INSTITUTE OF PHYSICS PUBLISHING METROLOGIA
Metrologia **40** (2003) S270–S288 PII: S0026-1394(03)61419-2

IGS/BIPM pilot project: GPS carrier phase for time/frequency transfer and timescale formation

J Ray[1] **and K Senior**[2]

US Naval Observatory, 3450 Massachusetts AV NW, Washington, DC 20392-5420, USA

E-mail: jimr@ngs.noaa.gov and Ken.Senior@nrl.navy.mil

Received 24 April 2002
Published 5 June 2003
Online at stacks.iop.org/Met/40/S270

Abstract

The development within the International GPS Service (IGS) of a suite of clock products, for both satellites and tracking receivers, offers some experiences which mirror the operations of the Bureau International des Poids et Mesures (BIPM) in its formation of TAI/UTC but some aspects differ markedly. The IGS relies exclusively on the carrier phase-based geodetic technique whereas BIPM time/frequency transfers use only common-view and two-way satellite (TWSTFT) methods. The carrier-phase approach has the potential of very high precision but suitable instrumental calibration procedures are only in the initial phases of deployment; the current BIPM techniques are more mature and widely used among timing labs, but are either less precise (common-view) or much more expensive (TWSTFT). In serving its geodetic users, the essential requirement for IGS clock products is that they be fully self-consistent in relative terms and also fully consistent with all other IGS products, especially the satellite orbits, in order to permit an isolated user to apply them with accuracy of a few centimetres. While there is no other strong requirement for the IGS timescale except to be reasonably close to broadcast GPS time, it is nonetheless very desirable for the IGS clock products to possess additional properties, such as being highly stable and being accurately relatable to UTC. These qualities enhance the value of IGS clock products for applications other than pure geodesy, especially for timing operations. The jointly sponsored 'IGS/BIPM Pilot Project to Study Accurate Time and Frequency Comparisons using GPS Phase and Code Measurements' is developing operational strategies to exploit geodetic GPS methods for improved global time/frequency comparisons to the mutual benefit of both organizations. While helping the IGS to refine its clock products and link them to UTC, this collaboration will also provide new time transfer results for the BIPM that may eventually improve the formation of TAI and allow meaningful comparisons of new cold atom clocks. Thus far, geodetic receivers have been installed at many timing labs, a new internally realized IGS timescale has been produced using a weighted ensemble algorithm, and instrumental calibration procedures developed. Formulating a robust frequency ensemble from a globally distributed network of clocks presents unique challenges compared with intra-laboratory timescales. We have used these products to make a detailed study of the observed time transfer performance for about 30 IGS stations equipped with H-maser frequency standards. The results reveal a large dispersion in quality which can often be related to differences in local station factors. The main elements of the Project's original plan are now largely completed or in progress. In major ways, the experiences of this joint effort can serve as a useful model for future distributed timing systems, for example, Galileo and other GNSS operations.

[1] Present address: NOAA, National Geodetic Survey, 1315 East-West Hwy, Silver Spring, MD 20910, USA.
[2] Present address: Naval Research Lab, Space Applications Branch, 4555 Overlook AV SW, Washington, DC 20375, USA.

1. IGS/BIPM pilot project overview

The 'International GPS Service/Bureau International des Poids et Mesures (IGS/BIPM) Pilot Project to Study Accurate Time and Frequency Comparisons using GPS Phase and Code Measurements' began in March 1998 with the participation of more than 35 groups [1]. The main goal of the pilot project is to investigate and develop operational strategies to exploit GPS measurements and geodetic techniques for improved global availability of accurate time and frequency comparisons. This is becoming more significant for maintaining the international UTC timescale as a new generation of frequency standards emerges with stabilities of 10^{-15} or better.

The respective roles of the IGS and BIPM organizations are complementary and mutually beneficial. The IGS and its collaborating participants bring a global GPS tracking network, standards for continuously operating geodetic-quality, dual-frequency GPS receivers, an efficient data delivery system, and state-of-the-art data analysis groups, methods, and products. The BIPM and its timing laboratory partners contribute expertise in high-accuracy metrological standards and measurements, timing calibration methods, algorithms for maintaining stable timescales, and the formation and dissemination of UTC. The progress of the project and related information are posted at http://maia.usno.navy.mil/gpst.html.

1.1. Potential of geodetic methods for timing

GPS-based techniques have been the basis for most high-accuracy time transfer for about two decades. The 'common-view' method [2], employed by the BIPM in the formation of the international UTC atomic timescale, uses single-frequency C/A pseudorange observations and relies heavily on cancellation of common-mode errors to achieve intercontinental comparisons between timing laboratories with uncertainties of a few nanoseconds at 5 day intervals. In contrast, high-accuracy geodetic methods using dual-frequency GPS carrier-phase observations achieve positioning repeatabilities at the sub-centimetre level (30 ps) for 1 day integrations. Assuming such positioning results can be realized also as equivalent light travel times, the potential of GPS carrier-phase and geodetic techniques for global time and frequency comparisons at the sub-nanosecond level is evident. This was the rationale in establishing a joint pilot project between the BIPM and the IGS to develop and demonstrate the operational capabilities.

GPS phase data are essential for modern geodesy. With the 'double differencing' technique, for example, which incidentally removes all clock-like effects, pseudorange data are not normally even included in the geodetic analysis, as these observables are about 100 times less precise than the phase measurements. However, to analyse undifferenced data and extract 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. See Larson and Levine [3] for a complete description of the GPS phase and pseudorange observation equations. For each tracking receiver–satellite pair, the quality of the clock estimates is maximized by ensuring the longest possible spans of continuous phase data free of cycle slips, thus minimizing the number of ambiguity parameters. High-performance geodetic receivers typically track individual satellite passes with arcs of up to three to four hours. Apart from viewing obstructions, the most problematic tracking is usually at the lowest elevation angles where the signal strength is most attenuated and the atmospheric path delay is greatest and most variable.

It is worth noting several distinctions between the observed estimates for clock parameters from a geodetic analysis and direct laboratory measurements of clock standards. In geodetic analyses, any effect which is bias-like will be absorbed into the clock estimates. The observability of clock parameters depends strongly on being able to model the geodetic problem adequately, with all other parameters of importance having partials that are much different from unity (or, more generally, from any constant value). In addition, to distinguish carrier-phase cycle ambiguities, pseudorange observables are required (as described above). These conditions are generally well satisfied for GPS data where the dynamics favour a very clear separation of clock effects from other parameters.

1.2. Satellite clocks

Satellite clock estimates are among the 'core' products of the IGS. Combined solutions for satellite clocks have been distributed with the IGS combined orbits since the founding of the service. It is essential that the clock information be as consistent as possible with the other IGS products to ensure that the precise point positioning (PPP) technique [4] can be applied by an autonomous user without degradation. In the PPP method, data from a single, isolated GPS receiver (pseudorange and phase at both frequencies) are reduced using the IGS orbits and satellite clock values instead of broadcast GPS navigation information. By using the carrier-phase data and precise IGS products, the instantaneous receiver position can be determined to the 5 cm level vertically with respect to the IGS reference system (better in the horizontal). Similarly, a precise determination of the receiver clock can be made relative to the underlying timescale of the IGS clocks. This approach is obviously well suited for simple, global dissemination of sub-nanosecond time and frequency information, limited principally by the quality of the timescale. Just as the IGS has adopted (and contributes significantly to) the International Terrestrial Reference Frame as its geodetic system, the timescale for its clock products should ideally be traceable to UTC.

1.3. Receiver clocks

Prior to this pilot project, the IGS did not produce clock information for its tracking receiver network. There are clear reasons for doing so, however, such as enabling comparisons between timing laboratories, a primary interest of the BIPM. Through the considerable efforts of the IGS Analysis Centers and Analysis Coordinator, new IGS clock products for a subset of the tracking network were initiated officially on 5 November 2000 [5]. A 'clock RINEX' file format was devised for the exchange of clock estimates and related information. By combining the analyses of at least two centres for each receiver, valuable integrity and quality checking have been implemented.

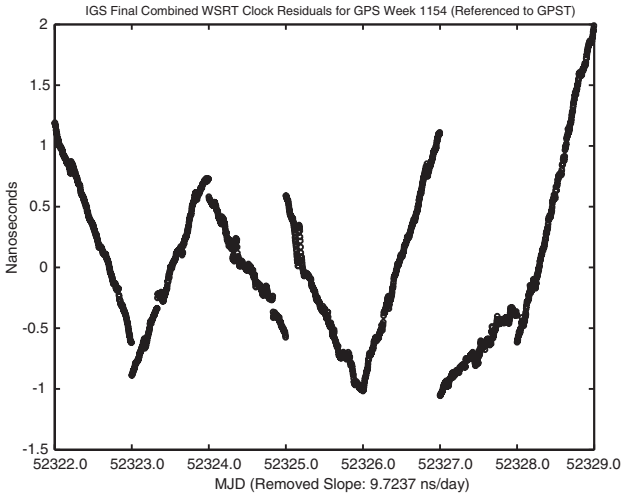

Figure 1. IGS combined clock estimates of a Sigma Tau hydrogen maser located at WSRT in Westerbork, The Netherlands. The estimates are referenced to a daily linear alignment to GPST. The plot shows the large day-to-day offsets in phase and frequency resulting from the GPST alignment strategy.

1.4. Timescales

It is necessary that all the IGS clock information, for satellites and receivers, be referenced to a common, consistent timescale. Historically, the IGS has used a simple linear alignment of the observed satellite clocks to broadcast GPS time (GPST) for each day separately. Unfortunately, the instability of GPST is sufficiently large that this procedure introduces large day-to-day discontinuities in the time and frequency of IGS clock products (see figure 1). There is no impact on the usefulness of the products for point positioning, but the utility for time/frequency dissemination is certainly limited. Consequently, a high priority has been the development of an internally realized frequency ensemble with improved stability as a reference for IGS clocks; this work is described below.

1.5. Instrumental calibration and comparison

In order to accurately relate the IGS estimates for receiver clocks to external standards, techniques must be developed to calibrate the instrumental delays within the tracking receiver hardware. From geodetic analyses of GPS data, the effective 'clock' of each receiver is determined for the ionosphere-corrected L3 phase centre of the antenna displaced by the electrical delay to the point in the receiver where the observables are measured and time-tagged. Calibration of the instrumental path delays is an essential requirement for time transfer uses. The requirements for frequency comparisons are much less stringent although the effects of environmental influences must be minimized. Considerable progress has been achieved in calibration of geodetic receivers, at least for the Ashtech Z-XII3T model [6, 7].

The topic of instrumental calibration is regarded as a separable issue being pursued elsewhere and will not be considered further here. Likewise, several authors have compared carrier-phase clock variations to independent techniques, so those results are not reviewed here. The interests of this report are the IGS clock analysis products, formation and alignment of the clocks to a stable timescale, and evaluation of the IGS clocks for their accuracy, stability, and utility for global time transfer.

2. IGS tracking network

As of early 2002, the IGS tracking network consists of nearly 300 permanent, continuously operating, geodetic stations globally distributed. Of these, almost 80 are equipped with external frequency standards, the remainder using internal crystal oscillators. The external standards in use are approximately 40 with H-masers, 25 with caesiums, and 15 with rubidiums. Not all tracking receivers are included in the IGS clock products due to limitations of the analysis capabilities, problems with data acquisition or transmission, or for other reasons. Generally, the IGS Analysis Centers give preference to stations with external frequency standards, so most of these are included in the clock products. Figure 2 shows the geographic distribution and type of frequency standard for receivers typically included in the IGS combined clocks. High-stability clocks are concentrated in Europe and North America, with lesser numbers in east Asia and Australia.

An important activity of the IGS/BIPM pilot project has been to promote deployment of additional IGS stations at timing laboratories. These must be high-quality geodetic receivers capable of continuously recording and rapidly transmitting dual-frequency pseudorange and carrier-phase observations. Generally, the IGS standards for configuration, operation, and documentation should be adopted. In addition, consideration should be given to electronic and environmental stability for the receiver equipment, low multipath surroundings for the antenna, high-quality cables, and other similar factors that might affect timing results. Some receiver models are available with options tailored for use in timing applications, such as the Ashtech Z-XII3T receiver. Currently about 18 timing labs host IGS stations, some with multiple installations (see table 1). Several other labs are in the process of installing IGS stations or evaluating equipment in preparation for doing so.

Documentation on all IGS stations, including station log files, is maintained at the IGS Central Bureau at http://igscb.jpl.nasa.gov/network/netindex.html. The IGS also publishes procedures for joining the tracking network and standards for station operation.

3. IGS timescales

3.1. Alignment to GPST

Starting with GPS week 1087, the IGS produces two sets of clock estimates, a rapid combination with a latency of less than 1 day, and a final set following about two weeks later. Both sets are tabulated at 5 min intervals. In addition to a subnetwork of approximately 90 stations, these include all the GPS satellites. Currently (March 2002) the GPS constellation consists of 28 spacecraft equipped with the following oscillators: 14 Block II/IIA with caesiums; 8 block II/IIA with rubidiums; 6 Block IIR with new-generation rubidiums.

