GPS Solut
DOI 10.1007/s10291-012-0286-4

ORIGINAL ARTICLE

Satellite clock bias estimation for iGPS

John Pratt · Penina Axelrad · Kristine M. Larson · Bruno Lesage · Richard Gerren · Nicholas DiOrio

Received: 30 May 2012 / Accepted: 22 August 2012

Abstract The High Integrity GPS program seeks to provide enhanced navigation performance by combining conventional GPS with a communications and ranging broadcast from the Iridium® Communications System. Through clock and message aiding, it would enable existing GPS receivers to acquire and track in more challenging environments. As is the case for standard GPS, accurate and precise timing is key to performance. An approach is presented for estimating the bias of each Iridium satellite clock using satellite-to-ground and satellite-to-satellite measurements. The satellite clock bias estimates are based on a Kalman filter that incorporates code-type observations from the measurements at 10 s intervals. Filter parameters are set based on the expected behavior of the clocks, allowing for discontinuous bias and frequency adjustments due to ground commands. Typical results show the current filter to be accurate to within 200 ns while always meeting the initial system specification of half a microsecond.

Keywords iGPS · Clock estimation · Kalman filtering · Clock ensemble

Introduction

The Office of Naval Research sponsored High Integrity GPS (iGPS) Technology Concept Demonstration Program is designed to improve the position, navigation, and timing performance for military GPS users by integrating the communications capability of the satellite network from Iridium Satellite LLC, hereafter referred to as Iridium. The Iridium constellation consists of 66 satellites. These satellites communicate with each other and the Iridium ground stations, or earth terminals, as well as users. With its network of satellites supplying coverage of the entire planet, Iridium provides global voice and data telecommunication services to both military and commercial customers with equipment and services targeting numerous markets such as maritime, aviation, defense/government, machine-to-machine communications, disaster response, and exploration/adventure (Foosa et al. 1998; Schuss et al. 1999).

The iGPS concept uses the Iridium communications capability to precisely transfer GPS time to properly equipped users in challenging environments such as natural and urban canyons, heavily wooded areas, and in the presence of intentional or unintentional interference. By establishing a robust means to provide this time to within 0.5 μs, the system will facilitate the acquisition of GPS and accelerate the time to first fix for properly authorized users in degraded environments. More information on using Iridium to augment GPS can be found in Joerger et al. (2009, 2010).

The Iridium satellites are in six orbital planes located in a low-Earth orbit altitude of ~780 km with a high inclination of ~86°. This leads to relatively short contact times with the ground, for 10 min or less, but higher received power signals than GPS. Each of the satellites is assigned a satellite vehicle (SV) number. The SV numbers are used to identify results shown in later sections.

The Iridium satellites use oven-controlled crystal oscillators onboard to generate the communications signals and maintain system time. Over short time intervals, these clocks are very stable, but over time spans larger than 100 s, the bias and drift of these clocks are less stable than

J. Pratt · P. Axelrad (✉) · K. M. Larson · B. Lesage · R. Gerren · N. DiOrio
Department of Aerospace Engineering Sciences, University of Colorado Boulder, 429 UCB, Boulder, CO 80309-0429, USA
e-mail: penina.axelrad@colorado.edu

Published online: 04 September 2012

the atomic clocks used by GPS satellites. The Iridium satellite clocks generally exhibit a flicker noise floor below 10^{-11} s/s from intervals of 0.1 to a few hundred seconds. At an interval of about 1,000 s and onward, the satellite clocks exhibit random walk behavior, still below 10^{-11} s/s (Fig. 1). Although this is the nominal behavior of the clocks, some do exhibit less stable behavior. Iridium issues commanded bias and frequency adjustments to each of the satellites at least twice per day to keep the clocks synchronized. Satellites with higher instabilities are updated more frequently, with no limit set by Iridium on how many times per day any satellite can be adjusted.

Figure 2 shows the general architecture of the iGPS system, which consists of reference stations and Iridium earth terminals that gather information from passing satellites and then relay that data to an operations center. To effectively utilize the Iridium constellation for ranging and augmentation of GPS, the position of the satellites must be known and the behavior of the satellite clocks must be estimated accurately and characterized with respect to GPS time. Each of the reference stations has a Rubidium clock calibrated to GPS time using an independent single-frequency L1 GPS receiver. Using the data collected from all the reference stations, the operations center determines the ephemeris and clock biases of the Iridium satellite constellation. The Iridium-augmented GPS reference stations are separate from the Iridium stations that are used for the standard constellation control and maintenance.

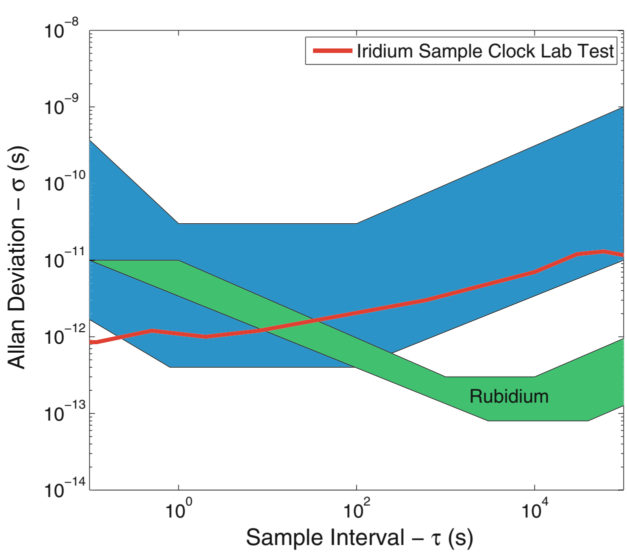

Fig. 1 Laboratory measurements of the Allan deviation for a sample clock expected to be representative of on-orbit performance of the Iridium satellite clocks. Sample clock measurements courtesy of Joseph White. Typical Allan deviation ranges of rubidium and quartz clocks are also shown for reference (Allan et al. 1997; Coates 2008; Vig 1992)

