

A review of time and frequency transfer methods

Judah Levine

Time and Frequency Division, National Institute of Standards and Technology, Boulder, CO 80305, USA

E-mail: jlevine@boulder.nist.gov

Received 25 June 2008

Published 5 December 2008

Online at stacks.iop.org/Met/45/S162

Abstract

I will discuss the three general methods that are commonly used to transmit time and frequency information: one-way methods, which measure or model the path delay using ancillary data, two-way methods, which depend on the symmetry of the delays in opposite directions along the same path, and common view, in which several stations receive data from a common source over paths whose delays are approximately equal. I will describe the advantages and limitations of the different methods including uncertainty estimates for systems that are based on them.

1. Introduction

In this tutorial I will discuss the various methods of transmitting time and frequency information. I will describe the advantages and limitations of the different methods, and I will provide an uncertainty estimate for each one. I will confine my discussion to the properties of the transmission medium, and I will not consider how a particular source of time information should be selected from among several nominally equivalent possibilities. Therefore, the focus of my discussion will exclude the significant portion of the implementation of the Network Time Protocol (NTP) [1] and similar time-transfer algorithms, which evaluate the possibility that the source of time might be broken.

Any time-transfer algorithm must also be designed with its purpose in mind. For example, if a time-transfer method is being used to discipline a local oscillator to a remote standard, the free-running stability of the local oscillator must be considered when designing the time-transfer procedure so that the resulting algorithm will make the best possible use of the calibration data. I will explain this point in a qualitative way in the following text.

Although the details vary from one situation to another, in general the free-running frequency stability of the local oscillator at short averaging times is better than the frequency stability of the remote device seen through the noisy calibration channel. The inherent stability or accuracy of the remote device is almost irrelevant to this conclusion, since the performance of the channel usually dominates the statistics of the time comparison. However, the time variance of

the channel delay is bounded, so that its frequency variance improves at sufficiently long averaging times. Since many oscillators have frequency variances that can be characterized as flicker or random walk of frequency at longer averaging times, the frequency calibration data from a quiet source usually will be better than the local device at longer averaging times as the contribution of the channel variance becomes less important. The situation for intermediate averaging times, where the free-running frequency stability of the local device is comparable to the stability of the calibration source seen through the channel, will vary from one situation to another, but the frequency stability of the oscillator may not be improved significantly in this regime.

There are similar considerations when the remote system is being used to calibrate the time of the local clock rather than its frequency. The time dispersion of the clock being calibrated is generally unbounded at sufficiently long averaging times, and it will therefore eventually exceed the finite channel noise, which is bounded. As in the previous discussion, a good calibration source, even if it can be seen only through a noisy channel, will be better than a poor local clock at sufficiently long averaging times. Depending on the details of the statistics of the channel, the time dispersion of the calibration data may improve as the averaging time decreases¹, and the time dispersion of the calibration can be further reduced by averaging if the contribution of the fluctuations in the channel delay can be characterized as white phase noise.

¹ The time dispersion caused by frequency fluctuations decreases with averaging time, whereas the time dispersion due to the measurement process generally does not.

These considerations become more complicated when the cost of a calibration (in computer time, network bandwidth, or some other quantity) becomes appreciable, so that the ratio of benefit to cost will also become a consideration. As a simple example, consider a measurement process that is dominated by a white noise process. The standard deviation of the measurements improves as the square root of the number of data observations, but the cost increases linearly with this number. The cost/benefit ratio therefore becomes increasingly unfavourable as more frequent calibrations are performed; if cost is the only consideration then the optimum strategy is to make a minimum number of measurements and to use as long an averaging time as possible. An additional consideration is the probability that the single measurement may be an outlier, and this concern may mandate multiple measurements on each calibration cycle or a shorter averaging time between cycles. Detecting outliers is always difficult because they represent a departure from the statistics of the process by definition so that it is difficult to assign an objective probability of their occurrence. When these considerations are important, the interval between measurements cannot be determined from statistical considerations alone.

2. Stability and accuracy requirements

In the following discussion, I will consider a local clock whose time is being compared with a remote device linked through some type of generalized communications channel. Depending on the context, I will refer to the local device as a receiver or as a client system and the remote device as a transmitter or as a time server.

If the purpose of the comparison is to calibrate the time of the local device, then the delay through the channel must be known, and any uncertainty in this delay will enter into the error budget for the overall calibration procedure. If, on the other hand, the purpose of the calibration is to evaluate the frequency of the local device, then the absolute magnitude of the delay through the channel is not important provided that it is a constant (or that its variation during the calibration procedure is small enough to be ignored). Although this may seem to be a weaker requirement, determining that the variation in the delay is sufficiently small so that it can be ignored may be almost as difficult as measuring its magnitude to begin with.

3. One-way methods

There are a number of systems where the delay from the transmitter to the receiver can be determined (either in whole or in part) by the use of ancillary data. The signals transmitted by the satellites of the global positioning system (GPS) are in this category—many contributions to the path delay can be estimated by the use of parameters transmitted from the satellite in the navigation message [2]².

It takes approximately 65 ms for a signal from a GPS satellite to reach a receiver on the surface of the Earth, and a first-order estimate of this delay can be computed by

combining the broadcast ephemeris of the satellite with the known position of the receiver. Although the transit time due to the geometrical path delay is relatively large, it can be determined relatively accurately so that its contribution to the error budget is much smaller than its magnitude. The use of the broadcast ephemeris, for example, will result in an uncertainty in the position of the satellite on the order of metres, so that the residual uncertainty in the geometrical path delay will be on the order of nanoseconds—a very small fraction of the geometrical delay itself. (I will not consider the more complicated situations when the position of the receiver is not known *a priori* or where it is moving fast enough so that the change in the time difference due to the motion of the receiver is comparable to the change due to the instability of its clock. These situations require data from additional satellites, since both the position of the receiver and the time offset of its clock must be determined simultaneously. The effects that we will consider in the following discussion are applicable to this situation as well, although there may be additional sources of uncertainty because of correlations between the solutions for the position of the receiver and the time of its clock.)