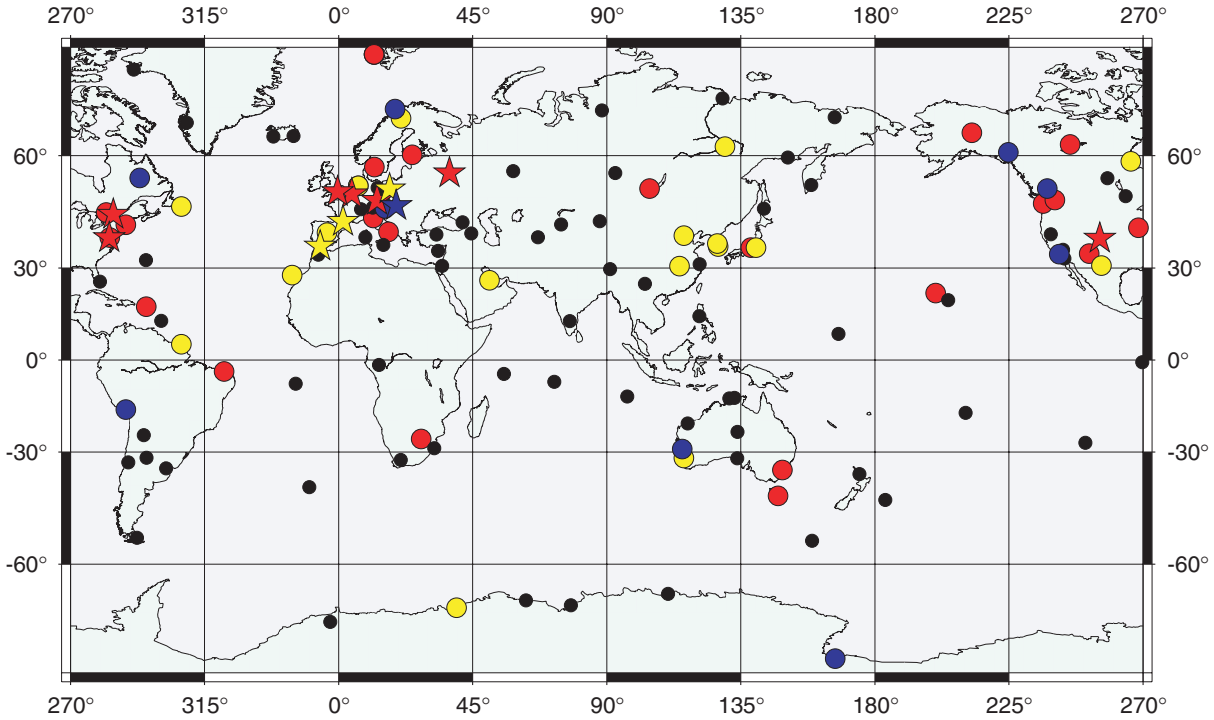

Figure 2. Geographical distribution of IGS stations currently (March 2002) contributing to the IGS combined clock products. Red denotes stations using hydrogen masers, yellow denotes caesiums, blue denotes rubidiums, and black denotes crystals. Stars indicate contributing IGS stations colocated at timing labs. Some stations with high-quality frequency standards or at timing labs are not yet included in the IGS clock products.
(This figure is in colour only in the electronic version, see `www.iop.org`)

Table 1. IGS stations colocated at BIPM timing laboratories (as of March 2002).

IGS site	Time lab	GPS receiver	Freq. std.	City
AMC2	AMC[a]	AOA SNR-12 ACT	H-maser	Colorado Springs, CO, USA
BOR1	AOS	AOA TurboRogue	Caesium	Borowiec, Poland
BRUS	ORB	Ashtech Z-XII3T	H-maser	Brussels, Belgium
KGN0	CRL[a]	Ashtech Z-XII3	Caesium	Koganei, Japan
MDVO	IMVP	Trimble 4000SSE	H-maser	Mendeleevo, Russia
MIZU	NAO	AOA Benchmark	Caesium	Mizusawa, Japan
NPLD	NPL[a]	Ashtech Z-XII3T	H-maser	Teddington, UK
NRC1	NRC[a]	AOA SNR-12 ACT	H-maser	Ottawa, Canada
NRC2	NRC[a]	AOA SNR-8100 ACT	H-maser	Ottawa, Canada
OBE2	DLR	AOA SNR-8000 ACT	Rubidium	Oberpfaffenhofen, Germany
OPMT	OP	Ashtech Z-XII3T	H-maser	Paris, France
PENC	SGO	Trimble 4000SSE	Rubidium	Penc, Hungary
PTBB	PTB[a]	AOA TurboRogue	H-maser	Braunschweig, Germany
SFER	ROA[a]	Trimble 4000SSI	Caesium	San Fernando, Spain
SPT0	SP	JPS Legacy	Caesium	Boras, Sweden
TLSE	CNES	AOA TurboRogue	Caesium	Toulouse, France
TWTF	TL[a]	Ashtech Z-XII3T	Caesium	Taoyuan, Taiwan
USNO	USNO[a]	AOA SNR-12 ACT	H-maser	Washington, DC, USA
WTZA	IFAG	Ashtech Z-XII3T	H-maser	Wettzell, Germany
WTZR	IFAG	AOA SNR-8000 ACT	H-maser	Wettzell, Germany

[a] Also participates in two-way satellite time transfer (TWSTFT) operations.

The daily linear alignment of the IGS clocks to GPS time, denoted here as GPST, results in large offsets in time and frequency between days (see figure 1). To overcome this limitation, a new system was recently developed to synthesize a more stable reference by forming a weighted ensemble of the IGS clocks. Corresponding to each of the rapid and final IGS combined clock products, both a rapid timescale, called IGRT, and a final timescale, IGST, are generated.

3.2. New integrated frequency scale

A number of issues regarding IGS clock data must be addressed in order to effectively ensemble the IGS clocks. For example, most IGS station receiver systems' internal clocks may reset by large amounts (e.g. from power cycling of the system), resulting in potentially large phase jumps in the clock estimates, and standard reductions of GPS carrier-phase data

often result in discontinuities of clock estimates at processing boundaries by up to 1 ns (for typical 1 day analysis arcs). Moreover, the overall clock offset between the IGS receivers and UTC is unknown since IGS receiver systems are usually uncalibrated. For these reasons as well as the White FM nature of IGS clock data for periods less than 1 day, it was decided to ensemble the IGS clocks in the form of an integrated frequency scale. That is, adjacent phase estimates for each clock are differenced to form 5 min frequency measurements for each clock where for example large phase jumps are easily identified and removed as outliers. A simple linear model (quadratic in phase) is then fit to the frequency measurements to detrend each clock against the frequency scale, which is formed on-the-fly in the filter using a standard weighting scheme discussed below. After iteration, the timescale is formed by taking the weighted average of the detrended frequency estimates of each clock and integrating back into a time series. This timescale is then steered to GPST with a long time constant. In this way, the new scale is loosely steered to GPST, resulting in a reference which is more stable than GPST in the short term, but maintaining a link to GPST in the longer term. Each IGS clock is then re-referenced to the new steered timescale.

3.3. Clock model

We briefly summarize the algorithm implemented by [8, 9] to generate the IGS timescales. The frequency (or rate) and frequency aging (or drift) of each clock are estimated using a two-state vector polynomial $\vec{p}$ driven by white noise processes. The usual state–space formulation (discrete-time) is

$$\vec{p}\,(t_{i+1}) = \varphi(\tau)\,\vec{p}\,(t_i) + \vec{e}\,(t_{i+1}),$$
$$\varphi(\tau) = \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix} \tag{1}$$

for $t_{i+1} = t_i + \tau$, where the first and second components of $\vec{p}$ are rate and drift respectively, and where the two components of $\vec{e}$ are independent white noise processes. Given frequency measurements $y(t)$ for a clock, the measurement model is

$$y(t) = \boldsymbol{h}\,\vec{p}\,(t) + n(t), \qquad \boldsymbol{h} = [1 \quad 0], \tag{2}$$

where n is a scalar white noise sequence which is assumed independent of the process noise $\vec{e}$.

For approximate (i.e. for small τ) modelling of random walk FM and random run FM process noise one would require that the covariance $\boldsymbol{q}$ for $\vec{e}$ be specified as [10–12]

$$\boldsymbol{q} \stackrel{\text{def}}{=} E[\,\vec{e}\;\vec{e}^{\,\mathrm{T}}] = \begin{bmatrix} a_{-1}\tau + a_{-2}\dfrac{\tau^3}{3} & a_{-2}\dfrac{\tau^2}{2} \\ a_{-2}\dfrac{\tau^2}{2} & a_{-2}\tau \end{bmatrix}$$

and for approximate modelling of White FM measurement noise the covariance r for n be given as

$$r \stackrel{\text{def}}{=} E[n^2] = \frac{a_0}{\tau},$$

where E is the expectation operator and a_0, a_{-1}, and a_{-2} are the noise spectral densitities of White FM, random walk FM, and random run FM processes, respectively, which may be empirically determined by analysing Allan variances for each clock. Only a minimal amount of random walk FM and random run FM, $a_{-1} = 10^{-3}\,\mathrm{ns^2\,day^{-3}}$ and $a_{-2} = 10^{-4}\,\mathrm{ns^2\,day^{-5}}$, noise is injected into the filter for each clock in order to maintain stability of the filter (cf [13], p 261). Future enhancements may consider improved values for these parameters.

However, values for the measurement noise White FM spectral densities, a_0 for each clock are determined empirically. These values are combined with the formal error sigmas, which are included with the clock estimates in the IGS combined clock product, to form the data weight matrix for the filter. The formal errors are a measure of the agreement among the contributing analysis centres' estimates and are therefore another way of identifying bad clock estimates. For each clock we set the measurement covariance as

$$r(t_{i+1}) = \frac{a_0}{\tau} + \frac{\sigma_x^2(t_{i+1}) + \sigma_x^2(t_i)}{\tau^2}. \tag{3}$$

The second term of (3) due to the formal errors typically dominates, but can sometimes be unrealistically small. The first term in (3), determined empirically, ensures a reasonable error floor for each clock.

Labelling the IGS combined clock phase estimate (referenced to GPST) for each clock k as x_k^{GPST}, we form frequency measurements y_k^{GPST} over each $\tau = 5$ min interval (time-tagged at the end of each interval) by simple linear differencing of adjacent measurements. This frequency measurement is then referenced to the new scale and used as data in (2) according to

$$y_k(t_{i+1}) \stackrel{\text{def}}{=} y_k^{\mathrm{GPST}}(t_{i+1}) - \sum_{j=1}^{N} w_j(t_{i+1})\big[y_j^{\mathrm{GPST}}(t_{i+1}) - (\varphi\,\vec{p}_{\,j}(t_i))_1\big], \tag{4}$$

where the weights w_j of each clock (discussed below) are required to sum to unity at each epoch t. Note that in the bracketed expression in (4) it is necessary to detrend each clock with respect to the rate estimate only (hence the subscript 1) since φ in (1) ensures that drift is already accounted for in the rate.

The standard frequency scale equation (4) detrends the data for each clock at $t = t_{i+1}$ using estimates of its rate and drift at $t = t_i$ with respect to the new scale (as propagated from t_i to t_{i+1} by φ). This is particularly useful in this application since at any 5 min epoch, up to a tenth of the clocks transition in or out of the filter because of data outages or editing. Equation (4) ensures that as clocks move in and out, the scale will be minimally affected.

In support of the continuing move of the IGS towards real-time production of its products, a standard Kalman filter is employed to solve (1)–(2), allowing one to process data sequentially in batches (currently, 1 day batches) and to model the noise processes mentioned earlier. Moreover, subject to the underlying assumptions, the Kalman approach yields optimal state estimates. Two iterations of the filter are performed in order to refine the weights. The filter accommodates the changing size of the state vector and covariance matrix as the number of clocks varies by updating only those subcomponents or submatrices of the state and covariance for which data are available.

3.4. Realigned output clocks

Aside from individual estimates of the rate and drift of each clock relative to the new frequency scale, the primary output of the Kalman filter implementation of (1)–(2) is a single series,

$$\sum_{j=1}^{N} w_j(t_i)\left(y_j^{\mathrm{GPST}}(t_i) - \vec{p}_{\,j,1}(t_i)\right) \stackrel{\mathrm{def}}{=} y_{\mathrm{IGS(R)T}}^{\mathrm{GPST}}(t_i) \tag{5}$$

of the new frequency scale relative to the original scale GPST. Integrating (5) yields its corresponding time series which we label as $x_{\mathrm{IGS(R)T}}^{\mathrm{GPST}}(t)$. Once this series is obtained, we may re-reference all of the original clock estimates in the IGS combined clock product by simple subtraction,

$$x_k^{\mathrm{IGS(R)T}}(t_i) = x_k^{\mathrm{GPST}}(t_i) - x_{\mathrm{IGS(R)T}}^{\mathrm{GPST}}(t_i), \tag{6}$$

so that relative clock-to-clock measures—and hence the use of the re-referenced clocks in positioning applications—are unaffected. However, before performing this re-referencing of the original clock estimates, we loosely steer the timescale to GPST (see section 3.7).

