A New Technique for Estimating Frequency from GPS Carrier-Phase Time Transfer Data

Christine Hackman
JILA, University of Colorado
Boulder, Colorado USA
chackman@jila.colorado.edu

Judah Levine and Thomas Parker
Time and Frequency Division
National Institute of Standards and Technology
Boulder, Colorado USA

Dirk Piester and Jürgen Becker
Physikalisch-Technische Bundesanstalt
Braunschweig, Germany

Abstract—GPS carrier-phase time transfer (GPSCPTT) offers good frequency stability at short averaging times, approaching a fractional frequency stability of 10^{-15} at an averaging time of 1 d. However, a discontinuity occurs in the time transfer estimates between the end of one processing batch (typically 1-3 d in length) and the beginning of the next. The conventional procedure for computing frequency from GPSCPTT data over a multi-day epoch is to remove the discontinuities and then to compute the frequency from the resulting continuous set of time transfer estimates. We present a new technique in which a frequency is computed from each batch solution and then these frequency values are averaged to obtain a mean frequency for the epoch of interest. The advantages of the new technique are (a) robustness in the event of a data outage and (b) the ability to compute frequency in the absence of the uncertainty introduced by the removal of the discontinuities.

The new method was tested by measuring the frequency difference between clocks located in Braunschweig, Germany and in Boulder, Colorado: frequency values obtained using the new technique were compared with those obtained using the conventional technique and those obtained from two-way satellite time and frequency transfer (TWSTFT). The frequency values obtained from the GPSCPTT data using the new method agreed with those obtained using the conventional method at 2-$3 \cdot 10^{-16}$; thus, it may be feasible to replace the conventional method with the new one. The frequencies obtained from the GPSCPTT data using either method agreed with those obtained from TWSTFT at a few to several parts in 10^{16}; therefore, both TWSTFT and GPSCPTT are approaching frequency-transfer performance sufficient for use during simultaneous fountain evaluations.

Keywords—carrier phase time transfer, frequency transfer, GPS, GPSCPTT, two way satellite time transfer, TWSTFT, TWSTT

I. INTRODUCTION

GPS carrier-phase time transfer (GPSCPTT) has shown great promise over the past several years. In an experiment conducted between receivers located at the United States Naval Observatory (USNO; Washington, D.C.) and at Schriever Air Force Base (Colorado Springs, Colorado), Larson *et al.* [1]

showed agreement between GPSCPTT and two-way satellite time and frequency transfer (TWSTFT) at a level of $\pm$ 1 ns over 96 d. Similarly, over the 6274 km baseline between USNO and the Physikalisch-Technische Bundesanstalt (PTB) in Braunschweig, Germany, Dach *et al.* [2] showed several-nanosecond-level agreement between GPSCPTT and TWSTFT over a period of about 180 d.

Unfortunately, GPSCPTT is limited in that a discontinuity occurs in the time transfer estimates between the end of one processing batch and the beginning of the next. A batch is typically 1-3 d in length; these steps generally range in size from 100 ps to 1 ns, depending on the level of pseudorange noise at the receiver sites. Though techniques exist for estimating and removing these steps, this introduces additional uncertainty into the results.

If the goal is to create a continuous time series that represents the phase difference between the two clocks being compared, then it is necessary to align the discontinuous results. This has become a common practice, even when the desired end is to compute the frequency difference between the clocks in question. However, the central goal of the NIST GPSCPTT program *is* to compare the frequency—not the time—of the world's primary frequency standards. This has raised the question of whether it is feasible to compute a frequency directly from each batch of data, and then to average (or otherwise combine) these frequencies in order to obtain the mean frequency over the epoch of interest.

There are two technical advantages to such an approach. The first is robustness in the event of a data outage: if GPSCPTT measurements are interrupted, there is no way to phase-connect (i.e., concatenate) the time transfer estimates obtained before the outage to those obtained after the outage. This makes computing the mean frequency over an interval containing the gap nearly impossible, as one needs to join the estimates in an arbitrary manner in order to do so. In contrast, if the frequency is computed directly from the batch results, one can continue doing so after the outage. Caution must still be used when computing the mean frequency over an interval that

encloses the data gap: if the relative frequency of the clocks measured is not stable, this loss of information can bias the estimate of the mean frequency. Nonetheless, it is possible.

The second advantage of this method is that it gives us the opportunity to estimate frequency in the absence of the uncertainty introduced by the phase-connection process.

The amount of uncertainty introduced by the phase-connection process is not well understood. If it is small enough to be negligible and if the predominant noise type is white FM, then a frequency computed from a concatenated set of time transfer estimates will have a smaller uncertainty than that computed from an average of individual batch-derived frequencies. That is because the uncertainty of the former will scale as $1/N\tau_o$, whereas the uncertainty of the latter will scale as $1/(\sqrt{N})\tau_o$, where τ_o represents the length of a batch and N represents the number of batches that are concatenated or averaged to obtain the result. However, if the phase-connection process introduces significant error, then the long-term frequency stability of the GPSCPTT solution may be compromised. That is because the uncertainty of a frequency computed from the concatenated time transfer estimates will have a component that increases as the number of connections within the epoch of interest increases. Estimating frequency from individual batches circumvents this problem.

To test this concept, we conducted a GPSCPTT experiment between frequency standards located at the National Institute of Standards and Technology (NIST) in Boulder, Colorado and at PTB in Braunschweig, Germany. TWSTFT was also performed between these clocks. Using the GPSCPTT data, we computed the frequency difference between the frequency standards in two ways: by phase-connecting the batch solutions and then estimating a frequency from the continuous set of time transfer estimates, and by computing a frequency directly from each batch solution and then averaging those estimates. We then compared these two frequency estimates with the value obtained from TWSTFT.

As we shall see, the results are promising: the frequencies obtained from all three methods agree at a few-to-several parts in 10^{16}. This implies that it may be feasible to replace the concatenation method with the new method. Furthermore, one of the fundamental goals of our GPSCPTT research is to attain a level of frequency-transfer performance sufficient for use during the simultaneous evaluation of fountain frequency standards. A fountain standard realizes the SI second with a frequency uncertainty of approximately $1\cdot10^{-15}$; any technique used to compare the frequency of fountain standards must have a frequency uncertainty considerably smaller than this. Because the frequencies obtained from GPSCPTT and TWSTFT—two independent methods—agree at a few-to-several parts in 10^{16}, it appears that the performance obtained from each of these methods is approaching that needed for this task.