In addition to the geometrical path delay, the refractivity of the path (the difference between the index of refraction of the propagation medium and its value in vacuum) also makes a significant contribution to the delay. For example, the refractivity of the troposphere near the surface of the Earth is about 3×10^{-4} at visible wavelengths, and has about the same value across much of the radio spectrum [3]. Therefore, the refractivity near the surface increases the effective path delay by about 1 ns per km of path. The effect is proportional to the density of the atmosphere and is therefore smaller at higher altitudes. The total additional delay to a navigation satellite due to the refractivity of the troposphere is typically about 6 ns in the zenith direction; the actual value at any location depends on the vertical profiles of the temperature and humidity. The refractivity is often estimated by means of meteorological data (primarily temperature and humidity) recorded at the receiving station. The validity of this approach depends on the assumption that the refractivity at any height above the station, and therefore the integrated zenith delay, is proportional to the temperature and humidity at the station.

The refractivity often varies with wavelength, and one way of estimating it is to measure this dispersion, which produces a variation in the path delay when different wavelengths are used in the time-transfer process. Both the refractivity and the dispersion usually have the same dependence on the parameters of the path (the column density of material, for example), so that the dispersion can be used to estimate the parameters of the material along the path, and these parameters can then be used to compute the refractivity. The dispersion is usually significantly smaller than the refractivity itself, so that estimating the refractivity in this way usually results in a decrease in the overall signal-to-noise ratio of the time-difference measurement. For example, although the index of refraction of the troposphere is about 3×10^{-4} for visible wavelengths [4], the red–blue dispersion of the troposphere [5] is only about 13×10^{-6} , so that the use of the dispersion to estimate the refractivity decreases the signal-to-noise ratio of

² The signals transmitted by the GPS satellites are described.

the measurement of the effective path delay by about a factor of 20. The dispersion of the troposphere at the microwave frequencies transmitted by satellites is essentially 0, so that the refractivity cannot be estimated in this way.

The refractivity of the ionosphere adds about 65 ns to the geometrical path delay at frequencies on the order of 1 GHz, which are the ones that are used in satellite navigation systems such as the GPS. This additional refractivity can be determined by measuring the dispersion between the signals at the L1 (1.5 GHz) and L2 (1.2 GHz) frequencies, since both the dispersion and the refractivity are proportional to the integrated electron density along the path, and both vary as the reciprocal of the square of the frequency. When the dispersion is used to estimate the refractivity, the degradation in the signal-to-noise ratio is about a factor of 3—a significant cost, but not as severe as in the previous example of the dispersion of the troposphere at visible wavelengths. (Single-frequency receivers cannot use this method for estimating the refractivity of the ionosphere and use a model of the ionosphere transmitted by the satellite.) Although the magnitude of the initial correction due to the refractivity of the ionosphere is much smaller than the geometrical path delay, the uncertainty in determining this correction is larger, so that the contributions of the two terms to the overall error budget are not very different.

Finally, there are contributions to the error budget of a one-way measurement that are small to begin with but that have a significant impact on the overall error budget because they are difficult to evaluate. In the case of signals from the GPS satellites (and similar systems, such as GLONASS and Galileo, which are intended primarily for navigation), these effects include the refractivity of the troposphere, which adds 5 ns to 10 ns of non-dispersive delay, and multipath reflections, in which the signal that reaches the antenna is a combination of signals that travel directly from the satellite to the antenna and those that reach the antenna after a reflection from some other object. (Impedance mismatches between the antenna and the receiver will produce additional reflections within the hardware itself. These reflections have similar effects as a reflection of the signal from a nearby object before it reaches the antenna.) Both of these effects always make the effective path length longer than the geometrical value, so that they introduce a bias into the measurements that is not removed by averaging. Since these effects are not easy to measure or to model, they often dominate the uncertainty of the measurement process even though they are relatively small to begin with.

If the delay through the troposphere is assumed to be homogeneous and azimuthally isotropic about the zenith direction, then its contribution to the refractivity can be estimated by the use of various models [6, 7]. The simplest model uses a constant delay at the zenith (on the order of 5 ns to 6 ns), which is determined from local meteorological data. The model assumes that the additional delay due to the refractivity of the troposphere depends only on the path length and is independent of azimuth and elevation. Therefore, when the signals from a particular satellite are used, the zenith delay is scaled by the ratio of the length of the slant path to the satellite through the troposphere to the length of the path if the satellite were at the zenith. In other words, the additional delay due

to the troposphere for any satellite is the product of the zenith delay and the reciprocal of the sine of the elevation angle of the satellite being tracked. A somewhat more sophisticated version of this model is based on the same assumptions but models the troposphere as a band of material confined to a range of altitudes. It is also possible to invert the model and solve for the zenith delay at each epoch from simultaneous observations of several satellites at different elevation angles. This method generally models the variation in the zenith delay as a random-walk process whose amplitude is constrained to be on the order of 10 ps h^{-1} . These methods are usually inadequate to model any variation in the contribution of the troposphere that varies with the azimuth of the satellite being observed, and they usually underestimate the effect of the troposphere at low elevation angles because the path through the troposphere is long and close to the surface so that the curvature of the Earth becomes significant.