3.5. Detecting outliers and frequency steps

Frequency data are rejected as outliers (e.g. due to phase jumps) in the filter by applying a 10-sigma test to the pre-fit residuals at each time step using the filter covariance. Discrete steps in frequency (e.g. hardware steering changes made at individual IGS stations) are flagged at each step by applying an optimal likelihood function, or smoothed chi-square statistic (see [14], pp 229–31). If such a step in frequency occurs for a specific clock, the weight for that clock is set to 0 for the following 12 h (and hence is removed from the scale for that duration) and additional noise is injected into the filter for the clock to allow the filter to establish its new rate and drift. Once the new rate and drift are determined, the clock is re-entered into the scale by returning its weight to its nominal value.

In addition to the above, particularly poor clocks whose root-mean-square (rms) of frequency calculated over the entire day's data are larger than 200 ns day^{-1} (e.g. crystal oscillators) are removed from the filter (and hence the scale) altogether for that day.

3.6. Weighting scheme

Nominal weights, $\tilde{w}$, for each clock are calculated daily as the inverse Allan variance of that clock sampled at $\tau_1 = 1.2\times10^3$ s, $\tau_2 = 1.02 \times 10^4$ s, and $\tau_3 = 4.32 \times 10^4$ s:

$$\tilde{w} = \frac{1}{\max\{\tau_1\sigma_y^2(\tau_1),\, \tau_2\sigma_y^2(\tau_2),\, \tau_3\sigma_y^2(\tau_3)\}}.$$

Two iterations of the filter are performed with respect to the weights for each clock. In the first iteration, only the current day's data are used to determine the weights where each clock is referenced to GPST. This iteration allows for the detection of unusual changes in the stability of clocks at the various IGS sites which occasionally occur and which are more slight than might be detected by the outlier and frequency step schemes mentioned above. In subsequent iterations, the nominal weights are determined using the current day's data as well as data from the previous six days. Also, in these subsequent iterations, clocks are referenced to the scale determined from the previous iteration.

Following an algorithm similar to that found in [15], nominal weights for each clock are limited to,

$$\max\left\{0.1,\, \frac{2.5}{\text{no of masers}},\, \frac{2.5}{\text{no of clocks}}\right\}.$$

To ensure the unity normalization constraint and the upper limit of weights at every epoch as clocks transition in/out of the ensemble, the actual weights used per epoch and per clock are determined iteratively from the nominal weights (see [8, 9] for details).

3.7. Steering to GPST

In order to utilize the short-term stability of the new IGS timescale while leveraging the longer-term stability of GPST, the scale is gently steered to GPST by applying a three-state (phase, frequency, and frequency drift) linear quadratic Gaussian (LQG) steering algorithm (see [16, 17], and for good references for control and filtering see [13, 14]). We do not include here the state–space formalism of the LQG algorithm, but instead refer the reader again to [9] for more detail. Empirically analysing the unsteered IGS timescale relative to GPST and applying the LQG formalism, we have set the steering gain as $G = [6.323\times10^{-6}\ \ 2.261\times10^{-5}\ \ 3.963\times10^{-4}]$. This steering choice effectively smooths variations in GPST over a period of roughly 30–40 day, consistent with the observed behaviour of GPST (see figure 6) reported in the BIPM *Circular T* [18].

4. Clock results using new timescale

The realigned IGS clock products are available, together with time domain and Allan deviation plots as well as other diagnostic information at the website http://timescales.nrl.navy.mil.

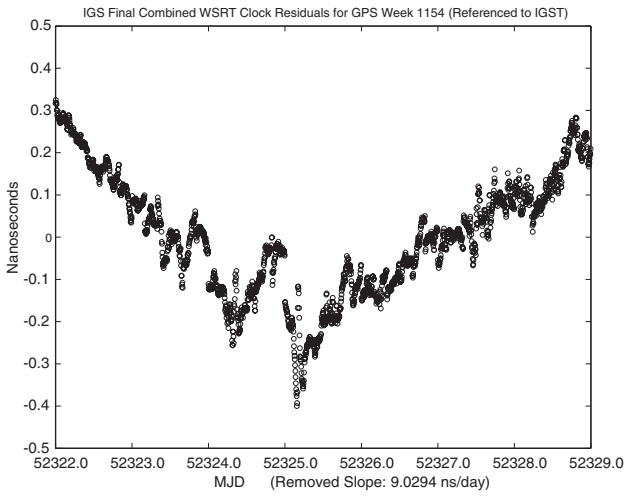

Figure 3. IGS combined clock estimates of a Sigma Tau hydrogen maser located at WSRT, Westerbork, The Netherlands, referenced against the new IGS timescale IGST. Comparison with figure 1 illustrates the improved timescale stability of IGST over GPST.

Though the IGS combined clocks currently contain approximately 40 hydrogen masers, many of these are not of sufficient quality to carry significant weight in the new IGS ensemble. Analysing the results from the beginning of the IGS combined clock product (November 2000) through March 2002 yields a timescale which is dominated by 10–15 hydrogen masers. Figures 3–6 display some representative results. Figure 3 shows the single time series WSRT-IGST of the IGS station located at Westerbork, The Netherlands referenced to the scale, which when compared with figure 1, provides an impression of the improved day-to-day stability of IGST (similarly for IGRT) compared against GPST. Figure 4 shows a time series of the 15 clocks (referenced to IGST) carrying the largest weight in the ensemble for one week of processing. The legend of the plot contains the average weight of each clock in the ensemble for this week. Figure 5 shows the corresponding Allan deviations for the same 15 clocks as in figure 4 and is representative of the stability performance of the IGS timescales. In particular, the deviation plots show that the stability for IGST is about 1×10^{-15} at 1 day, which is consistent with results for GPS carrier-phase time transfer found in [19–21]. Though figure 5 is typical of the stability of the IGS timescales generally, there are periods when the IGS combined clock products have included relatively few stations equipped with hydrogen masers. In some cases, only 15 min clock estimates are available for a few stations instead of the IGS standard of 5 min. These circumstances severely strain the internal ensemble method. However, we believe that the steering gain values handle such situations, but this is of course at the expense of short-term stability.

Figure 6 shows a comparison of the IGS timescales with both GPST as well as with the international timescale UTC for the period 5 November 2000–24 February 2002. In particular, the top plot shows IGST and IGRT against the IGS daily linear alignment to broadcast GPST as well as UTC against GPST via the BIPM *Circular T*. Assuming that GPST is equivalent to GPST via *Circular T*, we difference each of IGST-GPST and IGRT-GPST with UTC-GPST to obtain UTC-IGST and UTC-IGRT (middle plot). For this slightly more than one year of data, the steering strategy discussed earlier maintains the link between the IGS timescales and UTC to within ± 60 ns. The link is of course modulo any calibration biases since currently there are no calibrated receiver systems in the IGS network. Finally, the bottom plot compares IGST with IGRT. The two

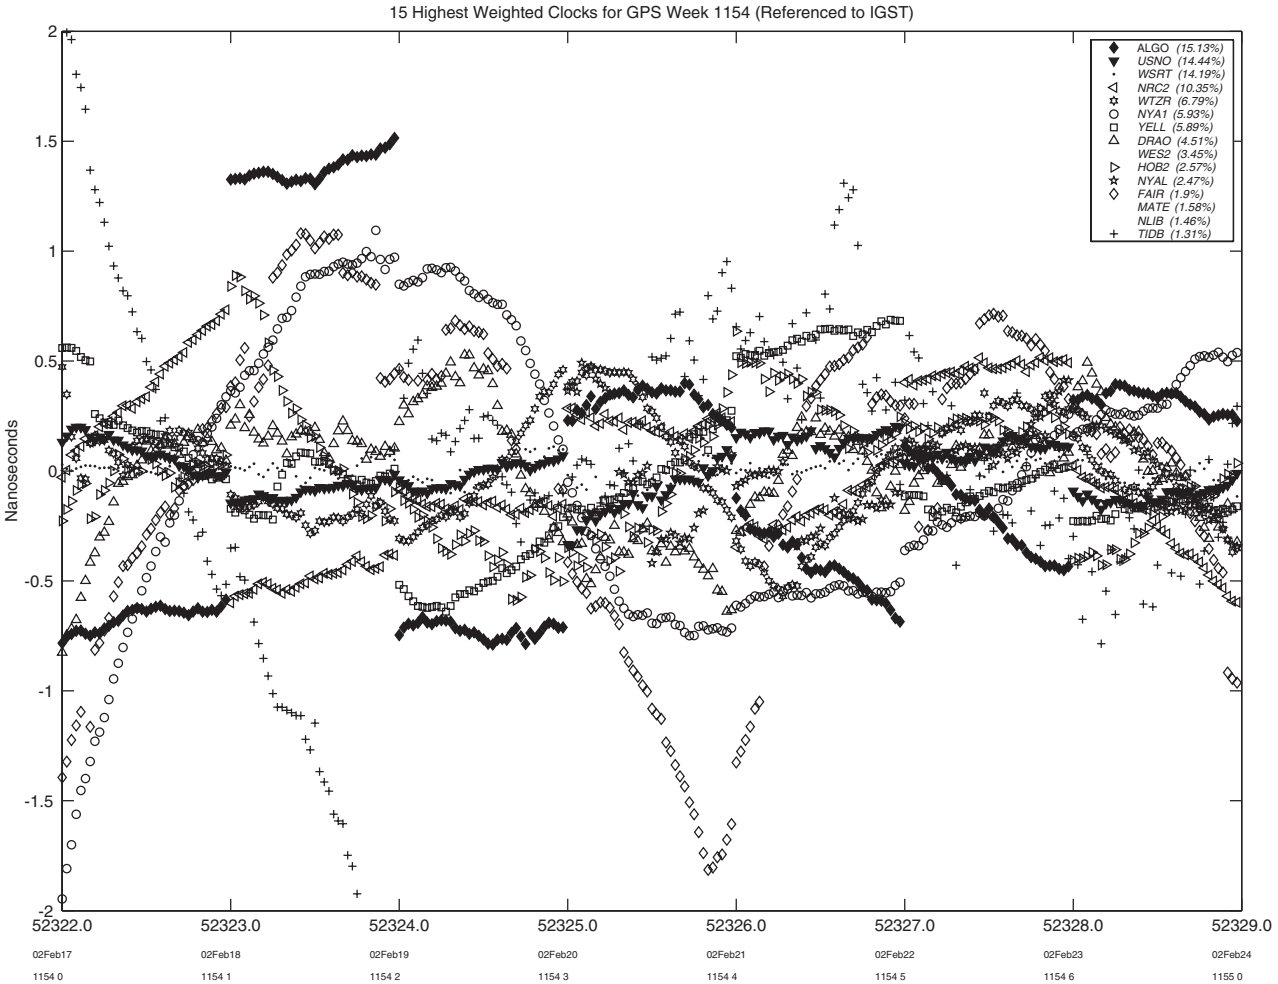

Figure 4. Phase plot of clocks carrying the highest weight in the IGS timescale IGST during GPS week 1154 (17–23 February 2002). Each clock estimate is referenced to IGST and a separate quadratic has been removed from each for plotting purposes. The legend indicates the average weight of each clock in the scale over this period. Only every tenth 5 min point is plotted. Note that three of the clocks listed in the legend are not shown in the plot because one or more phase breaks occurred during the period.

