New IGS Clock Products
A Global Time Transfer Assessment

Jim Ray National Geodetic Survey
Ken Senior Naval Research Laboratory

The International GPS Service (IGS) has a new suite of clock products available and is continuing to improve their usefulness for practical time and frequency transfer applications. In this month's column, Jim Ray and Ken Senior describe these IGS clock products, use internal repeatability analyses to assess their potential accuracy and stability limits, and compare them with the emerging requirements of the timekeeping community. They conclude that calibration of the internal delays in the GPS receiving equipment will probably continue to set the limit for time transfer accuracy, whereas frequency transfers can already achieve stabilities approaching 10^{-15} *over one-day intervals.*

Dr. Jim Ray *obtained a degree in space physics and astronomy from Rice University in Houston in 1980. After two post-doctoral fellowships, he worked for more than seven years as a member of the NASA Very Long Baseline Interferometry (VLBI) group at Goddard Space Flight Center, Greenbelt, Maryland. He later moved to the Geosciences Laboratory of the National Oceanographic and Atmospheric Administration (NOAA), where he remained for more than five years, first leading the VLBI analysis effort there, then joining the GPS analysis program. Afterwards, he was head of the Earth Orientation Department at the U.S. Naval Observatory (USNO) for five years. Recently he joined NOAA's National Geodetic Survey working on GPS and International GPS Service (IGS) projects. He also serves on the IGS Governing Board and acts as co-chair of the joint IGS/Bureau International des Poids et Mesures Time Transfer Pilot Project.*

Dr. Ken Senior *obtained a degree in applied mathematics from Bowling Green State University in Bowling Green, Ohio, in 1997. He served as an assistant professor of mathematics at the University of Tennessee, Chattanooga, for two years. Afterwards he was a member of the USNO clock development team in the Time Service Department for more than three years. His duties included GPS data analysis and mathematical research support. Recently he joined the Space Applications Branch of the Naval Research Laboratory in Washington, DC. He is also an associate member of the IGS.*

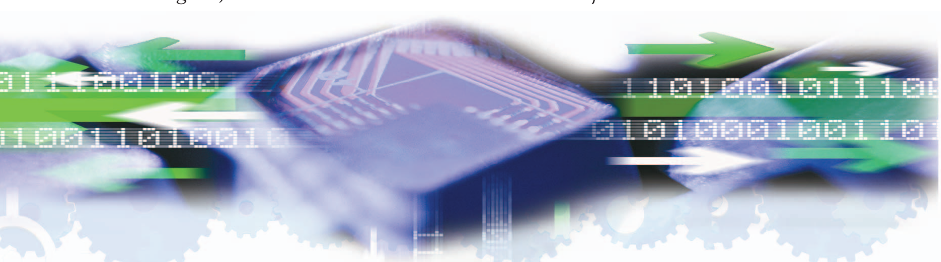

GPS has revolutionized the general accessibility of accurate global time and frequency transfer compared with prior terrestrial broadcast systems and the physical transport of clocks. Since the ending of Selective Availability (SA) on May 2, 2000, even users of inexpensive single-frequency, C/A-code-only GPS receivers now have access to GPS Time at a level around 100 nanoseconds, limited primarily by ionospheric propagation error. With a dual-frequency (usually codeless tracking) receiver, the performance improves to the order of 10 nanoseconds at a known location; the leading error sources in that case are inaccuracies in the broadcast GPS orbits and satellite clocks and pseudorange multipath. GPS Time itself differs from the international Coordinated Universal Time (UTC) time scale by similar amounts, up to 40 nanoseconds or so, ignoring leap-second differences. Point-to-point timing comparisons can be further improved to the few-nanosecond range, even for single-frequency users over intercontinental distances, by coordinating observing schedules, exchanging data sets, and forming single-differences to remove the common-mode satellite errors.

During the past two decades, the Bureau International des Poids et Mesures (BIPM) has relied primarily on this "common-view" method to form UTC based on comparisons between clocks at approximately 50 timing laboratories. Despite this important application, the common-view technique is not well suited for general time dissemination due to its coordinated and differential nature. Moreover, the emerging new generation of ultracold atomic frequency standards, having one-day stabilities of 10^{-15} or better, requires improving time transfer accuracy to the sub-nanosecond range.