There is no simple method for measuring or modelling multipath reflections. Directional antennas can reduce the impact of multipath, since they are much less sensitive to off-axis signals. Unfortunately, these antennas cannot be used with the GPS satellites or with similar systems intended for navigation, since the antenna must have a broad view of the sky if it is to track multiple satellites simultaneously. Choke rings (concentric rings spaced around a central antenna that produce destructive interference so as to attenuate signals arriving at low elevation angles) and ground planes provide some attenuation of reflected signals, especially if the reflectors are at a low elevation angle or are below the antenna. (A typical design will result in the attenuation of signals arriving at low elevation angles by about 10 dB relative to the response to a signal arriving at the zenith.) Even so, multipath reflections often are the largest contributor to the noise budget of time and frequency measurements that use signals from navigation satellites such as those of the GPS.

The geometrical relationship between the satellite, the receiving antenna and the reflectors that are the cause of multipath reflections has an approximate period of 1 sidereal day [8] (about 23 h 56 min), so that the amplitude of the multipath signal for each satellite also has this periodicity. The tracking schedules used for international time and frequency coordination exploit this periodicity by advancing the time of every track by 4 min every day [9], which results in an approximately constant multipath geometry from day to day for each satellite being tracked. Although this reduces the day-to-day variation in the multipath contribution, it converts the multipath variation into a systematic offset, which is different for each satellite that is being observed. This systematic offset changes slowly with time, since the geometrical periodicity is not exactly 23 h 56 min. This slow variation is difficult to detect, since it is usually masked by the long-period flicker and random-walk terms in the variance of the oscillator at the receiver. The multipath contribution can be observed (at least in principle) by comparing the data obtained by the use of satellites at different positions in the sky. Unfortunately, uncertainties in the ephemerides of the satellites contribute to this variation as well, and there is generally no easy way of separating the contribution of multipath from the contributions of other noise terms with similar periodicities.




Figure 1. Three days of consecutive one-second measurements of UTC(NIST) – GPS time via satellite 11. The data for consecutive days are displaced vertically for clarity and the time tags for the second and third days are advanced by 4 min and 8 min to account for the sidereal period of the orbit. The symbol ‘!’ shows a time interval of 13 min as discussed in the text.

The impact of multipath on GPS measurements can be demonstrated by measuring the variation in the time difference between a single GPS satellite and a local time scale with good stability so that the contribution of the clocks to the time difference does not vary during the time that the satellite is visible at the receiving location. Figure 1 shows the difference between GPS time and UTC(NIST) observed via a single satellite on consecutive days. The lowest, middle and upper traces show 1 s time-difference measurements for the first day, the second day and the third day, respectively. The vertical scale is nanoseconds in all cases, and the traces have been displaced vertically for clarity but not otherwise modified. In addition, the time tags associated with the second and third days have been advanced by 4 min and 8 min relative to the first day, respectively, to account for the sidereal periodicity discussed in the previous paragraph. It is clear from the figure that there is a significant part of the ‘noise’ in the data that has a sidereal-day periodicity. The data become noisy near the end of each trace because the satellite is very low in the sky at this time, and multipath reflections from nearby objects at low elevation angles (comparable to the elevation of the satellite itself) are particularly large. The symbol ‘!’ near the lower left-hand corner of the trace shows an interval of 13 min based on the x -axis scale of the figure. This time interval is significant because it is the averaging time used by receivers that implement the tracking schedule used for international time and frequency coordination described above [9]. It is clear from the figure that the 13 min average time difference measured by the receiver will be sensitive to where in the track the average is computed, and that the variation from one part of the track to another will introduce a bias that can be on the order of several nanoseconds. The bias is difficult to evaluate in general, since it is not clear from the figure which epoch is

associated with zero multipath. Many time transfer receivers do not track the same satellite for more than a single 13 min averaging time, which makes it essentially impossible even to evaluate the impact of multipath, much less to remove its effect.

The impact of multipath reflections can be attenuated by exploiting its sidereal-day periodicity. If τ is the length of a sidereal day (23 h 56 min), then the multipath contribution to measurements at times t and $t + \tau$ will be approximately equal, and the multipath contribution will cancel in the difference. Therefore, an estimate of the frequency based on data from a single satellite and an averaging time of τ :

$$y(\tau) = \frac{x(t + \tau) - x(t)}{\tau},$$

will be almost independent of multipath. This frequency can be averaged over the satellite track and integrated to yield a time difference that may have larger white noise than the original data but will be independent of multipath to first order. This method can also be used as a diagnostic tool to evaluate the significance of the multipath contribution to the measurements.

In spite of these limitations, one-way measurements are widely used because they are simple and often sufficiently accurate for the application. However, there are many applications where the residual uncertainties are too large to be acceptable, and I will now consider methods of attenuating these residual uncertainties.

The first method I will discuss is two-way time transfer, in which the delay from the transmitter to the receiver is estimated as one-half of the measured round-trip delay between the two stations. I will then discuss common view, which depends on the equality (or near equality) of two one-path paths between a single transmitter and two (or more) receivers.

4. Two-way time transfer

There are many different applications that use two-way time transfer, and all of them share some common characteristics. In all cases, both the local and remote stations are active and can transmit and receive messages. As we will see in the following discussion, the protocol depends on the fact that the time for a message to travel between the two stations is the same in both directions. If the symmetry of the delay can be realized, then its magnitude is not important and there is no need to understand the effects that contribute to it.

The message exchange is initiated when station 1 sends a message to station 2 at time T_{1s} , which is measured by the clock at station 1. The message is received at station 2 at time T_{2r} which is measured by the clock at station 2. At some other time, station 2 sends a message to station 1. This message is transmitted at time T_{2s} (measured by the clock at station 2) and is received at time T_{1r} (measured by the clock at station 1).

