GPS CARRIER-PHASE FREQUENCY TRANSFER ON THE NIMA MONITOR STATION NETWORK

B. W. Tolman and D. Munton
Satellite Geophysics Laboratory
Applied Research Laboratories: The Univ. of Texas at Austin
POB 8029, Austin, TX 78713-8029, USA

Abstract

The National Imagery and Mapping Agency (NIMA) maintains a worldwide GPS Monitor Station Network (MSN). We are in the process of studying carrier-phase frequency transfer as one method for understanding the overall network performance, and for potential use in testing future GPS receivers.

We have developed processing software that uses 30-second dual frequency data in a continuous sequential estimator, which allows unlimited forward processing as well as optional backwards (fixed-interval) smoothing. States in the estimation include tropospheric and ionospheric delays, corrections (assumed small) to the station positions, and double- differenced carrier-phase biases. Satellite positions are taken from either precise or broadcast ephemerides.

We have tested our process using both zero-baseline datasets and benchtop tests. Our initial results indicate that instabilities seen in the estimated time offsets are dominated by temperature variations coupling to the GPS receiver used.

INTRODUCTION

Over the last several years, GPS carrier-phase time and frequency transfer has been studied as a possible alternative to the standard common-view time transfer techniques [1-3]. The latter technique relies on the use of pseudorange data to GPS satellites in common view to produce estimates of clock offsets. The advantage of using the carrier phase lies in the reduced noise in the phase measurements when compared to the pseudorange measurements (a reduction of approximately 100-fold). The increased precision is at expense of introducing computational complexity, since the GPS carrier phase in known only modulo 2π, and the unknown integer ambiguities must be computed. In addition, there are a number of issues associated with stability of hardware delays which make full use of the GPS carrier phase for time transfer a challenging endeavor [4-9].

Our work has three main motivations: the desire to understand what new light examining the GPS carrier phase might shed on the behavior of an operational system (at a time resolution below that which is routinely examined), understanding and quantifying the effects hardware delays might have on measurements, and our desire to precisely monitor the behavior of a network of remotely distributed frequency standards.

THE NIMA GPS NETWORK

The National Imagery and Mapping Agency (NIMA) operates a worldwide network of Global Positioning System (GPS) Monitor Stations. These stations are located around the world, including sites in: Alaska, USNO, Washington, D.C., Ecuador, Argentina, Bahrain, England, Australia, South Africa, Tahiti, and Korea. Development stations (fully operational sites in the core configuration) are located in St. Louis, MO, and Austin, TX The locations of the stations are shown in Figure 1.

Applied Research Laboratories of the University of Texas at Austin (ARL:UT) developed the current station configurations and operational software, and provides lifecycle development and engineering for the NIMA MSN.

MSN STATION CONFIGURATION AND CHARACTERISTICS

There are currently two configurations of the stations: a fully redundant configuration (termed a "core" station), and an augmentation configuration, which trades redundancy for overall cost. Each core station is equipped with two Ashtech Z(Y)-12 GPS dual-frequency receivers, and all sites (with two minor exceptions) employ dual HP (now Agilent) 5071A cesium (Cs) frequency standards (with the high-performance tube option). The exceptions are that installation at USNO utilizes a hydrogen maser, steered to the USNO ensemble, and St. Louis employs a single common Cs standard. The augmentation stations are identically configured, but without the redundancy.

Each station is equipped with an Ashtech antenna assembly, including a Dorne-Margolin high-precision antenna, choke ring, and radome. Antenna cables used are RG-214 and RG-8A, except at USNO, where a phase-stabilized cable is used. Antenna cable runs at the sites range from approximately 30 to 60 m. The antenna signal is split at the top of the equipment rack, using a Minicircuits ZAPD-2 power splitter, with a DC block on each branch. A bias T is inserted on one branch to provide power to the antennas. Temperature, pressure, and humidity measurements are collected at each antenna.

Environmental control for the equipment at each site is dependent on the local facilities housing the station. Overall temperature stability, for example, ranges from very good, at USNO, to more widely varying at other sites. Daily temperature variations of 2° C are typical (inside the rack). In cases of power failures, variations greater than 5° C over several hours have been observed.