II. DATA COLLECTION

GPSCPTT and TWSTFT measurements are recorded at NIST and PTB on an ongoing basis. For this experiment, we analyzed measurements collected during MJDs 52926-46 and 53000-17 (October 14-November 3, 2003 and December 27,

2003-January 13, 2004, respectively). These periods were chosen because the PTB fountain, CSF1, was evaluated during MJDs 52929-44 and 52999-53014.

The same frequency standard is used to operate both the GPSCPTT and TWSTFT systems at NIST; this frequency standard is a hydrogen maser steered to UTC(NIST). Similarly, one hydrogen maser is used to operate both the GPSCPTT and TWSTFT systems at PTB. Thus, the frequency values obtained from GPSCPTT can be compared directly to those obtained from TWSTFT. The maser used at PTB during MJDs 52926-52946 was different than the one used during MJDs 53000-17; we refer to the former as H2 and the latter as H4.

TWSTFT measurements were made between NIST and PTB several days per week at approximately 14:50 UTC. The measurements were performed via an INTELSAT transponder using transmit and receive frequencies of 14.3 and 11.5 GHz, respectively. During MJDs 52926-46, measurements were taken Monday through Friday. During MJDs 53000-17, the intention was to collect TWSTFT measurements seven days per week; however, equipment failure prevented data collection on MJDs 53003-7 as well as on MJD 53014.

During each TWSTFT session, measurements were taken once per second for a total of 2 min. A quadratic fit was then made to the data obtained at each site. The midpoint of each of the quadratic fits was calculated, and then these midpoints were subtracted and divided by two. This yields the time transfer estimate of interest, i.e., the value of the time difference between H2 (or H4) and UTC(NIST).

The GPSCPTT data were collected on a continuous basis, as opposed to several days per week. Both NIST and PTB were equipped with dual-frequency geodetic-type receivers and choke-ring antennae. GPS measurements from all satellites in view were recorded every 30 s on both the L1 (1575 MHz) and L2 (1227 MHz) carrier frequencies. The data were analyzed as is described in the following section.

III. ANALYSIS OF THE GPSCPTT DATA

A. Estimation Procedures

The data were analyzed in 24-h batches using *GIPSY[1]* software provided by the Jet Propulsion Laboratory [3]. Each batch was started at 14:00 GPS time (GPST; 13:59:47 UTC) in order to facilitate comparison with the TWSTFT results. In order to phase-connect consecutive batches, the data were analyzed a second time in batches that started at 2:00 GPST. (The procedure used in phase-connecting the batches is described in Section IV.A.) Satellite orbits were obtained from the International GPS Service (IGS); earth orientation parameters were obtained from *IERS Bulletin B*.

When measurements are made of the GPS carrier-wavelength signal, i.e., on the L1 and/or L2 frequencies, it is not possible to record the integer number of wavelengths that initially lie between the tracking receiver and the GPS satellite. This unknown number of integer wavelengths that biases each

---

[1] A specific trade name is used for identification purposes only; no endorsement is implied.

arc of station-satellite data is known as the "carrier-phase ambiguity." It has been shown that if one can "resolve" these ambiguities, i.e., estimate the number of cycles that biases each arc and fix these numbers to their integer values, then both the geodetic positions and the receiver clock parameters estimated from GPS carrier-phase data will be strengthened [2, 4].

GIPSY uses double-difference techniques to resolve ambiguities, where a double-difference L_{ij}^{12} is defined as $L_{ij}^{12} = (L_i^1 - L_i^2) - (L_j^1 - L_j^2)$, where L_i^1 refers to a GPS measurement taken at site i of the signal from GPS satellite 1. Due to the limited amount of common-view satellite coverage, few double-differences can be formed directly between NIST and PTB. Therefore, to facilitate ambiguity resolution, we added GPS carrier-phase measurements recorded at the IGS stations ALGO (Algonquin Park, Canada), AMC2 (Colorado Springs, Colorado), NRC1 (Ottawa, Canada), and POTS (Potsdam, Germany) to our analysis. Like NIST and PTB, these sites are equipped with dual-frequency geodetic-type receivers and choke-ring antennae. The resulting network is shown in Fig. 1.

To begin the analysis, the data files from all six receivers were edited for cycle slips. The carrier-phase data were then decimated from a sampling rate of one data point every 30 s to a rate of one point every 5 min. The delay due to the ionosphere was then removed by forming the "ionosphere-free" linear combination of the L1 and L2 data [5]. Once this pre-processing was complete, the estimation was performed. Parameters estimated included (a) the coordinates of the antenna at each GPS site, (b) the delay of the GPS signal through the troposphere (again, evaluated at each site), (c) the time difference (i.e., the "receiver clock offset") between each of the frequency standards used to run the GPS receivers and a reference clock, and (d) the time difference ("satellite clock offset") between each of the clocks on the GPS satellites and the reference clock. These parameters were estimated as follows:

Antenna coordinates: We estimated one set of coordinates for the NIST GPS antenna and one set of coordinates for the PTB GPS antenna from each 24-h batch of data.

The IGS produces estimates of the coordinates of the ALGO, AMC2, NRC1 and POTS antennae for each GPS week; these are known as "SINEX" coordinates. We felt that these estimates would be superior to the estimates that we could produce using our six-station network; therefore, we would have preferred to fix the coordinates of these antennae to their SINEX values rather than estimating them. However, due to limitations in the *GIPSY* software, this was not possible. Therefore, we did estimate coordinates for the ALGO, AMC2, NRC1 and POTS antennae from each 24-h batch of data, but we constrained these estimates to be within 0.1 mm of the SINEX values. Thus, the coordinates of these antennae essentially did not change. (In contrast, the coordinates of the NIST and PTB antennae were estimated with *a priori* uncertainties of 10 m, i.e., the estimates of these coordinates were constrained only by the data.)

Troposphere delay: The amount by which a GPS signal is delayed as it passes through the troposphere can be expressed as the sum of two parts: a "wet" part caused by the dipole component of water vapor refractivity and a "hydrostatic" part
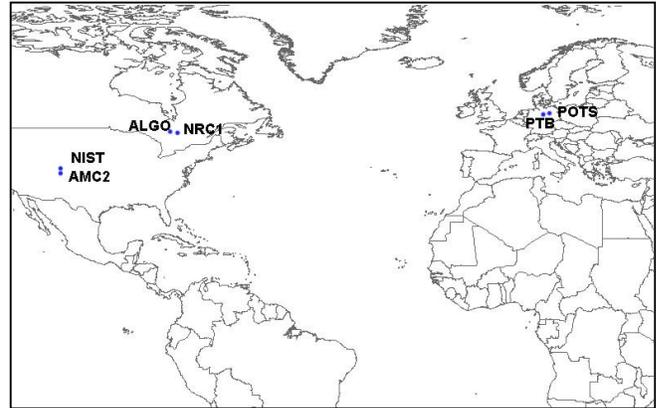
Figure 1. GPS stations used in computing the GPSCPTT results.

caused by dry gases and the non-dipole component of water vapor refractivity [6]. The wet part varies rapidly in time, whereas the hydrostatic part varies slowly.