Joint Pilot Project. The products of the International GPS Service (IGS) allow users to exploit GPS in an autonomous, non-differential mode (as well as differentially) to deliver user accuracies about 100 times better than the broadcast navigation system. (Users must additionally observe and analyze the pseudorange and carrier phase data at both GPS frequencies.) This means point positioning good to a few centimeters root-mean-square (r.m.s.) at each measurement epoch or to the sub-centimeter level for one-day integrations.

In principle, the corresponding time transfer errors could potentially be well below one nanosecond (note that a three-centimeter distance approximately equals 0.1 nanosecond light travel time). This recognition was the basis for establishing a joint pilot project between the IGS and the BIPM, starting in early 1998, to develop and demonstrate the operational capabilities for time transfer.

Hardware. In any GPS-based time transfer technique, the derived user clock readings apply to an internal point within the tracking system. Geodetic analyses of GPS pseudorange and phase data determine the effective receiver "clock", which is at the ionosphere-free phase center of the antenna but offset by the electrical delay to the point in the receiver tracking loop where the observables are measured and time-tagged. To accurately relate the internal clock values to external timing standards, which is essential for most practical applications, one must first determine the instrumental delays and biases within the receiving hardware chain. In general, these delays will not be constant under all circumstances, which greatly complicates matters.

Present methods to measure absolute instrumental delays are accurate only to a few nanoseconds. Differential calibration procedures, which compare a test receiver to an accepted standard unit, promise much better performance, possibly into the sub-nanosecond range. In 2001, the BIPM began a campaign to pursue this idea by circulating an absolutely calibrated receiver as a standard to differentially measure the biases of similar receivers deployed at timing labs. For frequency transfers (as distinct from time transfers), the instrumental requirements are less stringent in an absolute sense but demand that the GPS hardware be highly insensitive to, or well isolated from, environmental change. Such conditions are not unusual at timing labs but are rarely satisfied at typical geodetic facilities.

IGS Clock Products

Satellite clock values are among the "core" products of the IGS. Since its founding, the service has distributed combined solutions for satellite clocks together with the combined satellite orbits. In late 2000, the IGS also began distributing combined clock estimates for a subset of the global tracking network (see **Figure 1**), as well as for the satellites, both tabulated at five-minute intervals. As many as six independent analysis centers contribute clock determinations, at least two of which are required for each IGS combined value. This method assures quality control.

Required Consistency. The essential requirement for all the IGS clock products is that they be fully consistent at the centimeter level with the accompanying satellite orbits and the terrestrial reference frame (very closely tied to the International Terrestrial Reference Frame 2000 – ITRF2000). This condition ensures that IGS products can be used instead of GPS broadcast information to attain few-centimeter accuracy. An autonomous user can expect such performance using data from a single, isolated GPS receiver (observing dual-frequency pseudoranges and phases) for both position and clock determinations. The position will be expressed within the highly accurate ITRF2000 frame. On the other hand, the clock will be relative to the underlying IGS time scale, which has historically been only coarsely linked to GPS Time by a daily linear alignment to the broadcast satellite clocks. The time scale of the IGS clock products should ideally be accurately traceable to UTC, in the same way that the IGS geodetic reference frame conforms (and significantly contributes) to the ITRF. We have focused our recent efforts on improvements in this respect.

As with all high-accuracy GPS applications, the IGS relies primarily upon dual-frequency carrier phase observations, which are about 100 times more precise than the pseudorange data. For double-differencing analysis, pseudoranges are not normally even used except to aid in data editing. However, to analyze undifferenced data and determine clock estimates it is necessary to add the pseudorange data in order to permit separation of the otherwise indistinguishable clock offset and phase cycle ambiguity parameters. The quality of the clock estimates is maximized by ensuring the longest possible spans of continuous carrier phase data free of cycle slips, typically three to four hours for each receiver-satellite pair, thus minimizing the number of ambiguity parameters.

