

ORBIT PERIOD FREQUENCY VARIATIONS IN THE GPS SATELLITE CLOCKS

Everett R. Swift
Bruce. R. Hermann
Space and Surface Systems Division (K10)
Naval Surface Warfare Center
Dahlgren, Virginia 22448, USA

Abstract

A history of the GPS satellite clock behavior has been accumulated as a result of the weekly precise ephemerides produced at the Naval Surface Warfare Center (NSWC) under sponsorship of the Defense Mapping Agency (DMA). These ephemerides are produced using smoothed pseudorange data collected at a global set of ten tracking stations. The GPS satellite and station clocks are estimated simultaneously with the orbits. Time and frequency offset estimates are generated at one-hour intervals using a stochastic clock model. Studies using interferometric techniques to separate orbits and clocks have indicated that significant orbit period variations are present in the GPS satellite clocks. For PRN6/NAV3, which is currently operating on a rubidium frequency standard, these variations have amplitudes as large as 50 ns during the middle of eclipse season. For the four satellites operating on cesium frequency standards, the amplitudes of these variations are less than 15 ns. It is assumed that thermal cycling is the cause of these variations. The stochastic clock model has been tuned to allow the frequency offset state to track these variations. Starting with the first GPS week in 1988, this tuning has been used in NSWC's production processing. A brief description of the tuning experiments is given along with Allan variances computed for two satellite clocks based on the hourly estimates accumulated over a period of thirty weeks.

INTRODUCTION

The GPS precise ephemerides and clocks have been computed at NSWC since the beginning of 1986 using the OMNIS Multisatellite Filter/Smoothing (MSF/S) software system [1,2]. This system incorporates the square root information implementation of the Kalman filter and Rauch-Tung-Striebel fixed-interval smoother and processes data for all satellites and stations simultaneously. Current production processing consists of MSF/S fits using eight days of data for each week (the GPS week plus a half day on each end) so that adjacent fits overlap by one day. Smoothed ionospherically corrected pseudorange data collected at a global network of ten tracking stations are used in these fits. Five of the stations are the Air Force's Operational Control Segment monitoring stations in Colorado Springs, Ascension, Diego Garcia, Kwajalein, and Hawaii. Multiple STI receivers are deployed at each of these sites allowing all satellites in view to be tracked simultaneously. The other five stations are operated by DMA and are located in Australia, Argentina, England, Ecuador, and Bahrain. A TI4100 receiver

is deployed at each of these sites. Currently only one receiver is at each site so that only four satellites can be tracked simultaneously. Each station's receivers are driven by a Hewlett-Packard cesium frequency standard (HP 5061A-004).

Reference trajectories for all satellites are integrated using a truncated WGS 84 gravity field, point mass gravity fields for the Sun and Moon, solid earth tides, the ROCK4 radiation pressure model (with a y-axis acceleration), and a 5-minute integration step. The observations are processed in one-hour mini-batch intervals. This means that the Kalman filter process noise covariance matrix is only added in when propagating from one interval to the next and all observations within an interval are processed in one Kalman measurement update. An observation sigma of 75 cm is assigned to each smoothed pseudorange. The following parameters are estimated simultaneously for each weekly fit:

- For each satellite:
 - Orbital elements
 - Radiation pressure scale modeled stochastically
 - Y-axis acceleration modeled stochastically
 - Clock parameters (see next section)
- For each station except for the master station:
 - Clock parameters (see next section)
- Polar motion and rate of change of UT1-UTC modeled as random constants.

The WGS 84 station coordinates are held fixed in the estimation procedure. This procedure generates fitted trajectory files for each satellite and a clock file containing the time and frequency offsets between each satellite's clock and GPS time and frequency at one-hour intervals.