In GPS data analysis, a separate delay is not estimated for each of the satellite signals reaching a GPS receiver. Rather, the data recorded at that site are used to compute a hypothetical "zenith path delay" for the site; a mapping function is then used to convert the zenith path delay to a path delay value that is appropriate for each satellite given its elevation and azimuth.

The wet part of the zenith troposphere delay was estimated at each site as a time-varying random-walk parameter. New values were estimated every 5 min; consecutive estimates were constrained by 1.7×10^{-7} km/$\sqrt{s}$ (approximately 10.2 mm/$\sqrt{hr}$).

Because the hydrostatic part of the zenith troposphere delay varies slowly in time, it was not estimated. Rather, each site was assigned a fixed value based on height. Therefore, if the hydrostatic part of the troposphere delay changed during the day, that change was expressed as a change in the estimated value of the wet zenith troposphere delay.

The delay through the troposphere was assumed to be azimuthally symmetric; the Niell mapping function was used to model the elevation-angle dependence [7].

Satellite- and receiver-clock offsets: The hydrogen maser at NIST was chosen to be the reference clock for our system; therefore, the variations of all of the other receiver and satellite clocks in the system were estimated relative to it. Because this maser is steered to UTC(NIST), these clock estimates will be relative to UTC(NIST) as well.

The time difference between each satellite and/or receiver clock and UTC(NIST) was estimated once every 5 min as a white noise parameter. By "white noise" we mean that these parameters were modeled as if the value obtained at one epoch were completely independent of the value obtained at the next. A process-noise sigma of 1 s (effectively infinite) was used to reset the clock estimates at every 5-min epoch, so that the clock estimates were constrained only by the data.

This estimation process yields the quantity of interest, i.e., it yields a time series of time transfer estimates for H2 (or H4)–UTC(NIST). Each 24-h batch of data generates 288 time transfer estimates spaced at 5-min intervals.

B. Additions to the GPS Model

The time difference between H2 (or H4) and UTC(NIST) is estimated once every 5 min as a white noise parameter. This means that these time transfer estimates are unfortunately very good at absorbing the errors introduced by unmodeled effects. In addition, the GPS receivers at NIST and PTB are located 7532 km apart, which means that little common-mode error cancellation can occur. Hence, it is important to make the satellite-station range model as accurate as possible. Therefore, we have incorporated the following additions to standard GPSCPTT analysis technique:

Ocean loading: The motion of the ocean tides can cause a point on the earth's surface to move up and down with an amplitude of nearly a centimeter. This effect is known as "ocean loading"; on a long baseline such as NIST-PTB, near new moon or full moon, it can alter the frequency of a GPSCPTT-derived time series by nearly 10^{-15}.

To compensate for this, we turned on the ocean-loading model in *GIPSY* and supplied the model with coefficients of force for each site. The ocean-loading coefficients were computed using the GOT00.2 model [8] via the on-line program at the Onsala Space Observatory's website, http://www.oso.chalmers.se/~loading.

Center-of-mass-to-phase-center corrections: There is a spatial offset between the center of mass of a GPS satellite and the electrical phase center of the antenna from which it transmits. This offset must be accounted for in the modeling. The *GIPSY* model and the IGS use the same values for the offsets in Block I and Block IIA GPS satellites. However, in computing their orbits, the IGS assumes an offset of zero for the Block IIR satellites, whereas the *GIPSY* model uses a value of -63 cm along the z axis, where the z axis is defined as pointing from the center of mass of the satellite to the center of mass of the earth. (In other words, the "z" direction is exactly parallel to and opposite in sign of what we conventionally think of as the radial direction when it comes to orbits.)

In the past, it was believed that if one were estimating the time difference ("clock offset") between each of the satellite clocks and a reference clock, it wouldn't be necessary to correct for this discrepancy in the Block IIR modeling. That is because it was thought that the estimates of the Block IIR clock offsets would simply re-absorb this error. However, a satellite clock parameter cannot perfectly absorb what is essentially a radial orbit error because these effects have unit vectors that are not quite parallel: a clock parameter points along the line of sight between the station and the satellite, whereas a radial orbit error points along a line between the center of the earth and the satellite.

Because the GPS satellites are so far away from the earth, the two partial derivatives differ by a maximum of about 3 % when a minimum elevation angle of 15° is used. However, that is enough to noticeably change the frequency of the GPS-derived H2 (or H4)–UTC(NIST) time transfer estimates: when we tested the effect of turning the correction on and off, we observed an effect of over 120 ps/d, with similar results obtained for the NIST-POTS baseline. (Tests performed on the shorter baselines ALGO-NIST and NRC1-NIST showed a smaller effect, approximately 20 ps/d. This is expected, because the effect of a radial orbit error scales with the length of the baseline.) As a result, we decided to correct the Block IIR phase-center-to-center-of-mass discrepancy between the IGS orbits and *GIPSY*, because not to do so leaves a 63-cm radial orbit error unmodeled.

IV. COMPUTING FREQUENCY FROM THE GPSCPTT AND TWSTFT RESULTS

A. Concatenating Batches of GPSCPTT Time Transfer Estimates: GPSmerge

A 24-h batch of data generates a 288-point set of time transfer estimates that represents the time difference between the PTB maser and UTC(NIST). In this, the more-traditional method, we concatenate consecutive sets of time transfer estimates to form a single continuous time series. The frequency is then obtained from this continuous time series. The concatenation is performed as described below and as is shown in Fig. 2. For the remainder of this paper, we shall refer to this method as the *GPSmerge* method.

We wish to concatenate consecutive batches which start at 14:00 GPST. We shall refer to these batches as Day 1 and Day 2. The discontinuity between these batches occurs between the end of the first batch (13:55) and the beginning of the next (14:00). To estimate the size of the discontinuity—and thus remove it—we employ a third 24-h batch of GPS data that starts between these two batches at 2:00 GPST. This third time series straddles the discontinuity between Day 1 and Day 2 and provides a continuous set of time transfer estimates across it.

