TIME TRANSFER USING A MULTI-CHANNEL CODE-BASED GPS RECEIVER

Judah Levine
Time and Frequency Division
National Institute of Standards and Technology
Boulder, Colorado 80303

Abstract

We report on time-transfer experiments using a GPS receiver that we have constructed using a commercial GPS "engine" and a standard PC. The receiver measures the time difference between the local clock and a 1 pps signal that is synchronized to GPS time using data from up to eight satellites. The receiver also reports the difference between GPS time as estimated using each of the satellites being tracked and the composite output pulses that have a rate of 1 Hz (1 pps signal). These data can be used to construct the standard 13-minute tracks as defined in the BIPM standard; the same data can also be averaged in other ways that make better use of the multi-channel capabilities of the hardware. The 13-minute averages can be directly compared with standard time-transfer receivers using common-view analysis. The results of our tests suggest that we should re-examine the methods currently used for national and international time and frequency coordination, and we suggest an alternative approach based on multi-channel receivers that should be more flexible, simpler and easier to operate than the current system.

Introduction

Time transfer between timing laboratories is currently realized using single-channel GPS receivers. These receivers measure the time difference between the 1 pps output of a local clock and GPS time; these data are combined with similar measurements at other laboratories using the standard "common-view" method [1,2]. This method reduces the effects of radial orbit errors, errors due to the fluctuations in the satellite clock, perturbations in the transit time due to the atmosphere and the ionosphere and other common-mode effects.

The receivers that are currently used can only track one satellite at a time, and time-transfer between laboratories therefore requires a pre-defined tracking schedule. In fact, several different tracking schedules are defined. Each one is designed to optimize the time transfer in a geographical region (within the Western US, for example) or between two specified regions (between the

US and Europe, for example). This optimization is accomplished by choosing observation times so that the specified satellite is well above the horizon and roughly equidistant from all stations that will use the schedule. These tracking schedules are published periodically by the BIPM [3] and all major timing laboratories follow them.

A second important consideration in using GPS satellites for time-transfer is the selective availability ("SA") degradation of the signals imposed by the US Department of Defense. This degradation is generally accomplished by dithering the frequency of the clock in each satellite [4]. (The orbital parameters can also be dithered, but this has not been observed to date.) The dithering affects the receiver measurements of the time difference between the local clock and GPS time. However, its effect is the same at all stations and it therefore cancels in the common-view subtraction that is always a part of high-accuracy time transfer. (This cancellation in the common-view time difference would generally not be as effective if the orbital parameters were dithered instead of the clock.)

The cancellation of the effects of SA in the common-view time differences depends on the fact that all stations process the measurements in exactly the same way. In addition to the tracking schedules, the International Bureau of Weights and Measures (BIPM) has therefore defined an analysis method for averaging the observations at each station [5]. Almost all participants in GPS time-transfer use this averaging method in analyzing their time-difference data. This method is often implemented as part of the internal firmware of the receiver. (The BIPM definition also includes a specification of the output format that is to be used to report the data, but this is a secondary issue that will not concern us here.)

The data that result from implementing both the tracking schedule and the specified averaging algorithm are usually called tracks. Each track reports the average time difference between the local clock and GPS time computed using 13 minutes of data. A 24-hour day could accommodate 110 13-minute tracks in principle, but the existing receivers require some time between tracks for synchronization and housekeeping, and generally can handle only 96 tracks or less as a result. The track times advance by 4 minutes every day starting from the origin times defined in the published schedule, since the satellite

284

geometry is synchronous with the sidereal day. This advance is normally handled automatically by the receiver firmware.

In this paper we suggest that we should reconsider both the concept of tracking schedules and the details of the averaging method, given the advent of relatively inexpensive multi-channel GPS "engines" that can be used to construct general purpose time-transfer receivers. In order to support this idea, we have constructed a simple GPS receiver based on such an engine, and we have conducted a series of tests both at NIST and between NIST and other stations.