Operationally, each receiver continuously collects 1.5 sec and 30 sec GPS dual-frequency pseudorange and carrier-phase measurements. The 1.5 sec measurements are used to produce a 15-minute smoothed pseudorange product. The 1.5 s data from the core stations flow back to St. Louis in near real-time. The 30 s data are archived at each site and retrieved on occasion. Overall data collection statistics show that the core stations routinely collect greater than 95% (usually 98%+) of all possible smoothed points. The robustness of this system, the routine high availability of geodetic quality dual-frequency pseudorange and phase data, and the geographic dispersion of the stations make the data collected ideal for testing new analysis tools.

Parenthetically, we note that the Ashtech receivers discussed here are distinct from the Ashtech Z12T receivers that have been studied in the timing literature. The Z(Y)-12's used here predate those receivers, and in addition have been augmented with security modules.

CARRIER-PHASE PROCESSING

The software process that we have developed uses dual-frequency GPS carrier-phase and code data to estimate the clock offsets, ionospheric and tropospheric corrections, and several other quantities. The data processing consists of a preprocessing step, which is applied to each receiver's data independently, followed by a double-difference estimation processor that combines the data from two or more receivers to produce a network solution. Finally, the relative clock bias between different receivers is constructed from the single-difference equation using all the final estimated quantities, along with an (arbitrary) initial value. When processing a large dataset in separate blocks, this single arbitrary bias in the final clock offset may be preserved between blocks, as long as there is physical continuity in the data stream and the clocks themselves.

The preprocessing program does several things: it edits anomalous data, finds and marks cycle slips and clock jumps, and computes an autonomous position and receiver clock solution from the pseudoranges. Data for satellites below 10° elevation are excluded at this point. Finally, the preprocessor applies the computed receiver clock offset to a polynomial fit of the carrier phase to synchronize all the data with GPS time. This eliminates the need for "quality checking" of the data within the estimation processor, and allows things like initial estimates of the phase ambiguities and selection of the reference satellite to be conducted outside the estimator.

The estimation processor is a sequential (Kalman) filter with optional fixed-interval smoothing, implemented using the square root information filter (SRIF) formulation. Our method is loosely based on the pioneering work of Blewitt **0**. The states to be estimated consist of receiver position coordinates, tropospheric and ionospheric zenith delays, and phase ambiguities. The estimator is linear, except that the position states are linearized. The estimator may be used to compute either a single baseline or a network solution, using datasets from any number of sites; the position of any or all of the receivers may be fixed via an input option.

The troposphere is modeled as a zenith value times a mapping function; any one of several different tropospheric mapping functions may be chosen. An input option allows the user to estimate only relative troposphere or ionosphere (fixing the values at one receiver), or to omit these from the state altogether (e.g. for work with short baselines). The position coordinates and phase ambiguities are modeled as random constants, and given very large process noises. The tropospheric and ionospheric zenith delays are modeled as exponentially time-correlated processes, with time constants of several hours. Temperature, pressure, and humidity means (differing by site) are used to compute an initial zenith value.

The measurement data consist of double differences of both the wide lane and ionosphere-free linear combinations of the dual frequency phases, as well as the wide lane pseudorange minus phase. As an option, the geometry-free linear combination may also be used. The measurement noise covariance, for phases and pseudoranges, is constructed from the measured signal to noise ratio in the standard way, using the tracking loop bandwidth. Typically, this gives the ratio between code and carrier uncertainties of 10^5; this ratio can also be specified and fixed by an input option.

There are different strategies for resolving the phase ambiguities. The wide lane ambiguities are resolved first, and rather easily **0**, because of the long wavelength and relatively low measurement noise of this linear combination of the phase. The remaining ambiguities (actually the L1 integer bias) may be resolved either manually, by inspection of the post-fit residuals, or automatically within the estimator via a search algorithm.

ZERO-BASELINE RESULTS

Our initial analysis focuses on St. Louis and USNO, because each of these stations uses only one frequency standard to drive both receivers, resulting in a less complex system to work with. In addition, daily temperature variations at USNO are significantly smaller than those at other sites. Previous analysis involving the NIMA receivers at USNO **0**, suggested a long-term variation between the two receivers, possibly due to cable multipath. We were specifically interested in examining the origin of this effect.

The data used for this study consisted of 30-second GPS data in Ashtech binary file format. The actual data rate is not critical and the 30-second data were chosen as the best compromise between data rates (1.5 second data were also available) and data volume/processing time. Similarly, the format of the data used is not critical, but was chosen as a matter of convenience.