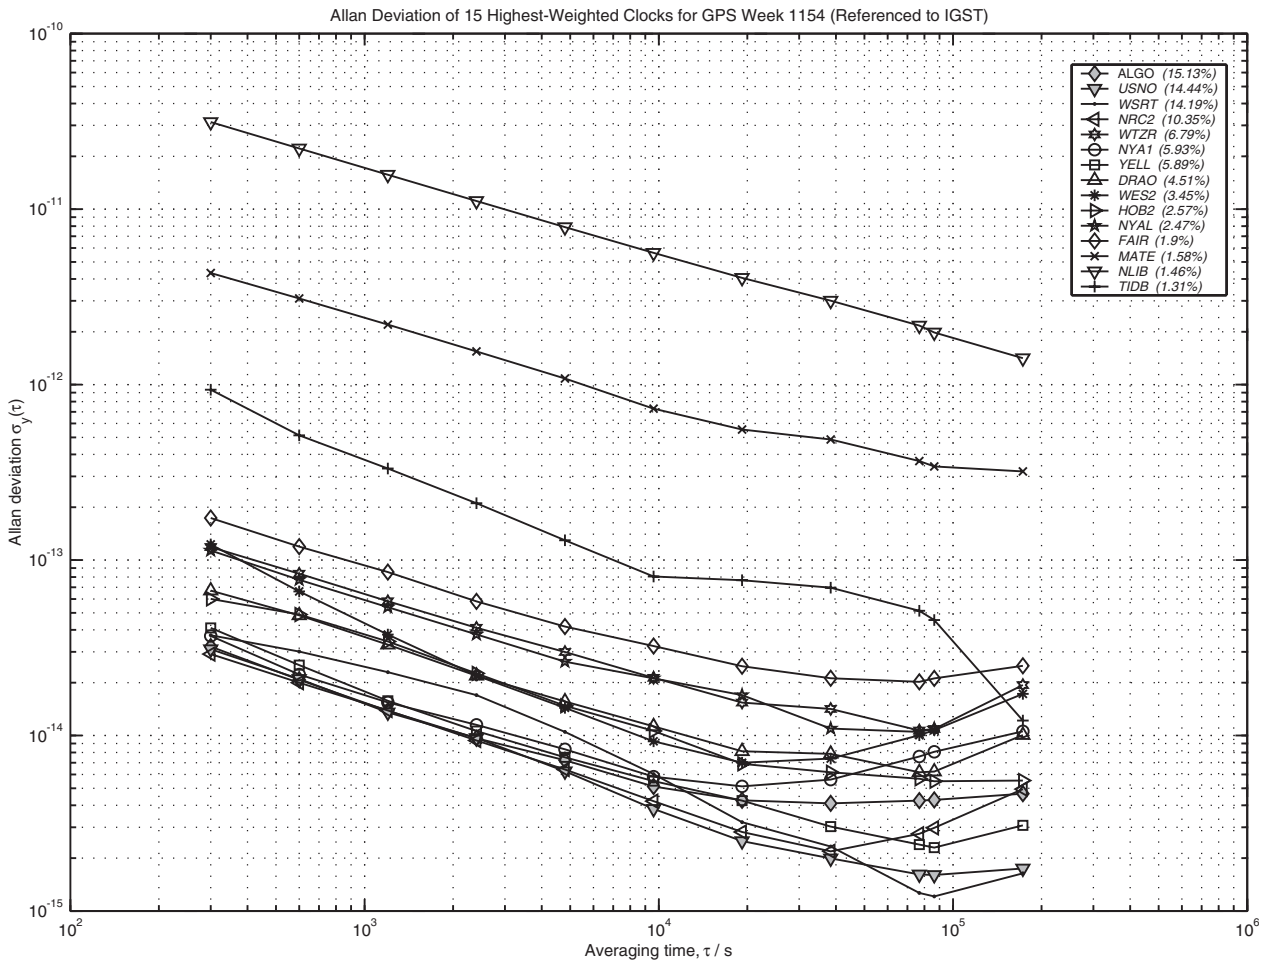

Figure 5. Allan deviation plot of the 15 highest-weighted clocks in the IGS timescale IGST for GPS week 1154 (17–23 February 2002). The clocks are referenced to IGST and no detrending has been performed. Phase and frequency breaks have been removed from each clock for the creation of this plot.

largest anomalous excursions (as well as several smaller ones) of IGRT-IGST on 26 January 2001 (MJD 51935) and 24 May 2001 (MJD 52053) are due to poor alignment of the IGS rapid clock estimates to GPST during the combination process and are not attributable to IGRT or IGST.

In addition to the tracking receiver clocks, the satellite clocks also participate in the ensemble timescale. Some of the newer generation Block IIR satellites, which are equipped with improved rubidium standards, carry weights of about 1%. The improved stability of the IIR rubidiums is illustrated in figure 7, which compares these six spacecraft to the average performances of the rubidium and the caesium clocks in prior GPS generations. Up to about a day, all the IIR clocks are more stable except for SVN44 (PRN28). SVN51 (PRN20) has the best stability at all averaging intervals for the month shown. Also shown in figure 7 is the Allan deviation of GPST as observed by the IGS, for intervals longer than 1 day. (The daily linear alignment of the IGS clocks to GPST makes intervals less than 1 day meaningless.) Interestingly, three Block IIR satellites (PRN20, PRN14, and PRN13) have multi-day stabilities better than GPST itself. This observation demonstrates that the stability of GPST is not inherently limited by the quality of the satellite frequency standards, as is generally assumed.

5. Accuracy and precision of carrier-phase time/frequency transfer

The 'absolute' accuracy of geodetic GPS clock estimates (neglecting instrumental calibration biases) is determined solely by the pseudorange data. More properly, it is the average over the chosen analysis span of all the pseudorange data involving a particular clock that determines the accuracy of the estimates for that clock. For an analysis arc of 24 h using a global tracking network sampled at 5 min intervals, the formal error of each clock estimate is typically about 100 ps to 125 ps, assuming each pseudorange observation has an uncertainty of 1 m and using standard analysis strategies and *a priori* variances. The thermal noise figure for pseudorange data is only a few centimetres, but the effects of multipath are much larger and account for the widely adopted 1 m value. On the other hand, the far more precise carrier-phase data determine the detailed time variation, and hence precision, of the clock estimates. Several experiments seem to bear out the expectation of timing precisions (or frequency stabilities) better than the formal errors, down to the level of 10^{-15} or better at 1 day intervals [22, 23]. But these tests have only been conducted for zero-length or very short baselines where an independent ground truth can be established. Due to

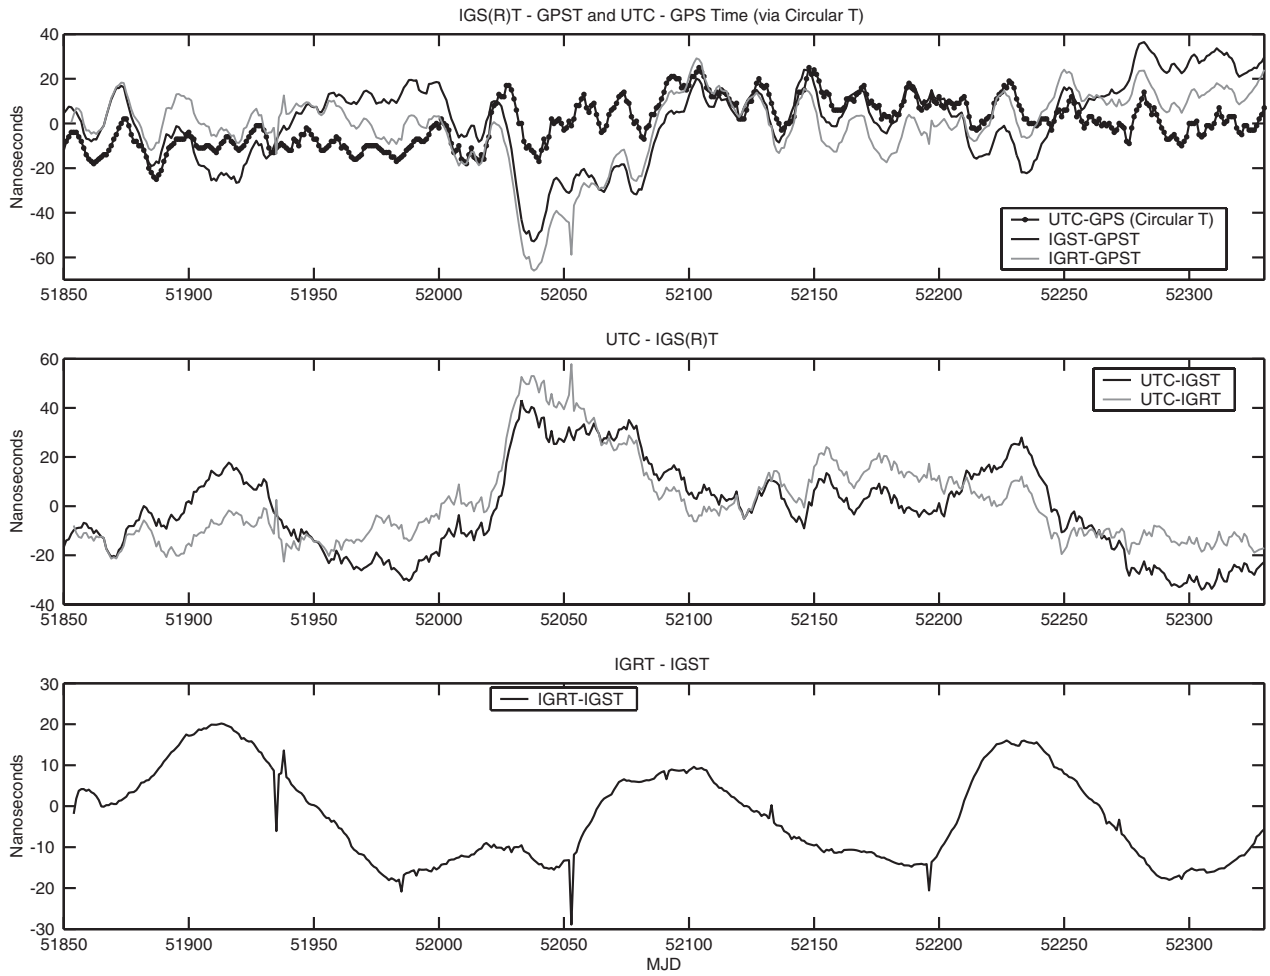

Figure 6. Comparison of the IGS timescales IGST and IGRT against GPST and against UTC (modulo leap seconds) for the period 5 November 2000–24 February 2002. The top plot shows UTC minus GPST via the BIPM *Circular T* as well as the IGS timescales IGST and IGRT minus GPST. The middle plot shows each of UTC-IGST and UTC-IGRT obtained by differencing the time series in the top plot which assumes that GPST via *Circular T* is equivalent to GPST. This assumption can lead to errors at the 10 ns level or greater because of the differing strategies used to track GPST. Finally, the bottom plot shows IGRT minus IGST.

differences in common-mode error cancellation, such results cannot be safely extrapolated to longer baselines where it is more difficult to determine the true time transfer precision, especially for averaging intervals much less than a day.

An internal test of actual clock measurement accuracy can be performed by comparing independent clock estimates at the boundaries between consecutive 1 day analysis arcs. This is analogous to the classic geodetic repeatability test using a time series of position determinations. To ensure that the extrapolation error due to instabilities in the station clocks is minor, only those receivers equipped with H-maser external frequency standards are included. Alternatively, one could consider the clock differences during overlapping intervals, which would not require high-stability oscillators. Larson *et al* [20] have done this for a single long baseline to remove the discontinuities at arc boundaries. As an estimate of measurement accuracy, however, that approach suffers from the unknown correlation between the overlapping arcs. For this study we prefer to avoid that difficulty by using statistically independent clock estimates from adjacent, non-overlapping analyses. Similar, preliminary results were reported by Ray *et al* [24] based on clock solutions posted by the US Naval Observatory (USNO). Those were for pairs of stations, however. Here individual stations are analysed from the combined clock results of the IGS for a 67 week period. Preliminary results for a shorter span were reported by Senior and Ray [25]. Time series of day-boundary clock discontinuities are obtained for 30 tracking receivers, together with computed rms deviations. From these we can infer bounds on the true accuracy and stability of the clock estimates.

5.1. Observed day-boundary clock discontinuities

Our raw data set consists of the combined clock estimates of the IGS for a subnetwork of the global tracking network. The IGS final clock series is preferentially used here starting 29 October 2000 for the 15-month period through 2 February 2002. Since the IGS rapid clock series contain some stations not included in the finals, such as the two timing lab sites BRUS and NPLD, these have been used when finals data are unavailable or when the rapids yield 25% or more day-boundary clock jumps. The clock estimates for 30 tracking receivers, all equipped with H-maser frequency standards, have been analysed. All

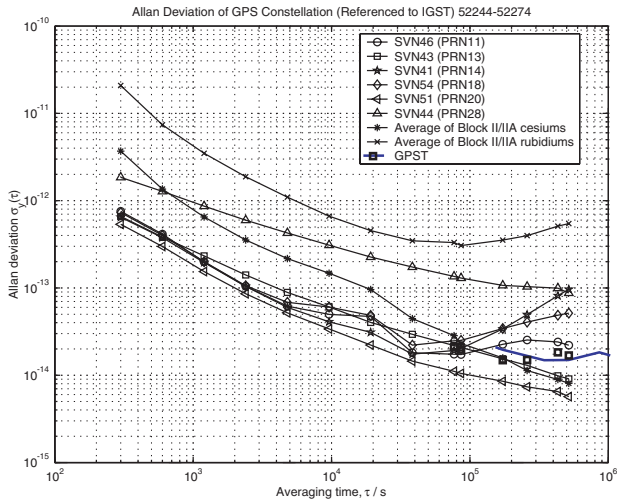

Figure 7. Allan deviation plot of each Block IIR satellite clock as well as a plot of the average Allan deviation of the Blocks II/IIA cesium and II/IIA rubidium satellites during December 2001. All clocks are referenced to IGST and no data editing has been applied.

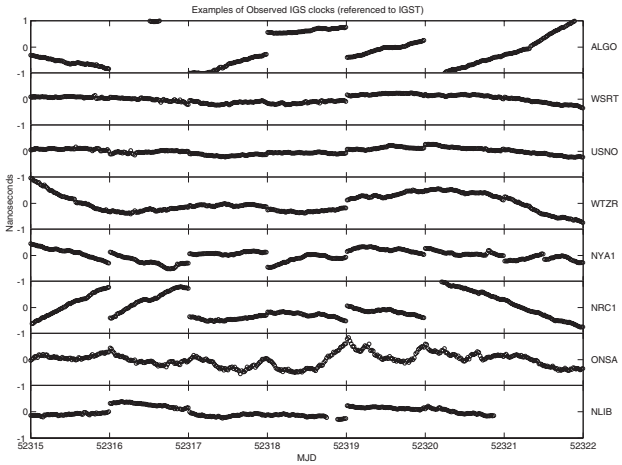

Figure 8. Time series of IGS clock estimates (after internal timescale realignment) for a sample of eight stations equipped with H-maser frequency standards to illustrate day-boundary clock discontinuities. The period shown is GPS week 1153 (10–16 February 2002). A separate quadratic has been removed over the entire interval from each series for the purpose of plotting. Only every fifth 5 min point is plotted.

were referenced to the new timescales described above, IGST or IGRT.