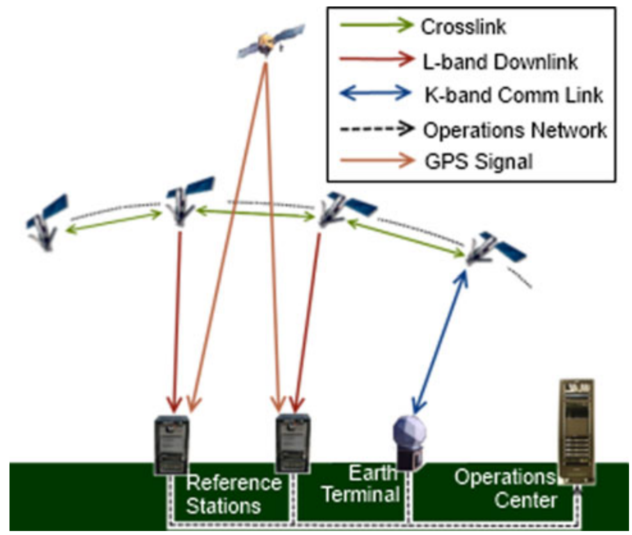

Fig. 2 iGPS elements and signals. iGPS uses several reference stations that receive and process signals from both the GPS and Iridium satellites. Each reference station uses GPS information to determine the bias in its reference clock. The central estimator at the operations center uses the Iridium downlink and crosslink signals to determine the Iridium satellite clock biases

We describe an approach to globally estimate the bias for each Iridium satellite clock using inter-satellite and satellite-to-ground ranging signals. The basic methodology is first presented, followed by some example results.

Measurement models and characteristics

There are two measurement types available to the Iridium global satellite clock estimator. The first type, which is similar to GPS pseudorange, is an L-band downlink ranging measurement between an Iridium satellite and a reference station. The second measurement type is a K-band ranging signal between neighboring Iridium satellites, referred to as a crosslink measurement. The basic structure of each is developed below.

In order to use the downlinks or crosslinks to determine clock bias, the locations of the Iridium satellites must be well known. The locations of the Iridium satellites used for this research were calculated from an ephemeris set generated by the Iridium operations group. These are currently accurate to better than 20 m. Also necessary for the use of the downlink measurements are the locations of the reference stations. These have been determined using GPS precise point positioning with an expected accuracy of better than 10 cm.

The currently implemented satellite to reference station downlink signal provides an unambiguous measurement of the satellite clock relative to the reference station clock with a noise level of about 600 ns. The bias and drift of the reference station clock are solved for separately using GPS signals. The uncertainty of these estimates is typically between 10 and 20 ns, much smaller than the Iridium satellite clock errors and measurement noise present in the typical downlink. The timing error in the downlink signal therefore is chiefly a result of the satellite clock bias and the measurement processes involved in the downlink itself, as seen below. The downlink range measurement, D, can be modeled as:

$$D = \frac{R}{c} + \text{tof} + B_{\text{ref}} - B_{\text{sat}} + \varepsilon_{\text{ref}} - \varepsilon_{\text{sat}} + \gamma + \eta, \quad (1)$$

where R is the geometric range, tof is the time-of-flight correction, B_{ref} is the reference station clock bias, B_{sat} is the satellite clock bias, ε_{ref} is the reference station receiver hardware bias, ε_{sat} is the satellite transmitter hardware bias, γ includes the ionospheric and tropospheric delays, and η represents all other error sources.

Using the satellite and reference station position information, the geometric range and time-of-flight correction are removed from the downlink signal. The reference station bias, as calculated via GPS, is also removed. Because of the relatively coarse ranging requirements, the tropospheric delay is simply modeled as a site-specific constant divided by the sine of the elevation angle of the satellite relative to the reference station. This should remove about 75 % of the associated error leaving residual troposphere errors of less than 2 ns. The ionosphere and remaining troposphere error, residual ephemeris error, multipath, and tracking noise are combined in the variable η' and are expected to be less than 50 ns. The downlink residual model is then given by:

$$\delta D = -B_{\text{sat}} + \varepsilon_{\text{ref}} - \varepsilon_{\text{sat}} + \eta' \quad (2)$$

A downlink pass lasts between 3 and 10 min with a data rate of one measurement per second. Figure 3 shows the residuals for a typical pass. At intervals of less than 10 min, the Iridium clocks have an expected time deviation (Riley 2008) of less than 1.1 ns. Thus, for the purpose of determining the Iridium satellite clock biases to better than 0.5 µs, the high-rate downlink range measurements can be reduced to a single bias and drift value for the pass, by simply applying a linear least squares fit to the downlink residuals. The bias value estimated at the midpoint of the pass is then input as a measurement to the global satellite clock estimator. The measurement uncertainty is determined from the covariance of the least squares fit.