The protocol does not depend on any particular relationship among these times, but different implementations may impose a relationship between T_{2r} and T_{2s} . For example, the NIST Automated Computer Time Service (ACTS) [10] assumes that station 2 will echo the message from station 1 with negligible delay, so that $T_{2s} = T_{2r}$. The NTP assumes that station 2 will reply to a message received from station 1, so that $T_{2s} > T_{2r}$. The delay between the query and the response is not important provided that it is measured accurately and that both clocks are well behaved during this interval. In both these protocols, station 2 is passive—it replies to a query but does not initiate a message exchange. This is not a requirement of the protocol, however, and completely symmetric configurations are also used. There is no specific relationship between the times in two-way satellite time transfer—each station typically transmits a message each second as measured by its local clock or time scale.

Since the message exchange is symmetrical, there is no loss of generality in limiting the discussion to the case where station 1 transmits the first message in the exchange and $T_{2s} \geq T_{2r}$. If d_{12} is the time it takes a message to travel from station 1 to station 2, then the time difference between the clocks at stations 1 and 2, Δ_{12} , is given by

$$\Delta_{12} = (T_{1s} + d_{12}) - T_{2r}. \quad (1)$$

The first term is the time of clock 1 when its request reaches station 2, and the second term is the reading of clock 2 at that instant. The clock at station 1 is fast with respect to station 2 if this time difference is positive. In a similar manner, we can compute a second estimate of this time difference from the second message, which travels in the opposite direction. As in the previous equation, the first term is the time when the reply is received at station 1 and the second term is the reading of the clock at station 2 at that instant:

$$\Delta_{12} = T_{1r} - (T_{2s} + d_{21}). \quad (2)$$

Combining these two equations, we obtain

$$\Delta_{12} = \frac{(T_{1s} + T_{1r}) - (T_{2s} + T_{2r}) + (d_{12} - d_{21})}{2}, \quad (3)$$

$$D = d_{12} + d_{21} = (T_{1r} - T_{1s}) - (T_{2s} - T_{2r}), \quad (4)$$

where D is the measured round-trip travel time.

All two-way methods assume that the path delay is symmetric, so that the third term in the numerator of equation (3) is 0, and equation (4) can be used to find the one-way path delay in either direction. Thus, assuming a symmetric delay,

$$d_{12} = d_{21} = \frac{D}{2}. \quad (5)$$

If the path delays are not the same in both directions, then equation (4) still gives the round-trip path delay, D , but equation (5) is no longer correct and the third term in the numerator of equation (3) is no longer 0. We can characterize the symmetry of the path delay by means of a parameter k , which is a fraction between 0 and 1:

$$d_{12} = kD, \quad (6)$$

$$d_{21} = (1 - k)D,$$

where the assumption of symmetry in the delay implies that $k = 0.5$. For other values of k , the third term in the numerator of equation (3) introduces a time offset whose magnitude is

$$\varepsilon = (k - 0.5)D, \quad (7)$$

so that the effect of the asymmetry becomes more important as the path delay itself increases. In other words, the effect of any asymmetry can be minimized by keeping the path delay itself small. This comment is relevant to the NTP and similar systems where the path delay can be changed by the users by choosing a nearby (in the sense of network delay) source of timing information. (This choice is important only if the path delay is asymmetric, or if there is a concern that it might be so. If the delay is the same in the two directions, then keeping the delay small is not important, since its effect will be removed no matter how large it is.)

The real problem with a path asymmetry is that it is usually not a constant. (A constant path asymmetry would introduce a timing offset but would not be a problem for a system that was distributing frequency.) And the spectrum of its fluctuations is often similar to the variance of the clock in the receiver, so that it is difficult to separate the two. This is particularly serious with NTP and other methods that use the Internet for time distribution. We will discuss this point in more detail below.

In addition to the obvious requirement that the path delay must be symmetric, there is a more subtle requirement that the clocks be well behaved during the interval in which the messages are exchanged. Assuming a perfectly symmetric path delay, so that the third term in the numerator of equation (3) is 0, the time difference can be thought of as comparing the time of one clock at the epoch corresponding to the midpoint between the time when the initial request was sent and when the reply was received with the epoch corresponding

to the midpoint of the time between when the request was received at the remote end and when the reply was transmitted. However, the actual message exchanged spanned a time from $D/2$ before this epoch to $D/2$ after it, and the protocol depends on the fact that the time difference was well behaved over this interval. In practice, this requirement will be satisfied if the time difference between the two clocks varies no faster than linearly with epoch and has no discontinuities. (That is, the time difference results from a time offset combined with a simple frequency difference.) The algorithm will produce a value even if this requirement is not satisfied, but the computed time difference will not be correct.

Another less obvious requirement is that the reference plane for the two-way measurement process must be the same as the reference plane for the applications that use the clock that has been calibrated. This is not always true. For example, consider a computer whose clock is being compared with the clock on a remote server by the use of messages exchanged in the NTP format. In order to minimize the jitter in the network delay measurements, some implementations of NTP use some form of hardware-assist to measure the transmission and reception times that are needed by the protocol. (Other methods of digital time distribution may also use hardware-assist for the same reason. For example, the time-transfer protocol IEEE 1588 [11] is generally implemented in this way.) The point at which the time stamp is applied by the hardware (or by low-level driver software) is the implicit reference plane for the measurement, and delays that are internal to the measurement loop between this point and the equivalent point on the remote server are estimated by the protocol. However, an application running on the same hardware and applying time stamps to events must request the time from the system each time an event occurs. There will be a varying delay in satisfying this request, and this additional delay is not part of the two-way process and is not estimated at all. On the other hand, if NTP (or whatever calibration method is being used) were also running as an application (rather than as a low-level process in the hardware driver), then it would also have to request the time from the system when NTP messages were received, and this system delay would then become part of the two-way process so that it would be estimated by the two-way algorithm. Although the network delay measured in this way might have greater jitter, it would more accurately reflect the time that was actually used by the application. Therefore, the goal of improving the statistics of the measurement algorithm might actually result in poorer performance of an application that depended on the clock whose time difference was being calibrated.

