the snapshot estimates, but not necessarily any more accurate. The filter avoids some of the spikes that affected the snapshot solution, but also introduces some lag to the estimate. The lag comes about because there are no actual measurements of velocity being incorporated into the filter. Instead it is estimating velocity by differencing position estimates. Such a process always introduces latency to the velocity estimate. The situation may be improved by incorporating true velocity or acceleration estimates into the filter. GNSS measurements are often combined with measurements from inertial sensors in more advanced Kalman filter implementations. Such implementations are beyond the scope of this book, but the interested reader will find many such examples in the literature, for example [refs].

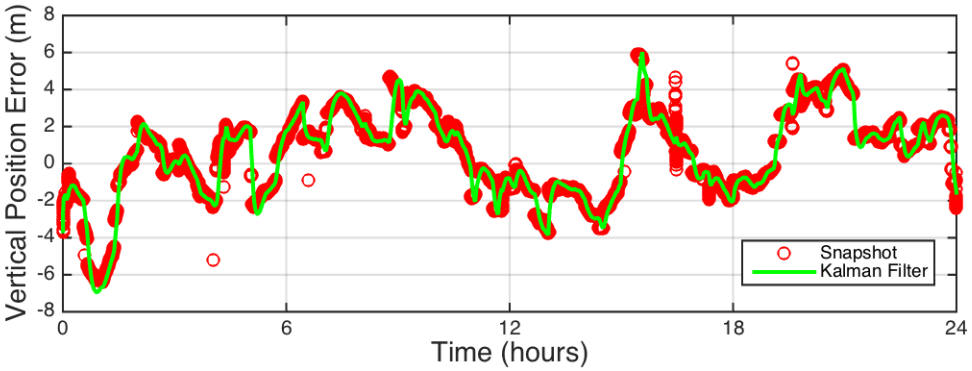

Figure 9-x Dual-frequency, Iono-free vertical position errors for snapshot and Kalman filter solutions

The velocity estimates from our Kalman filter implementation are shown in Figure 9-x. Since the receiver was stationary, these estimates also represent the velocity estimation error. The spikes in the vertical velocity estimate correspond to jumps and spikes in the snapshot position estimate as shown in figure 9-x. Note that it is on the order of our assumed value of 5 mm/s2. Had we assumed a much smaller value in **Q**, it would have led to a smoother estimate with even greater lag. Had we assumed a much larger value, the Kalman filter output would have more closely matched the snapshot estimate. This selection of **Q** (relative to $\mathbf{W}^{-1}$) is critical to the tuning of the Kalman filter. If **Q** is too small, the data may be ignored in favor of the model, leading to smooth, but inaccurate results. If **Q** is too large, the model has little influence and the filter is providing little benefit relative to the snapshot solution.

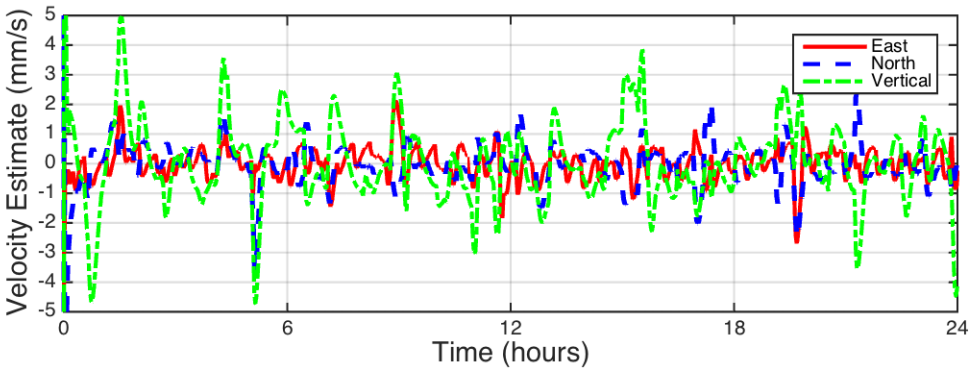

Figure 9-x Dual-frequency, Iono-free velocity errors for the Kalman filter solutions