The downlink example in Fig. 3 shows a pass of ~7.5 min. The original downlink observations have 1σ noise of 600 ns. For typical downlink passes, the resulting midpoint bias formal uncertainty is reduced to between 25 and 45 ns, depending on the length of the downlink pass. As a result of hardware biases and other unmodeled

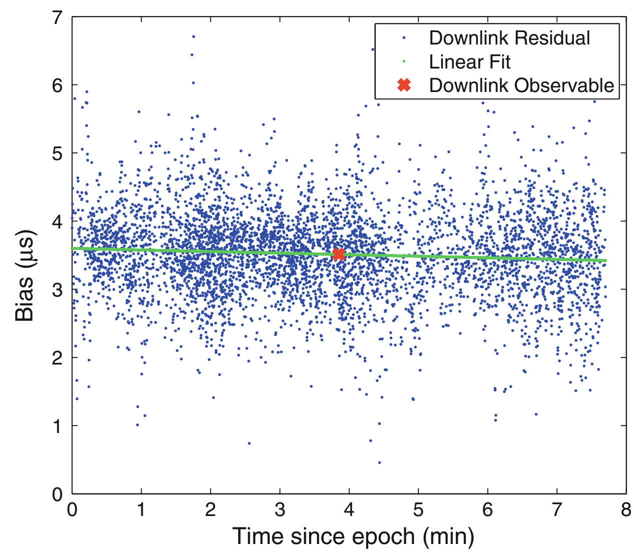

Fig. 3 Preprocessing of downlink measurements to create downlink observable. The individual downlink measurement residuals received by the reference stations currently have a 0.6 µs level of noise. Over short time intervals, the clock exhibits mostly linear behavior. By fitting a line and using a single point from the midpoint as a representative aggregate observable, the measurement noise input to the filter and computational load are reduced

systematic errors, the actual error in the downlink measurements is estimated to be ~100 ns. The results shown are typical of the downlink preprocessing.

The crosslink measurement set includes an unambiguous but noisy range measurement known as a variable receive offset (VRO) measurement, and a precise but ambiguous range measurement known as a unique word phase (UWP) measurement. As the current design favors simplicity and accuracy over precision, only the VRO measurements are currently being used. Each satellite has the potential to measure four crosslink ranges with satellites in front and behind it in the same orbital plane and with one satellite in each neighboring orbital plane. The specific cross-plane crosslinks vary depending on whether the satellite is ascending or descending in its orbit. The satellites in the first and sixth plane, which are moving in opposite directions, do not communicate with each other.

The in-plane crosslink range measurements have lower noise than the cross-plane crosslinks, but both are useful for connecting clock measurements across the satellite network. In general, the precision of the in-plane crosslink measurements is on the order of 2 ns. The overall accuracy of the measurements though is closer to 20 ns, again due to uncalibrated biases and unmodeled systematic effects (Fig. 4). The cross-plane crosslink observations have errors that are two to three times larger, depending on the satellites. Despite this, the cross-plane crosslinks are very important for keeping all satellites clock biases observable when an orbital plane lacks downlink measurements. It

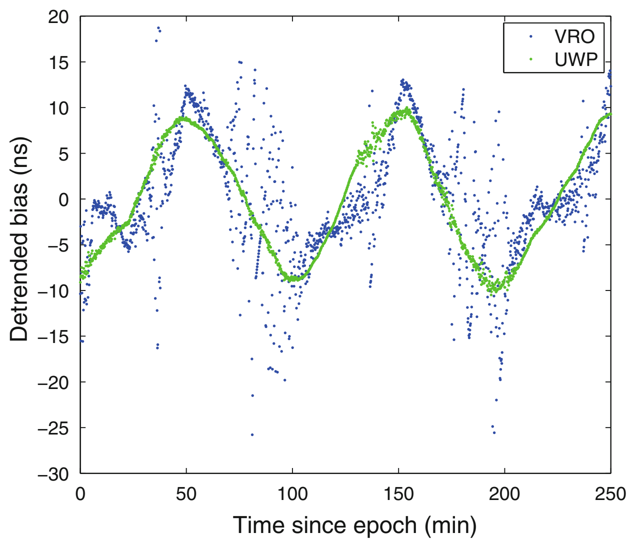

Fig. 4 Comparison of the noise in the in-plane crosslink VRO measurement, which is being used in the filter, and the more precise UWP measurement. There are systematic errors in the VRO measurement, especially when the satellite is in high-latitude regions

should be noted that cross-plane measurements are not made by the satellites when they are near the poles of the Earth, because of the potentially short ranges between neighboring satellites in the region where the orbital planes cross. The accuracy of the cross-plane measurements is also better near the equator.

The crosslink range, R, between the two satellites is measured by each of the satellites involved in the link. The one-way link from satellite A to satellite B is referred to as χ_{AB}. Differencing two one-way measurements for the same link removes the geometrical range and common error sources including the minimal atmospheric effects present at an altitude of 780 km. The two-way crosslink measurement is referred to as X_{AB}.

$$\chi_{AB} = \frac{R_{AB}}{c} + B_A - B_B + \varepsilon_{recA} - \varepsilon_{trB} + \eta \quad (3a)$$

$$\chi_{BA} = \frac{R_{BA}}{c} + B_B - B_A + \varepsilon_{recB} - \varepsilon_{trA} + \eta \quad (3b)$$

$$\begin{aligned} X_{AB} &= \chi_{AB} - \chi_{BA} \\ &= 2(B_A - B_B) + \varepsilon_{recA} - \varepsilon_{trB} - \varepsilon_{recB} + \varepsilon_{trA} + \eta \end{aligned} \quad (4)$$