Figure 8 illustrates the day-boundary clock jumps for eight stations during a representative week. The algorithm for measuring day-boundary clock jumps must recognize and ignore a variety of data anomalies. First, each 2 day span is detrended using a linear fit. The interval between clock observations of adjacent days closest to midnight must not exceed 30 min to avoid excessive interpolation error due to oscillator instability. The formal error of the clock difference at the midpoint between the two days must be less than 500 ps to ensure adequate observability. The rms variation of the clock values for each individual day must be less than 150 ps, compared with typical H-maser variations of about 45 ps over 24 h, to avoid confusion by other data problems unrelated to time transfer. Any clock jump greater than 5 ns is rejected as an outlier. The following results are not sensitive to these specific selection criteria. Statistics for the resulting day-boundary clock jumps are given in table 2, together with the receiver and antenna model used at each site. Antennas equipped with a radome cover are so indicated.

The results in table 2 are listed in order of increasing rms clock jumps. Not all stations are equally well sampled and results for some are clearly not robust, particularly NOT1 and KOKB. Several stations show distinct changes in clock-jump behaviour with time: WES2 (Westford, MA, USA), MATE (Matera, Italy), YELL (Yellowknife, NWT, Canada), ALGO (Algonquin, Ontario, Canada), HOB2 (Hobart, Tasmania, Australia), and NRC1 (Ottawa, Ontario, Canada). For each of these cases except HOB2, day-boundary statistics are also computed for the different periods of performance that can be distinguished. The variations at HOB2 are sufficiently complex that this is not done; this station is considered in detail below. For WES2, two different receiver types have been used and each period is considered separately. The MATE clock performance appears to sharply degrade following a brief data gap but recovered later after the receiver was replaced. At YELL, ALGO, and NRC1, the clock jumps are much greater during winter than during other times of the year, by factors of two to three. Figure 9(*a*) shows the day-boundary jumps at NRC1 for the first nine months, but the winter increase does repeat in the most recent data. Of the two receivers at Ny Alesund, Norway (NYA1 and NYAL), which are both connected to the same H-maser frequency standard but which use different antennas, NYA1 is affected by a period of larger clock jumps mostly during May–June 2001. A similar pattern is not evident in the NYAL data. Likewise, AMC2 (Colorado Springs, CO, USA) has very large outliers (about 3 ns jumps) at various times throughout the period. Because the numbers of outliers are small compared to the total data sets, these have not been subdivided. Instead, a second set of statistics has been computed for each station after removing outliers greater than three times the unedited rms. In all other cases, the day-boundary clock jumps do not show clear variations with time, although some are not well sampled.

5.2. Analysis of time transfer accuracy

The mean values for the clock jumps tabulated in table 2 are always comparatively small, except for the sparsely sampled cases of FORT and KOKB. This is shown graphically in figure 10. The FORT station is plagued with poor data quality due to its older generation receiver (AOA TurboRogue SNR-8000), which cannot track L2 under conditions of high ionospheric activity, causing the data coverage to be fragmentary most of the time during this period near solar maximum. The overall smallness of the mean clock jumps 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. To examine the distribution of clock jumps, histograms for the well sampled USNO and WSRT stations are shown in figure 11 compared with Gaussian distributions of like mean and standard deviations. All points outside the plot range are placed in the endpoint bins. The WSRT distribution matches a Gaussian quite closely, whereas

Table 2. Day-boundary clock discontinuities using IGS combined clock products (finals for GPS weeks 1086–1151; rapids for GPS weeks 1087–1153).

IGS site	Mean jump/ps	Rms of jumps/ps	Number of jumps	Receiver model	Antenna model	Remarks
NOT1	66	80	2	TRIMBLE 4000SSI	TRM29659.00	
GODE	4	173	34	AOA SNR-12 ACT	AOAD/M_T dome	
WTZR	−5	184	301	AOA SNR-8000 ACT	AOAD/M_T	
ONSA	−54	186	66	ASHTECH Z-XII3	AOAD/M_B dome	
BRUS R	−36	207	140	ASHTECH Z-XII3T	ASH701945B_M	
WSRT	−46	222	414	AOA SNR-12 ACT	AOAD/M_T dome	
TIDB	−51	222	172	ASHTECH Z-XII3	AOAD/M_T	
CRO1	73	248	165	ASHTECH Z-XII3	AOAD/M_T	
USNO	7	252	392	AOA SNR-12 ACT	AOAD/M_T	
USUD	−97	256	75	ASHTECH Z-XII3	AOAD/M_T	
NPLD R	−65	280	92	ASHTECH Z-XII3T	AOAD/M_T	
WES2	−43	361	242	Mixed types	AOAD/M_T	All data
	−74	400	143	AOA ROGUE SNR-8000	AOAD/M_T	Until 29 Jun 2001
	2	295	99	ASHTECH UZ-12	AOAD/M_T	Starting 27 Jul 2001
IRKT	−117	375	123	AOA ROGUE SNR-8000	AOAD/M_T	
NYAL	−57	378	209	AOA BENCHMARK ACT	AOAD/M_B dome	
NLIB	161	411	90	AOA ROGUE SNR-8000	AOAD/M_T dome	
FAIR R	−97	422	125	AOA SNR-8100 ACT	AOAD/M_T dome	
PIE1	100	442	290	AOA ROGUE SNR-8000	AOAD/M_T	
MATE	−7	444	271	TRIMBLE 4000SSI	TRM29659.00	All data
	−71	605	95			18 Apr–24 Sep 2001
	27	325	176			All other times
NYA1	−81	456	364	AOA BENCHMARK ACT	ASH701073.1 dome	Has several large outliers
	−63	323	356			Edit eight very large outliers
ALBH R	32	534	222	AOA BENCHMARK ACT	AOAD/M_T dome	
DRAO	37	551	406	AOA BENCHMARK ACT	AOAD/M_T	
YELL	−30	613	382	AOA BENCHMARK ACT	AOAD/M_T	All data
	−63	810	155			Winters (1 Dec–22 Mar)
	−7	431	227			Summers (all other times)
MEDI R	−94	680	40	TRIMBLE 4000SSI	TRM29659.00	
ALGO	29	691	389	AOA BENCHMARK ACT	AOAD/M_T	All data
	14	1097	124			Winters (1 Dec–22 Mar)
	36	373	265			Summers (all other times)
AMC2 R	−40	718	156	AOA SNR-12 ACT	AOAD/M_T	All data
	−8	329	150			Edit six very large outliers
HOB2	−79	816	359	AOA ICS-4000Z ACT	AOAD/M_T	Distinct time-variable behaviour
FORT	373	858	61	AOA ROGUE SNR-8000	AOAD/M_TA_NGS	
NRC1	−28	1009	418	AOA SNR-12 ACT	AOAD/M_T	All data
	−29	1514	159			Winters (1 Dec–22 Mar)
	−27	485	259			Summers (all other times)
METS	53	1105	64	ASHTECH Z-XII3	AOAD/M_B	
KOKB R	227	1223	18	AOA SNR-8100 ACT	AOAD/M_T dome	

R with the site name designates results from IGS rapid clocks; all others from IGS finals. Rapids clocks are used only when a station is unavailable in the finals or when >25% more jumps are available in the rapids.

the USNO data are overly peaked near the mean. The means in both cases are insignificantly different from zero.

The most striking feature of the rms clock jumps listed in table 2 is the very wide range of values observed, from about 170 ps to 1200 ps. Since each day-boundary difference involves two measurements, the implied error in any single clock measurement will be smaller by $\sqrt{2}$ if the correlations between successive days are negligible. Interpreting the results in terms of lower limits on the time transfer accuracy, the performance of the technique is evidently highly station-specific, ranging from about equal to the formal error expectation to nearly an order of magnitude greater. Generally, the choice of receiver and antenna model does not seem to be the determining factor, judging from table 2. However, the best performance of any station equipped with an older TurboRogue (IRKT) is poorer than either the newer AOA (ACT models) or Ashtech receivers, though possibly for other reasons. In addition, the WES2 performance did improve after the TurboRogue receiver was replaced, though the rms change of about 25% is not conclusive. Use of antenna radomes does not appear to affect time transfer performance based on these results, although any static bias would not be detectable.

To investigate possible sources of station-specific time transfer error, it is instructive to consider those cases of clear temporal changes in performance. The HOB2 clock discontinuities are shown in figure 12(*a*). Following an interruption in the time series during April–May 2001, the HOB2 day-boundary clock jumps became much noisier. The gap in the time series is caused by frequent receiver clock resets beginning on 16 April (MJD 52015), which make it impossible to determine reliable day-boundary clock jumps. Inspection of IGS Mail messages and the site log does not indicate any particular event that might be responsible for the change in performance then. Correspondence with station

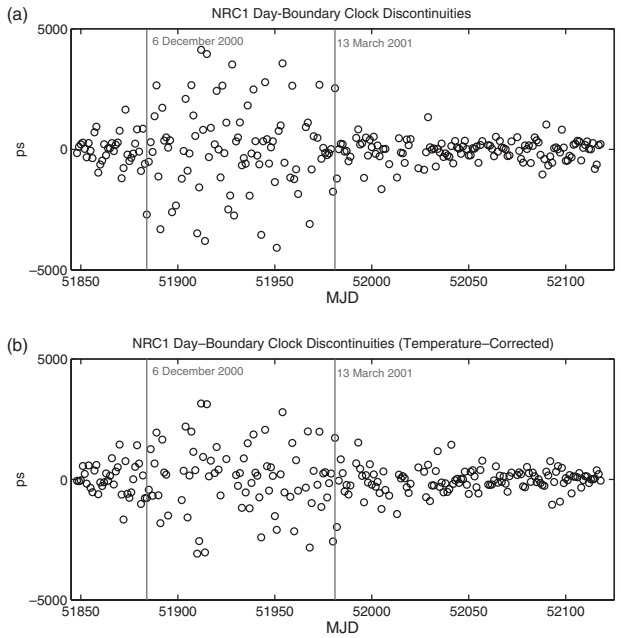

Figure 9. Time series of day-boundary clock discontinuities for NRC1 (Ottawa, Canada). Upper panel (*a*) shows the raw results, with much larger variations during winter than at other times of the year. Lower panel (*b*) shows the same results after removing a linear temperature-dependent trend (see table 3). While the level of discontinuities during winter has been reduced, they remain significantly larger than at other times.

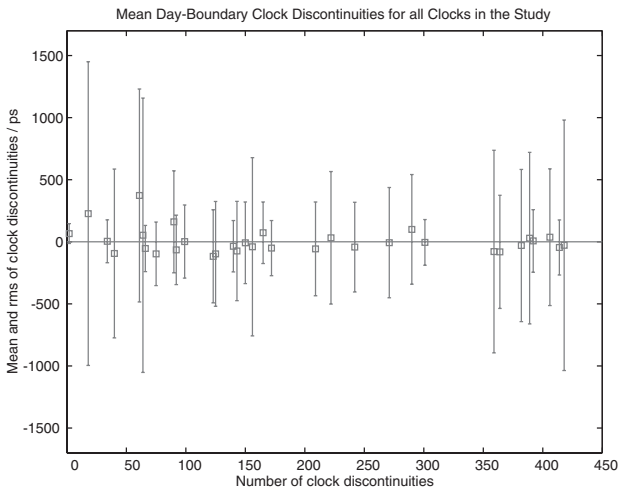

Figure 10. The mean and rms values for the day-boundary clock discontinuities of all stations in the study as a function of the number of discontinuities available.

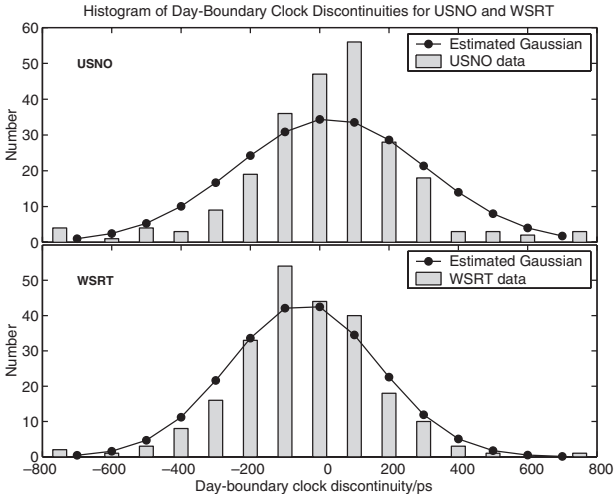

Figure 11. Histograms of the day-boundary clock discontinuities for USNO and WSRT compared with Gaussian distributions for each.