CLOCK MODEL

Nominal clock offsets for both the satellites and stations, which consist of second order polynomials in time with possibly step changes, are removed from the smoothed pseudorange data residuals before they are used in the orbit/clock estimation procedure. Corrections to this nominal for each satellite clock are modeled as the outputs of a third order linear system with white noise inputs. The three corrections are called frequency drift $\delta\ddot{r}$, frequency offset $\delta\dot{r}$, and time offset $\delta\tau$. The model in current state continuous variables is given by:

$$\begin{pmatrix} \delta\ddot{r} \\ \delta\dot{r} \\ \delta\tau \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \delta\ddot{r} \\ \delta\dot{r} \\ \delta\tau \end{pmatrix} + \begin{pmatrix} w_1 \\ w_2 \\ w_3 \end{pmatrix} \quad (1)$$

The model is actually implemented in discrete form using pseudoePOCH state variables as described in [2]. The frequency drift state is modeled as a random constant for the fit, i.e., no white noise drives this state so its spectral density, q_1 , is set to zero. The white noise, with spectral density q_2 , driving the frequency offset state defines the random walk frequency noise variations for the model. The white noise, with spectral density q_3 , driving the time offset state defines the white frequency noise variations. No explicit flicker noise parameter is present in this model. The Allan variance corresponding to this model is then given by

$$\sigma_y^2(\tau) = \frac{q_3}{\tau} + \frac{q_2\tau}{3} \quad (2)$$

where τ is the averaging time in seconds. Figure 1 shows this relationship in graphical form. The same model, but without the frequency drift state $\delta\ddot{r}$, is used for the station clock corrections. The clock for the master or reference station is constrained to be the assumed offset between this clock and GPS time. Therefore all clock solutions are estimated relative to GPS time. Using these models for the satellite and station clocks, the orbit, clock, and polar motion parameters are estimated simultaneously for all satellites using the MSF/S system.

CLOCK MODEL TUNING RESULTS

During each year there are two time spans when a given GPS satellite enters the Earth's shadow on every revolution. These eclipse seasons cause larger thermal variations within the spacecraft than for the rest of the year, and consequentially may cause larger clock frequency variations. Three weeks were selected for use in the tuning experiments based on which satellites were or were not experiencing eclipses. These were GPS weeks 403 (September 27-October 3, 1987), 413 (December 6-12, 1987), and 419 (January 17-23, 1988). PRN6/NAV3, PRN9/NAV6, and PRN12/NAV10 were in the middle of their eclipse season during week 403. The other four satellites were in the middle of their eclipse season during week 419. No satellites were being eclipsed during week 413. Weeks related to eclipse season were selected because previous work done by Aerospace Corporation, IBM, and NRL [3] has indicated that thermal cycling of the rubidium clocks on PRN6/NAV3 and PRN9/NAV6 was worst during eclipse season. All the other satellites except PRN8/NAV4 were operating on cesium frequency standards during these three weeks. PRN8/NAV4 was operating on a quartz crystal standard.

To tune the stochastic clock model a way of separating the clock and orbit estimates was needed. The interferometric processing approach was adopted to do this separation. By differencing two simultaneous pseudorange measurements for the same satellite from a pair of stations, the satellite clock effects are removed from the measurement. The globally distributed ten-station network provides almost continuous two-station tracking that allows this differencing for the entire orbit. Measurement differencing to remove satellite clock effects is equivalent to solving for independent time offsets for each satellite at each measurement time. To test this equivalency, single difference measurements using the smoothed pseudorange data for week 403 were derived and used for orbit determination. This differencing introduced measurement noise correlations and redundant measurements that were not properly handled in the orbit estimation procedure due to limitations of the software. For comparison, the MSF/S system was configured to solve for independent satellite time offsets for each measurement time along with the orbits. The orbits determined by these two approaches differed by less than 0.5 m root-mean-square (RMS) in each component. The second method is preferred because it fully accounts for the correlations and additionally provides satellite time offset estimates. These estimates do contain some very high frequency noise, which is not the actual clock behavior but due to aliasing of the measurement noise into the estimates.

Figure 2 contains a plot of the time offset corrections to the nominal for week 403 for PRN6/NAV3 derived using the second interferometric method. The aliased noise is present in the plot but is hard to see because of an orbit period, nearly sinusoidal variation of amplitude 50 ns. Figure 3 contains a plot of the time offset corrections for PRN3/NAV11 for week 413. The noise is now more obvious because the orbit period variation is less than 15 ns in amplitude. For the three weeks processed using this technique, the observed variations were worst during the mid-eclipse week for the two rubidium clocks (PRN6/NAV3 and PRN9/NAV6) with 35 ns amplitude worst case variations for PRN9/NAV6. The

cesium clock variations were always less than 15 ns in amplitude (worst case was for PRN3/NAV11 for week 413) and no correlation with eclipse season was observed.