The ranges from A to B, R_{AB}, and from B to A, R_{BA}, include the time-of-flight difference between the two ranges. Time-of-flight would not be removed through differencing, because the value is different for each of the two one-way links. Therefore, a time-of-flight-adjusted range is removed from each one-way link before differencing. The range adjustment is calculated based on ephemeris information. The two-way residual error due to inaccuracy in the ephemeris is negligible. The largest remaining errors in the two-way crosslink are the transmitter and receiver hardware biases (ε). Though the common-mode hardware bias is removed from all satellites, there is some variability between satellites that has not yet been extensively characterized. Based on the consistency of initial comparisons between crosslinks and downlinks, the mean hardware bias difference is estimated to be less than 100 ns.

Although the downlink measurements exhibit a larger overall noise than the crosslink measurements, they are necessary to tie the satellite clocks to a known reference. Without the downlink measurements, the satellite clocks are not fully observable and the bias estimates would drift. Thus, both downlink and crosslink ranges are needed for accurate and stable global estimates of the Iridium constellation clock biases. More precise, Iridium downlink measurements, similar to GPS carrier phase, are currently being tested for clock estimation, which show promise to improve future performance of the system once they become available.

In addition to the measurements being used in the filter, Iridium has provided two measurement types to assist in the validation of the filter clock estimates. First, Iridium has independent estimates of the satellite clock biases. These estimates are formed using a set of K-band feeder link measurements that are separate from the Iridium downlink measurements used by the filter. The two measurement systems have a known mean offset of 3.58 µs, which is removed in the subsequent comparisons. The Iridium estimates have a stated accuracy of 1 µs, though in practice the accuracy seems to be better. There is no set interval for these estimates, but they are generally made at least once every 2 h. The limited accuracy of the Iridium values means that they can only be used to verify the global clock estimates in a rough sense.

Iridium also provides the values of the commanded adjustments to the satellite clocks. These provide a very precise measurement of the change in clock bias at the time of clock adjustments and demonstrate how well the filter can respond to the clock maintenance. Both of the Iridium measurement types are marked on the plots shown in the results section.

Clock filters

To determine the bias of the Iridium satellite clocks based on downlink measurements, the biases in the reference stations relative to GPS time must be estimated. Each reference station clock is an atomic frequency standard, primarily rubidium but also may be a cesium or a hydrogen maser standard depending on local availability of the higher-stability source. The stability of the Iridium satellite clocks has already been described. The two clock sets will be estimated separately using different configurations as detailed below.

The ground station clock errors are determined based on point solutions from an independent GPS receiver component in the reference station. A Kalman filter is used to improve the estimates based on a classic model as described by Stein (1988, 1989). The filter consists of three states: bias (b), frequency offset (f), and frequency drift (d). A separate Kalman filter is run for each ground station so that the ground reference clock estimates are completely independent of each other.

$$X = \begin{bmatrix} b \\ f \\ d \end{bmatrix} \tag{5}$$

At each epoch, the filter estimates the state of the clock referenced to GPS time. The estimates are propagated from one epoch to the next using a simple linear state transition matrix in (6) and (7).

$$\varphi = \begin{bmatrix} 1 & \Delta t & \frac{\Delta t^2}{2} \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{bmatrix} \tag{6}$$

$$X(t) = \varphi(t - t_0)X(t_0) \tag{7}$$

The covariance of the state is also propagated using the standard Kalman filter equation in (8)

$$P(t) = \varphi(t - t_0)P(t_0)\varphi(t - t_0)^T + Q_d, \tag{8}$$

where Q_d represents the process noise in the dynamic model of the clock. The process noise is defined based on the spectral noise densities of the clocks for the ground station as defined below.

$$Q_\mathrm{d} = \begin{bmatrix} S_0\Delta t + \frac{S_2\Delta t^3}{3} + \frac{S_4\Delta t^5}{20} & \frac{S_2\Delta t^2}{2} + \frac{S_4\Delta t^4}{8} & \frac{S_4\Delta t^3}{6} \\ \frac{S_2\Delta t^2}{2} + \frac{S_4\Delta t^4}{8} & S_2\Delta t + \frac{S_4\Delta t^3}{3} & \frac{S_4\Delta t^2}{2} \\ \frac{S_4\Delta t^3}{6} & \frac{S_4\Delta t^2}{2} & S_4\Delta t \end{bmatrix} \tag{9}$$

The ground station clock filter is a typical filter used for clock applications and similar filters have been validated in several different papers including Senior et al. (2003). The satellite filter, detailed below, uses this standard model for a basis while modifying key parts to improve performance with the Iridium system.