Design of the NIST multi-channel receiver

A number of inexpensive GPS "engines" are now available commercially. The typical device produces an output pulse that is synchronized to GPS time (possibly with a varying time offset that is specified in a nearly simultaneous data message). This pulse is then connected to one port of a separate time-interval counter, which measures the time difference between this pulse and the corresponding one from the local clock. These time-differences are assigned a time-tag that can be specified either in UTC or in GPS time and they are then stored for further processing.

The timing of the output pulse from the GPS engine is usually derived using data from several satellites—up to 12 in some cases. As discussed above, the time differences recorded by this receiver will fluctuate because of the dithering of the satellite clocks as a result of the implementation of SA. (The effects of SA are not canceled by averaging data received at the same time from different satellites.) Since receivers at different locations will generally be tracking a different group of satellites, a simple common-view subtraction of the two data sets will not cancel the effects of SA. The only way to do this is to subtract the data satellite by satellite, and the receiver must provide enough information to support this computation. In other words, the receiver must provide the offset between the time message received from each satellite and the time of the composite average that is used to produce the physical tick. Although it is not absolutely necessary, it is also highly desirable if the time of the output pulse is specified with nanosecond level resolution.

Using these criteria as a guide, we have constructed a GPS receiver using a commercial GPS engine, a standard PC and a time-interval counter board that connects to the PC bus. The hardware is controlled and configured using the standard Graphical User Interface of the PC operating system, and the data are written to the local hard disk in a standard system format. The GPS engine produces an output pulse every second that is synchronized to GPS time using data from up to eight satellites. The time-interval board measures the time difference between each of these

pulses and the corresponding pulse from the local clock. The receiver also provides the ancillary data that are needed to compute the contribution of each satellite in view to the timing of the composite pulse; these data are transmitted in serial format and are received using one of the COM ports of the PC. The data from the receiver are combined with the time intervals measured by the counter to produce measurements of the time difference between the local clock and GPS time for every satellite that is being tracked. The output data stream therefore consists of a UTC time-tag and up to 8 pairs of numbers. Each pair consists of the number of the satellite being tracked and the time difference between the local clock and GPS time using the data from that satellite. When all 8 channels are active, the output data rate is somewhat more than 100 bytes per second.

The PC also has additional software to average these data, to store them on the local disk or to transmit them to distant users using standard electronic-mail techniques and formats. All of these programs are managed by the operating system of the PC; in all cases the interface to the user is via standard graphical techniques (i.e., "point and click" using a mouse). All of the software is written in a standard high-level language, and no special modifications to the hardware or base software of the PC are necessary.

Common-view experiments

Our first test was to demonstrate that our receiver design is at least as good as the existing single-channel systems. We installed the receiver at NIST and averaged the one-second measurements of the time difference between a local clock that was traceable to UTC(NIST) and GPS time. We used the BIPM method for averaging the data and observed the satellites as specified in the BIPM track schedule for the Western US that was current at the time. We compared these data against a standard single-channel GPS receiver that made the same measurements, averaged them in the same way and followed the same track schedule. (This reference receiver is the one that is used for time-comparisons between NIST and other laboratories.)

Since the reference clock is the same for both receivers and since the paths from the receivers to the satellite are nearly identical, variations in the measurements due to any of these sources should cancel in the common-view difference between these two data sets. The only exceptions are differences in cable lengths, which would produce a constant time difference between the two data sets, and effects such as multipath reflections, which do not affect the two systems in exactly the same way (even though the antennas are only a few meters apart).

We conducted this experiment for about 30 days, during which time we measured 1238 common-view tracks. (The published track schedule for the Western US

contains substantially fewer than the maximum number of tracks that could be implemented in principle.) The results of this experiment are shown in fig. 1. Based on a direct measurement of the cable delays, the expected static time difference between the measurements of the two receivers should have been -24 ns; the mean of the data was measured to be -25.54 ns, in good agreement with this value. The spectrum of the data can be approximately characterized as white phase noise at short periods with an amplitude of 2.7 ns RMS; this is comparable to what would have been observed using two standard single-channel timing receivers in the same configuration. As we will show below, this is not totally due to receiver noise. Differences in multipath reflections contribute to the time-difference data at periods of a few minutes and longer, and a differential sensitivity to fluctuations in the ambient temperature may also make a smaller contribution.