Figures 2 and 3 provide a comparison of the phase offset between the two receivers at USNO and at the St. Louis site, computed from data collected from these sites. The data for USNO cover a period from day 138 to 178 of 2001 (MJD 52047-52087), with one data gap of approximately one-quarter day, while the data for St. Louis cover days 151 – 165 of 2001 (MJD 52060-52074).

The two clearly exhibit different behavior. Daily variations in St. Louis are on the order of 1 ns, while variations at USNO are of order 200 ps, with smaller variations visible. The difference in the behavior is at least partly explained by the fact the St. Louis station (which serves as a test and training station) is located in an area with substantial foot traffic, and does not have the same level of temperature control as USNO.

Figures 4 shows the time deviation (TDEV) for the USNO and St. Louis data sets, with 95% confidence error bars. For averaging times greater than approximately 900 seconds, the USNO results show evidence of white FM noise. However, the St. Louis TDEV indicates that flicker FM is present, possibly due to the environmental variations.

Our belief that environmental variations are the cause of this behavior is based on experience with the MSN sites. In order to ensure that environmental variations are the root cause, and in order to establish which equipment is most vulnerable to these variations, we performed a series of benchtop tests that are now described.

BENCHTOP TESTS

A zero baseline test was performed at ambient temperature in an assembly room, using the same hardware as a core MSN station, including two Ashtech Z(Y)-12 receivers and an HP5071 frequency standard. Room temperature was recorded at 5-minute intervals on a HOBO environmental probe, which has a temperature resolution of about 0.4° C. GPS data were recorded at 30-second intervals. A Minicircuits ZAPD-2 power divider (with isolation 28 dB at L1 and 40 dB at L2) was used to divide the GPS signal approximately 1 meter from the receivers. Data were collected for approximately 5 days, with one break after approximately 2 days. At this point, the power splitter used was swapped for a GPS Networking power splitter, which provides approximately 20 dB of additional isolation.

The results of this test are shown in Figures 5-7. Figure 5 shows the phase and temperature recorded over the 5-day test. There is a clear correlation between temperature and phase, which persists regardless of the power splitter used. Figure 6 shows the TDEV plot for these data and St. Louis for comparison. The results of this test are consistent with the results from St. Louis, and suggest that the power splitter plays little role in the results. Separate analysis of the data taken with each splitter shows similar behavior, although at short time intervals the higher isolation splitter seems to yield slightly better TDEV values. Longer tests would be required to be decisive in whether there is a small difference between the two.

Figure 7 shows a plot of phase against temperature for the data shown in Figure 5. Since the temperature is recorded only once every 5 minutes, each phase value has been assigned the temperature of the closest measurement. A linear relationship is suggested, with a coefficient of ~ 750 ps/C.

We also conducted a second test, placing both receivers into a temperature control chamber. The power splitter was left outside the temperature control chamber. Temperature was recorded both inside the temperature chamber and at the location of the power splitter. Three temperature changes were introduced during this test: one on starting the measurements, one after 2 elapsed days, and the final after 3 elapsed days. The results from this test are shown in Figure 8. In all cases of temperature change, it is clear that the system responds, due to changing path delays. Notable is the fact that the final temperature change does not induce a phase change of the size that might be expected from the first two. No correlation between the external temperature measurements and the computed phase difference was found.

Figure 9 shows the TDEV for this test (computed from periods when the temperature was constant) and what is computed from the latter half of the period of the USNO data. The similar structure of these plots suggests that removing rack temperature variations significantly improves the observed behavior of the system.

FUTURE WORK

Our preliminary work has been limited in scope in order to concentrate on the effects that might be reproducible in a benchtop environment. We plan to continue the work studying the potential use of phase-stabilized cables and the possible effects of cable multipath, with the goal of identifying potential modifications to make at each MSN site. We have planned several combined short- and zero-baseline tests, utilizing four or more receivers, in order to capture the affects of multipath on our results. In the longer term, we will be performing longer continental baseline studies on a routine basis and considering how this process might be made to operate in near real-time.

CONCLUSIONS

The zero-baseline tests, when analyzed, have yielded reasonable results, given our knowledge of MSN hardware and environment. Also, our limited tests suggest that our software process is working correctly.