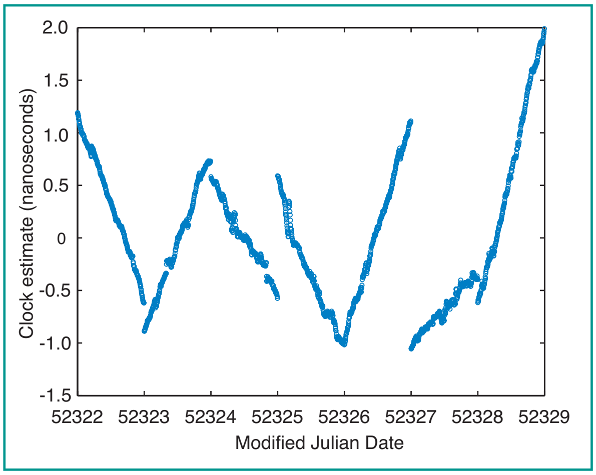

FIGURE 2 The original IGS combined clock estimates for the WSRT station (located in Westerbork, The Netherlands, and equipped with an H-maser frequency standard) for GPS week 1154 (February 17-23, 2002) referenced to a daily linear alignment to GPS Time. We have removed an overall linear trend of 9.7237 ns/day. The large day-to-day discontinuities in offset (time) and rate (frequency) illustrate the limitations of GPS Time as a stable reference time scale.

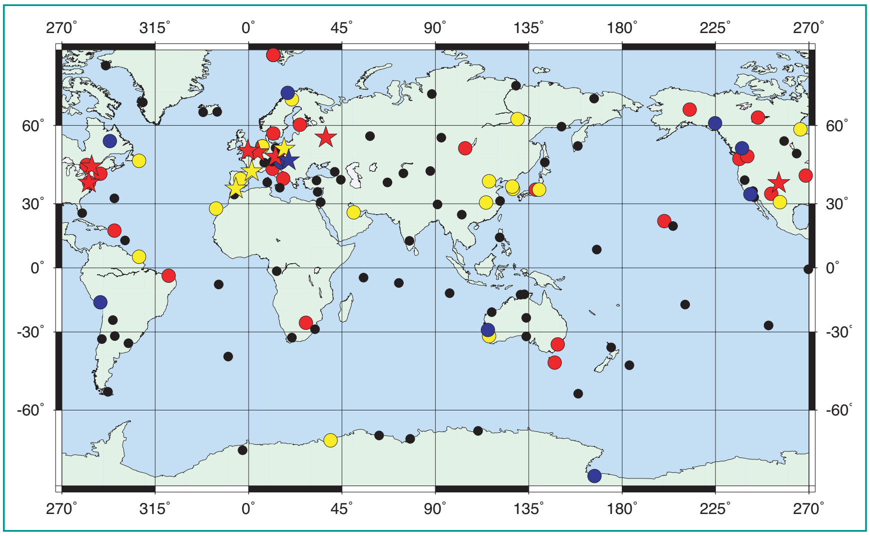

FIGURE 1 The geographical distribution of stations included in the IGS combined clock products. The larger, colored symbols denote stations equipped with external frequency standards: H-masers (red), cesiums (yellow), rubidiums (blue). The smaller black dots indicate stations using internal crystals. IGS stations co-located at timing labs are shown as stars.

Improving the IGS Time Scale

The IGS practice of daily linear alignment of its clocks to GPS Time is perfectly satisfactory for all geodetic applications, since changing every clock value at any particular epoch by a common offset has no impact on positioning determinations. From this point of view, the turning off of SA was inconsequential since at each data epoch, the IGS estimates the satellite clock offset, including any SA contribution. Therefore, the large amount of dithering applied by SA to the satellite clock values did not affect the precision of the IGS clock determinations. (Interpolation to epochs between those tabulated by the IGS is greatly improved now, though.)

Still, a more stable IGS time scale has clear advantages for time/frequency applications. For intervals of one day and longer, GPS Time is stable to about two parts in 10^{14}, while IGS stations equipped with hydrogen maser (H-maser) frequency standards can be up to an order

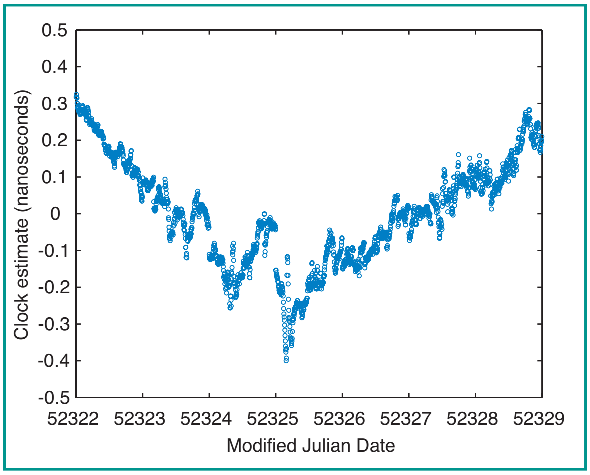

FIGURE 3 The same WSRT clock data shown in Figure 2 after realignment to the new IGS time scale. We have removed an overall linear trend of 9.0294 ns/day. The observed timing performance is now consistent with that expected of a high-quality H-maser frequency standard.

of magnitude more stable. **Figure 2** illustrates the situation for the WSRT (Westerbork, The Netherlands) station, which uses an H-maser. The large discontinuities in clock offsets and rates between successive days are an artifact of the daily alignment to the less stable GPS Time scale.