temperature of the hardware. The fluctuations in the ambient temperature were, of course, common to both receivers, and their effects would therefore have canceled to first order in our calibration procedures.
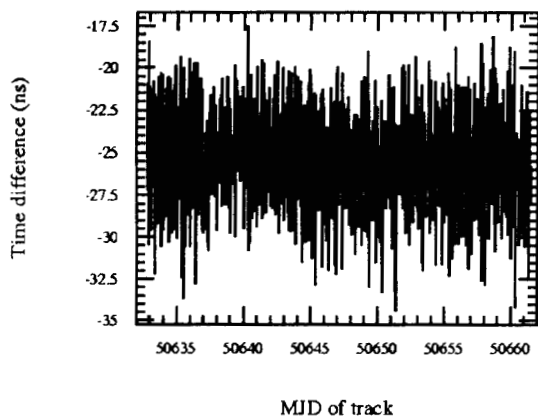

Fig. 1. A short-baseline common-view experiment between a multi-channel receiver whose data are averaged according to the BIPM method and a single-channel timing receiver. Both receivers are located at NIST and both use a reference clock that is traceable to UTC(NIST). The static-time difference between the two receivers is due to a difference in cable delays as described in the text.

Fig 2. The time difference between USNO and NIST. (a—upper figure) Measured using standard single-channel receivers and the BIPM track schedule. (b—lower figure) Measured using a multi-channel receiver that recorded data from all satellites in view. Both figures use the same horizontal and vertical scales.

We have repeated the measurement of the delay through this receiver a number of times over a period of nearly one year, and we have performed similar tests using other receivers based on the same type of hardware. In every case, the time-difference data measured using these short-baseline common-view experiments have agreed with the measured cable delays to better than the RMS uncertainty in the measurements.

The receivers used in these tests were located in a standard laboratory environment with no special temperature controls. Our hardware design uses heat sinks to attenuate the effects of temperature changes on the receiver, but we did not measure or actively stabilize the

The accuracy and stability of the receiver calibration was confirmed in an experiment between NIST and the US Naval Observatory (USNO). We measured the time difference UTC(USNO)–UTC(NIST) using a multi-channel receiver at each of the sites, and compared these data with the same time difference measured using standard single-channel hardware. The multi-channel receivers did not use a tracking schedule but instead simply recorded data from all satellites in view. The 5-minute average common-views were computed afterwards using all possible satellites; the standard single-channel receivers realized the standard BIPM track schedule and analysis

method. The experiment started on MJD 50770 (18 November 1997) and is still in operation. Figures 2a and 2b show a 25-day portion of both data sets. The two plots use the same scale for both axes to facilitate a comparison of the two results. The noise level in both data sets is about 4 ns RMS (4.2 ns for the standard one-channel data and 3.6 ns for the 5-minute multi-channel averages), and the disagreement between the two data sets is less than 1 ns RMS over the full period of observation. The multipath effects at the USNO are not as serious as they are at NIST, and most of the contribution to the variance at intermediate periods came from the NIST stations.

The RMS noise in both of these data sets is larger than what we observed in the corresponding short-baseline common-view tests (Fig. 1 above and Fig. 3 below). This is to be expected—the variations in the troposphere and the ionosphere are not as well correlated over the longer baseline and their contributions do not completely cancel in the common-view differences over longer distances. In addition, most errors in the satellite ephemeris affect the paths to the two stations by different amounts and are therefore not perfectly cancelled by the common-view difference.

The USNO–NIST comparison has no adjustable parameters or offsets. The multi-channel receiver that we installed at USNO was calibrated at NIST using data analogous to the short-baseline common-view measurements presented in Fig. 1 above. This calibration was combined with the cable delays in the 1 pps distribution system at the USNO to arrive at the overall calibration for the receiver. Both at USNO and at NIST, the cables used to interface the multi-channel receiver to the local clocks are not the same as the ones that are used for the single-channel receivers. (In fact, the two receivers are NIST are driven from two different physical clocks whose times are only related to each other through the NIST AT1 time scale hardware.) The agreement between the two data sets suggests that all of these independent measurements have been done correctly.