Our tests also strongly support the idea that it is the temperature, coupling to the MSN receivers, and not the level of isolation provided by the splitter, that is the dominant source of periodic perturbations in this type of analysis. The pattern in the results indicates that the effect is localized at the receiver, although we cannot currently distinguish between cable multipath effects induced by a changing reflection at the receiver, and changes in internal path delays. The overall delay instabilities seem to contribute a variation of approximately 750 ps/^{o}C between pairs of receivers. The limited data available do not suggest that the type of splitter plays an immediate role in the overall environmental dependence seen.

Finally, the Ashtech receiver is now nearing the end of its lifecycle, and the procurement of a replacement receiver/antenna system is now well beyond initial stages. We are developing a process and a set of standard tests, initially performed with the Ashtech, which we hope to eventually use for any newly procured receiver. The Ashtech data will provide a baseline for judging any future receiver, and, we hope, provide some indications about the suitability of any future receiver for time and frequency transfer applications.

REFERENCES

[1] K. M. Larson, and J. Levine, 1998, *"Time Transfer Using the Phase of the GPS Carrier,"* **IEEE Tranactions on Ultrasonics, Ferroelectrics, and Frequency Control, UFFC-45,** 539-540.

[2] K. M. Larson and J. Levine, 1999, *"Carrier-Phase Time Transfer,"* **IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, UFFC- 46,** 1001-1012.

[3] K. M. Larson, J. Levine, L.M.Nelson, and T. E. Parker, 1999, *"Assessment of GPS Carrier-Phase Stability for Time-Transfer Applications,"* **IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, UFFC-47,** 484-494.

[4] T. Schildknecht and T. Springer, 1999, *"High Precision Time and Frequency Transfer Using GPS Phase Measurements,"* in Proceedings of the 30th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 1-3 December 1998, Reston, Virginia, USA (U.S. Naval Observatory, Washington, DC), pp. 281-292.

[5] F. G. Ascarrunz, T. E. Parker, and S. R. Jefferts, *"Group-Delay Errors Due to Coherent Interference,"* in Proceedings of the 1999 Joint Meeting of the European Frequency and Time Forum (EFTF) and the IEEE International Frequency Control Symposium, 13-16 April 1999, Besançon, France (IEEE Publication 99CH36313), pp. 198-202.

[6] M. A. Weiss, F.G.Ascarrunz, T.Parker, V.Zhang, and X.Gao, *"Effects of Antenna Cables on GPS Timing Receivers,"* in Proceedings of the 1999 Joint Meeting of the European Frequency and Time Forum (EFTF) and the IEEE International Frequency Control Symposium, 13-16 April 1999, Besançon, France (IEEE Publication 99CH36313), pp. 259-262.

[7] M. A. Weiss, 2000, *"Long Term Effects of Antenna Cables on GPS Timing Receivers,"* in Proceedings of the 2000 IEEE/EIA International Frequency Control Symposium and Exhibition, 7-9 June 2000, Kansas City, Missouri, USA (IEEE Publication 00CH37052), pp. 637-640.

[8] E. Powers, 1999, *"Hardware Delay Measurements and Sensitivities in Carrier Phase Time Transfer,"* in Proceedings of the 30th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 1998, pp. 293-306.

[9] E. Powers, P. Wheeler, D. Judge, and D. Matsakis, 1999, *"Hardware Delay Measurements and Sensitivities in Carrier-Phase Time Transfer,"* in Proceedings of the 30th Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 1-3 December 1998, Reston, Virginia, USA (U.S. Naval Observatory, Washington, DC), pp. 293-305.

[10] G. Petit, C. Thomas, Z. Jiang, P. Uhrich, and F. Taris, 1999, *"Use of GPS ASHTECH Z12T Receivers for Accurate Time and Frequency Comparisons,"* **IEEE Transactions on Ultrasonics,**

Ferroelectrics, and Frequency Control, UFFC-46, 941-949.

[11] G. Blewitt, 1989, *"Carrier Phase Ambiguity Resolution for the Global Positioning System Applied to Geodetic Baselines up to 2000 km,"* **Journal of Geophysical Research, 94, B8,** 10187-10203.