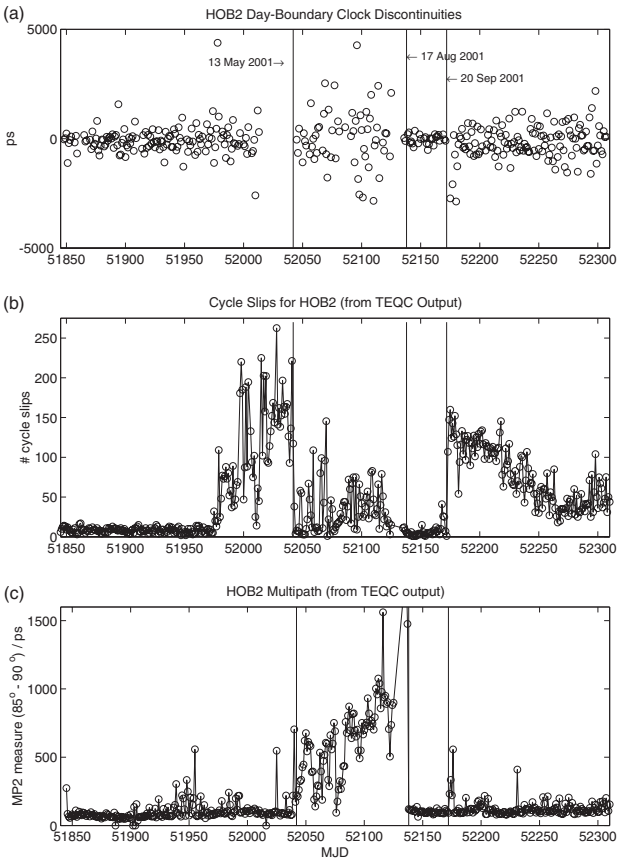

Figure 12. Time series of results for HOB2 (Hobart, Australia). Upper panel (*a*) shows the raw day-boundary clock discontinuities illustrating the large increase in variations after 13 May 2001. Middle panel (*b*) shows the daily number of carrier-phase cycle slips as reported by the quality-checking utility TEQC. Lower panel (*c*) shows the TEQC multipath index MP2 for observations in the elevation angle bin from 85° to 90°.

personnel (P Digney, private communications) revealed that the clock reset period corresponds with a large increase in phase cycle slips. This can be seen in figure 12(*b*), which plots the daily number of cycle slips as determined by the quality-checking utility TEQC [26]. The increase in cycle slips was traced to a sharp drop in the signal-to-noise ratio (SNR) due to corrosion and moisture in the connection of the Heliax 50 Ω underground cable to the antenna. Judging from figure 12(*b*), this problem seems to have begun around 7 March 2001. The connection was repaired, but not until 17 August 2001 (MJD 52138; see IGS Mail no 3477) when the day-boundary jumps and the cycle slips both sharply improved. Apparently, some temporary corrective action was taken earlier, on about 14 May (MJD 52043), after which day-boundary clock jumps were

again measurable but with rms values increased by more than a factor of two. Figure 12(*b*) shows that the frequency of cycle slips improved in mid-May, but not to the level seen earlier in the year before the problem began. Around 20 September 2001 (MJD 52172; see IGS Mail no 3534) further maintenance work was done on the cables at HOB2. This coincides with another abrupt increase in the cycle slips and a moderate increase in the rms clock jumps. We have consulted another diagnostic output from TEQC to check for correlations with the variations in time transfer performance. MP1 and MP2 report the rms variations of the pseudorange multipath at the L1 and L2 frequencies, allowing for unknown biases for each satellite and assuming the effect of phase multipath is negligible [26]. Aggregate MPi statistics are computed by TEQC, as well as values for individual satellites and for 5° elevation angle bins. Inspection of MPi measures for HOB2 fails to reveal any patterns clearly correlated with either the day-boundary clock jumps or the cycle slip trends except for MP2 at high elevation angles. Figure 12(*c*) plots MP2 for the 85° to 90° bin as a function of time. In mid-May 2001 (around MJD 52041) the MP2 (85°–90°) measure begins a sharp increase closely matching the greater clock jump rms in figure 12(*a*) and coincident with a dramatic drop in the number of cycle slips (figure 12(*b*)). The pattern reversed in August 2001 when the antenna cable was repaired.

We hesitate to assert that high-elevation pseudorange multipath at the L2 frequency is directly or causally linked to clock measurement accuracy. For one thing, the correlation seen for HOB2 is not seen for other sites with time-varying day-boundary clock jumps. More likely, the high-elevation MP2 index reflects some other underlying error source, which also affects the number of cycle slips and leads to poorer time transfer accuracy. Without further information it is not possible to isolate the ultimate source. However, the sequence of events at the station suggests that perhaps some attempt in mid-May 2001 to ameliorate the cable corrosion problem, prior to its ultimate repair in August, may have improved the SNR but inadvertently introduced some other signal problem, such as an impedance mismatch giving rise to internal multipath at the L2 frequency. After the problems were temporarily resolved in August 2001 further cable work a month later apparently introduced another type of signal degradation associated with greater cycle slips but no large increase in MP2.

The case with MATE is less informative because no other quantitative measure seems to correlate with the degradation in time transfer accuracy beginning in late April 2001. There is no public record of any specific events that could have affected the data quality, but correspondence with station personnel (D Del Rosso, private communication) reveals that the frequency synthesizer of the H-maser was adjusted on 26 April 2001. This inadvertently caused a tracking problem in the GPS receiver that was corrected on 2 May. These events caused a gap in our time series of day-boundary clock jumps, but it is unclear how the degraded time transfer accuracy afterwards could have resulted. The receiver was swapped on 25 September 2001, to a new Trimble 4000SSI unit (see IGS Mail no 3538), but it was not stated whether this action was taken to correct any specific problem. An older firmware version was installed in the new receiver, possibly to cope with a known problem in version 7.29. (The Trimble 4000 firmware problem was brought to our attention by D Hogarth, private communications.) Whatever the reason for the receiver change, the MATE time transfer performance improved by nearly a factor of two afterwards.

Interpretation of the NRC1, ALGO, and YELL clock measurement results is complex. The indication of a seasonal trend in the day-boundary jumps (see figure 9(*a*)) suggests checking for temperature dependences. Of the stations in our study set, meteorological data are available for 11. We have extracted daily mean temperatures during the period 29 October 2000 to 28 July 2001, and tested for correlations between the day-boundary clock jumps and differences in mean daily temperatures. The results are given in table 3. Only ALGO and NRC1 exhibit temperature dependences that are significant at a level greater than 5 sigmas, while WTZR, DRAO, and ALBH show effects at about 4 sigmas. None of the remaining six stations shows any correlation of day-boundary clock variations with temperature.

Before examining the ALGO and NRC1 results more carefully, let us first consider the implications of those stations with no evidence of temperature sensitivities. Various studies [22, 23, 27] have demonstrated the effects of temperature variations on individual components of the tracking systems. Typical sensitivities for the receivers alone are of order $\pm 100\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}$ with large variations among individual units, including for the same model [28]. Common RF antenna

Table 3. Day-boundary clock jumps versus temperature changes (GPS weeks 1086–1124).

IGS site	Correlation coefficient	Rms residual/ps	Slope/ps $^{\circ}\mathrm{C}^{-1}$	Slope sigma/ps $^{\circ}\mathrm{C}^{-1}$	Number of points	Range of temperature jumps/$^{\circ}\mathrm{C}^{-1}$
WTZR	0.3615	164	25.3	5.5	145	13.4
FAIR	0.0068	232	0.5	9.7	61	17.6
USNO	0.0435	293	3.7	6.4	177	20.1
WES2	0.1137	377	12.2	9.1	139	21.0
ALBH	−0.4167	456	−127.3	35.0	65	8.3
MATE	−0.0520	503	−12.6	19.8	152	12.2
DRAO	−0.2608	555	−54.6	14.6	193	16.8
YELL	0.0671	617	8.6	9.8	172	32.0
ALGO	−0.5973	545	−101.3[a]	9.6	202	26.4
NRC1	0.5224	919	155.9[a]	16.7	235	22.9
METS	0.1359	1221	80.7	131.6	22	8.6

[a] Slope determination significant by more than 5 sigmas.

cables have thermal sensitivities around $1\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}\,\mathrm{m}^{-1}$ and cable runs frequently exceed 50 m. Much better cable types are available, having temperature coefficients of about $0.03\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}\,\mathrm{m}^{-1}$ [29], which are coming into use at timing labs though not so commonly at other geodetic installations. Trenching and other means of environmental isolation are sometimes used to minimize the effects of temperature cycling on antenna cables. We have previously shown [30] that the short-term (diurnal) temperature stability of AOA Dorne Margolin choke–ring antennas is better than $2\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}$. However, that study could not exclude the possibility of longer-term temperature-induced effects due to sensitivity in the daily average of the pseudorange observations.

The results in table 3 for USNO, which show a long-term null temperature sensitivity of $3.7 \pm 6.4\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}$, can now be used to extend our earlier antenna results. As documented by Ray and Senior [30], the USNO cable and receiver systems are very well isolated from environmental variations. The only significant uncontrolled component of the station hardware is the outside antenna unit, an AOA Dorne Margolin T choke–ring (see table 2). The uncertainty in our estimate of the long-term USNO station thermal coefficient therefore places an upper limit of $10.1\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}$ (1 sigma) on any possible pseudorange-induced thermal variation due to this type of antenna.

Returning to ALGO and NRC1, their inferred temperature sensitivities are very significant, about 10 sigmas. Kouba (2001, unpublished report) has documented diurnal temperature dependences in the raw pseudorange data from ALGO and DRAO, with the effects being larger at ALGO by a factor of two to three. The antenna cables in both cases are underground and should not be especially sensitive to diurnal changes. The cable type at ALGO is Andrew LDF4 (Caissy, private communications), which has a temperature coefficient between $0.027\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}\,\mathrm{m}^{-1}$ and $0.061\,\mathrm{ps}\,^{\circ}\mathrm{C}^{-1}\,\mathrm{m}^{-1}$. For this reason, Kouba concluded that the thermal variations are caused by either the antennas or the receivers. Our results above for USNO exclude any significant effect due to the Dorne Margolin antennas, assuming that individual units behave similarly. Thus, the temperature sensitivities at ALGO, NRC1, and to a lesser extent at DRAO, point to inadequate isolation of the receivers from external environmental changes. Station personnel for ALGO and NRC1 (M Caissy, private communications) indicate that this possibility cannot be excluded. Efforts are underway to monitor temperature variations near those receivers. However, other components of these systems must also be considered, such as the connectors and the antenna power splitter used at NRC1.

Regardless of the underlying source of the temperature sensitivity, this effect alone cannot fully account for the magnitude of the clock jumps at ALGO and NRC1, and is evidently unrelated to the seasonal variation at YELL. That can be seen in the rms day-boundary clock residuals after removing temperature-dependent trends (table 3), which are reduced by only minor amounts compared with the original rms values (table 2). Even more important is the persistence of apparent seasonal trends after removal of the temperature dependences. Table 4 gives the day-boundary statistics for ALGO and NRC1 after accounting for the temperature trends in table 3. The rms clock jumps remain two to three times greater in winter compared with summer. Of the diagnostic metrics we have checked for data quality and station performance (such as those discussed above for HOB2), only the multipath MPi measures from TEQC show any correlations with the ALGO or NRC1 day-boundary clock jumps. However, the daily multipath variability is much smaller (roughly 10% to 50% between seasons compared with factors of two to three for the clock jumps), and the multipath indices are smaller in winter than in summer. Despite these inconclusive results, we speculate that greater wintertime multipath might occur due to signal reflection off snow-covered surfaces near the antennas, leading to poorer time transfer accuracy during those periods. It should be noted that MPi measures multipath variations with periods from 1 min (IGS data sampling is 30 s) to about 3 h (the typical satellite pass) and is not sensitive to longer wavelength or quasi-static biases. The ALGO and YELL antennas are both mounted over concrete pillars with the L1 and L2 phase centres being 21 cm and 22.8 cm above the tops of the pillars, and with a space of 13.5 cm between the bottom of each antenna choke ring and its pillar. This configuration is likely to be prone to standing-wave reflections, especially during winter when the pillar tops are covered by snow or ice. Elosegui *et al* [31] described serious back-scattering problems for a similar antenna configuration, although not specifically involving snow cover. In a detailed study of phase multipath in a regional network in Germany, Wanninger and May [32] found marked changes in the daily multipath behaviour caused by snow but not by rainfall.

Table 4. Day-boundary clock discontinuities after removing temperature-dependent linear trends (GPS weeks 1086–1124).

IGS site	Mean jump/ps	Rms of jumps/ps	No of jumps	Remarks
ALGO	0	545	202	All data
	−31	913	48	Winter
	10	355	154	Summer
NRC1	0	919	235	All data
	35	1406	77	Winter
	−17	532	158	Summer

Without further information, the causes for the general dispersion of time transfer accuracies among the various stations in table 2 cannot be further isolated. We should expect essentially any factor that can degrade the pseudorange and/or carrier-phase observables to be considered a candidate. Because multipath is usually the dominant observational error, it deserves special attention. Doing so will require well constructed models tailored for the individual stations, a task beyond the scope of this paper. Other local factors to be considered include temperature sensitivities (examined above for some stations), RFI and electromagnetic environment, internal impedance mismatches, receiver firmware, and so forth.