Multi-channel receiver noise measurements

A comparison with a standard timing receiver (as in the USNO–NIST experiment) does not make the best use of the multi-channel capabilities of the GPS engine, since we are constrained by the single-channel limitation of the standard receiver. In addition, short baseline, common-clock experiments are better suited to evaluate the noise performance of the receiver itself, since most external effects are cancelled in the common-view differences.

Figure 3 shows a common-view test between two identical multi-channel receivers whose antennas were located a short distance apart. The RMS noise in the difference between the two receivers was about 1.7 ns,

with longer-period fluctuations that are primarily due to multipath.


MJD of track

Fig. 3. A short-baseline common-view experiment between two multi-channel receivers. The two receivers are connected to the same clock. The antennas are about 1 m apart. The average time-differences between the two receivers are computed for all satellites in view as described in the text.

As in the multi-channel experiment between USNO and NIST, the data shown in Fig. 3 were processed by computing 300-second averages of the time-differences using all satellites that were in common-view at any time. Each receiver recorded time-differences using all of the satellites that it could see, and the common-views were computed after the fact in software. Since there were about 6 satellites visible at any time, each data point represented an average of about 1800 one-second measurements. Each point in a standard common-view comparison (such as Fig. 1), on the other hand, represents the average of 780 one-second measurements (where the term "average" actually means the more complicated procedure defined by the BIPM, which produces a single time difference output from 780 one-second measurements.). Assuming that both data sets were limited by the white phase noise in the measurement process itself, we would expect the data in Fig. 3, acquired using a multi-channel receiver, to have a RMS amplitude of about 66% of the data in Fig. 1, which was obtained using the standard 13-minute analysis. The observed ratio is 63%—in good agreement with this expectation. (In fact, each 5-minute data point from the multi-channel receiver contains somewhat less than 1800 points, since the averaging algorithm used in the multi-channel receivers does not use the full 300 seconds of data from each satellite. The relative performance of the multi-channel receiver is therefore somewhat better than would be expected based on this simple proportionality.)

Figure 4 presents the comparison between the different types of receivers using the TVAR statistic, $\sigma_x(\tau)$. The points plotted with the symbol "+" illustrate the common-view time-difference between two single-channel receivers operating at NIST. The points plotted with the symbol "*" show the same measurements between two multi-channel receivers. In both cases, the two receivers use the same clock as a reference. The straight line between the two sets of points has a slope of -1/2, the slope that is characteristic of white phase noise on a TVAR plot. The data from the single channel receivers are acquired using the standard BIPM track schedule and are therefore not equally spaced in time. We have used simple linear interpolation to repair this deficiency, recognizing that this procedure may distort the variance calculation to some extent. The data from the multi-channel receivers use 5-minute averages of all satellites that are in common-view at any epoch and are used as is. Both data sets are well characterized as white phase noise, except possibly at the longest time intervals. Based on this analysis, we would conclude that the multi-channel receiver data has an RMS noise level of about 40% of the single-channel results, an advantage that is somewhat greater than the value calculated above.
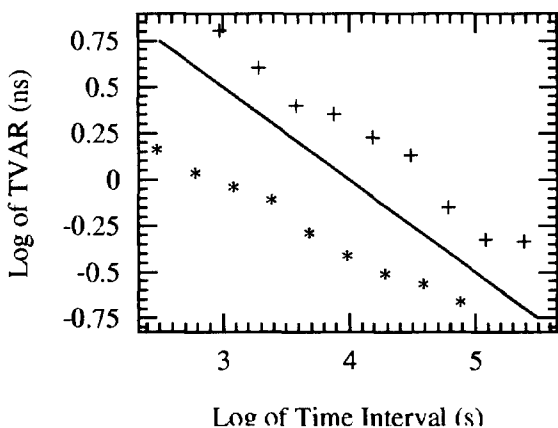

Fig. 4. TVAR computed using short-baseline common-view time differences. The points plotted using "+" show data from single-channel receivers; those with "*" show data from multi-channel receivers. Each data set is processed as discussed in the text.