Two strategies were tested for specifying the two white noise spectral densities available in the adopted model. These strategies were designed to allow the model to accommodate the orbit period clock variations with the one-hour mini-batch interval processing used in production. The interferometric results described above were considered "truth". A 12-hour period sinusoid of amplitude 10 ns results in a peak of at most 1.31×10^{-12} in the square root of the Allan variance at an averaging time of 6 hours. Both strategies were based on the assumption that the Allan variance of the model at 6 hours must be able to accommodate the variations observed using the interferometric technique. The first strategy involved increasing the white noise driving the time offset state (white frequency noise) much more than the white noise driving the frequency offset state (random walk frequency noise) so that the white frequency noise part of the Allan variance curve was above the required level at 6 hours. The second strategy involved increasing the white noise driving the frequency offset state and leaving the white noise driving the time offset state at a realistic level so that the random walk frequency part of the Allan variance curve was above the required level at 6 hours. This second approach is consistent with the assumption that the clock frequency is actually changing due to thermal effects and the pseudorange measurements include the integrated effect of these changes.

The RMS differences over all satellites between the orbits derived interferometrically and those derived using spectral densities based on the first strategy were 0.6 m in the radial direction, 1.2 m in the along-track direction, and 0.6 m in the cross-track direction. Using spectral densities based on the second strategy these were 0.2 m, 0.5 m, and 0.3 m respectively. The RMS clock differences were 3.0 ns for the first strategy and 2.0 ns for the second. Figure 4 contains a plot of the time offset corrections for PRN6/NAV3 for week 403 using white noise spectral densities based on the second strategy. Comparing this plot with the plot in Figure 2 indicates that the time offset estimates agree with the interferometric results but do not contain the high frequency noise. Figure 5 gives the corresponding frequency offset corrections. A sinusoidal frequency variation of amplitude 1.46×10^{-12} gives, when integrated, a sinusoidal time offset variation of amplitude 10 ns. Therefore integrating the frequency offsets in Figure 5 should approximate the time offsets in Figure 4. For the first strategy the peak-to-peak frequency offset variation was approximately 3×10^{-12} less than in Figure 5. Therefore the frequency offset state did not track the total variation. Figures 6 and 7 give the time and frequency offset corrections for PRN3/NAV11 for week 413 corresponding to Figure 3. The second strategy resulted in better agreement with the interferometric results and allowed the frequency offset state to track the assumed thermal effects.

ALLAN VARIANCE COMPUTATIONS

The white noise spectral densities for equation (2), determined using the second strategy above, were adopted for the MSF/S production runs beginning with the first full week of 1988 (GPS week 417). In a plot similar to Figure 1, the resulting Allan variance models for the satellites are shown in Figure 8. The procedures required to compute Allan variances from the clock solutions derived using these statistics are described in this section.

The Allan variance processing begins with the MSF/S satellite clock offset file which is the result of an eight day fit centered on each GPS week. The data consists of the mini-batch time of week, the time offset, and the frequency offset for all satellites. In order to use these files to compute Allan variances, the weekly files must be joined together and corrected for any clock jumps that might have occurred. This requires performing the steps listed below for each satellite.

1. Remove duplicate estimates of offsets at each end of week.

2. Eliminate all clock jumps from the offsets.
3. Compute fractional frequencies from each offset.
4. Fit and remove a first degree polynomial from the resulting fractional frequencies.
5. Compute and plot the Allan variances.

The 24-hour overlap region at the beginning/end of each week was removed by using a weighted average to form a smooth transition. As time progresses in the overlapping region, the sinusoidal weighting function gradually shifts from full weight for the first week to full weight for the second week.