[12] D. Matsakis, L. Schmidt, K. Senior, E. Powers, and J. DeYoung, 2001, *"Comparison of High-Precision Time Transfer Techniques,"* in Proceedings of the International Conference on Time and Frequency, 6-7 February 2001, New Delhi, India (Metrology Society of India and NPL, New Delhi), pp. 115-132.

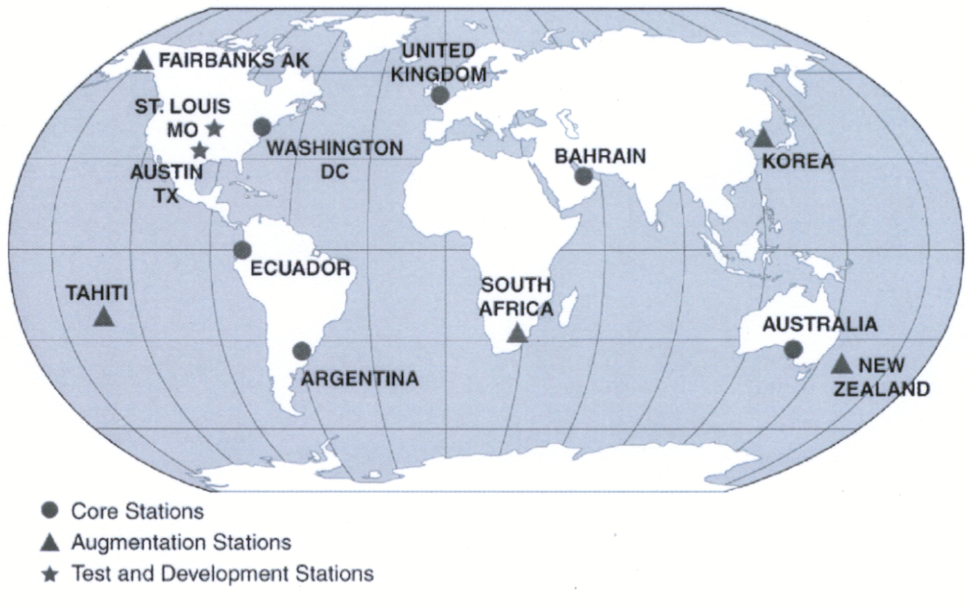

Figure 1. The network of NIMA GPS Monitor Stations.

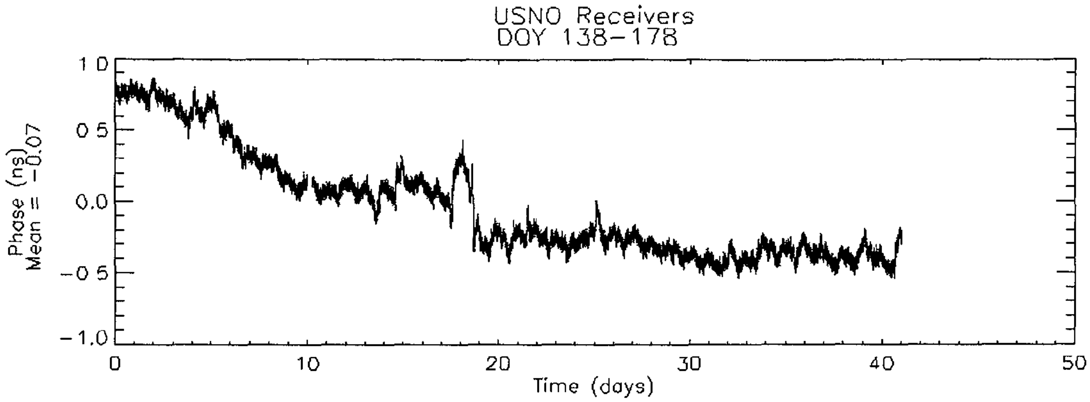

Figure 2. Phase difference between receivers at USNO for days 138-178 of year 2001.