The size of the discontinuity is computed as follows: let d_1 represent the constant offset between the time transfer estimates of Day 1 and those of the "transfer" batch which starts at 2:00. Let d_2 represent the constant offset between the time transfer estimates of Day 2 and those of the transfer batch. The difference $d_2–d_1$ estimates the size of the discontinuity between the time transfer estimates of Day 1 and Day 2.

d_1 is computed by subtracting the time transfer values obtained from the 2:00 batch from those obtained from the Day 1 batch over the interval 5:00-10:55 GPST and then computing the mean of these values. d_2 is computed by subtracting the time transfer values obtained from the 2:00 batch from those obtained from the Day 2 batch over the interval 17:00-22:55 GPST and then computing the mean of those values.

Once the value of $d_2–d_1$ is obtained, the time transfer estimates of Day 2 are shifted by the appropriate amount. This process is repeated forward through the data set, with each computation joining an unconnected 288-point time series to the ever-growing phase-connected time series preceding it. Thus, we obtain a continuous set of time transfer estimates. Note that the above procedure assumes that there is a negligible difference in frequency between the time transfer estimates obtained from Day 1 and those obtained from the transfer batch over the hours 5:00-10:55 GPST; the same assumption is made for the time transfer estimates obtained from the Day 2 batch and those obtained from the transfer batch over the hours 17:00-22:55 GPST.

Estimate of discontinuity = $d_2 - d_1$.

Day 1

Day 2

Original time transfer
results for Day 2

14:00

14:00 GPST   5:00   11:00

Time transfer
results for Day 1

d_1 d_2 Phase-connected time
transfer results for Day 2

2:00 GPST   17:00   23:00

Batch run from 2:00-1:55 GPST used
to connect Day 1 and Day 2 results
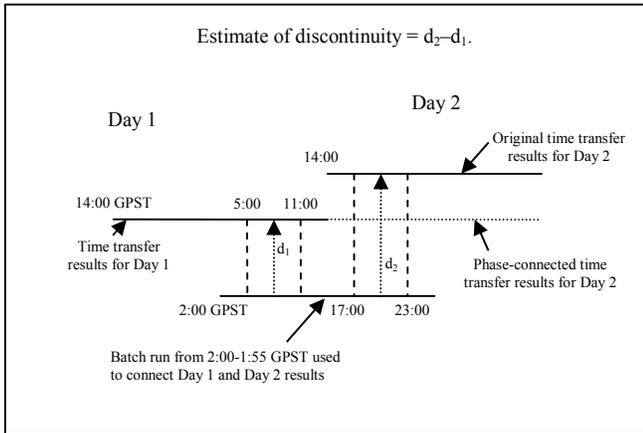
Figure 2. Concatenation method used to phase-connect consecutive batches of GPSCPTT data starting at 14:00 GPST. Each solid line represents a set of time transfer values obtained from a 24-h batch of GPS processing. A "transfer" batch is run from 2:00-1:55 GPST in order to connect the batches starting at 14:00. d_1 is computed by subtracting the time transfer values obtained from the 2:00 batch from those obtained from the Day 1 batch over the interval 5:00-10:55 GPST and then computing the mean of these values. d_2 is computed by subtracting the time transfer values obtained from the 2:00 batch from those obtained from the Day 2 batch over the interval 17:00-22:55 GPST and then computing the mean of these values. Finally, the Day 2 values are shifted by the quantity $d_2 - d_1$. This process is repeated forward through the entire data set.

Once the continuous time series is obtained, the computation of a frequency becomes straightforward. The simplest way to compute a mean frequency over a time interval—and the method that is optimal for white FM noise—is to subtract the time transfer estimate at the beginning of the interval from the time transfer estimate at the end of the interval and to divide this quantity by the time interval between the points. We perform a procedure that is almost identical. Our time transfer estimates are separated by 5-min intervals. As we shall show in *Results and Discussion*, the phase stability of our results is slightly better at an averaging time of 10 min than it is at an averaging time of 5 min. Therefore, to compute the mean frequency over an epoch, we average two time transfer estimates at the beginning of the epoch to form one endpoint and two time transfer estimates at the end of the epoch to form the other endpoint. We then use these endpoints to compute the frequency as described above.

B. Estimating Frequency Directly from Daily GPSCPTT Batches: GPSave

This, our new method, is simple in comparison. We shall refer to it as the *GPSave* method.

Each 24-h batch of data produces a 288-point set of time transfer estimates representing the time difference between H2 (or H4) and UTC(NIST). To compute a frequency from this batch, we average the first two time transfer estimates together to form one endpoint and the last two time transfer estimates together to form another endpoint. We subtract these endpoints and divide by the time between them to obtain the frequency for that batch. A frequency derived in this manner will have an averaging time of 23 h 50 min.

This process is performed for each 24-h batch that starts at 14:00 GPST. To obtain a mean value for the frequency of H2 (or H4)–UTC(NIST) over a multi-day interval, we average the individual frequency values contained within that interval.

C. Computing Frequency from TWSTFT Results

To compute the mean frequency of H2 (or H4)–UTC(NIST) over a given interval, we subtract the time transfer estimate at the beginning of the interval from that at the end of the interval and divide by the time between the points.

V.   RESULTS AND DISCUSSION

Fig. 3 shows the frequency of the maser at PTB relative to UTC(NIST) as computed from GPSCPTT and TWSTFT. The GPSCPTT values were obtained using the *GPSave* method. The TWSTFT values were obtained by subtracting adjacent TWSTFT time transfer estimates and dividing by the interval between them. As Fig. 3 shows, good agreement is obtained between the GPSCPTT and TWSTFT results. It is interesting that both GPSCPTT and TWSTFT show a sharp change in the relative frequency of H2-UTC(NIST) at about MJD 52934.

As Fig. 3a shows, there is a gap in the GPSCPTT values for MJDs 52942-43. The ionosphere was extremely active during these days, causing outages at some of the IGS sites and rendering the GPSCPTT results nearly unusable. We did not phase-connect the GPSCPTT data past MJD 52941.4 due to a complete loss of results for MJDs 52941.8-52942.0 and a jump/data loss at MJD 52941.4. However, the data became usable again during MJD 52943, so we computed daily frequency values for the remainder of the measurements.

Fig. 4 shows the phase stability (TDEV) of the results obtained, and Fig. 5 shows the frequency stability as measured by the Allan deviation (ADEV). The TDEV plot confirms the decision to average adjacent 5-min time transfer estimates when creating endpoints. The ADEV plot shows that it is appropriate to use the endpoint method when computing a frequency from each of the 24-h batches, because the noise type at one-day averaging times is consistent with white FM.