New System. To overcome this limitation, we have recently developed a new system to synthesize a more stable timing reference by forming a dynamically-weighted ensemble of the IGS clocks themselves. Since all the satellite clocks are candidates for use in this new scale and since numerous more stable clocks are available in the IGS tracking network, the result is guaranteed to be more stable than GPS Time. The new time scale algorithm is actually formulated in the frequency domain using a two-state (clock rate and drift) Kalman filter model, then integrated back into the time domain. The weights of contributing clocks are adjusted dynamically to favor the most stable ones while allowing for continuous changes in individual clock performance and data availability.

One must exercise great care to detect and reject outliers, such as those due to receiver clock resets, swaps in frequency standards, and other types of data problems. In practice, the time scale is overwhelmingly dominated by the H-maser stations although some of the more stable cesium clocks (aboard satellites and at tracking stations) contribute with small weights. We estimate that the stability of the new IGS scale is about 1 part in 10^{15} over a one-day averaging period, roughly 20 times better than GPS Time.

Long-term Stability. Because the IGS clocks currently have no direct linkage to UTC or other absolute frequency standards, the long-term stability of the new scale cannot be assured without further provision. This is accomplished by gently steering the ensemble scale to GPS Time with an effective time constant of roughly 30 to 40 days, consistent with the observed behavior of GPS Time as reported by the BIPM. In this way, the short-term stability of the new IGS time scale is assured while leveraging the longer-term stability of GPS Time, which itself is maintained through continuous monitoring and steering of the system to UTC(USNO) kept by the U.S. Naval Observatory.

To demonstrate the improved performance of the new IGS scale, **Figure 3** shows the clock data from WSRT, previously shown in Figure 2, after realignment. In comparing these figures, note the reduced vertical plot scale (from 3.5 to 1.0 nanoseconds) and the visibly improved day-to-day stability in both offset (time) and rate (frequency). Similar results are obtained for other stations equipped with H-masers, where the limitations of GPS Time instability are most evident.

Accuracy Assessments

Even after realignment to the new time scale, small discontinuities in the clock results of the most stable stations are usually apparent at day boundaries. This is a manifestation of the finite measurement accuracy of the IGS clock values themselves. Recall that the GPS-based clock accuracy (ignoring the instrumental calibration biases) is determined solely by the pseudorange data averaged over each analysis arc. The formal errors for one-day analyses of five-minute data are typically about 100 to 125 picoseconds (ps), assuming each pseudorange observation has an uncertainty of 1 meter. Comparing independent clock estimates at the boundaries between consecutive days allows an objective, albeit internal, assessment of the true clock measurement accuracy. This is analogous to the classical geodetic repeatability test using a time series of position determinations.

We have studied the day-boundary clock discontinuities over a 67-week period for 30 IGS stations equipped with H-masers. (The variations of less stable frequency standards at other stations overwhelm the clock measurement errors.)