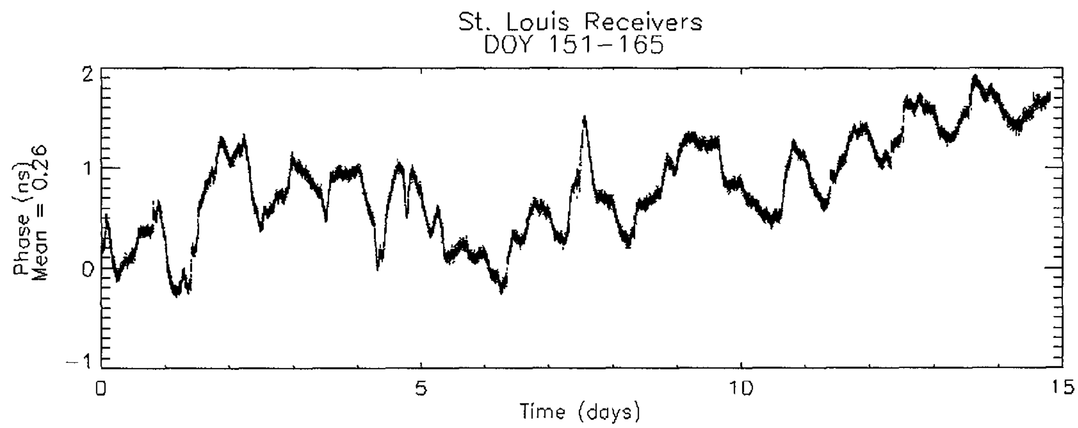

Figure 3. Phase difference between receivers at St. Louis for days of year 151-165 of 2001.

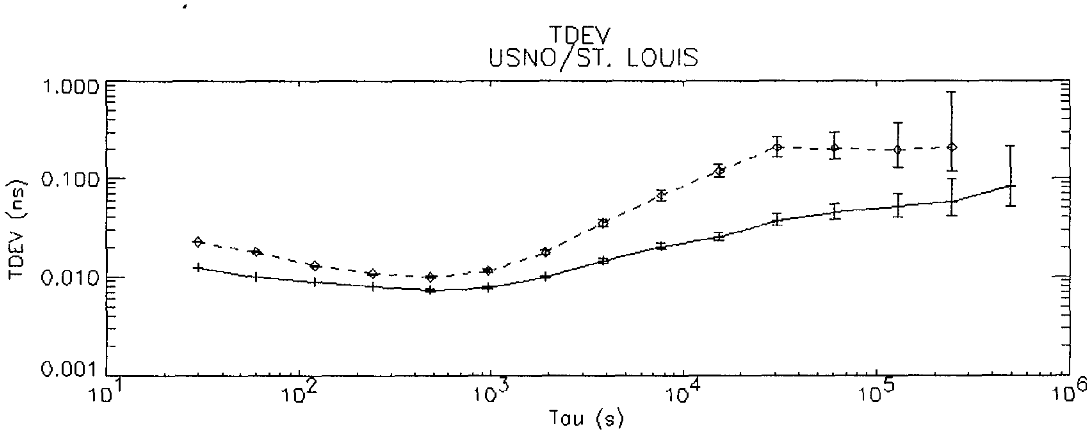

Figure 4. TDEV for the data in Figures 2 and 3 (USNO, full line, and St. Louis, dashed), showing flicker FM at times greater than 1000 s.

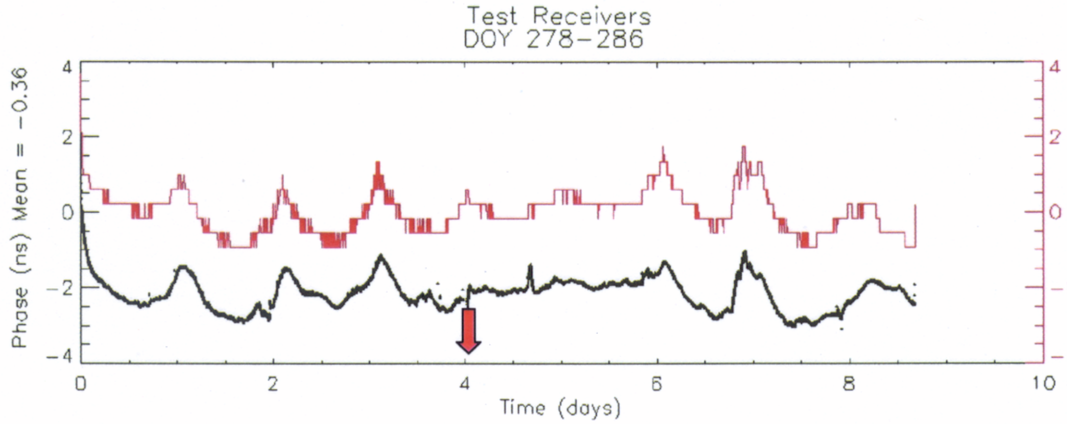

Figure 5. Zero-baseline comparison of phase offset between two Ashtech receivers at room temperature. The arrow marks the time at which the power splitter was swapped out. Temperature is the upper curve.