9.6 Differential Positioning

In later chapters we will learn about the ground-based augmentation system (GBAS) and the satellite-based augmentation system (SBAS). Rather than performing stand alone positioning as described in this chapter, these systems utilize differential positioning. In differential positioning, there is one or more nearby reference stations that are affected by similar errors. These reference stations can either provide their measurements to the user or predict the user errors by exploiting the fact that they already know their true position. Let us first examine the case where a very close by reference station makes its own pseudorange measurements to the satellite and provides them to the user. The user can then subtract these from their own measurements and replace (9.18) with

$$\delta\rho^{(i)} = \rho_{user}^{(i)} - \rho_{ref}^{(i)} - \Delta b_0 \tag{9.x}$$

If the user and reference station are sufficiently close together, the ionospheric errors, tropospheric errors and satellite clock offsets will all effectively cancel. This $\boldsymbol{\delta\rho}$ placed into (9.22) will solve for the $\boldsymbol{\Delta x}$

position between the user and the reference station rather than the user's absolute position. However, if the reference station location is precisely known, the user can combine it with their position difference estimate and obtain their absolute position. Typically differential positioning has much better accuracy as it significantly reduces the largest error sources. However, the receiver noise and multipath from the reference station will affect the positioning accuracy and these error sources will not be cancelled for the user. Greta care must be taken to minimize these contributions from the reference receiver.

Usually it is more convenient to provide the user with estimated pseudorange errors rather than the full pseudorange measurement. The reference station subtracts its known range and estimated clock to provide pseudorange corrections that are added to the user's measurements to reduce the effects of ionospheric errors, tropospheric errors and satellite clock offsets. The corrections from a single reference station can be obtained from:

$$-\left(\rho_{ref}^{(i)} - r_{\text{ref}}^{(i)} - \hat{b}_{\text{ref}} + B_{ref}^{(i)}\right) \tag{9.x}$$

The satellite clock is removed to reduce the effects of any synchronization error between the reference station and the receiver. Sometimes ionospheric and tropospheric models are also included for both the user and the reference. These models are used to reduce the effects of spatial variation of those error sources. The corrections (9.x) are added to either (9.18) or (9.24), depending on whether the propagation models are used, to obtain the pseudorange residuals for (9.22). Providing corrections is better because less information needs to be transmitted to the user and it is easier to combine multiple reverence station corrections together. However both approaches are equivalent mathematically.

9.7 Positioning accuracy

The accuracy of the position solution is affected by the errors on the pseudorange residuals and the distribution of the satellites about the user. This latter effect is usually referred to as geometry. Reducing the magnitude of the pseudorange errors will improve the positioning accuracy. Similarly, having more ranging measurements or ranging measurements with a better distribution about the user also improves accuracy. If all of the satellites were clustered in one small part of the sky, the position accuracy will be worse than if there are satellites in each part of the sky. The effect of geometry is illustrated in Figure 9-8. We use a simple two-dimensional model that also assume synchronization of the satellite and the user (no need to estimate a clock offset). The range measurement indicates that the user is a fixed distance from the satellite. In two-dimensions, the possible set of user locations can be illustrated as a circle around the satellite. The dashed line indicates the range measurement and the fat gray circle indicates the measurement uncertainty. The position estimate lies at the intersection of the two circles. Note that there can be an ambiguity as there are two possible locations where the circles intersect. Similar ambiguity exists for the full four-dimensional GPS position solution and there does exist another potential solution out in space beyond the satellite orbits. This image solution is avoided by initiating the non-linear search, described in Section 9.4, on or near the Earth. Three example scenarios are depicted in Figure 9-8. In the first scenario, the satellites are close together. The position estimate is at the point of intersection and the darker shaded area, where the two gray circle overlap, represents the position uncertainty. The uncertainty is smallest in the direction of the lines of sight. However, in the direction perpendicular to the lines of sight, the position uncertainty is much greater. The middle scenario has the lines of sight orthogonal to each other, which results in the smallest overall position uncertainty. We can see that the possible error is well constrained in both dimensions. The last scenario has the satellites on opposite sides of the user. This is similar to the first scenario where the position solution is constrained in the direction of the lines of sight and is much larger perpendicular to them.