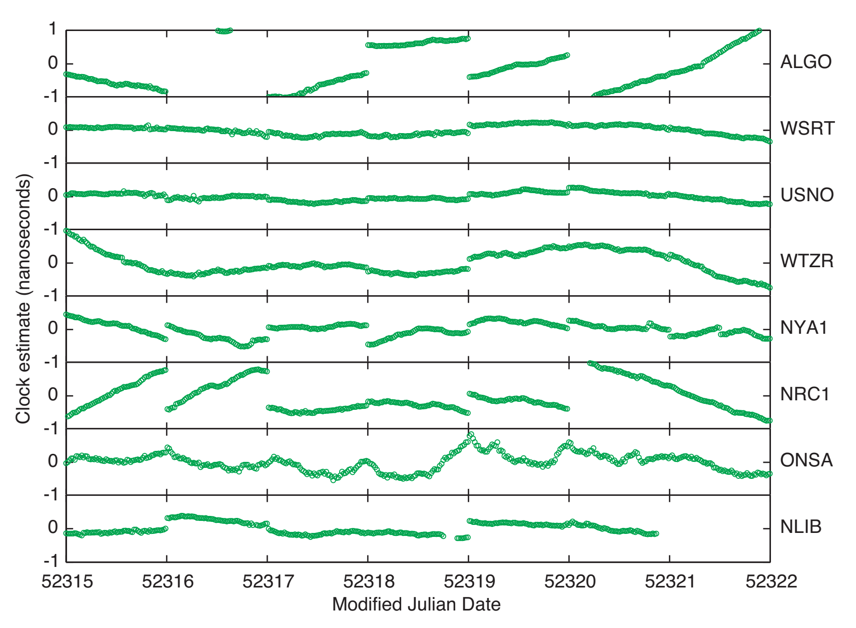

FIGURE 4 Time series of IGS clock estimates for GPS week 1153 (February 10-16, 2002), after realignment to the new time scale, for a sample of eight stations equipped with H-masers. This illustrates the occurrence of day-boundary discontinuities, which are much more distinct at some sites than at others. (For plotting purposes we have removed a separate quadratic trend over the entire week from each series; this figure shows only every fifth 5-minute point.)

Figure 4 displays a sample of the type of data used for a representative week. When we examine the r.m.s. statistics for each of the day-boundary series we find several significant results.

Clock Jumps. The mean values for the clock jumps are always small compared with the r.m.s. of each distribution, except for a couple of stations which are only sparsely sampled. This observation confirms that the geodetic time transfer technique does not induce spurious long-term drifts into the clock estimates, at least not at a significant level. In addition, examination of the distributions of clock jumps about their means shows that the behavior is very close to Gaussian for well-sampled stations.

The most striking feature of the r.m.s. clock jumps in our analysis is the very wide range of values obtained, from about 170 to 1,200 picoseconds depending on the station. Since each day-boundary difference involves two independent measurements, the implied error in any single time transfer measurement will be smaller by $\sqrt{2}$ if the correlations between successive days are negligible. In other words, we infer that time transfer accuracy varies from about 120 picoseconds for the best stations (consistent with the formal errors) up to at least 850 picoseconds. Such a large dispersion among stations in timing performance is

Allan Variance and Clock Stability

The Allan variance is a widely-used statistic for assessing the performance of oscillators and clocks over a specified time interval. By "performance" we mean how well an oscillator maintains a particular frequency or how well a clock keeps time. If we have a time series of measurements of the relative frequency difference, or frequency deviation, of two oscillators,

$$y_i = \frac{\Delta f_i}{f}$$

then the Allan variance, or 2-sample variance, is defined as

$$\sigma_y^2(\tau) \equiv \frac{1}{2}\left\langle \left(y_{k+1} - y_k\right)^2 \right\rangle$$

where τ is the time interval over which the frequency deviation is measured and the angled brackets indicate the statistically expected value which could be obtained, theoretically, by averaging an infinite number of measurements. The factor "1/2" causes this variance to equal the classical variance if the y_i constitute a set of random uncorrelated values – in other words, "white noise."

In practice, the Allan variance is estimated from a limited number of measurements, *m*, by

$$\sigma_y^2(\tau) = \frac{1}{m}\sum_{j=1}^{m}\frac{1}{2}\left(y_{k+1} - y_k\right)_j^2$$

But why not simply use the classical variance to describe oscillator performance? It turns out that the classical variance diverges for some commonly observed oscillator noise processes; that is, the variance increases with an increasing number of data points.

In comparing two clocks, the clocks' time differences can be related to the underlying frequency differences and the Allan variance can be computed directly from the time series of *clock* differences.

The square root of the Allan variance is the Allan deviation. A plot of the Allan deviation versus the averaging interval, τ, portrays the relative stability of a pair of oscillators or clocks over different intervals. Various factors affect the stabilities (or, more precisely, the instabilities) of oscillators and clocks, including the physical properties of the devices and environmental factors. Furthermore, when a pair of devices is compared, the comparison process itself may contribute noise to the measurements. If a "test" clock is compared to one which is vastly superior in performance, say a particular reference clock, or to a "composite" clock derived from the comparison of a large number of clocks (such as that represented by the IGS's new time scale), then the Allan deviation of the relative measurements approximates the performance of the test clock.