To accommodate the frequent bias and frequency adjustments made by Iridium to the satellite clocks, the satellite clock filter is designed to estimate the bias only. This type of filter emphasizes stability and quick convergence over accuracy. Unlike the reference station clock filters, all the satellite clock biases are estimated together to take full advantage of the crosslink measurements.

$$X = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_{66} \end{bmatrix} \tag{10}$$

Based on the expected short-term stability of the satellite clocks, each clock state is assumed constant between measurement updates, which generally occur every 10 s. At this interval, the biases of the clocks show a variation of about 0.01 ns. Therefore, the constant assumption is reasonable. The covariance of the state, P, however, is still propagated using the formula in (8) using the identity matrix for φ. The process noise is defined in (11) with the spectral noise densities of all clocks assumed to be identical.

$$Q_\mathrm{d} = \begin{bmatrix} S_0\Delta t + \frac{S_2\Delta t^3}{3} + \frac{S_4\Delta t^5}{20} & 0 & \cdots & 0 \\ 0 & & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & S_0\Delta t + \frac{S_2\Delta t^3}{3} + \frac{S_4\Delta t^5}{20} \end{bmatrix} \tag{11}$$

The adaptation of the filter estimates to the clock adjustments is made by the manipulation of the state covariance matrix. After the filter is updated with a set of measurements, measurement residuals are compared with the residual or innovation variance. If the residuals are outside the 3σ bounds of the residual covariance matrix, it is assumed that a clock adjustment has occurred. In the case of a clock adjustment, the covariance of the adjusted clock is reset to an uncertainty of 3 µs and all the cross correlation terms are zeroed. The measurement set is then reprocessed. The large covariance on the adjusted clock allows it to quickly converge to an adjusted estimate within a few time steps. If the flagged adjustment turns out to be an anomalous measurement, the clock state returns to normal also within that time frame. Generally, anomalous measurements exhibit little effect as nearly all satellites have more than one crosslink and most have four.

Filter results

The satellite clock filter has been tested on approximately 10 days of data from the iGPS development testbed. Results are presented here for 25.5 h starting on December 16, 2011 at 20:49:19 UTC. This day is fairly representative of the other data sets. As there is no absolute truth reference, the validation was done by comparing the filter estimates with Iridium estimates and checking their consistency with commanded clock adjustments. The initial goal of the filter is to maintain an estimate of the clock bias with an accuracy of 0.5 µs. This goal has been set due to the current accuracy limits imposed by the length of the Iridium native TDMA downlink burst. Advanced signal designs, currently under test and evaluation, are expected to improve the accuracy in later implementations.

Figures 5 and 6 show example estimates of the clock bias for SVs 18 and 80. The SV 18 clock is typical of the Iridium constellation. SV 80 is a less stable, rapidly drifting clock that requires five bias adjustments over the 24 h period instead of two. The figures show the bias estimate with its 2σ uncertainty, the downlink measurements color-coded by reference station, the available Iridium estimates over the time period, and any bias adjustments that were performed during the measurement period. The bias adjustments are indicated by a vertical line with text specifying the magnitude.

Figures 7 through 10 provide additional information to illustrate the accuracy and precision of the satellite filter design. Figures 7 and 8 show a comparison of the filter estimates to the estimates provided for Iridium for both the clock bias and the adjustments made to the clock bias by the Iridium control station. Figure 9 presents the residuals of all downlink measurements from the sample data set. Finally, Fig. 10 shows the clock bias estimates for all of the satellites in orbital plane 6.

Overall, the estimates from the filter are consistent to within 0.5 µs of the Iridium estimates, which is below the Iridium estimate uncertainty. A comparison of the Iridium and iGPS estimates shows the differences have a variation of about 200 ns. This is true for all satellites and indicates that the filter is performing within the accuracy available from the Iridium estimates. A 200-ns positive bias in the residuals, which was removed from the data shown in

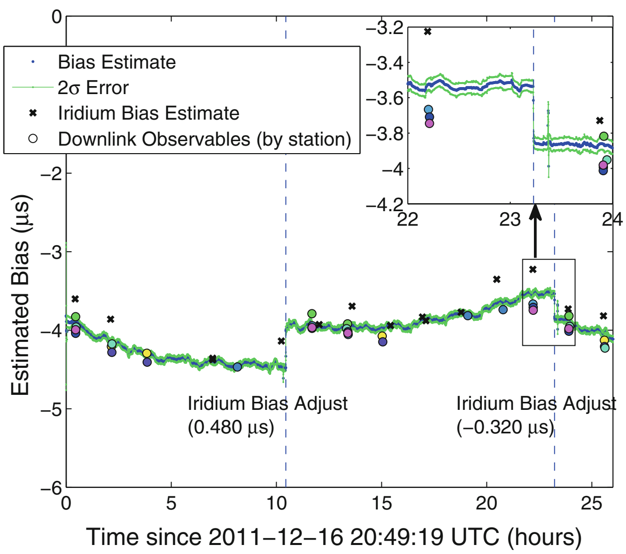

Fig. 5 SV 18 clock bias estimate with 2σ deviation, downlinks, and Iridium estimates marked

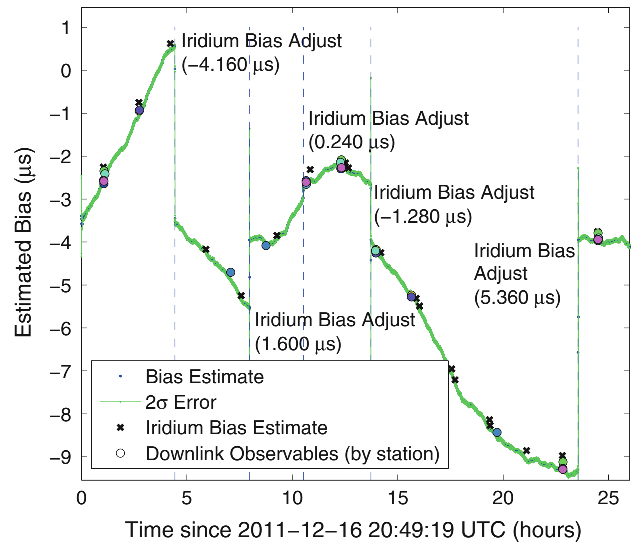

Fig. 6 SV 80 clock bias estimate with 2σ deviation, downlinks, and Iridium estimates marked