The time of all clock jumps are listed in the clock file. This information is relied on to identify and eliminate each occurrence of a jump from the time or frequency offsets. The removal is performed in several steps. The first step consists of a least squares fit of a four coefficient function to the last 24 offsets before the jump. This function consists of a straight line plus a sinusoid. A second function of the same type is then used to fit the first 24 offsets after the jump. The two fitted functions are then evaluated at the midpoint of the interval in which the jump occurs. The difference between the two results at the midpoint is the jump amplitude and is added to all succeeding time offsets. Subsequent jumps from the same satellite in later weeks are handled in the same manner. Removing jumps by this method destroys the true offset magnitude, but this is of no consequence to the Allan variance computation.

Fractional frequencies can be computed from both the time offsets and the frequency offsets. In fact, the frequency offsets are fractional frequencies directly. The time offsets can be converted to fractional frequencies by computing the difference between two consecutive time offsets and dividing by the one-hour time interval between them.

The Allan variance is computed from the corrected fractional frequencies. Before beginning this computation, an independent least squares straight line fit is made to fractional frequencies from both the time offsets and frequency offsets to remove the frequency drift from the data. The Allan variance can then be computed using the residual fractional frequencies from the time offsets and/or from the frequency offsets.

Representative Allan variance plots are presented for the GPS weeks 417-446 (January 3 through July 30, 1988). A plot for a satellite with a rubidium frequency standard (PRN6/NAV3) is shown in Figure 9. The pronounced fluctuations seen in this Allan variance are due to the periodic oscillations that are present in the time offsets (Figure 4) and frequency offsets (Figure 5). A portion of Figure 9 is enlarged in Figure 10 to show this effect in detail. The first null is close to the orbital period of 43082 seconds. The true performance of this rubidium without the periodic component can be estimated by connecting the nulls in Figures 9 and 10.

An Allan variance plot for a satellite (PRN3/NAV11) with a cesium frequency standard is presented in Figure 11. The effects on the Allan variance of the periodic signals in Figures 6 and 7 is evident.

Allan variances computed for any satellite based on clock files from the MSF/S system derived using the model statistics shown in Figure 8 will always be below the curves drawn in this figure. In order for the clock solutions to represent the true performance of the frequency standards, the time and frequency offsets must not be artificially constrained by the model. The plots show that the computed Allan variances are well below the curves of Figure 8 and consequently are representative of the satellite clock performance.

CONCLUSIONS

Significant orbit period frequency variations are present for the GPS Block I satellites, especially for those operating on rubidium frequency standards. The results of a tuning experiment for the stochastic clock model used in the MSF/S orbit/clock determination system indicate that these variations can be accurately estimated. The adopted model statistics allow realistic Allan variances for the satellite frequency standards to be computed from these satellite clock estimates.

Tracking the performance of the GPS frequency standards is of interest to many users in the time transfer community. The satellite clock estimates derived using the MSF/S system provide a continuous history of satellite time and frequency offsets as a by-product of the precise ephemeris computations.

References

1. Swift, E. R., "NSWC's GPS Orbit/Clock Determination System," Proceedings of the First Symposium on Precise Positioning with the Global Positioning System, April 15-19, 1985, Rockville, Maryland.
2. Swift, E. R., Mathematical Description of the GPS Multisatellite Filter/Smoothing, NSWC TR87-187, Naval Surface Warfare Center, Dahlgren, Virginia, October, 1987.
3. Presentations at GPS Control Segment Performance Analysis Working Group Meetings, 1985-1987.