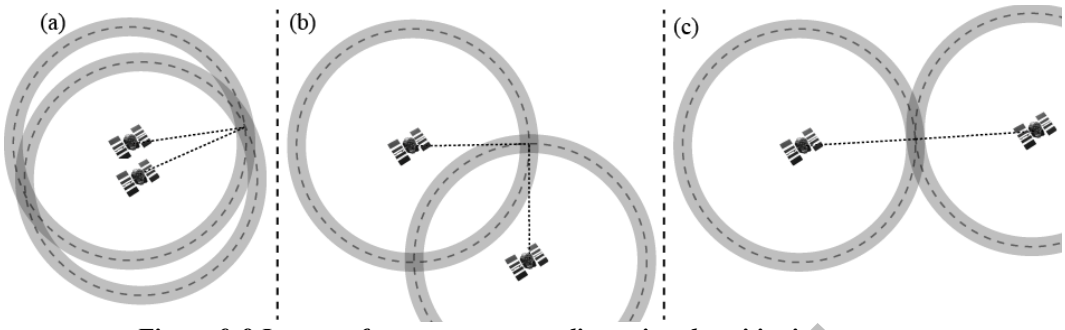

Figure 9-8 Impact of geometry on two-dimensional positioning accuracy

Good geometry is obtained when the lines of sight span all dimensions. When also estimating the clock offset, it is also better to span those dimensions from opposite directions. To see why, consider the first and last scenarios depicted in Figure 9-8. When ignoring the clock, as described above, they appear to create similar magnitudes of uncertainty. However, now imagine that there is also a clock state to be estimated. In the first scenario, the clock term would be difficult to resolve, as it would cause the radii of the two circles to grow or shrink by the same amount. This causes the region of uncertainty to now grow in the same direction as the lines of sight. In contrast, under the last scenario, growing or shrinking the radii does not cause more uncertainty in the direction of the lines of sight, but would rather move the position estimate in the perpendicular dimension. In reality, a third satellite is needed to resolve the position in two dimensions with an unknown clock state and therefore, this example is a bit of an oversimplification. Nevertheless, it remains true that having two satellites on opposite sides of the user results in better geometry than having the two satellites on the same side. The ideal situation is to be surrounded in all three dimensions. We will better quantify geometry below, but first we will look at the overall position accuracy before separating out the effect of geometry by itself.

The expected position error covariance matrix, **P**, for the solution obtained in (9.22) was given by (9.x). **P** is a four by four matrix whose diagonal elements describe the expected variance of the position and clock estimates. If the variances in **W** are most-likely pseudorange variance estimates, then the above covariance is also a most-likely estimate. However, if the pseudorange variances are safety-of-life overbounds, then the position estimate variances are also overbounding values that are likely much larger than the observed sample variances. The **G** matrix is expressed in a particular coordinate frame. It is most commonly either: an Earth-centered Earth-fixed (ECEF) frame aligned with the prime meridian and the North pole; or a local East, North, Up (ENU) frame specific to the users location. When looking at the user's accuracy we will use the local ENU frame.

The effects of geometry by itself are best seen by assuming that the weighting matrix can be replaced by the identity matrix, that is, that all satellites are equally accurate. We then obtain terms only related to the geometry that are referred to as dilution of precision (DOP). For positioning, we most frequently examine the horizontal, vertical and position DOPs (HDOP, VDOP, and PDOP respectively) as given by

$$\begin{aligned} \text{HDOP} &= \sqrt{\left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1}_{1,1} + \left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1}_{2,2}} \\ \text{VDOP} &= \sqrt{\left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1}_{3,3}} \\ \text{PDOP} &= \sqrt{\left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1}_{1,1} + \left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1}_{2,2} + \left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1}_{3,3}} \end{aligned} \tag{9.x}$$