4.1. Systems based on the two-way method

The two-way method is widely used to transfer time and frequency information, since the method does not depend on the actual path delay provided only that it is symmetrical. I will discuss three applications of the two-way method in this section. All of them implement the method as explained in the previous section, although the details vary somewhat from one system to the other.

The NTP is very widely used to transfer time information between computer systems linked over a network such as the Internet. The two-way exchange of messages is typically initiated by a client system, and the calculation of the path delay and the time offset is also done by that system based on the time tags provided by the remote server as described in the previous section.

A typical network path will have a delay of 50 ms to 100 ms, and both the magnitude of the delay and its symmetry can vary significantly from one query to the next one. As we showed in the previous section, the impact of an asymmetry in the measured round-trip delay is proportional to the delay itself, so that messages with shorter round-trip delays will tend to have smaller time errors due to path asymmetry. Many implementations of the NTP exploit this relationship by weighting a measured time difference by the round-trip path delay measured during the exchange of messages under the assumption that longer path delays are more likely to have a larger asymmetric component. Another version of this idea treats the delay between two stations as having a minimum value that is based on the average of some number of recent exchanges³. This minimum delay is assumed to be symmetric. If the round-trip delay measured during a subsequent message exchange exceeds this minimum value, then the additional contribution is assumed to be caused by an asymmetry in the delay, and the measured time difference can be corrected based on this assumption. The method is most useful when the sign of the asymmetry is known from other data. (This might be the case when a network path is known to be much more heavily loaded in one direction than in the other one.) This assumption is based on experience with the Internet and is not a fundamental feature of the two-way protocol, which is not sensitive to the magnitude of the delay if it is symmetric.

If the asymmetry is assumed to be a zero-mean process then it is possible to reduce its impact by averaging several requests, but such averaging cannot remove a static asymmetry or one that has significant flicker or random-walk characteristics. As a result, it is difficult to realize timing accuracies of better than 5 ms to 10 ms over wide-area networks. If the network is particularly congested, the asymmetry might be as high as 80% of the total round-trip delay in one direction and 20% in the other. If the path delay is of order 100 ms, then, from equation (7), an asymmetry of this magnitude will bias the measured time difference by about 30 ms. From equation (7) the maximum possible bias due to an asymmetry is one-half of the round-trip delay. This is usually less than 100 ms in almost all cases, so that it is often small enough to be ignored in many applications.

The NIST ACTS is similar to the NTP in that it transmits time information by means of a two-way protocol. The message consists of a string of characters giving the time followed by an 'on-time marker' character, which is initially '*'. The transmission medium is the voice telephone network. The receiver echoes the on-time marker back to the transmitter with a very small delay (which is assumed to be zero). The transmitter evaluates the one-way path delay (as one-half of the

³ This idea is often called the 'huff'n'puff' filter in the NTP literature. See [1, p 54, ff].

round-trip delay) as I described above and uses this estimate to adjust the time advance of a subsequent transmission of the on-time marker so that it will arrive on time at the client system. The transmitter changes the on-time marker from ‘*’ to ‘#’ to indicate that the advance has been computed in this way. Although the one-way path delay is comparable to the network delays on the Internet, the timing accuracy is significantly better primarily because the delays are more nearly symmetric and both the magnitude of the delay and its symmetry have a much smaller variation with time. After the first few seconds of a telephone connection (during which time the delay may vary due to adaptive equalizers on the line), the measured delays usually vary by less than 0.5 ms. The variations can be characterized as white phase noise, so that averaging consecutive measurements can be effective in reducing the variance of the estimated time difference. The asymmetry in the delay is typically no more than 1% or 2%, so that the protocol can transmit time information with a residual error of 1 ms to 2 ms, and it is quite common to achieve timing accuracies of better than 1 ms by averaging three or four consecutive one-second measurements. The performance of the ACTS system is discussed in greater detail in [12].

Finally, two-way satellite time transfer is used to compare the clocks and time scales of timing laboratories [13]. The method uses a modem to generate a pseudo-random code synchronized to the 1 Hz ticks of the local clock. This code is used to modulate a microwave carrier at the ‘up-link’ frequency of about 14 GHz, and this signal is transmitted to a geostationary satellite. The satellite re-transmits the modulation on the ‘down-link’ frequency, which is typically about 12 GHz. The signal is received at the remote end, the pseudo-random code is extracted and a modem converts the code back to 1 Hz ticks by maximizing the cross-correlation of a locally-generated copy of the pseudo-random code with the received version. The system is full-duplex and transmits signals in both directions simultaneously. Although the path delays are of the same order as those for the NTP and ACTS systems described previously, the timing accuracy is typically better than 1 ns, which is orders of magnitude better than the accuracies of the NTP and ACTS systems. The higher accuracy is due to both the higher frequency that is used to transmit the information and the very much more symmetric path delay, which is a prerequisite for the two-way method.

4.2. Disadvantages of the two-way method

Just as all the time-transfer systems that use the two-way method share the common advantage of being not sensitive to the path delay provided only that it is symmetric, all of them share a number of common disadvantages, which can limit the usefulness of the method in some situations.