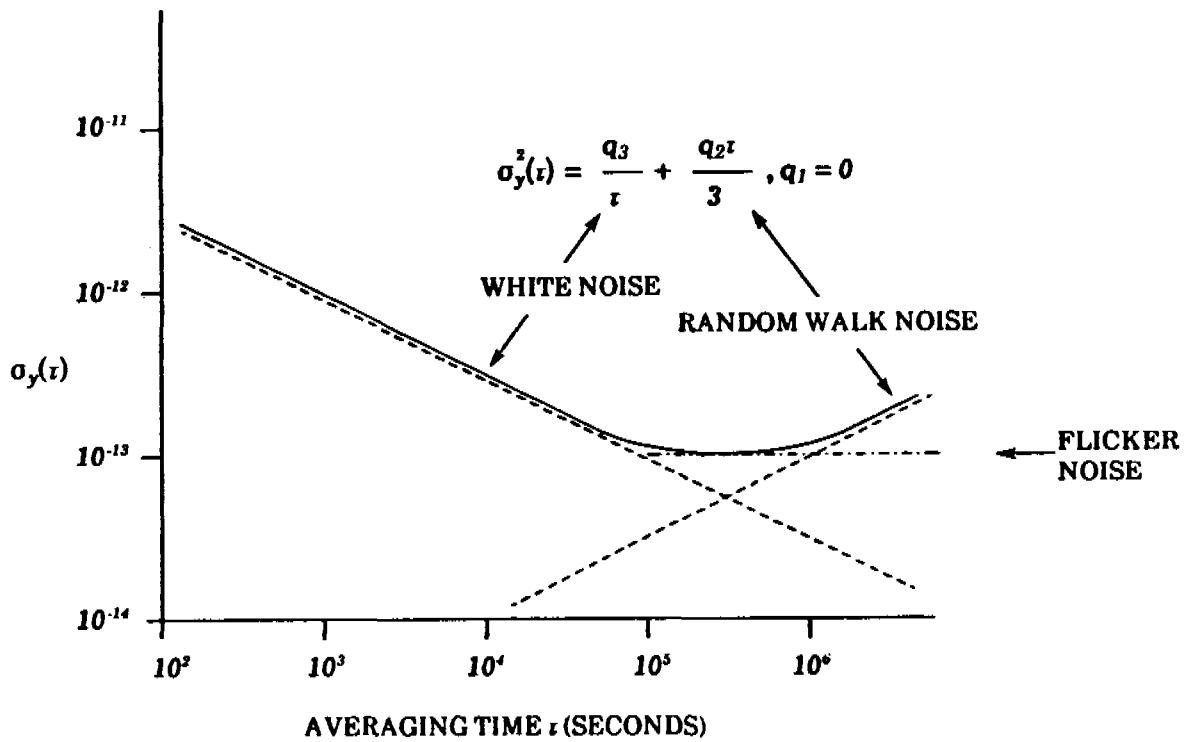


FIGURE 1. CORRESPONDENCE BETWEEN CLOCK MODEL SPECTRAL DENSITIES AND ALLAN VARIANCE

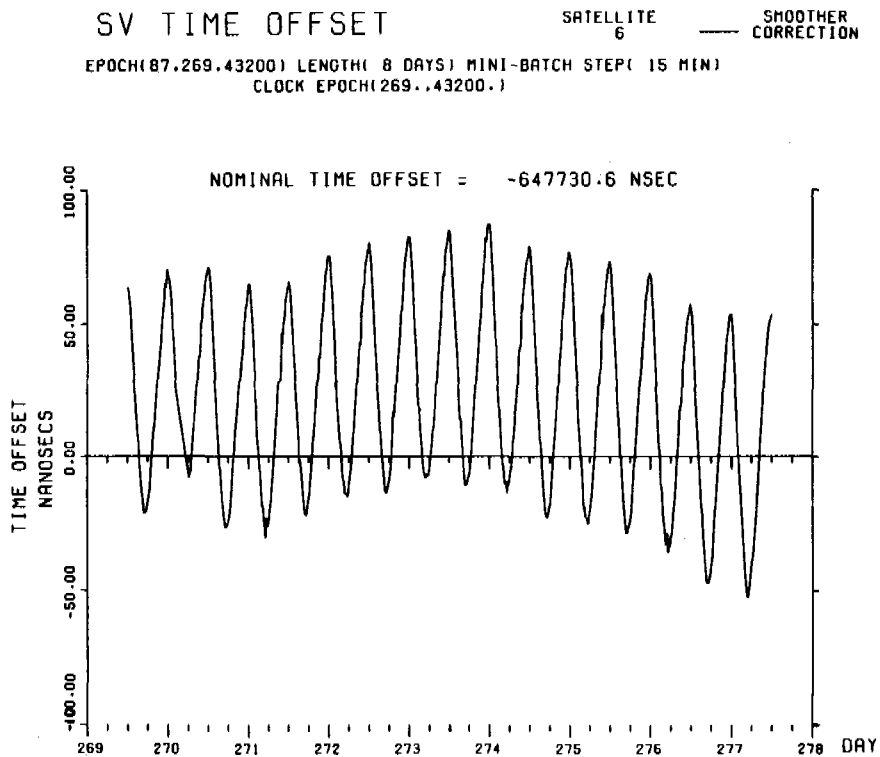


FIGURE 2. INTERFEROMETRIC TIME OFFSET CORRECTIONS FOR PRNG/NAV3 FOR WEEK 403

SV TIME OFFSET

SATELLITE 3 — SMOOTHER CORRECTION