The multi-channel receiver is clearly more efficient in using the satellite constellation, and it can obviously make measurements more rapidly than a single-channel device. However, this experiment shows that the noise in both designs is roughly proportional only to the number of elemental one-second data points that are used to compute each time-difference. The analysis method used in the standard receivers does not provide superior performance over the much simpler averaging scheme we have used for

the multi-channel results shown in Fig. 3. (As we pointed out above, when the noise in the common-view data is normalized by the number of elemental one-second measurements that are used to compute it, the BIPM analysis method is somewhat worse than a simpler average.) This suggests that, at the very least, the computationally expensive analysis method used in the current single-channel receivers does not do a better job than simpler methods, and therefore might not be worth the computing power needed to realize it.

This result should not be terribly surprising. Figure 5 shows the one-second time-difference data between the local clock and GPS time measured by a single receiver observing a single satellite. These data are not noise in the usual sense of the word—to see this, compare the vertical scale on this figure, which is much larger than those on the previous ones. The variance in this figure is completely dominated by the degradations in the satellite transmissions due to Selective Availability. The upper left hand corner of the figure also shows two horizontal lines corresponding to averaging periods of 5 minutes (identified by the symbol "/") and 13 minutes (identified by the symbol "!"). Both of these averaging times are very short compared to the typical period of the SA-induced variations shown in the figure. Neither period will "average" the SA in the commonly understood sense of that term no matter what averaging algorithm is chosen. In other words, any algorithm operating using either of those averaging times will produce a result that is not stationary. The improvement that is realized in the common-view method arises solely because of the fact that both stations are processing the data in the same way and then subtracting the results—not because the details of the averaging algorithm in the receivers are well matched to the received data.


Days from MJD 50714

Fig. 5. The time difference between the local clock and GPS time measured using data from satellite 9 only. The figure shows 1 s measurements with no averaging, and the observed variation is due to the effects of selective availability (SA).

Multipath reflections often make a significant contribution to the variance of common-view data. Figure 6 shows the one-second common-view time differences measured by two receivers whose antennas were about 1 m apart. Each trace shows the time-difference data as a function of UTC hour on a single day using only satellite 9. The time-tags of all of the observations on each day after the first one have been advanced by 4 minutes relative to the previous day, and each trace has been offset vertically for clarity. Although the details of the structure vary as the satellite moves through the sky, the same pattern repeats after one sidereal day when the satellite returns to the same position in the sky relative to the reflectors. As on the previous figure, the two horizontal lines near the top of the figure have lengths corresponding to averaging intervals of 13 minutes (identified by the symbol "!") and 5 minutes (identified by the symbol "|"). The periods that characterize the variation in the data due to the multipath reflections vary depending on the exact position of the satellite in the sky with respect to the reflectors, but neither averaging interval is really appropriate for dealing with these variations.
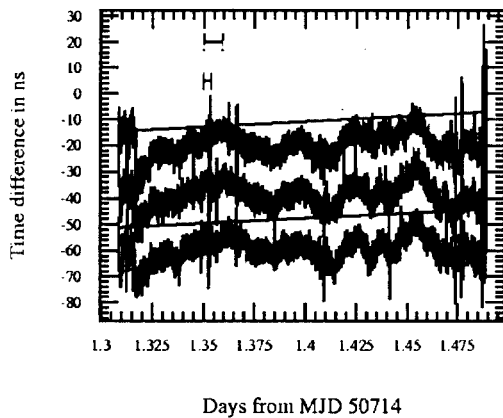

Days from MJD 50714

Fig. 6. A short-baseline common-view experiment using data from satellite 9 only. Time differences from successive days are offset by 4 minutes and are displaced vertically by an arbitrary amount for clarity. The figure shows 1 s measurements with no averaging.

The data plotted in Fig. 6 show the *differential* effect of multipath on two antennas that are only about 1 m apart. The multipath effect on each individual receiver can be estimated using satellite 15, which has no SA. Figure 7 shows the time difference between a local clock and GPS time measured using this satellite only. The multipath effects here are comparable to those in the previous figure, and the same comments with respect to averaging times apply. (The fluctuations in the time difference between the local reference clock and GPS time are very small on this time scale.)