In the first place, the two-way method requires that both end stations be able to transmit and receive timing information. The hardware at the stations must be more complicated and more expensive than for a simple one-way receive-only system, and a license to transmit signals may be necessary in some circumstances. The delays through the transmitter and receiver portions of the hardware must

be carefully balanced to avoid introducing asymmetries, and this balance must be independent of variations in the local environmental parameters such as temperature and supply voltage. (The reference planes for the two measurements also must be carefully matched as discussed above.)

Second, satellite time transfer does not support casual or unplanned message exchanges—both end stations must agree in advance on the format of the data exchanges and the schedule that will be followed. No two-way system can support an anonymous association—a capability that can be supported automatically by one-way and common-view methods. Finally, both stations must devote internal resources to support the message exchange. This is a significant issue for the NTP, where a single server may be required to support several thousand simultaneous two-way exchanges with different client systems. The server must maintain state information for each of these exchanges while it is in progress. This requirement is not shared by simple one-way systems or by the common-view system to be described in the following section. A telephone-based system (such as the NIST ACTS servers) has the same problem—a telephone line and its associated modem must be dedicated to each caller for the duration of the call.

5. The common-view method

In the simplest version of the common-view method, a number of stations receive time signals transmitted from a single source. Each station measures the time at which a particular signal arrives at its location. The stations then compare these measurements and subtract them. If the one-way path delays from the transmitter to the two stations are equal, then both the delay and the characteristics of the transmitter cancel in the difference, so that the time difference between the two receiving sites can be computed without knowing anything about the source or the path delay. Even fluctuations in the path delays will be cancelled to the extent that they are common to the paths to the two stations.

In practice it is difficult to find a transmitter such that the path delays to both receivers are exactly equal. This has two consequences: (1) since the path delays are not exactly equal, the difference between them affects the measured time difference in first order. The difference in the two path delays must be measured or modelled in some way, and fluctuations in the path delays and errors in the estimate of the difference degrade the accuracy of the measured time difference. (2) Signals received simultaneously by the two receivers were emitted at different times by the transmitter, so that the measured time difference depends on the stability of the transmitter clock between the two transmission times. (This requirement is analogous to a similar requirement in a two-way method.) If the transmitter is moving (as is the case for signals from navigation satellites such as those of the GPS), then the difference in transmission times translates into a difference in the two positions of the transmitter when the signals were emitted, so that any uncertainty in the ephemeris of the source also contributes to the error budget. Conversely, signals transmitted simultaneously by the satellite are received




Figure 2. Three days of consecutive one-second common-view measurements between two receivers whose antennas were located about 1 m apart. The two receivers measured UTC(NIST) – GPS time via satellite 19 for consecutive days. The data for consecutive days are displaced vertically for clarity and the time tags for the second and third days are advanced by 4 min and 8 min to account for the sidereal period of the orbit. The symbol ‘!–!’ shows a time interval of 13 min as discussed in the text.

at different times by the two receivers, so that the measured time difference now depends on the relative frequency offset of the clocks at the receiving stations and on the stability of this parameter.

Although the choice between an analysis based on a common transmission time from the satellites or a common reception time at the receivers would seem to be arbitrary, the common reception time is more easily implemented for multi-channel receivers, which can process the data from several satellites simultaneously and which usually apply a common time tag to the data received from all of the satellites being tracked. Even for single-channel receivers, most measurements apply a time tag derived from the clock in the receiver to the received signals. In the case of timing laboratories, this time tag is usually derived from the UTC maintained at that site.

The receivers in a common-view time transfer must cooperate by choosing the transmitter and the details of the measurement procedure. However, there need be no connection or cooperative relationship between the receivers and the transmitter, and there is no requirement that the transmitter even know that it is being used for a common-view algorithm. Different groups of receivers, with a very different path delay in each group, can use the same transmitter, provided only that the delays from the transmitter to the receivers in each group are equal (or nearly equal), and that the differences in the path delays within each group can be measured or modelled with an uncertainty that is small enough to satisfy the requirements of the application. A single transmitter can support an arbitrary number of common-view associations, so that the common-view method has much more favourable scaling properties than two-way time transfer.

There are three situations where the common-view method does not work well. (1) The method is not effective if the path delays from the transmitter to the receivers are very unequal and when there is no good way of measuring or modelling this difference. (As we pointed out above, when the path delays are very unequal there is also the implicit requirement that the transmitter clock be well behaved over the time interval corresponding to the difference in the time delays along the two paths.) This limitation can be particularly important for the common-view observation of a ground-based transmitter, since the path delays in different directions can be very different (even if the geometrical lengths are similar), the fluctuations in these delays are often uncorrelated (even if the delays themselves are the same on the average), and one station can be much closer to the transmitter than the other one. The time delay in the receiver (and fluctuations in this delay) and local problems such as multipath reflections are also in this category, since they are usually hard to model and have fluctuations that are unique to each site. Multipath reflections can be particularly troublesome, since they can have significant variation over surprisingly short distances. This point is illustrated in figure 2, which shows the common-view time difference between two independent receivers which measure the time difference between a common clock and the time transmitted by a single GPS satellite. The antennas for the two receivers are at the same height and are separated horizontally by about 1 m. The figure shows the common-view time differences of consecutive 1 s measurements of the two receivers. The three traces show the measurements obtained on consecutive days, with the data obtained on the second and third days advanced by 4 min and 8 min, respectively, to account for the sidereal period of the global positioning

satellites as discussed above in connection with figure 1, above. Although the details of the figure are not the same as for figure 1, there are clear fluctuations that have a period of one sidereal day, and these fluctuations have amplitudes that are not significantly smaller than the data shown in figure 1. The symbol ‘!–!’ shows a time interval of 13 min based on the same x -axis scale as the rest of the figure; as in the previous figure, the exact placement of this averaging time will introduce a systematic offset into the common-view time difference that is difficult to evaluate and remove. (2) Even if the path delays are equal in an average sense, the method is not effective if the fluctuations in the two path delays are not correlated, since these uncorrelated fluctuations will not cancel in the time differences. For example, the contributions of the ionosphere and the troposphere to the path delays from a single satellite to several receivers on the ground tend to become uncorrelated as the distance between the two receivers increases. (3) The method depends on the fact that all the receivers can receive the signals from a single physical transmitter. Receivers that are too far apart cannot benefit from this simple form of common view.