We now examine how well the frequencies obtained from the *GPSave*, *GPSmerge* and TWSTFT techniques agree for averaging times of 4 d and longer.

We do not wish to interpolate between the TWSTFT time transfer estimates. Therefore, for a given data set (e.g., MJDs 52926-41), we list all subsets of the TWSTFT data that have a given averaging time. For example, for MJDs 52926-41, there are five sets of TWSTFT time transfer estimates that have an averaging time of 4 d: 52928-32, 52929-33, 52932-36, 52935-39 and 52936-40. We then compute the frequency y(H2–UTC(NIST)) for each of these epochs using TWSTFT, *GPSave* and *GPSmerge*. For example, for MJDs 52929-33, the value of y(H2–UTC(NIST)) obtained from TWSTFT is $-53.51 \cdot 10^{-15}$, that obtained from *GPSave* is $-53.88 \cdot 10^{-15}$, and that obtained from *GPSmerge* is $-54.00 \cdot 10^{-15}$. We then compute the difference between the frequencies obtained from each of the methods for a given epoch. For example, for MJDs 52929-33, y(*GPSave*)–y(TWSTFT) = $(-53.88 + 53.51) \cdot 10^{-15} = -0.37 \cdot 10^{-15}$. Similarly, y(*GPSmerge*)–y(TWSTFT) and y(*GPSave*)–
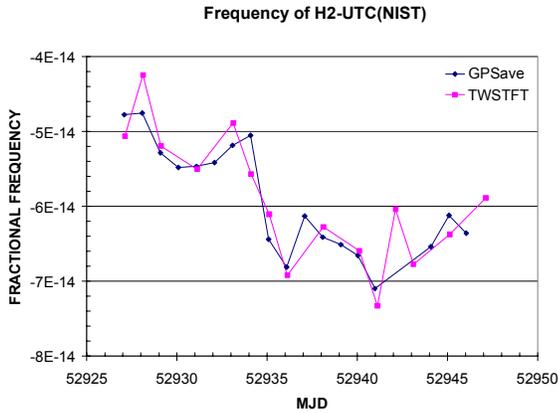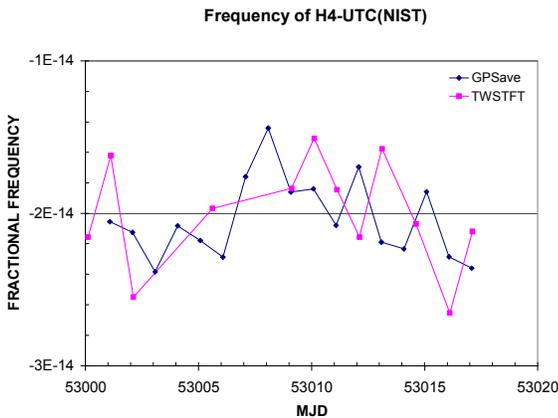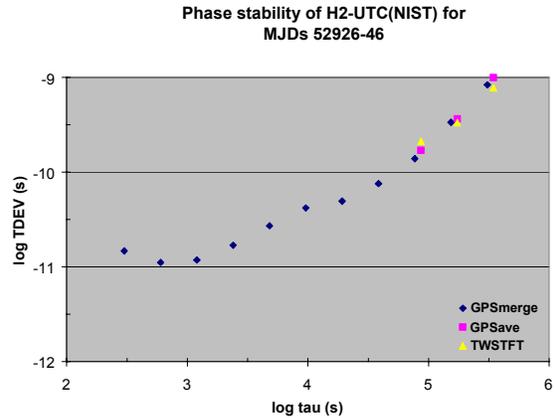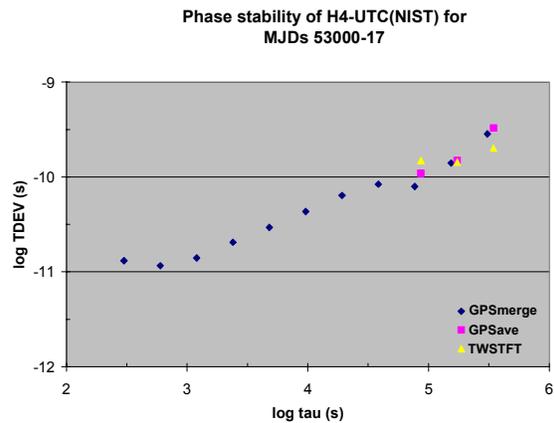
Figure 3. Frequencies derived from GPSCPTT and TWSTFT. The GPSCPTT values were computed using the *GPSave* method, i.e., each frequency shown was computed from a single 24-h batch of data. The TWSTFT values were obtained by subtracting adjacent time transfer values and dividing by the interval between them.

y(*GPSmerge*) = -0.49·10^{-15} and 0.12·10^{-15}, respectively. This yields one set of estimates for the agreement among the three methods for an averaging time of 4 d; the process is then repeated for the other four epochs with averaging times of 4 d, i.e., for MJDs 52928-32, 52932-36, etc. Thus, for MJDs 52926-41 and an averaging time of 4 d, we obtain five values of y(*GPSave*)–y(TWSTFT), five of y(*GPSmerge*)–y(TWSTFT) and five of y(*GPSave*)–y(*GPSmerge*). To obtain a final value for the agreement of the methods for this averaging time (4 d) and this data set (MJDs 52926-52941), we compute the simple mean of the five values of y(*GPSave*)–y(TWSTFT), and do the same for the five values of y(*GPSmerge*)–y(TWSTFT) and of y(*GPSave*)–y(*GPSmerge*). This is then repeated for averaging times of 5 d, 6 d, etc.

As Table 1 shows, for MJDs 52926-41, the frequencies obtained using the two GPSCPTT techniques agree at approximately 2·10^{-16} for all averaging times. For this set of data, the *GPSave* technique shows better agreement with the results from TWSTFT than does the *GPSmerge* technique, with the *GPSave* and TWSTFT results agreeing at roughly 3·10^{-16} for all averaging times. While the frequencies obtained from the *GPSmerge* technique do not agree quite as well with the TWSTFT results as do those obtained from the *GPSave*

Figure 4. Phase stability of the GPSCPTT results. The values obtained for the TWSTFT results are included for completeness; however, they should not be used due to the amount of interpolation required to obtain them.

technique, the agreement between these values and TWSTFT is still good: at worst, we see agreement at the 6·10^{-16} level.