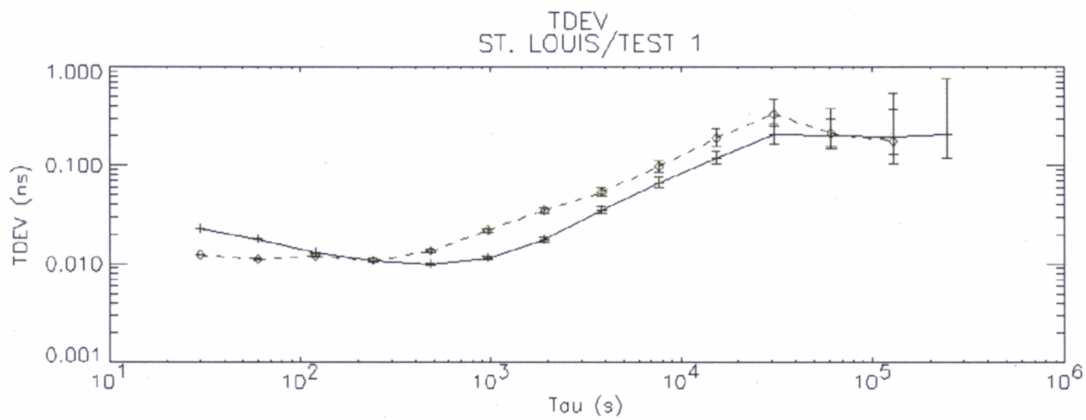

Figure 6. Comparison of TDEV for St. Louis (dashed, same as Figure 4) and data from Figure 5.

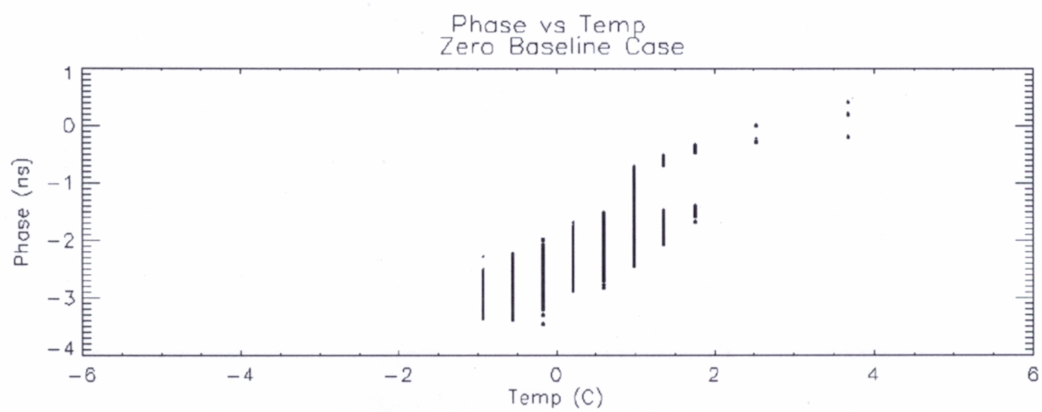

Figure 7. From Figure 5, an estimate of the phase as a function of temperature.