EPOCH(87.339.43200) LENGTH(8 DAYS) MINI-BATCH STEP(15 MIN)
CLOCK EPOCH(339..43200.)

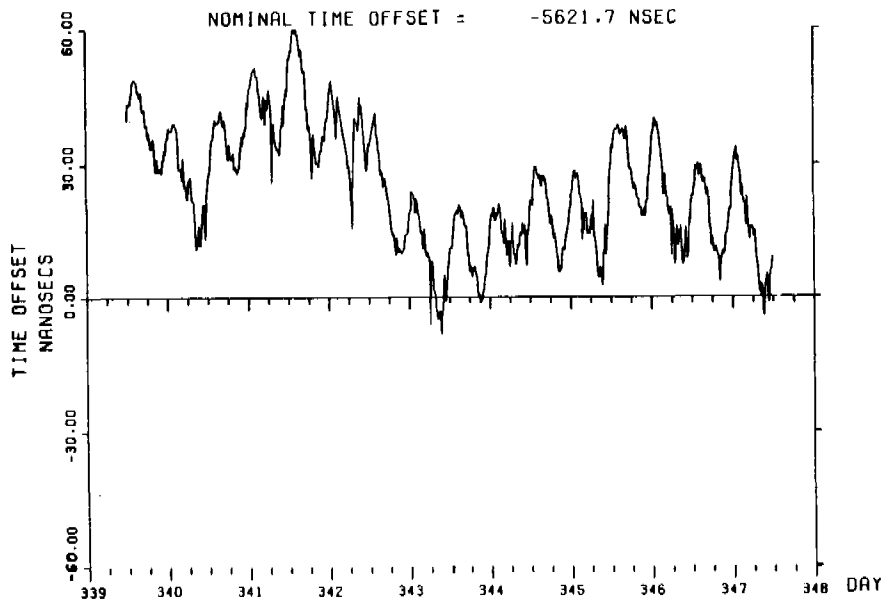


FIGURE 3. INTERFEROMETRIC TIME OFFSET CORRECTIONS FOR PRN3/NAV11 FOR WEEK 413

SV TIME OFFSET

SATELLITE 6 — SMOOTHER CORRECTION

EPOCH(87.269.43200) LENGTH(8 DAYS) MINI-BATCH STEP(60 MIN)
CLOCK EPOCH(269..43200.)