5.3. Analysis of time transfer stability

The precision of clock estimates within a given analysis arc (and hence the frequency stability) is generally expected to be much better than either the formal errors or the absolute accuracy measures because the relative clock values are determined mostly by the carrier-phase data (under normal

circumstances). However, the actual performance has not been well characterized experimentally. For the few long baselines that have been well studied, the observed frequency stabilities are rarely better than about 2×10^{-15} at 1 day intervals [20]. Also, the question remains whether the noise of the carrier-phase time transfer process behaves as white noise (Allan deviation as τ^{-1}) or random walk ($\tau^{-0.5}$) in the time domain. If the time transfer errors were perfectly white, the formal clock measurement uncertainties would imply 1 day stabilities at the level of 1.3×10^{-15} assuming no other effects are significant (such as intrinsic clock instability), but a very poor 3.8×10^{-13} at 5 min intervals.

Given the very large dispersion in accuracy performance among IGS stations, it is natural to consider whether short-term stability varies similarly. Figure 13 shows the average Allan deviations at 300 s intervals compared with the rms day-boundary clock jumps for 23 of the IGS stations in table 2 (through GPS Week 1124 only). For each station and day, Allan deviations were computed from the realigned IGS clock products for intervals of 300 s, 10^3 s, 10^4 s, and 3 h . Then average values were computed for each station and each interval rejecting as outliers any values greater than $(8 \times 10^{-13})\sqrt{300\,\mathrm{s}/\tau}$. For ALGO, HOB2, MATE, and NRC1 with variable day-boundary behaviours, separate points are plotted for each sub-interval and those points are joined by lines. Five stations (ONSA, TIDB, CRO1, IRKT, and METS) have much poorer short-term stabilities than the rest. Inspection of the time domain plots for ONSA and METS shows clear diurnal variations, likely to be related to environmentally induced effects, while the CRO1 results appear 'noisy' at high frequencies, for unknown reasons. In each of these five cases, aspects of the local station configurations are more likely to be responsible for the poor short-term stability than is the time transfer method. For this reason, these stations have been excluded from a fit of the average Allan deviations versus rms day-boundary discontinuities. Our objective is to determine an approximate lower envelope of the short-term stability regime. The existence of a correlation between these distinct measures of long-term accuracy and short-term stability is reinforced by the observation that the two separate periods for ALGO, HOB2, and NRC1 parallel the general trend of the other included stations; only MATE behaves somewhat differently.

By extrapolating the trend in figure 13 to an intercept of zero rms discontinuities (i.e. presumably perfectly accurate time transfers), we can infer a limiting value for the short-term stablity due to the time transfer method. Repeating for

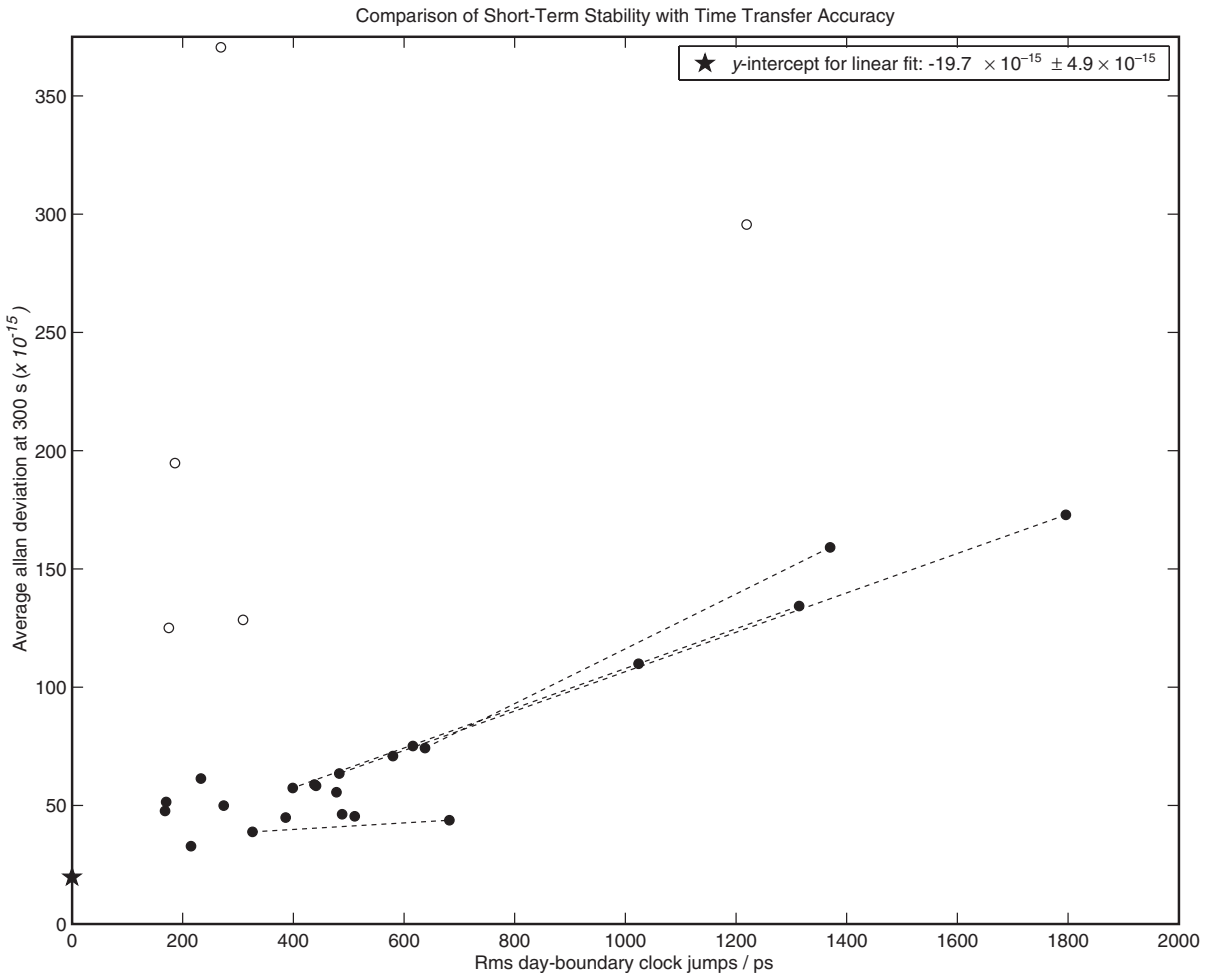

Figure 13. The average Allan deviations computed at 300 s intervals versus the rms day-boundary clock discontinuities for all stations. Two connected points are shown for each of the stations ALGO, HOB2, MATE, and NRC1, for the intervals of markedly different behaviour. The stations with local timing instabilities (ONSA, TIDB, CRO1, IRKT, and METS, indicated by open circles) are not included in the linear fit for the general trend.

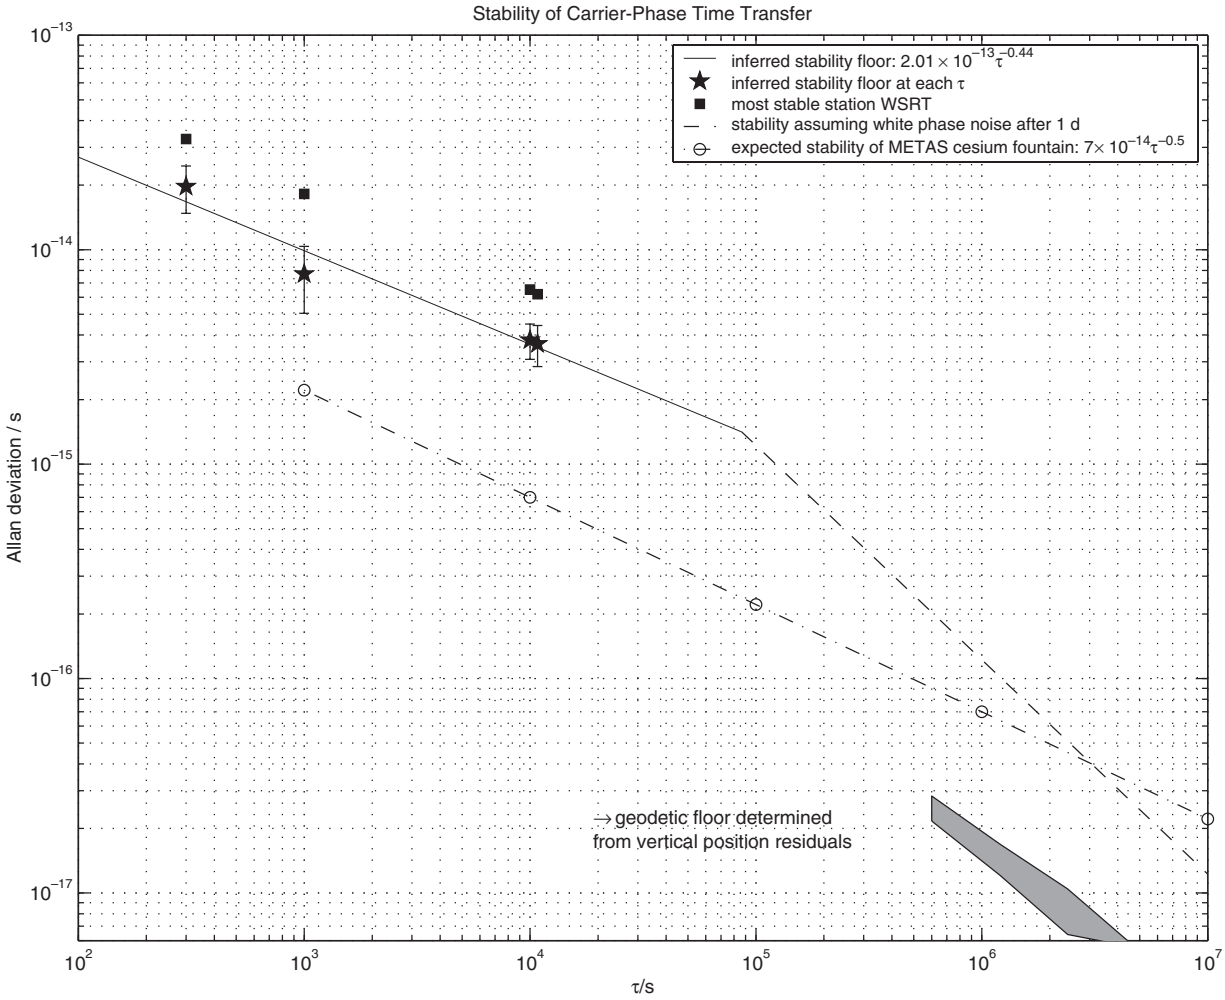

Figure 14. The inferred floor for Allan deviation stability of the carrier-phase time transfer method (——) compared with the average stability for WSRT, a station among the best in terms of overall stability. Also shown (- - - -) is the expected behaviour if the time transfer stability behaves as a white noise process with an instantaneous uncertainty equal to the typical measurement formal errors (115 ps), which is consistent with the observed time transfer accuracy for the best performing stations. For reference, the target stability for the METAS cesium fountain (— · —) [34] is also shown, together with an estimate of the fundamental geodetic stability floor derived from the vertical position residuals for ALGO, NRC1, USNO, and WSRT (see text).

10^3 s, 10^4 s, and 3 h gives the values plotted in figure 14 and an overall stability limit power law of $2.01 \times 10^{-13}\tau^{-0.44}$. The formal error for the determination of the power law exponent is 0.07, so the behaviour is not significantly different from $\tau^{-0.5}$, as expected for a random walk error process. At an interval of 1 day, the inferred stability is 1.4×10^{-15}. Also shown in figure 14 are the average Allan deviations for WSRT, a station among the best in terms of overall short-term stability and long-term accuracy. Our estimate for the limiting time transfer stability is about half that of WSRT.

We conclude that the short-term time transfer stability is indeed much better than implied by the instantaneous formal errors (as good as 2×10^{-14} versus 3.8×10^{-13} at 300 s), at least for intervals shorter than 1 day (see figure 14). The benefit of implicit carrier-phase 'smoothing' to improve the frequency stability, which must introduce important temporal correlations, probably explains the observed random walk rather than white noise behaviour. In this regard, it is noteworthy to recall that GPS carrier-phase observations can be regarded as equivalent to an integrated Doppler data type. At 1 day intervals the inferred stability is approximately equal to the level expected from the formal time transfer measurement errors and a white noise process (1.4×10^{-15} versus 1.3×10^{-15}). We expect, but cannot demonstrate, that the time transfer stability for intervals longer than 1 day will follow a white noise trend for the appropriate station accuracy. This will continue until ultimately limited by longer-term geodetic errors, such as due to unmodelled geophysical effects, and especially by longer-term variations in pseudorange multipath, which are not reflected in the geodetic positioning results. For reference, we have shown an estimate of the fundamental geodetic stability floor in figure 14. This estimate has been derived using the local vertical positioning residuals for the IGS weekly station coordinate combinations for ALGO, NRC1, USNO, and WSRT (results kindly provided by R Ferland, private communications).