Common-view time transfer is usually implemented with electromagnetic signals and atmospheric paths. The requirement that the transmission path be well defined and easy to characterize favours the use of short-wavelength signals, since these signals travel in very nearly straight lines and can be modelled by means of the approximations of geometrical optics. On the other hand, it is impossible to realize global coverage from a single transmitter by the use of short wavelengths. The all-in-view method addresses this limitation.

5.1. The all-in-view version of common view

One way of realizing global coverage is to implement common view by means of multiple physical transmitters whose signals are synchronized to a common reference, and this idea is the basis for the ‘all-in-view’ or ‘melting pot’ algorithms. For example, if the signals from all of the satellites of the GPS could be traced back to a common time reference, then the common-view principle could be realized with respect to this common time reference even for stations that received time signals from different physical satellites. The links between each physical satellite and the common timing reference add an additional layer to the common-view algorithm, and the success of all-in-view algorithms depends on the fact that the uncertainties introduced by this layer are small compared with the other contributions to the error budget. Although the broadcast ephemerides of the global positioning satellites have always provided a link back to a common GPS system time, the uncertainties in this link and in the real-time orbital parameters that are needed to estimate the geometrical path delay are too large to support an all-in-view algorithm. However, this is not the case for post-processed ephemerides and satellite clock solutions, and these products can support common-view time comparisons with the all-in-view method.

The choice between common view based on a single physical transmitter and all-in-view based on multiple

transmitters that are synchronized to a single time reference is complicated, and it depends on an analysis of the error budget of the measurement process in each case [14]. I will consider a number of limiting cases in the following discussion.

The all-in-view method is particularly useful for very long paths (between the United States and Australia, for example), since there is no physical transmitter that can be received simultaneously at locations so far apart. It is possible to use traditional common view by use of a third station that is midway between the two sites, and this method has used a station in Hawaii, which can support common view measurements both to the United States and to Australia by the use of a (generally different) common satellite in each case. However, double common view is not always possible, and the characteristics of the clock at the midpoint station will enter into the measurement if the two common views are not measured simultaneously.

Conversely, the two methods are likely to be equivalent over very short paths, since both stations will see exactly the same group of satellites, so that the all-in-view method is effectively the same as common view.

It is difficult to provide an unambiguous comparison for intermediate paths, since the different contributions to the error budget of the measurement must be evaluated in detail. I will outline the considerations in the following discussion. A more detailed discussion can be found in [15].

- (1) If the receiver has several channels so that it can receive signals from multiple satellites simultaneously, and if the measurement noise makes a significant contribution to the error budget, and if this noise can be approximated as white phase noise, then the all-in-view method has an advantage, since the all-in-view method can use more data. The signal-to-noise ratio of the measurement process will be improved by a factor proportional to the square root of the number of satellites that are observed.
- (2) If the receiving antenna does not have a clear view of the sky in all directions, then all-in-view may continue to provide data even when a satellite that is suitable for common view is obstructed.
- (3) Neither all-in-view nor common view provides any advantage when the error budget is dominated by local effects such as multipath reflections or uncertainties in the hardware delay of the receiver. All-in-view may have some advantages in detecting and estimating the impact of multipath, since the effect of multipath will tend to be different for the different satellites that are observed simultaneously, so that the scatter among simultaneous measurements may provide an estimate of the level of multipath.
- (4) If the all-in-view data from multiple satellites received at a single site are simply averaged together, then it can be difficult to detect outliers caused by a single bad measurement. However, more sophisticated comparisons are possible and can be very useful, since the data from the different satellites represent a redundant array of independent time measurements (T-RAIM) of the same time difference of the local clock with respect to the common system reference time.

- (5) The full advantage of all-in-view can be realized only by employing post-processing with precise ephemerides and satellite clock solutions. Many satellite receivers use the broadcast ephemeris to report the time difference between the local clock and satellite time and do not provide the details of this computation. The contribution of the broadcast ephemeris must be ‘backed out’ of these data, and this computation must be done separately for each satellite that contributes to the all-in-view average. This may not be easy to do, since different stations may have used different broadcast ephemerides for the same satellite.
- (6) The all-in-view method has an additional complication if the delay through the receiver varies with the satellite that is being observed. For example, the GLONASS satellites transmit on different frequencies [16], and the delay through the RF portion of the receiver may vary from one satellite to another so that the overall receiver calibration will be a function of exactly which satellites are observed at any epoch [17].

5.2. Reference stations for common view

The orbits of navigation satellites and the characteristics of their clocks are determined from data obtained at a number of monitor stations. In some sense, time transfer based on data from a navigation satellite is therefore a common-view measurement between the user’s clock and the clocks at the ensemble of monitoring stations. However, the data from the GPS monitor stations are not available directly, but instead are used to compute the ephemeris message that is uploaded to each satellite and is broadcast along with the timing information.