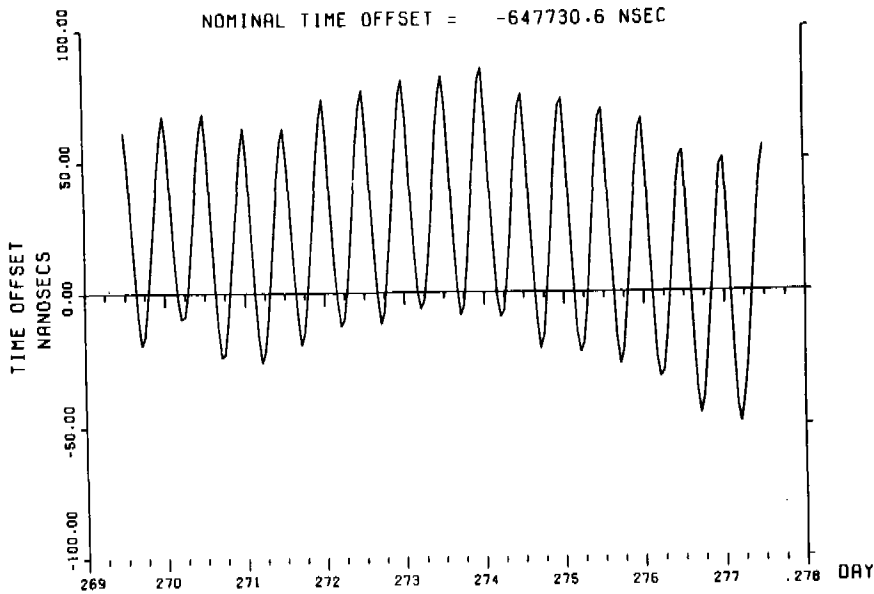


FIGURE 4. STRATEGY 2 TIME OFFSET CORRECTIONS FOR PRN6/NAV3 FOR WEEK 403

SV FREQ. OFFSET

SATELLITE
6

SMOOTHER
CORRECTION

EPOCH(87.269.43200) LENGTH(8 DAYS) MINI-BATCH STEP(60 MIN)
CLOCK EPOCH(269..43200.)




FIGURE 5. STRATEGY 2 FREQUENCY OFFSET CORRECTIONS
FOR PRN6/NAV3 FOR WEEK 403

SV TIME OFFSET

SATELLITE
3

SMOOTHER
CORRECTION

EPOCH(87.339.43200) LENGTH(8 DAYS) MINI BATCH STEP(60 MIN)
CLOCK EPOCH(339..43200.)

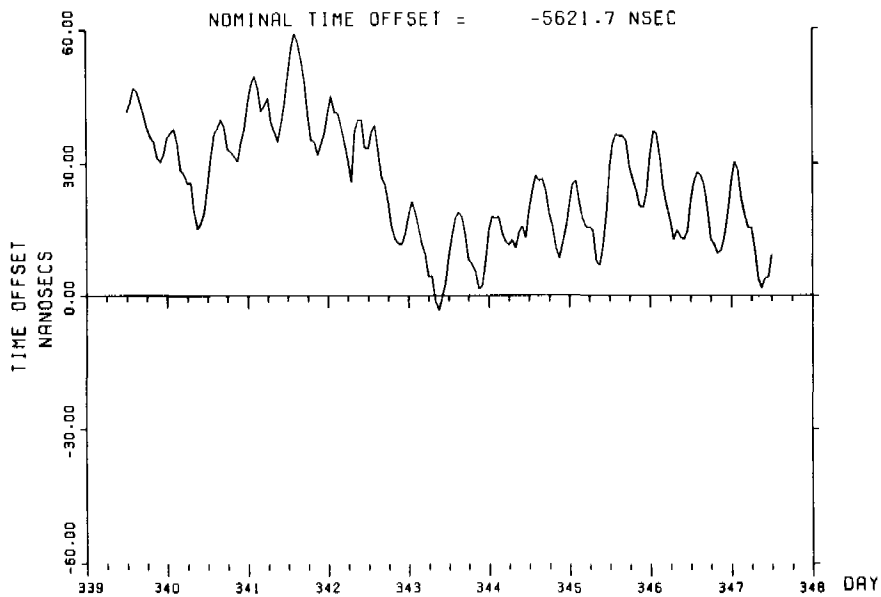


FIGURE 6. STRATEGY 2 TIME OFFSET CORRECTIONS
FOR PRN3/NAV11 FOR WEEK 413

SV FREQ. OFFSET

SATELLITE 3 — SMOOTHER CORRECTION

EPOCH(87.339.43200) LENGTH(8 DAYS) MINI BATCH STEP(60 MIN)
CLOCK EPOCH(339..43200.)