Table 2 shows the results for MJDs 53000-17. The frequencies obtained from the two GPS techniques agree with each other at about 3·10^{-16} for all averaging times. In this case, the frequencies obtained from the *GPSmerge* technique show better agreement with those obtained from TWSTFT than do those obtained from the *GPSave* technique. The agreement between the *GPSmerge* and TWSTFT results is remarkable, with the average agreement dropping below 10^{-16} for some averaging times. Frequencies obtained from the *GPSave* technique agree with those obtained from TWSTFT at approximately 4-6·10^{-16}.

While Tables 1 and 2 show good agreement among the various methods of computing frequency, they also show that further work is needed to eliminate systematic errors. For MJDs 52926-41, on average, the frequency obtained from TWSTFT tended to be slightly higher than that obtained from the *GPSave* method, and the frequency obtained from the *GPSave* method tended to be slightly higher than that obtained from the *GPSmerge* method. During MJDs 53000-17, the frequency obtained using TWSTFT tended to be slightly higher than that obtained using the *GPSmerge* technique, and the

Frequency stability of H2-UTC(NIST) for MJDs 52926-46

5a

Frequency stability of H4-UTC(NIST) for MJDs 53000-17
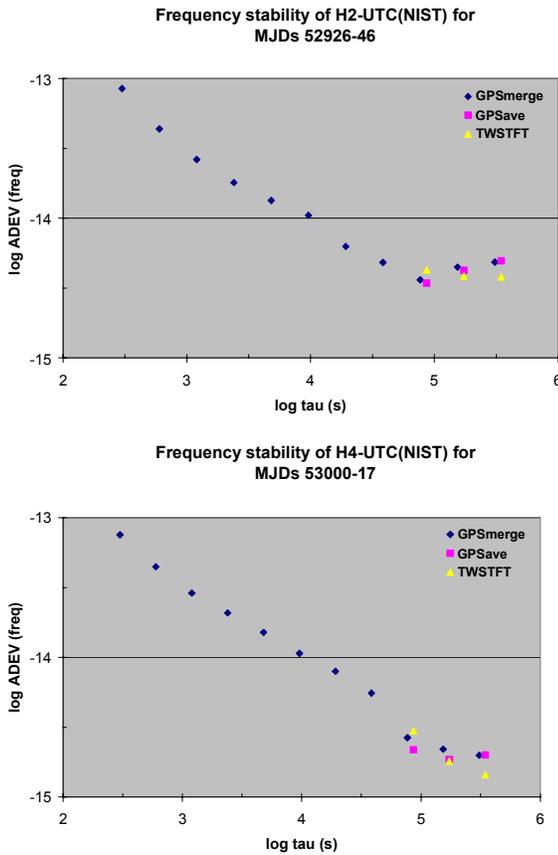
5b

Figure 5. Frequency stability of the GPSCPTT results. The values obtained for the TWSTFT results are included for completeness; however, they should not be used due to the amount of interpolation required to obtain them.

TABLE I. AGREEMENT OF FREQUENCY VALUES OBTAINED FROM *GPSave*, *GPSmerge* AND TWSTFT FOR MJDs 52926-52941

averaging time (d)	number of values averaged	average value, y(*GPSave*)-y(TWSTFT) (units: 10^{-15})	average value, y(*GPSmerge*)-y(TWSTFT) (units: 10^{-15})	average value, y(*GPSave*)-y(*GPSmerge*) (units: 10^{-15})
4	5	-0.250	-0.410	0.159
5	5	-0.380	-0.619	0.240
6	6	-0.291	-0.511	0.220
7	6	-0.198	-0.348	0.150
8	4	-0.234	-0.392	0.157
9	2	-0.407	-0.575	0.168
10	2	-0.125	-0.252	0.127
11	2	-0.145	-0.270	0.125
12	2	-0.383	-0.542	0.159
13	2	-0.438	-0.623	0.185
14	1	-0.328	-0.542	0.214

TABLE II. AGREEMENT OF FREQUENCY VALUES OBTAINED FROM *GPSave*, *GPSmerge* AND TWSTFT FOR MJDs 53000-53017

averaging time (d)	number of values averaged	average value, y(*GPSave*)-y(TWSTFT) (units: 10^{-15})	average value, y(*GPSmerge*)-y(TWSTFT) (units: 10^{-15})	average value, y(*GPSave*)-y(*GPSmerge*) (units: 10^{-15})
4	5	-0.490	-0.114	-0.376
5	4	-0.630	-0.420	-0.210
6	4	-0.394	-0.062	-0.332
7	5	-0.416	-0.093	-0.323
8	5	-0.450	-0.193	-0.258
9	3	-0.563	-0.220	-0.343
10	3	-0.542	-0.224	-0.318
11	3	-0.610	-0.295	-0.314
12	2	-0.459	-0.198	-0.261
13	2	-0.805	-0.500	-0.305
14	2	-0.444	-0.094	-0.350
15	3	-0.476	-0.148	-0.328
16	2	-0.352	-0.045	-0.307
17	1	-0.530	-0.235	-0.295

frequency obtained using the *GPSmerge* technique tended to be slightly higher than that obtained from the *GPSave* technique. These effects are small, about $1\text{-}3\cdot10^{-16}$ apiece, but persistent.

It is possible that when analyzing the data from MJDs 52926-41, a systematic error in frequency was introduced by phase-connecting the data, and that the *GPSave* method produced frequencies that showed better agreement with those obtained from TWSTFT because it sidestepped this error. Fig. 6 shows the evolution in time of the difference between the concatenated GPS time transfer estimates, i.e., those used in the *GPSmerge* method, and those obtained from TWSTFT. If the phase-connection process were perfect and if GPSCPTT and TWSTFT were measuring the same quantities, we would expect these plots to have slopes of zero. As expected, there appears to be little overall trend in the data for MJDs 53000-17. However, the plot for MJDs 52926-41 shows an obvious slope, which equates to a frequency offset between the *GPSmerge* and TWSTFT methods. If this offset were caused by errors in the phase-connection process, then it would make sense that the *GPSave* method produced frequency values that showed superior agreement with those obtained from TWSTFT.

It is also worth noting that during both MJDs 52926-41 and 53000-17, the frequency derived from TWSTFT tended to be slightly higher than the frequency derived from either of the

GPSCPTT methods. Phase-connection errors aside, it is possible that during these periods, there was a small frequency offset between the two time transfer systems, and that whichever of the two GPSCPTT methods yielded a higher frequency would be the one that showed closer agreement with TWSTFT.