The fact that the multipath effects have a period of one sidereal day is exploited in the design of the BIPM tracking schedule. The start time of each track is advanced by 4 minutes every day so that it always starts at essentially the same sidereal time on consecutive days. This synchronous sampling of the multipath effect does not remove it, but rather converts it from a time-varying effect to a static offset in first order. Based on the data shown in Figs. 6 and 7, the magnitude of this offset can approach 10 ns peak to peak. This offset actually changes slowly with time because the observation periods are not precisely synchronized with a sidereal day and because of perturbations on the orbits of the satellites. This slowly varying multipath offset may masquerade as a response to long-period seasonal temperature fluctuations.

It is possible to reduce the effects of multipath reflections using choke-ring antennas; at least at NIST, these do not work as well as one might like, and residual multipath is still a problem even when they are used.
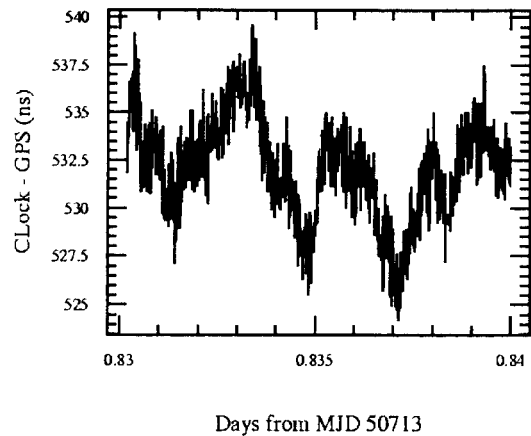

Days from MJD 50713

Fig. 7. The time difference between the local clock and GPS time measured using data from satellite 15 only. This satellite has no SA, and the observed fluctuations provide an estimate of the effects of multipath and receiver noise as discussed in the text.

Discussion of the data

Our analyses used three kinds of data—raw one-second measurements as measured by the hardware and two different averages of these values. These results can provide some guidance in the design of an operational time-transfer system based on these multi-channel receivers.

Averaging data is justified from a statistical point of view only when the data can be characterized (at least approximately) as white phase noise. Although the average of a time series will always exist in a formal sense even when the data cannot be even approximately characterized as white phase noise, its behavior is not

always well defined in these cases. The best way to satisfy this requirement would be to postpone averaging GPS measurements until after the common-view time differences had been computed because the common-view subtraction cancels or attenuates many systematic contributions to the data variance. In other words, common-view data are about as white as time differences ever get.

On the other hand, this strategy maximizes the amount of data that must be transmitted and stored. Using the rough estimate that an 8-channel receiver produces data at the rate of about 100 bytes per second, the world-wide network of timing laboratories might easily produce a data base of 300 Mbytes per day. This figure includes only those laboratories that contribute data directly to the BIPM—it does not include the many other users of these measurements.

This is not an unthinkable amount of data given the size of mass storage devices that are now generally available, but the real objection to storing and transmitting this much data is that their actual information content is much smaller than this volume would indicate. If it were not for effects such as SA and multipath reflections, averaging and then decimating the measurements to one point per hour would be more than enough for most applications.

The problem would become more tractable if we did not try to solve it by transmitting all of the data to a central location for the common-view computations, and if we were prepared to relax the notion that this central archive must support a direct common-view calculation between *any* two contributing laboratories. Suppose, instead, that we defined two "hub" laboratories in each region. These hub laboratories would make a commitment to publish electronically their measurements of the difference between their local clock and GPS time with little or no averaging and on a near real-time basis. Other laboratories in the region would download these data and compute the common-view time differences between themselves and the hub. These common-view differences would then be averaged to something on the order of 1 point per hour, and these decimated data could be archived locally or transmitted to a central location for archival storage. Depending on the exact format that was used, the size of the system archive might be somewhat smaller than the current common-view data from the laboratories that contribute data to the BIPM.