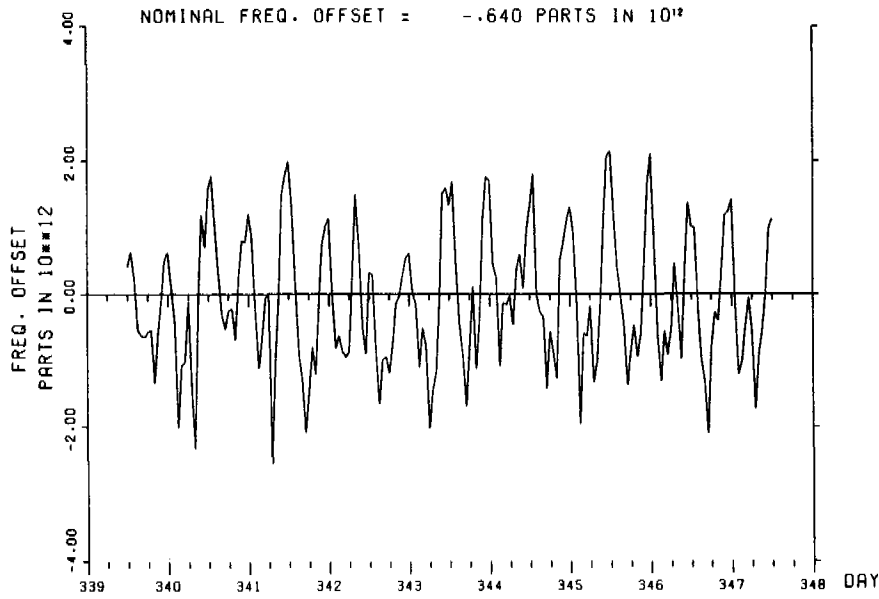


FIGURE 7. STRATEGY 2 FREQUENCY OFFSET CORRECTIONS FOR PRN3/NAV11 FOR WEEK 413




FIGURE 8. MODEL ALLAN VARIANCES FOR 1988

ALLAN VARIANCE PLOT
 FROM OMNIS TIME OFFSETS
 GPS WEEKS 417-446, PRN 6, NAV 3




FIGURE 9. ALLAN VARIANCE FOR PRN6/NAV3
 FOR WEEKS 417-446, 1988

ALLAN VARIANCE PLOT
 FROM OMNIS TIME OFFSETS
 GPS WEEKS 417-446, PRN 6, NAV 3




FIGURE 10. ENLARGEMENT OF A PORTION OF FIGURE 9

ALLAN VARIANCE PLOT
FROM OMNIS TIME OFFSETS
GPS WEEKS 417-446, PRN 3, NAV 11




FIGURE 11. ALLAN VARIANCE FOR PRN3/NAV11
FOR WEEKS 417-446, 1988

QUESTIONS AND ANSWERS

MARK WEISS, NIST: What is the difference between the broadcast ephemeris and your precise ephemeris in, say, meters?

MR. SWIFT: That is a difficult question. We believe that our ephemeris is good to 3-5 meters.

MR. WEISS: How far off is the broadcast ephemeris?

MR. SWIFT: I think that it could be off fifteen meters. The rubidium ones are probably worse than the cesium ones.

MR. WEISS: Is all the variation in the clock put into the ephemeris, or is some put into the broadcast ionospheric model as well?

MR. SWIFT: I don't think that any of it gets into the ionospheric model. It all goes into the ephemeris.

DAVID ALLAN, NIST: Did you say that for PRN3 this was worst case and that it wasn't during an eclipse?

MR. SWIFT: That's right, this was worst case of all the cesiums for those three weeks.

MR. ALLAN: Why would it not be worse during an eclipse?

MR. ALLAN: I don't know.

DR. GERNOT WINKLER, USNO: A sinusoidal variation in time can be interpreted as an integrated variation of frequency or as a phase modulation. If you interpret that as a phase modulation, which can come from any temperature sensitive point in the circuitry, the statistics would look quite different. Have you considered that possibility?

MR. SWIFT: I have been assuming, based on all the information that I have obtained, that the general consensus is that it is a frequency variation due to the temperature sensitivity of the clock.

JIM SEMLER, INTERSTATE ELECTRONICS: I will vouch for Everett's numbers for the precision ephemeris. We have been using it at Interstate pretty successfully for about two years now. I have question for you concerning the interferometric processing of those differences. Did you use your Kalman filtering for that? How did you process the data?

MR. SWIFT: Yes, we used the Kalman filter. We essentially gave it infinite Q's and processed it with 15 minute mini-batches, which is our measuring interval. That is how we implemented the independent estimate.