As stated previously, the PTB fountain CSF1 was evaluated during MJDs 52929-44 and 52999-53014. It is thus possible to use our results to compute the frequency y(CSF1–UTC(NIST)) for these periods, and then to compare these values to those that can be obtained from *BIPM Circular T* [9-12].

[H2-UTC(NIST)](GPSmerge) – [H2-UTC(NIST)](TWSTFT)   6a

[H4-UTC(NIST)](GPSmerge) – [H4-UTC(NIST)](TWSTFT)   6b
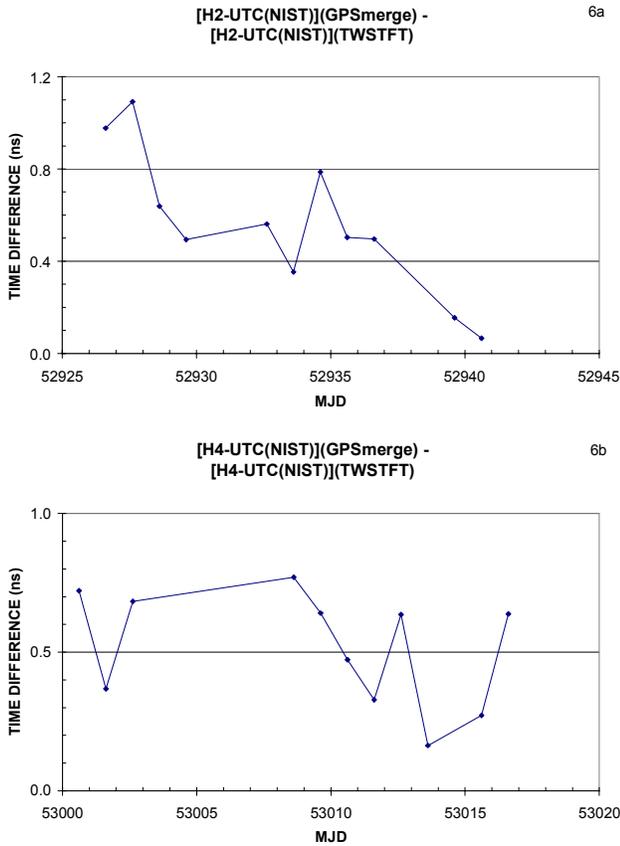
Figure 6. Evolution in time of the difference between the concatenated GPS time transfer estimates, i.e., those used in the *GPSmerge* method, and the time transfer estimates obtained from TWSTFT. The plot for MJDs 53000-17 has little or no trend, as one would expect; however, the plot for MJDs 52926-41 shows a non-zero slope. This may indicate that the concatenation process has introduced a frequency offset into the GPSCPTT results; see text for details.

To compute y(CSF1–UTC(NIST)) from our data, we use the equation

y(CSF1–UTC(NIST)) =
  y(CSF1–H2(or H4)) + y(H2 (or H4))–UTC(NIST)).   (1)

The first quantity on the right-hand side is measured by PTB during the fountain evaluation: for MJDs 52929-44, y(CSF1–H2) = $71.4 \cdot 10^{-15}$, and for MJDs 52999-53014, y(CSF1–H4) = $31.6 \cdot 10^{-15}$ [13]. The second quantity on the right-hand side is the mean frequency we obtain from GPSCPTT or TWSTFT.

To compute y(CSF1–UTC(NIST)) from *BIPM Circular T*, we use

y(CSF1–UTC(NIST)) =
  y(CSF1–UTC) + y(UTC–UTC(NIST)).   (2)

y(CSF1–UTC) is equal to the scale interval "d" reported in section 4 of *BIPM Circular T*; the value of d reported for the 52929-52944 evaluation was $11.5 \cdot 10^{-15}$ and the value reported

for the 52999-53014 evaluation was $13.2 \cdot 10^{-15}$ [9, 10]. y(UTC–UTC(NIST)) can be calculated using the values of the time series UTC–UTC(NIST) reported in section 1 of *Circular T*. Using the values of UTC–UTC(NIST) reported for MJDs 52929 and 52944, we obtain an average frequency y(UTC–UTC(NIST)) of $1.6 \cdot 10^{-15}$; using the values for MJDs 52929 and 52939 yields y(UTC–UTC(NIST)) = $0.5 \cdot 10^{-15}$ [9, 11]. Using the values of UTC–UTC(NIST) reported for MJDs 52999 and 53014, we obtain y(UTC–UTC(NIST)) = $0.2 \cdot 10^{-15}$ [10, 12].

Combining these values from *Circular T*, we obtain estimates of y(CSF1–UTC(NIST)) = $13.1 \cdot 10^{-15}$, $12.0 \cdot 10^{-15}$ and $13.4 \cdot 10^{-15}$ for MJDs 52929-44, 52929-39 and 52999-53014, respectively. Because the value of y(CSF1–UTC) alone has an uncertainty of about $2.5 \cdot 10^{-15}$, we can assume that the frequencies we have derived from *Circular T* will be uncertain by at least this amount.

Table 3 shows the agreement of y(CSF1–UTC(NIST)) as computed from GPSCPTT and TWSTFT with that computed from *Circular T* for MJDs 52929-52944. Because we could not phase-connect the data past 52941.4, we used the mean frequency computed for the interval 52928.6-52940.6 and compared that to the *Circular T* result obtained for 52929-39. (The ".6" part of the MJD is present because the TWSTFT measurements were made/the GPS batches were started at approximately 14:00 UTC.) As Table 3 shows, our values agree with the values from *Circular T* at 2-$6 \cdot 10^{-16}$, well within the uncertainty of $2.5 \cdot 10^{-15}$ stated above.

We were not able to compute frequency values for the 24-h GPS batches that started at MJDs 52941.6 and 52942.6. However, we were able to compute a frequency value from the 24-h batch that started at 52943.6. Therefore, we computed the mean value of y(H2–UTC(NIST)) for MJDs 52928.6-44.6, despite the fact that we were missing frequency information from the middle of the epoch. We used TWSTFT to compute an average frequency for MJDs 52928.6-52943.6. Table 3 shows the agreement of these values with the value from *Circular T* for MJDs 52929-44. These results display a slightly worse agreement with *Circular T*, about -2.1×10^{-15}, than those obtained for MJDs 52928.6-52940.6. However, this is still within the uncertainty of the value from *Circular T*.