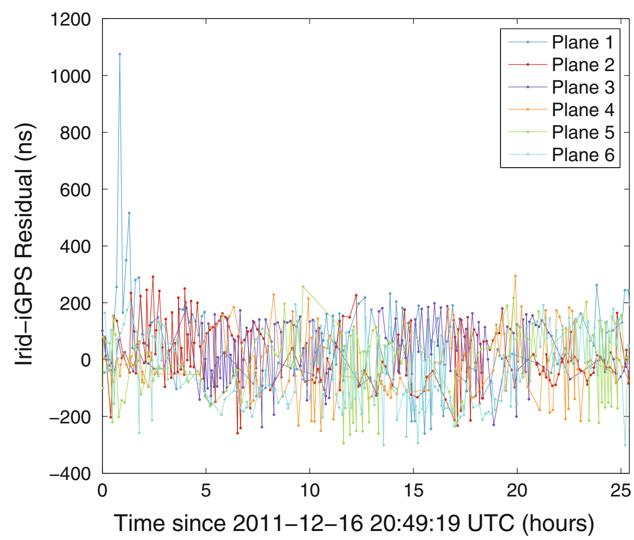

Fig. 7 Difference between the Iridium and new global clock estimates. Expected accuracy of the Iridium estimates is 1 µs, though in practice they appear to be more accurate. A 200-ns mean offset between the estimates has been removed to center the results

Fig. 7, indicates that there is an offset between the two systems beyond the 3.58 µs expected value given earlier.

A more powerful indication of accurate filter performance is the agreement between the commanded bias adjustments, provided by Iridium, and the change in the estimates given by the filter at the same time (Fig. 8). The 2σ variation in the difference between estimated bias adjustments and the commanded values is about 30 ns for the sample set. This indicates that the filter estimates have a very high internal accuracy. Any offsets seen in the system comparisons due to hardware biases do not appear to

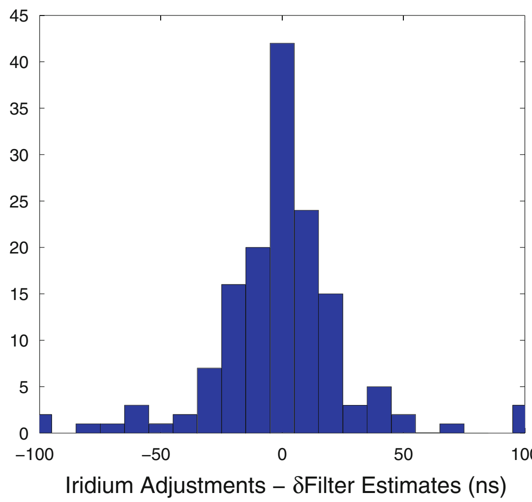

Fig. 8 Histogram of the difference between the commanded Iridium adjustments and the change in the filter estimates at the same time (for all 66 Iridium satellites). The plot is based on 148 clock adjustments that occurred within the data set. A total of five outliers, with values in excess of 100 ns occurred. They are represented by the bins at the far left and right edges of the plot

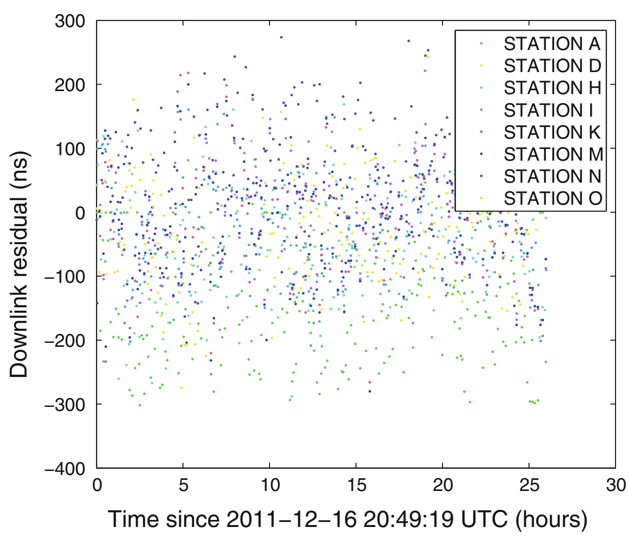

Fig. 9 Downlink residuals for the filter estimates from all 6 planes (colored by reference station). The 2σ deviation is about 210 ns. Each of the stations has a nonzero mean bias offset. These results are consistent with the expected downlink variances

change over time. Thus, once these hardware biases have been determined, it will be possible to reliably remove them from the measurements.