Typically, the stability of an oscillator or a clock can be characterized by a power-law dependence on the averaging interval. If the logarithm of the Allan deviation is plotted versus the logarithm of τ, the resulting plot will be approximated by a series of connected lines of different slopes. The slopes of the line segments indicate the particular noise process afflicting the device. The following table lists some common noise types and their dependency on τ:

Noise type	Slope
white phase noise	−1
flicker phase noise (approx.)	−1
random walk phase noise	−1/2
white frequency noise	−1/2
flicker frequency noise	0
random walk frequency noise	1/2

An oscillator or clock might be very precise but may not be accurate because its mean frequency departs from that of some standard. In fact, due to physical changes in an oscillator, there is a usually a systematic change in its frequency with time, a phenomenon called aging. Because they are so common, linear frequency drifts due to aging and other factors are sometimes removed before computing the Allan variance. – R.B.L.

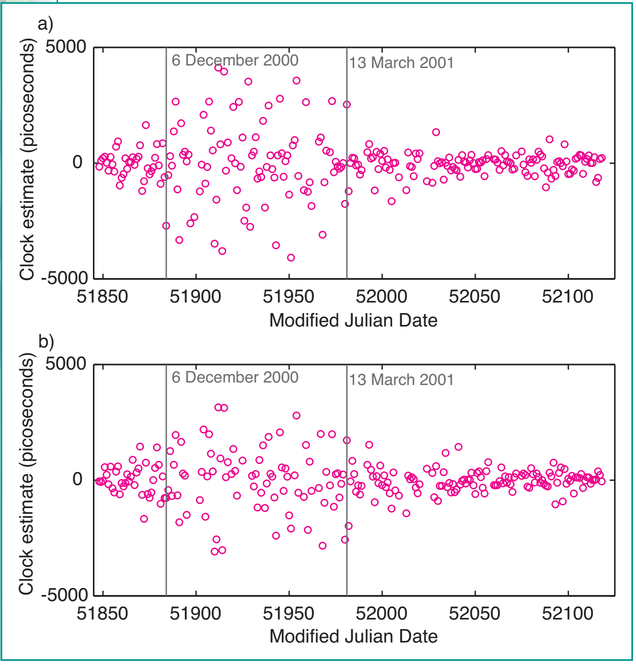

FIGURE 5 Time series of day-boundary clock discontinuities for NRC1 (Ottawa, Canada). Upper panel (a) shows the raw results, with much larger variations seen during winter than at other times of the year. Lower panel (b) shows the same results after correction by removing a linear temperature-dependent trend. While the resulting level of discontinuities during winter has been somewhat reduced, they remain significantly larger than at other times.

surprising given that differences in geodetic accuracy are normally smaller.

There is no clear correlation between timing accuracy and the choice of antenna or receiver model. Also, use of antenna radomes does not appear to affect the timing performance, although any static bias (which is likely) would not be detectable. Without more detailed studies of the individual stations, we cannot account for the general dispersion in performance. In principle, one must consider any factor that can degrade the pseudorange and/or carrier phase observables. Because multipath is usually the dominant observational error, especially for the pseudoranges, it deserves special attention. One must also evaluate such other local factors as temperature sensitivities, radio-frequency interference and the electromagnetic environment, internal impedance mismatches, and receiver firmware.

Station-Dependent Performance. We can obtain some further insight into station-dependent performance among the several stations that show distinct temporal changes in day-boundary behavior: ALGO (Algonquin Park, Ontario, Canada), HOB2 (Hobart, Tasmania, Australia), MATE (Matera, Italy), NRC1 (Ottawa, Ontario, Canada), and YELL (Yellowknife, NWT, Canada). The changes at HOB2 can be linked to variations in signal strength at the receiver, which affect the frequency of phase cycle slips, apparently due to problems with the antenna cable. At MATE, the time transfer accuracy improved by nearly a factor of two after the receiver was swapped with a new unit of the same type but with a different firmware version. The three Canadian stations at ALGO, NRC1, and YELL display much greater clock jumps during winter than during other times of the year, by factors of two to three. **Figure 5a** shows the day-boundary jumps at NRC1 for a nine-month period (though the same pattern repeats in more recent data). This evidence of seasonal trends suggests checking for temperature dependencies. Of the stations in our study set, meteorological data are available for 11.

Testing the daily mean temperature differences for correlations with the day-boundary clock jumps reveals very significant temperature sensitivities for ALGO (-101.3 ± 9.6 ps/°C) and NRC1 (+155.9 ± 16.7 ps/°C), but not for YELL (+8.6 ± 9.8 ps/°C). Five of the other stations display no correlation with temperature, while three stations have marginally significant trends. Various studies have demonstrated the effects of temperature variations on individual components of the tracking hardware. Typical sensitivities for the receivers alone are of the order of ±100 ps/°C with large variations among individual units. Common RF antenna cables have thermal sensitivities around 1 picosecond per degree Celsius per meter and cable runs often exceed 50 meters. On the other hand, a widely deployed choke-ring antenna has no detectable temperature effect on timing.