Table 4 shows the agreement with *Circular T* of the values obtained for MJDs 52999-53014. It is unfortunate that during this interval, the GPS values showed some of the worst agreement with TWSTFT of this entire experiment; thus, the mean frequency values obtained for this epoch have a spread of $0.8 \cdot 10^{-15}$. Nonetheless, all three values of y(CSF1–UTC(NIST)) are within $1.8 \cdot 10^{-15}$ of the value from *Circular T*.

VI.   FUTURE WORK

The frequency values obtained using GPSCPTT and TWSTFT agree well. However, the uncertainty with which cesium fountains realize the SI second is constantly improving. Therefore, it would be useful to obtain even better agreement between the frequency values obtained from GPSCPTT and those obtained from TWSTFT. Toward that end, it would be helpful to understand why the mean frequency values obtained from GPSCPTT tended to be slightly lower than those obtained from TWSTFT. This issue could be addressed by analyzing

A. y(CSF1–UTC(NIST)) from Circular T = 12.0·10⁻¹⁵ for MJDs 52929-39. B. y(CSF1–UTC(NIST)) from Circular T = 13.1·10⁻¹⁵ for MJDs 52929-44. *Parentheses in column 5 indicate which value was used in the comparison.*				
method	averaging epoch	y(H2– UTC(NIST)) (units: 10^{-15})	y(CSF1– UTC(NIST)) (units: 10^{-15})	y(method)– y(Circular T) (units: 10^{-15})
TWSTFT	52928.6 - 52940.6	-58.8	12.6	0.6 (A)
GPSave	52928.6 - 52940.6	-59.0	12.4	0.4 (A)
GPSmerge	52928.6 - 52940.6	-59.2	12.2	0.2 (A)
TWSTFT	52928.6- 52943.6	-60.5	10.9	-2.2 (B)
GPSave	52928.6- 52944.6	-60.4	11.0	-2.1 (B)

y(CSF1–UTC(NIST)) from Circular T = 13.4·10⁻¹⁵ for MJDs 52999-53014				
method	averaging epoch	y(H4– UTC(NIST)) (units: 10^{-15})	y(CSF1– UTC(NIST)) (units: 10^{-15})	y(method)– y(Circular T) (units: 10^{-15})
TWSTFT	53000.6 - 53013.6	-19.2	12.4	-1.0
GPSave	53000.6 - 53013.6	-20.0	11.6	-1.8
GPSmerge	53000.6 - 53013.6	-19.8	11.8	-1.6

additional sets of data to see whether this frequency offset persists. If it does, then its source should be investigated.

It would also be useful to know why the mean frequency values obtained using the *GPSave* method were consistently higher or lower than those obtained using the *GPSmerge* method. To begin, the discontinuities between batches should be examined to see if there is anything systematic or cumulative about them. Perhaps the method by which the time transfer estimates are concatenated can be improved. Finally, we may need to optimize the method by which the frequency is computed from each of the 24-h batches.

VII.   CONCLUSIONS

A new technique for computing frequency from GPSCPTT data is proposed and tested. In the conventional method, consecutive batches of GPSCPTT-derived time transfer estimates are concatenated by estimating and removing the discontinuities between them; frequency values are then computed from the resulting continuous time series. In the new method, a frequency value is computed from each batch; to obtain the mean frequency over a multi-day interval, the frequency values contained within that interval are averaged.

The new technique was tested by comparing the mean frequency values it yields with those obtained using the conventional method and those obtained using TWSTFT. These tests were performed over MJDs 52926-46 and 53000-53017; the clocks compared were located at PTB in Braunschweig, Germany and at NIST in Boulder, Colorado.

Although further study is needed to investigate the source of low-level systematic errors, the results obtained were promising. The frequency values obtained from the GPSCPTT data using the new technique agreed with those obtained using the conventional technique at 2-3·10^{-16}. Thus, it may be feasible to replace the conventional technique with the new technique. Furthermore, the frequency values obtained from the GPSCPTT data using either of the two techniques agreed with those obtained from the TWSTFT data at a few to several parts in 10^{16}. One would expect little common-mode error cancellation between GPSCPTT and TWSTFT. This implies that both GPSCPTT and TWSTFT are approaching a level of frequency-transfer performance sufficient for use during simultaneous fountain evaluations.

REFERENCES

[1] K.M. Larson, J. Levine, and L.M. Nelson "Assessment of GPS carrier-phase stability for time-transfer applications," *IEEE Trans. UFFC*, vol. 47, pp. 484-493, 2000.

[2] R. Dach, G. Beutler, U. Hugentobler, S. Schaer, T. Schildknecht, T. Springer, G. Dudle, and L. Prost, "Time transfer using GPS carrier phase: error propagation and results," *J. Geodesy*, vol. 77, pp. 1-14, 2003.

[3] F.H. Webb and J.F. Zumberge, editors, *An Introduction to GIPSY/OASIS-II: Precision Software for the Analysis of Data from the Global Positioning System*, JPL D-11088 (internal publication), Jet Propulsion Laboratory, Pasadena, CA, July 1997.

[4] G. Blewitt, "Carrier-phase ambiguity resolution for the Global Positioning System applied to geodetic baselines up to 2000 km," *J. Geophys. Res.*, vol. B94, pp. 10,187-10,203, 1989.

[5] K.M. Larson and J. Levine, "Carrier-phase time transfer," *IEEE Trans. UFFC*, vol. 46, pp. 1001-1012, 1999.

[6] Y.E. Bar-Sever, P.M. Kroger, and J.A. Borjesson, "Estimating horizontal gradients of tropospheric path delay with a single GPS receiver," *J. Geophys. Res.*, vol. B103, pp. 5019-5035, 1998.

[7] A.E. Niell, "Global mapping functions for the atmospheric delay at radio wavelengths," *J. Geophys. Res.*, vol. B101, pp. 3227-3246, 1996.

[8] R.D. Ray, "A global ocean tide model from TOPEX/POSEIDON altimetry: GOT99.2," *NASA Technical Memorandum 209478*, 1999.

[9] *BIPM   Circular   T   191*; published   on-line   at ftp://62.161.69.5/pub/tai/publication/cirt.191

[10] *BIPM   Circular   T   193*; published   on-line   at ftp://62.161.69.5/pub/tai/publication/cirt.193

[11] *BIPM   Circular   T   190*; published   on-line   at ftp://62.161.69.5/pub/tai/publication/cirt.190

[12] *BIPM   Circular   T   192*; published   on-line   at ftp://62.161.69.5/pub/tai/publication/cirt.192

[13] S. Weyers, personal communication, 2004.