In practical terms, however, internal investigations near 1 day and longer are limited by the stability of the underlying timescale, its steering, and the frequency standards themselves [8]. To explore the actual time transfer stability performance

over those intervals will require comparisons between stations equipped with more stable frequency standards, such as the new cold atom standards now under development. While our observational results here accord remarkably well with the expected performance predicted by Petit and Thomas [33] based on a purely theoretical error budget analysis, an experimental verification will require direct comparison between two or more caesium fountain clocks. To illustrate, we have plotted in figure 14 the target stability for the METAS Cs fountain (see [34], p 95). For stations equipped with such standards, the white noise character of carrier-phase time transfer should become evident between about 1 day and 35 days, if other errors have not intervened by that point.

5.4. Discussion of accuracy and precision results

Our central finding is that the accuracy of carrier-phase clock estimates is, in the best cases, about equal to the formal errors of approximately 115 ps for 1 day arcs. However, the performance can be much poorer, by nearly an order of magnitude, and seems to be highly site-dependent. Since several different antenna and receiver models in common use appear to give similar results, other site-dependent factors are probably responsible for the variable performance. Our study has identified only a few of the local error sources, notably temperature-induced variations, antenna cable problems, and receiver failures. Standing wave multipath effects are also likely. Generally, any effect which degrades the quality of the pseudorange or carrier-phase data must be considered. Multipath errors, which have been well studied for their effects on phase data, deserve particularly close scrutiny for pseudorange and clock effects.

Some authors who have compared carrier-phase results with other time transfer techniques have chosen to remove the day-boundary discontinuities analysed here. Larson *et al* [20], for example, use half-day overlaps between successive 4.5 day analysis arcs to estimate and remove the discontinuities. They found a discontinuity rms of 222 ps over a 236 day period for the baseline between USNO and AMC2. This is somewhat smaller than our own rms estimate for 1 day arcs of 252 ps for USNO and 329 ps for AMC2 (after outlier removal). The difference could conceivably be related to the longer analysis arcs used by Larson *et al*, which might yield slightly more accurate time transfer results, although there are also many other differences in the analyses. However, the approach of successively removing discontinuities presents some drawbacks that should be considered. If, for example, the statistical relationship between the time transfer errors of successive analysis arcs can be described by a Gauss–Markov process, then the accumulated error in the adjusted clock time series will have a variance that grows with time as

$$\sigma^2(t) = 2T\sigma_0^2[t - T + T\mathrm{e}^{-t/T}],$$

where T is a characteristic correlation time constant for the clock errors between successive analysis arcs and σ_0 is the standard error for clock estimates of an individual arc (see [13] for background). Such an accumulation error could degrade the frequency stability if the correlation time is not very short (i.e. less than 0.5 day) and should be checked whenever this procedure is applied; our data set is not suited to analyse the autocorrelation properties. Dach *et al* [35] have used simulations to illustrate other problems that can be introduced by long-term accumulation of systematic errors.

Another approach to remove discontinuities between processing arcs, by simply continuing the analysis forward without allowing any interruption in the time series [36, 37], raises other questions. While this will give the appearance of 'error-free' clock estimates, no physical measurement process can actually be free of uncertainty. Potentially, only the information needed to objectively quantify the measurement errors is removed. It could well be that longer analysis arcs give more accurate clock estimates, as expected for the formal errors [37], but an error 'floor' always intervenes in any real physical process to eventually deny $1/\sqrt{N}$ improvements. Whether the floor is significant at 1 day intervals or for longer periods remains to be established. It should also be noted that the time transfer stability could suffer for intervals longer than about 1 day due to the continuation of carrier phase-based correlations that would not otherwise occur with 1 day analysis arcs. That is, the expected transition from random walk to white noise stability at about 1 day (and at a stability of roughly 1.3×10^{-15}) for 1 day analysis arcs could very well lead to random walk behaviour extending to n days for an n-day analysis arc.

We have also seen that the short-term stability of carrier-phase time transfer results varies linearly with the inferred 1 day accuracy. This conclusion is not necessarily expected, since the pseudorange data largely determine the accuracy whereas the carrier-phase data dominate the stability. This result implies that the error sources that affect pseudorange data and accuracy are highly correlated with the quality of phase data. Certainly that should be true for multipath errors, but other effects should be studied as well.

6. Concluding remarks and future prospects

The joint efforts of the IGS, BIPM, and their collaborating groups has produced a rich return of global satellite and tracking receiver clock results. The usefulness of these products was already well demonstrated for geodetic positioning and other applications. The utility of this project for global time and frequency comparisons is now developing rapidly. An internally realized and stable timescale has been synthesized from the IGS clock products, steered loosely to GPST and thereby traceable with good approximation to UTC. The internal timing accuracy and stability of these clock products has been quantified for a large number of the current tracking receivers, including some at timing labs. The results show a very wide range of performances, but the best stations approach the formal error level of about 115 ps. This implies 1 day frequency stabilities of about 1.4×10^{-15}. The method of using day-boundary discontinuities as a quality assessment index has been shown to provide high sensitivity to a range of station-specific factors that can degrade timing performance. This approach is strongly recommended for the evaluation of new station installations and for general diagnostic purposes.

Meanwhile, parallel efforts have developed methods to calibrate the instrumental delays in geodetic receivers, currently to the level of a few nanoseconds, and these are

being applied to field receivers. When sufficient numbers of calibrated IGS stations are colocated at timing labs, then direct links of the IGS timescales can be made to TAI. This should also permit carrier-phase clock results to be used, together with other techniques, in the formation of TAI. The fast turnaround of the IGS rapids could be especially useful in realizing a predicted form of TAI.

Acknowledgments

We are indebted to all components of the IGS without whose accomplished efforts this work would not have been possible. We also thank Mark Caissy, François Lahaye, Paul Digney, and Domenico Del Rosso for their assistance in diagnosing some of the station-related effects, as well as Kristine Larson and Jan Kouba for helpful discussions. Doug Hogarth pointed out the firmware change made at MATE in September 2001 and Remi Ferland generously provided time series of IGS coordinate residuals.

References

[1] Ray J and Petit G 1999 IGS/BIPM pilot project to study time and frequency comparisons using GPS phase and code measurements—interim report to the CCTF *Presented to the 14th Meeting of the Consultative Committee for Time and Frequency (BIPM, Sèvres, France)*

[2] Allan D and Weiss M 1980 Accurate time and frequency transfer during common-view of a GPS satellite *Proc. 1980 IEEE Freq. Control. Symp. (Philadelphia, PA, 1980)* pp 334–56

[3] Larson K M and Levine J 1999 Carrier-phase time transfer *IEEE Trans. Ultrason., Ferroelect., Freq. Control.* **46** 1001–12

[4] Zumberge J, Heflin M, Watkins M and Webb F 1997 Precise point positioning for the efficient and robust analysis of GPS data from large networks *J. Geophys. Res.* **102** 5005–17

[5] Kouba J and Springer T 2001 New IGS station and satellite clock combination *GPS Solutions* **4** 31–6

[6] Petit G, Jiang Z, White J, Beard R and Powers E 2001 Absolute calibration of an Ashtech Z12-T GPS receiver *GPS Solutions* **4** 41–6

[7] Petit G, Jiang Z, Moussay P, White J, Powers E, Dudle G and Uhrich P 2001 Progresses in the calibration of 'Geodetic Like' GPS receivers for accurate time comparisons *Proc. 15th European Frequency and Time Forum* pp 164–6

[8] Senior K, Koppang P, Matsakis D and Ray J 2001 Developing an IGS timescale *Proc. 2001 IEEE Freq. Control. Symp. (6–8 June 2001)* pp 211–18

[9] Senior K, Koppang P and Ray J 2003 Developing an IGS time scale *IEEE Trans. Ultrason., Ferroelect., Freq. Control.* **50** (6) at press

[10] Ahn I and Brown R 1986 Assessing the validity of suboptimal 2-state clock models *Proc. 18th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting (Washington, DC, 2–4 December 1986)* pp 421–36

[11] Barnes J and Allan D 1985 Time scale stabilities based on time and frequency Kalman filters *Proc. 39th Annual Frequency Control Symp. (Philadelphia, PA, 29–31 May 1985)* IEEE 85CH2186-5, pp 107–12

[12] Stein S and Filler R 1988 Kalman filter analysis for real time applications of clocks and oscillators *Proc. 42nd Annual Symp. on Frequency Control (Baltimore, MA, 1–3 June 1988)* IEEE 88CH2588-2, pp 447–52

[13] Brown R and Hwang P 1997 *Introduction to Random Signals and Applied Kalman Filtering* 3rd edn (New York: Wiley)

[14] Maybeck P 1994 *Stochastic Models, Estimation, and Control, Volumes 1–3, Mathematics in Science and Engineering* vols 141-1, 141-2, 141-3 (Arlington, VA: Navtech Book and Software Store)

[15] Thomas C and Azoubib J 1996 TAI Computation: study of an alternative choice for implementing an upper limit of clock weights *Metrologia* **33** 227–40

[16] Koppang P and Leland R 1995 Steering of frequency standards by the use of linear quadratic Gaussian control theory *Proc. 27th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting (29 November–1 December 1995)* vol 3334 (San Diego, CA: NASA Conference Publication) pp 257–67

[17] Matsakis D, Miranian M and Koppang P 2000 Steering strategies for the master clock of USNO *Proc. ION GPS 2000 (Salt Lake City, UT, 19–22 September 2000)* pp 933–6

[18] Bureau International des Poids et Mesures *Circular T* Pavillon de Breteuil, F-92312 Sèvres, Cedex

[19] Young L, Jefferson D, Lichten S, Tjoelker R and Maleki L 1996 Formation of a GPS-linked global ensemble of hydrogen masers and comparison to JPL's linear ion trap *Proc. 1996 IEEE International Frequency Control Symp.* pp 1159–62

[20] Larson K, Levine J, Nelson L and Parker T 2000 Assessment of GPS carrier-phase stability for time-transfer applications *IEEE Trans. Ultrason., Ferroelect., Freq. Control.* **47** 484–94

[21] Matsakis D, Senior K and Breakiron L 1999 Analysis noise, short-baseline time transfer, and a long-baseline GPS carrier-phase frequency scale *Proc. 31st Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting (Dana Point, CA, 7–9 December 1999)* pp 491–504

[22] Bruyninx C, Defragne P and Sleewaegen J-M 1999 Time and frequency transfer using GPS codes and carrier phases: onsite experiments *GPS Solutions* **3** 1–10

[23] Petit G, Thomas C, Jiang Z, Uhrich P and Taris F 1998 Use of GPS Ashtech Z12T receivers for accurate time and frequency comparisons *Proc. 1998 IEEE International Frequency Control Symp.* pp 306–14

[24] Ray J, Arias F, Petit G, Springer T, Schildknecht Th, Clarke J and Johansson J 2001 Progress in carrier phase time transfer *GPS Solutions* **4** 47–54

[25] Senior K and Ray J 2001 Accuracy and precision of GPS carrier phase clock estimates *Proc. 33rd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting (Long Beach, CA, 27–29 November 2001)* pp 199–217

[26] Estey L and Meertens C 1999 TEQC: the multi-purpose toolkit for GPS/GLONASS data *GPS Solutions* **3** 42–9

[27] Overney F, Schildknecht T h, Beutler G, Prost L and Feller U 1997 GPS time transfer using geodetic receivers: middle-term stability and temperature dependence of the signal delays *Proc. 11th European Frequency and Time Forum* pp 504–8

[28] Schildknecht Th and Dudle G 2000 Time and frequency transfer: high precision using GPS phase measurements *GPS World* **11** 48–52

[29] Powers E, Wheeler P, Judge D and Matsakis D 1998 Hardware delay measurements and sensitivities in carrier phase time transfer *Proc. 30th Precise Time and Time Interval Meeting* pp 293–305

[30] Ray J and Senior K 2001 Temperature sensitivity of timing measurements using Dorne Margolin antennas *GPS Solutions* **5** 24–30

[31] Elosegui P, Davis J L, Jaldehag R T K, Johansson J M, Niell A E and Shapiro I I 1995 Geodesy using the global positioning system: the effects of signal scattering on estimates of site positions *J. Geophys. Res.* **100** 9921–34

[32] Wanninger L and May M 2001 Carrier-phase multipath calibration of GPS reference stations *Navigation* **48** 113–24

[33] Petit G and Thomas C 1996 GPS Frequency transfer using carrier phase measurements *Proc. 1996 IEEE International Frequency Control Symp.* pp 1151–8
[34] Bureau International des Poids et Mesures 2001 Consultative Committee for Time and Frequency *Report 15th Meeting* Pavillon de Breteuil, F-92312 Sèvres, Cedex
[35] Dach R, Schildknecht Th, Springer T, Dudle G and Prost L 2002 Continuous time transfer using GPS carrier phase *IEEE Trans. Ultrason., Ferroelect., Freq. Control* **49** 1480–90
[36] Petit G, Jiang Z, Taris F, Uhrich P, Barillet R and Hamouda F 1999 Processing strategies for accurate frequency comparison using GPS carrier phase *Proc. 1999 Joint European Frequency and Time Forum & 1999 IEEE International Frequency Control Symp.* pp 235–8
[37] Senior K, Powers E and Matsakis D 1999 Attenuating day-boundary discontinuities in GPS carrier-phase time transfer *Proc. 31st Precise Time and Time Interval Meeting* pp 481–9