In the cases of ALGO and NRC1, we have insufficient information to isolate the specific components that contribute to the observed thermal variations. Regardless of the underlying sources, however, temperature sensitivity alone cannot fully account for the seasonal variation of the clock jumps. That can be seen when the raw day-boundary clock jumps are "corrected" using the observed linear temperature sensitivity. **Figure 5b** shows the results for NRC1 and the persistence of seasonal variations. The null temperature sensitivity for YELL despite its seasonal variation supports this conclusion. Instead, we suspect, but cannot demonstrate, that the Canadian stations are probably also affected by standing-wave back-reflections off the tops of the concrete pillars over which the antennas are mounted. During winters, when snow or ice often covers the pillar tops, back-reflections are likely to be most pronounced.

Stability Considerations

Generally, one expects the precision of clock estimates (and hence the frequency stability) within a given analysis arc to be much better than either the formal errors or the accuracy measures indicate, because the time variation of the clock values is mostly determined by the very precise carrier phase data. For the few long baselines that have been well studied, the observed frequency instabilities approach two parts in 10^{15} over an interval of one day and are consistent with a random walk ($\tau^{-0.5}$, where τ is the averging period) for clock variations over intervals shorter than a day.

We have examined our day-boundary data set for any relationship with Allan

Further Reading

For an introduction to timekeeping, time transfer, and GPS Time, see

- "Time, Clocks, and GPS" by R.B. Langley in *GPS World*, Vol. 2, No. 10, November/December 1991, pp. 38-42.

For background information on IGS clock products, see

- "New IGS Station and Satellite Clock Combination" by J. Kouba and T. Springer in *GPS Solutions*, Vol. 4, No. 4, 2001, pp. 31-36.

For a discussion of time transfer using GPS carrier phase, see

- "Time and Frequency Transfer: High Precision using GPS Phase Measurements" by T. Schildknecht and G. Dudle in *GPS World*, Vol. 11, No. 2, February 2000, pp. 48-52.
- "IGS/BIPM Pilot Project: GPS Carrier Phase for Time/Frequency Transfer and Time Scale Formation" by J. Ray and K. Senior in *Metrologia*, Vol. 40, Spring 2003 (in press).

For access to the realigned IGS clock products, plots, and related data files, see

- "IGS Time Scales" at <http://timescales.nrl.navy.mil>.