The residuals from the downlink measurements, shown in Fig. 9, have a 2σ deviation of about 210 ns. This is about what was expected, based on the calculations shown in the measurement section. There is an individual offset for each station from nominal due to hardware bias. The filter incorporates each downlink measurement with equal

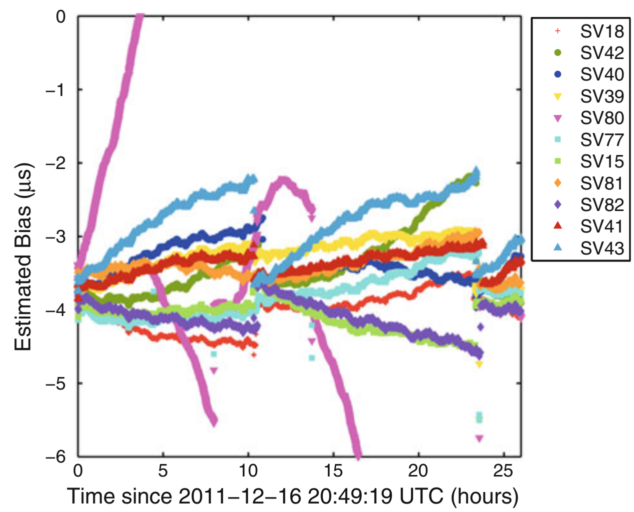

Fig. 10 Clock estimates for all satellites in orbital plane 6. The full range of the SV 80 satellite clock bias can be seen in Fig. 6

weighting so that the overall downlink residual mean is zero.

The time histories of the bias estimates for satellites in plane 6 (Fig. 10) show a high level of correlation between the biases. This correlated structure indicates that some systematic errors are present in the filter estimates, because it is not physically realistic for the clock biases to be related in this manner. A frequency analysis of the estimates identified a strong once-per-orbit signal present in the estimates of the clocks. Because these orbital variations occur in all clock errors simultaneously, it indicates that the problem is likely due to inconsistencies in specific cross-plane crosslinks, which then are quickly spread across all clock estimates.

A key strength of the current filter design is that it leverages both the downlinks, to keep the system observable, and the crosslinks, to frequently update the filter with high-precision measurements. The cross-plane crosslinks are of particular interest because of their ability to link clocks between neighboring orbital planes, thus extending the influence of a downlink measurement, and their potential negative impact of due to higher errors compared with the in-plane links. To study this, the filter was applied to the same data sets described above, with the cross-plane crosslink measurements removed. Figure 11 shows a comparison of the two results for SV 18. Figure 12 provides an expanded view of a 30-min portion of the comparison, around 22 h, where there is a notable difference between the two estimates.

Without the cross-plane data, the estimates of SV 18's clock bias are smoother, but have a larger uncertainty seen in the 2σ deviation. Furthermore, these estimates drift more between downlink updates, compared with the nominal case. Small jumps in the no-cross-plane estimates can be

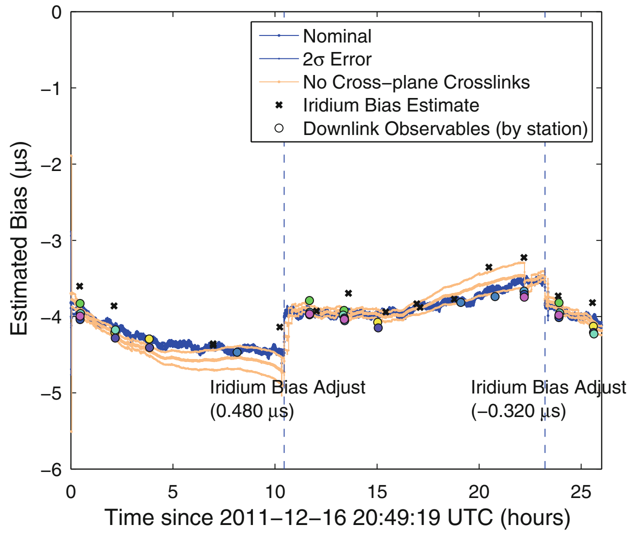

Fig. 11 SV 18 comparison of nominal filter to one without cross-plane crosslinks

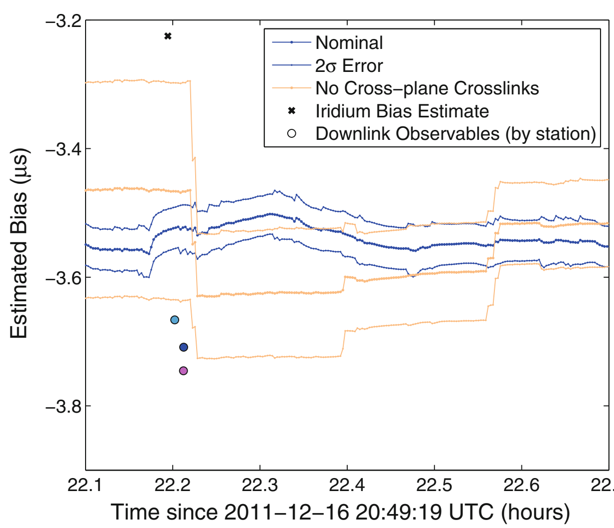

Fig. 12 Example of bias jumps in SV 18 comparison of nominal filter to one without cross-plane crosslinks

seen when downlink measurements from different satellites are incorporated by the filter. This is particularly noticeable in Fig. 12. In each case, the filter jumps in response to a downlink measurement bring the no-cross-plane estimates closer to the nominal filter estimates; thus, we can conclude that the nominal approach is more accurate.

Conclusions

An Iridium-augmented GPS system offers a unique opportunity for military users to access a full constellation of 66 satellites with high-power signals to assist in positioning via GPS. The fact that the Iridium constellation is a communication system and was not originally designed for navigation or time synchronization leads to interesting challenges in the implementation of an Iridium-based augmentation system. In particular, the lack of atomic clocks aboard the Iridium satellites necessitates real-time filtering to be able to provide accurate estimates of the clock bias for each of the satellites. The global clock estimator approach determines biases for all satellite clocks in the constellation concurrently, by incorporating downlink measurements with a 2σ error of 100 ns after preprocessing and crosslink measurements with noise levels of 20 ns (2σ) for the in-plane crosslinks and 40 ns (2σ) for the cross-plane crosslinks. It is designed to provide accurate estimates that are robust to large, discrete, commanded adjustments to the satellite clock bias and drift that can occur several times per day on each satellite.