Figure 9-9 shows coresponding DOP values for the data presented in Figures 9-5 and 9-6. It is obvious that the East and North dimensions have better geometry than the vertical dimension. This is because satellites may be to the East and the West as well as to the North and the South of the user, but they can only be observed above the user and not below the user due to the Earth blocking the signals. Further, because the inclination of the satellites only allows them to go between 55° S and 55° N, there are typically fewer satellites

to the North of a user in the Northern hemisphere. Thus, the user is more likely to have good geometry in the East-West dimesnion than the North-South (except near the equator). This too can be seen in Figure 9-9 which corresponds to a user located at 37° N. The larger errors in figure 9-6 often correspond to larger DOP values in Figure 9-9. The reverse is not always true as a large uncertainty does not guarantee a large error. Larger errors are more likley with larger uncertainty but not necessarily all the time. There is less correspondence with Figure 9-5 as the ionospheric errors are more significant and have a larger differences from satellite to satellite. This variation makes the equally weighted assumption of the DOP less representative of reality. The dual frequency case of Figure 9-6 still has differences from satellite to satellite, but they are not as significant as in the single frequency case.

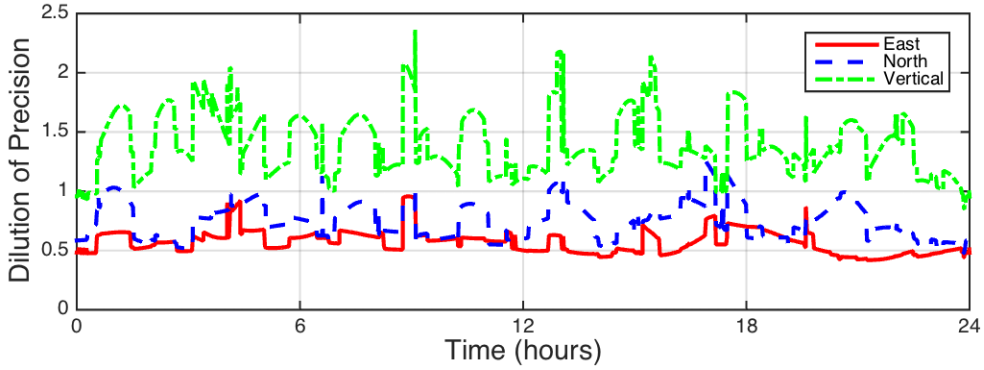

Figure 9-9 Single-constellation dilution of precision values

The previous figure was for the single constellation case. Adding a second constellation further improves the geometry as can be seen in Figure 9-10 which corresponds to the data in Figure 9-7. Because GLONASS has a higher inclination angle it does a better job of reducing the DOP in the North-South direction. Unfortunately, the DOP values overstate the expected improvement, because, as mentioned previously, the GLONASS ranging accuracy is worse. The equally weighted assumption used to create DOPs is a poor model in this instance. However, the DOP improvement does indicate the type of positioning accuracy improvement that would occur in the second constellaiton's ranging accuracy was comperable to the first.