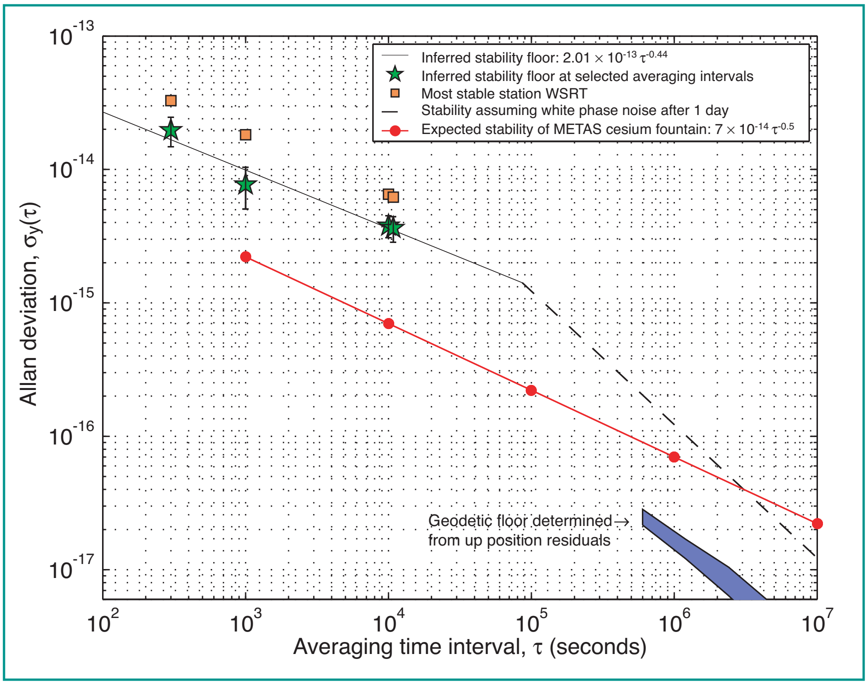

FIGURE 6 Comparison of the inferred Allan deviation stability floor for the geodetic time transfer method (solid line) with the average stability for WSRT, among the best stations in terms of overall performance. This figure also shows (continued dashed line) the expected behavior for intervals longer than one day if the time transfer stability behaves as a white noise process with an average daily uncertainty equal to the typical formal errors (115 picoseconds). For reference, we also show the target stability for the METAS cesium fountain under development in Switzerland (connected red dots) together with an estimate of the fundamental geodetic stability floor derived from the vertical position residuals for ALGO, NRC1, USNO, and WSRT. The geodetic floor may not be attainable due to long-term variations in pseudorange multipath or other factors which affect timing results but not positioning.

deviation measures (see **"Allan Variance and Clock Stability" sidebar**) at sub-daily intervals. Ignoring a few stations with clearly substandard environmental controls, we indeed did find good correlations. We have fit these to linear trends and extrapolated to zero r.m.s. clock discontinuities (corresponding to perfectly accurate time transfers averaged over one day) to infer a limiting value for the short-term stability due to the geodetic time transfer method. **Figure 6** shows the results, which closely follow the expected random walk behavior up to one day. The instability of WSRT, among the best IGS stations in terms of overall short-term stability and long-term accuracy, is about double our inferred stability floor. We expect, but cannot yet demonstrate, that the time transfer instability for intervals longer than the one-day analysis period will follow a white noise trend (τ^{-1}) for the appropriate station accuracy (see Figure 6). This will continue until ultimately limited by longer-term geodetic errors and pseudorange multipath variations, which may not be reflected in positioning results.

Prospects for Time Transfer

The carrier phase geodetic technique and the realigned IGS clock products appear to be satisfactory for frequency and stability comparisons between the best standards in current use, namely actively tuned H-masers. Over all intervals beyond five minutes, the time transfer method need not seriously limit the observed stability performance, although the actual behavior is evidently highly station-dependent and not all stations reach their full potential. However, for the new generation of frequency standards being developed with one-day stabilities better than 10^{-15}, current geodetic methods may not be sufficient for comparisons over intervals of less than a few weeks. When such clocks become available, we will need to study this regime further.

Comparisons of time in an absolute sense are limited by the finite errors of hardware delay calibration and this is likely to remain unchanged for the foreseeable future, although differential calibration methods offer promise. Only at the very poorest performing IGS stations do the time transfer errors approach the level of the calibration uncertainties.

Among the areas requiring further work is the development of closer links between the IGS clock products and UTC. Improved traceability and better long-term steering of the IGS time scale will allow the clock products to be useful for general dissemination of accurate and easily accessible UTC. Accomplishing this will require calibrated IGS receivers co-located at a sufficient number of timing labs, a development that is well underway. In addition, user access to the new timing capabilities would benefit from enhancements in the design of geodetic receivers.

Acknowledgments

We are indebted to all components of the IGS without whose accomplished efforts this work would not have been possible. The IGS Analysis Centers and the Analysis Coordinator deserve special thanks. We also appreciate the friendly assistance and invaluable cooperation of the BIPM and its staff. This article is based, in part, on a paper to appear in the journal *Metrologia*. We thank the journal for permission to use some of the figures which appear in the paper.

Manufacturers

The absolutely calibrated receiver circulated by the BIPM is an **Ashtech Precision Products** (Sunnyvale, California) Z-XII3T.

WSRT's H-maser is a **Sigma Tau** model MHM - 2010 now manufactured by **Datum Incorporated**'s Timing, Test and Measurement Division (Beverly, Massachusetts).

The choke-ring antenna deployed at many IGS stations is a **Dorne & Margolin** (now **EDO Corporation**, New York, New York) Model T distributed by **Allen Osborne Associates, Inc.** (Westlake Village, California).

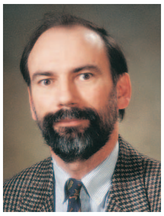

"Innovation" is a regular column featuring discussions about recent advances in GPS technology and its applications as well as the fundamentals of GPS positioning. The column is coordinated by Richard Langley of the Department of Geodesy and Geomatics Engineering at the University of New Brunswick, who appreciates receiving your comments as well as topic suggestions for future columns. To contact him, see the "Columnists" section on page 4 of this issue.