The satellite clock filter has been shown to provide estimates that meet or exceed the initial design requirements. The measurement residuals and comparison with Iridium's independent estimates show the filter to be well within 0.5 µs for all the clock estimates. The 2σ accuracy of the filter estimates is currently about 200 ns, and the 2σ precision of the estimates is about 30 ns. Although the cross-plane crosslink measurements appear to add noise to the overall system, they are indispensable in keeping the estimates from drifting between downlinks. The systematic nature of the cross-plane noise indicates that there may be a way to reduce it in the future.

The next step in the development of the Iridium global clock estimator is to focus on improving the overall accuracy and short-term precision of the results. Downlink biases can be better estimated by increasing the number of ground stations which will allow for an accurate estimate of the bias of each ground station relative to the mean. The mean bias can then be calibrated through hardware measurements at one or two of the ground stations. Clock drifts can be better estimated by incorporating high-precision phase measurements from the crosslinks (UWPs) and the downlinks. Again, special attention is required to provide the best possible accuracy with rapid response to discrete adjustments of a satellite clock bias or frequency. Downlink phase measurements are currently being made and incorporated into the reference receiver which will support this effort. Once they are reliably established, improvements to the clock drift estimates can be integrated with the bias measurements to further improve the performance of the clock bias filter. Future efforts will assess the impact of the global clock estimator on Iridium-augmented GPS user performance.

Acknowledgments This work was sponsored by Coherent Navigation under a prime contract from Boeing awarded by the Naval Research Laboratory (NRL). The authors are grateful to William

Bencze, Clark Cohen, Isaac Miller, and Tom Holmes of Coherent Navigation; Misa Iovanov (Boeing) and Mark Nelson (Kinetx) supporting Iridium Communications; Brian Patti (via Boeing subcontract) and John Rice of Iridium Communications; Peter Fyfe, Don Tong, Rick Gerardi, and Phil Stranahan of Boeing Defense, Space & Security; and Joseph White, Ken Senior, and High Integrity GPS Program Manager Jay Oaks of NRL for their efforts in facilitating this work.

References

Allan DW, Ashby N, Hodge C (1997) The science of timekeeping. Hewlett Packard, Geneva (Appl. Note 1289, p 60)

Coates P (2008) Frequency measurement. In: Access science. Retrieved from http://www.accessscience.com/content.aspx?searchStr=Frequency+Measurement&id=272000

Foosa C, Raines R, Gunsch G, Temple M (1998) An overview of the Iridium low Earth orbit satellite system. In: Aerospace and electronics conference, pp 152–159

Joerger M, Neale J, Pervan B (2009) Iridium/GPS carrier phase positioning and fault detection over wide areas. In: Proceedings of the 22nd international technical meeting of the Satellite Division of the Institute of Navigation, pp 1371–1385

Joerger M, Gratton L, Pervan B, Cohen C (2010) Analysis of Iridium-augmented GPS for floating carrier phase positioning. Navigation 57(2):137–160

Riley WJ (2008) Handbook of frequency stability analysis. NIST 1065:26

Schuss J, Upton J, Myers B, Sikina T, Rohwer A, Makridakas P, Froncois R, Wardle L, Smith R (1999) The IRIDIUM main mission antenna concept. IEEE Trans Antennas Propag 47: 416–424

Senior K, Ray J, Koppang P (2003) Developing an IGS time scale. IEEE Trans Ultrason Ferroelectr Freq Control 50:585–593

Stein SR (1988) Kalman filter analysis for real-time applications of clocks and oscillators. In: 42nd annual frequency control symposium, pp 447–452

Stein SR (1989) Kalman filter analysis of precision clocks with real-time parameter estimation. In: 43rd annual frequency control, pp 232–236

Vig J (1992) Introduction to quartz frequency standards. IEEE Ultrason., Ferroelectr., and Freq. Control Society Tutor. Retrieved from http://www.ieee-uffc.org/frequency\control/teaching.asp?vig=vigcomp

Author Biographies

John Pratt is a PhD candidate in Aerospace Engineering Sciences at the University of Colorado Boulder. He previously received his BS and MS in Electrical Engineering from Utah State University.

Penina Axelrad is Professor and Chair of Aerospace Engineering Sciences at the University of Colorado Boulder. She has been involved in GPS-related research since 1986 and is a Fellow of the ION and the AIAA.

Kristine M. Larson is a Professor of Aerospace Engineering Sciences at the University of Colorado Boulder. She has been involved in GPS-related research for over 25 years and is an AGU Fellow.

Bruno Lesage graduated with an MS in Aerospace Engineering Sciences in May 2012 from the University of Colorado Boulder. He completed his BS in Aerospace Engineering at the University of Texas at Austin.

Richard Gerren received his BS, MS, and PhD from University of Colorado Boulder in Aerospace Engineering Sciences. He is currently a Professional Research Associate with the department.

Nicholas DiOrio is currently a graduate student at Colorado State University. In May 2012, he completed his BS in Aerospace Engineering Sciences at the University of Colorado Boulder.