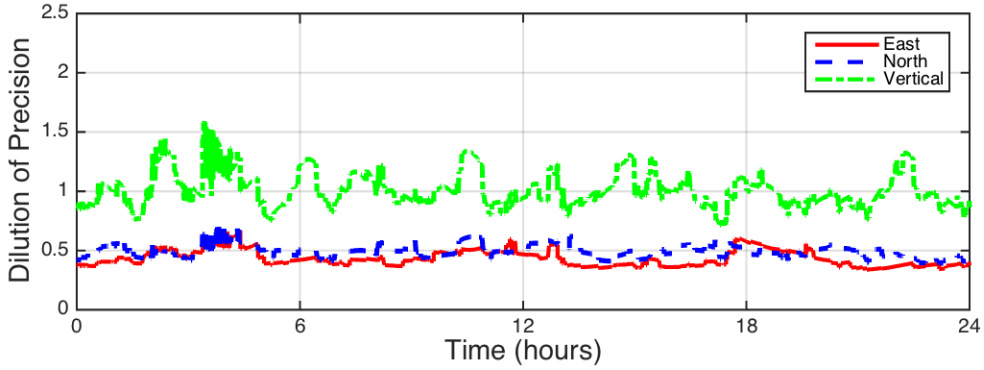

Figure 9-10 Dual-constellation dilution of precision values

The DOP is a dimensionless quantity. It can be multiplied by the mean expected ranging accuracy to obtain an approximation of the expected positioning accuracy. However, it is better correctly model the expected weights as is done in (9-x). Instead of isolating the geometry influence, this covariance estimate provides the combined influence of geometry and ranging uncertainty on the positioning accuracy. Figures 9-11 and 9-12 present these 1-σ confidences for the single and dual frequency cases respectively. Because the sigma values from Section 9.2.7 are very conservative, the predicted position accuracy is also very conservative relative to the actual errors observed in Figures 9-6 and 9-7. This positioning confidence term

serves as the basis for safety-of-life confidence bounds used by the augmentation systems as will be described in later chapters. The user does not know what their true error is, but by providing an upper bound on its possible value, the user will know whether it is safe to use for their intended operation.