The differences between the hub laboratories would be computed in the same way. A receiver at the BIPM might act as the central hub for the entire network. The exact location of this central hub is not important, of course, since its clock would drop out in any calculation of the time-differences between two contributing laboratories. In fact, the location of any of the hubs need not be at a timing laboratory at all, although such locations are more likely to have the infrastructure necessary to support the

commitment that is implied in the concept. Since the hubs represent single points of failure, it is important to have more than one site in each region functioning in this capacity.

Some laboratories will probably choose to keep their existing single-channel receivers for some time to come, and it would be most useful if the new system remained compatible with the averaging method used in this equipment. There are two ways of achieving this. The first would be for the hub laboratories to publish the raw one-second time-difference measurements, and the second would be to implement the hub system using simple 15-second averages of the raw data. The one-second data could be used to construct tracks that conformed exactly to the currently defined BIPM procedure, and laboratories that kept their old receivers could down-load the one-second data and construct standard BIPM tracks in software for comparison with the data produced by the existing equipment. Constructing tracks using the 15-second averages would be different in principle from the BIPM algorithm, because that algorithm specifies a quadratic least-squares fit over the 15-second data rather than a simple average. However, both our simulations and our tests with real data show that the difference between the two schemes would be too small to measure in almost all cases. We think that the decrease in the amount of data that would have to be processed in this case would more than offset the slight loss caused by the change in the way the 15-second blocks were computed.

Conclusions

We have constructed a multi-channel GPS receiver using a commercial GPS engine, and we have used this receiver in a number of time-transfer experiments. The receiver can be used to realize the averaging format and track schedule specified by the BIPM for international time coordination. These algorithms do not make the best use of the multi-channel capabilities of the hardware, and we have used these capabilities to compute common-view time differences using all satellites in view. These measurements are more efficient than the standard method in that they make better use of the GPS constellation. They can also be used to implement multi-hop common-view measurements between two stations that have no satellites in common-view. To do this, the stations at the two end-points would use a mid-point hub whose clock drops out of the final time differences.

In addition to making better use of the GPS constellation, multi-channel receivers can track most satellites from horizon to horizon. This continuous tracking can be used to evaluate the contribution of multipath effects to the measured time-differences. This is more difficult to do with single-channel receivers, since the receiver cannot be dedicated to a single satellite for a long

period of time. In addition, the 4-minute advance in the BIPM track schedule aliases the multipath variation and converts it to a slowly varying offset that may appear to be a sensitivity to long-term temperature fluctuations or similar effects.

For all of these reasons, the increase in the volume of data that is available from multi-channel hardware suggests that the traditional methods for computing the time-differences among timing laboratories should be re-examined. We have suggested one alternative approach that has a number of advantages over the current method based on single-channel receivers and tracking schedules. This method would support near real-time estimates of the time-differences between contributing laboratories. Using the software that we have already developed in our pilot program (or some other equivalent implementation), the network we envision would be almost totally automated and would require almost no manual intervention in normal operations. It should substantially simplify the tasks associated with international time and frequency coordination and the distribution of time and frequency information from national timing laboratories to their customers.

References

1. Allan, D.W. and M.A. Weiss (1980), "Accurate Time and Frequency Transfer During Common-view of a GPS Satellite," *Proc. 34th Symposium on Freq. Control*, Philadelphia, PA, pp. 334-346.

2. W. Lewandowski (1993), "GPS Common-View Time Transfer," *Proc. 25th Precise Time and Time Interval Planning Meeting*, pp. 133-148.

3. *Annual Report of the BIPM Time Section*, Vol. 9, Sevres, France: BIPM, pp. 69-73, 1996.

4. Allan, D.W., M. Granveaud, W.J. Klepczynski and W. Lewandowski (1990), "GPS Time Transfer with Implementation of Selective Availability," *Proc. 22nd Precise Time and Time Interval Planning Meeting*, pp. 145-154.

5. Allan, D.W., and C. Thomas (1994), "Technical Directives for Standardization of GPS Time Receiver Software," *Metrologia*, Vol. 31, pp. 69-79.