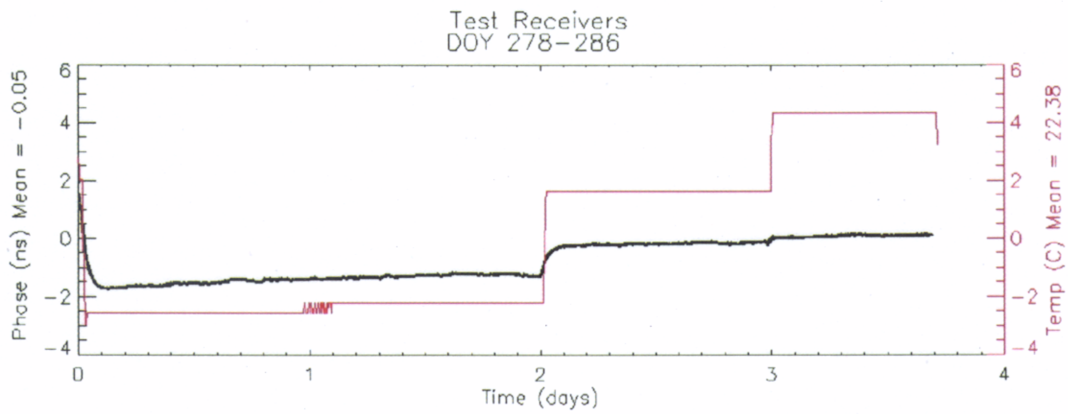

Figure 8. Relative clock phase in the zero-baseline test with both receivers in the temperature controlled chamber. Temperature (with mean removed) is the curve that goes from lower left to upper right.

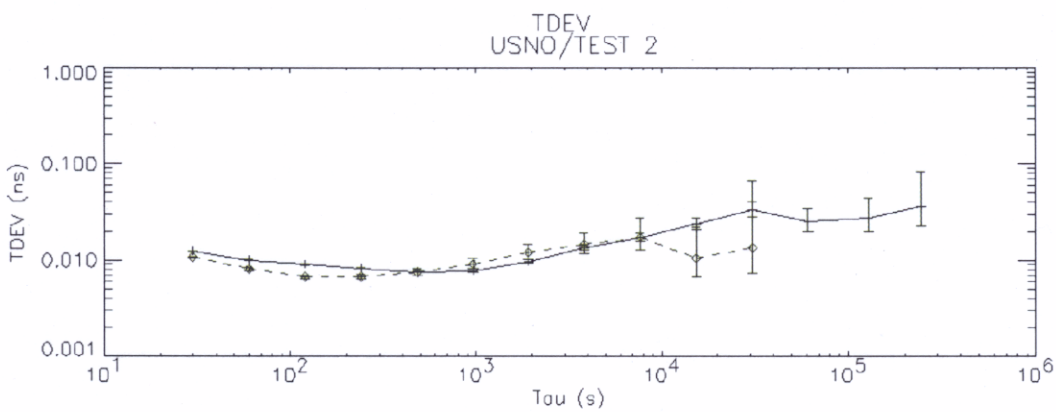

Figure 9. Plots of TDEV for USNO (full line; compare Figure 4) and zero-baseline case (dashed) with both receivers temperature controlled. Noise levels approach 7-8 picoseconds.

QUESTIONS AND ANSWERS

VICTOR REINHARDT (Boeing Satellite Systems): I've heard the term before, but exactly what do you mean by "wide-lane phase?"

BRIAN TOLMAN: It is a linear combination of the L-1 and L-2 phase. I cannot remember it off the top of my head what it is, but it is just a linear combination that has an effective wavelength of 86 centimeters. So it is a way of combining the two phases that remove some properties and add some others. It has a long wavelength, but relatively high noise.

REINHARDT: So it is a way to get rid of ambiguity?

TOLMAN: It is a strategy for finding the ambiguities in the phase. You have two ambiguities that you have to find, but you can form linear combinations and find one linear combination, and then the orthogonal one. And that's what we do.

REINHARDT: Thank you.

TOM CLARK (NASA/Goddard and Syntonics): In terms of Victor's question, you construct an observable at L-1 minus L-2 as a frequency, which ends up being 386 megahertz, which corresponds to the 86-centimeter wavelength that he mentioned. The question I was going to ask is when you put these things in the chamber, you are obviously seeing some thermal effects. I, for one – and I think many of the others in the community – would like to know what is it inside the receiver that is actually causing this? In other words, understand the physics, not just figure out how to make a Band-Aid.

TOLMAN: Absolutely.

CLARK: It would be very interesting to see a little more detailed information on the time constants when you made temperature steps. And also, to see if there is asymmetry when you make the step up from when you make it down. So I think there is a lot to be done there. It would also be very nice if you would come up with some way to apply, rather than using a large environmental oven, to be able to heat up specific things in the receiver – crystal oscillators, RF amplifiers, and things like that would be the logical thing to suspect; probably power supplies are not something that would be an effect.
You've made a good start. I, for one, want to see some more.

TOLMAN: We are continuing to work on it. I do not think we will get inside the receiver, but we are going to continue to pursue this. And you are right; I would love to know what it is inside the receiver that is doing this.