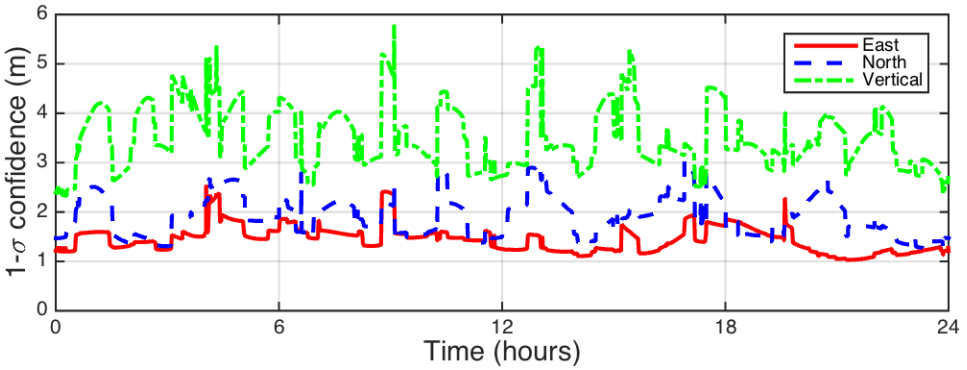

Figure 9-11 Single-constellation 1-σ position uncertainty

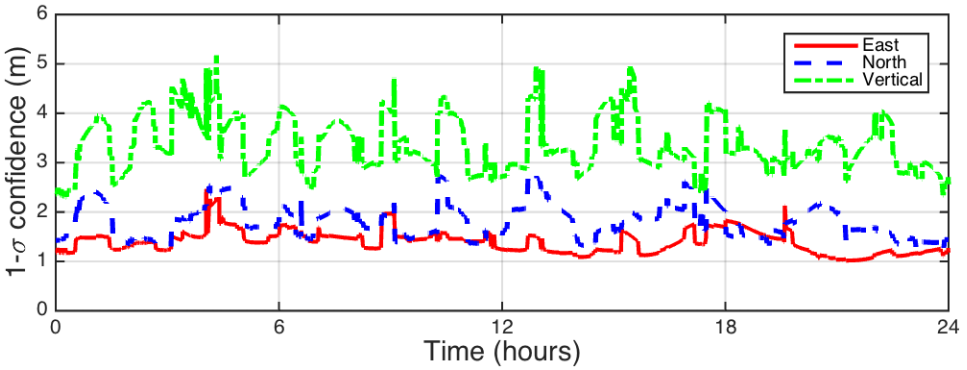

Figure 9-12 Dual-constellation 1-σ position uncertainty

plots including snapshot vs. Kalman

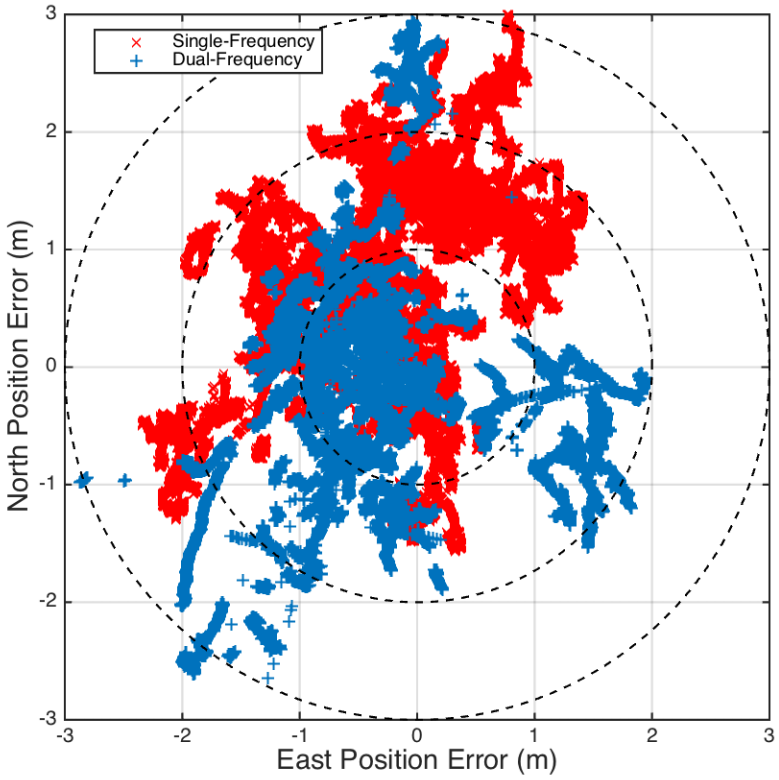

References

Misra, P. and Enge, P., **Global Positioning System: Signals, Measurements, and Performance,** 2nd Edition, Chapters 5 and 6, Ganga-Jumana Press, 2001.

Axelrad, P. and Brown, R. G., "GPS Navigation Algorithms," in **Global Positioning System: Theory and Applications,** Vol I, Parkinson, B., Spilker, J., Axelrad, P., and Enge, P., (eds), AIAA, 1996.

Klobuchar, J. A., "Ionospheric Effects on GPS," in **Global Positioning System: Theory and Applications,** Vol. I, Parkinson, B., Spilker, J., Axelrad, P., and Enge, P., (eds), AIAA, 1996.

Spilker, J. J., "Tropospheric Effects on GPS," in **Global Positioning System: Theory and Applications,** Vol. I, Parkinson, B., Spilker, J., Axelrad, P., and Enge, P., (eds), AIAA, 1996.

Braasch, M. S., "Multipath Effects," in **Global Positioning System: Theory and Applications,** Vol. I, Parkinson, B., Spilker, J., Axelrad, P., and Enge, P., (eds), AIAA, 1996.

Langley, R. B., "GPS System Receiver Noise," Innovations Column in GPS World, June, 1997 accessed via http://gauss.gge.unb.ca/papers.pdf/gpsworld.june97.pdf

Brown, R. G. and Hwang, P. Y. C., **Introduction to Random Signals and Applied Kalman Filtering**, 4th Edition, John Wiley & Sons Inc., 2012.

Hatch, R., "The Synergism of GPS Code and Carrier Measurements," in Proceeding of 3rd International Symposium on Satellite Doppler Positioning, Las Cruces, NM, 1982.

GPS Directorate, “Navstar GPS Space Segment/Navigation User Interfaces,” IS-GPS-200H, September, 2013. Available at http://www.gps.gov/technical/icwg/IS-GPS-200H.pdf

GPS Directorate, “Navstar GPS Space Segment/ User Segment L5 Interfaces,” IS-GPS-705D, September, 2013. Available at http://www.gps.gov/technical/icwg/IS-GPS-705D.pdf

RTCA, “Minimum Operational Performance Standards for Global Positioning System/ Satellite-Based Augmentation System Airborne Equipment,” DO-229D, Change 1, December, 2006.