The ephemeris and clock parameters in the broadcast message are an extrapolation based on previous data from the monitor stations, so that the accuracy of the implicit common view is limited because the model used to compute the extrapolation cannot fully account for stochastic influences. These effects usually have a flicker or random-walk character, so that the accuracy of the message degrades as the broadcast estimates age.

The limitations of the broadcast ephemeris can be overcome with post-processing with the precise ephemerides computed by the International GNSS Service⁴. Precise point positioning (PPP) is one way of realizing this idea [18]. This method uses both the time codes transmitted by the GPS satellites and the phase of the carrier [19], which is generated from the same oscillator onboard the satellite. Use of the phase of the carrier [20] improves the resolution of the measurements, since the carrier has a period much shorter than that of the code. Since the carrier phase measurements are ambiguous modulo an integer number of cycles of the carrier, this integer must be determined in some way. A typical method determines the integer from the transmitted code, and the accuracy of the carrier phase data will be no better than the code measurements when this technique is used [21]. In

⁴ See the website igs.cbl.jpl.nasa.gov for information about the service and its products.

addition, all carrier phase analyses must cope with possible slips of an integer number of cycles in the data due to the finite signal-to-noise ratio of the measurements. Detecting these cycle slips becomes more complicated when the refractivity of the ionosphere is estimated using the dispersion between the two frequencies transmitted by the satellite, since the signals at the two frequencies may not experience a cycle slip at the same time.

If data from several stations are available, it is also possible to solve for the orbits and clocks of the satellites and the clocks of the ground stations by means of a simultaneous least-squares solution. One (or more) of the ground stations are used as the references for this computation. The GIPSY⁵ and Bernese [22] software packages operate in this way, and the method is being used for time and frequency comparisons between national laboratories [23].

Common-view systems based on ground-based transmitters are subject to the same kinds of problems. The propagation delay is a function of the properties of the ionosphere and the troposphere, the conductivity of the ground along the path and other factors. These factors vary both in time and in azimuth and they are difficult to model accurately.

One method that can be used to improve the real-time performance of both ground-based and satellite-based transmitters is to use reference stations with very good clocks at various locations and to publish (or broadcast) the data from these reference stations in near real time. The data from these reference stations are then used in real-time common view. The density of these reference stations is determined from the correlation distance of the perturbations in the path delay.

Many timing laboratories (including NIST) support this technique by publishing measurements of various timing sources with respect to the local realization of UTC. These data are not available in absolute real time for administrative rather than technical reasons. For example, the NIST measurements for the time signals received from the GPS satellites with respect to UTC(NIST) are published with a delay of about 24 h⁶.

There are also stations that measure the times of GPS satellites and broadcast corrections in real time. The US Coast Guard operates a network of stations which support ‘differential GPS’ corrections⁷. These data are transmitted on a low-frequency carrier (about 285 kHz). The signals can be received and processed by many commercial GPS receivers. The ground-based augmentation system provides a similar service⁸. These transmissions can also be used to inform the users of a bad satellite.

Reference stations are particularly useful for ground-based low-frequency transmitters such as those of the LORAN

⁵ GIPSY was developed at the Jet Propulsion Laboratory. See <http://www.jpl.nasa.gov> for details.

⁶ The NIST data can be obtained with a delay of about 24 hours from the NIST web site tf.nist.gov or from <ftp://clock.bldrdoc.gov/pub/time/nist.gps>. The ftp file contains the measurements made at NIST and averaged for 13 minutes as defined in the technical directives as described in the reference above.

⁷ The US Coast Guard service is described at <http://www.tfhrc.gov/its/ndgps/02072.htm>.

⁸ Many other augmentation services are described at <http://www.beaconworld.org.uk/dgps.htm>.

navigation system⁹. This system transmits periodic pulses with a carrier frequency of 100 kHz. The shape of the pulses is very carefully controlled so that a receiver can interpolate to a fraction of the full width of the pulse. A receiver determines its position by measuring the difference in the arrival times of pulses from several synchronized stations. Each such difference locates the receiver on a hyperbola, and several pairs of measurements are used to fix the exact location. Since the navigation solution depends on timing differences, the system is less sensitive to variations in the propagation characteristics of the path. (This advantage is basically the same as the common-view method described above, where several receivers observe a single transmitter. Here a single receiver observes several transmitters. In both cases fluctuations in the delays that are common to both paths cancel in the difference.) The cancellation is not complete, however, because the stations that transmit the signals used to compute the difference are not located at the same place. The range of the system is improved as the separation between the transmitting stations is increased, but this increase in separation can degrade the advantage of the differencing technique.

Unfortunately, the differencing technique that is used for a navigation solution cannot usually be carried over to a common-view timing application because of the variability of the propagation delay in time, in distance, and, especially, in azimuth. One way of mitigating this variability is to compute 'additional secondary factors (ASF)' which are measurements of the propagation delay from a transmitter to a reference receiver which has a very good local clock. By the use of these factors, a user located not too far away from one of these reference stations could correct the received data for estimating the measured delay to that region. As with a satellite augmentation system, the spacing of these reference stations would depend on the correlation distance of the perturbations in the delay. This distance can be as short as tens of kilometres, so that a relatively dense network of reference stations may be needed at some regions where the spatial variability of the delay is large and accurate measurements are needed. These reference stations may be needed for a navigation solution in rivers and narrow channels where the cancellation of the hyperbolic method does not provide adequate attenuation of the fluctuations in the path delay.

Various enhancements to LORAN have been proposed, and a common thread in these proposals is a method for determining and transmitting these local corrections¹⁰. These enhancements are collectively referred to as e-LORAN¹¹. A significant number of reference stations will be required to provide wide-area coverage for these correction factors, since the correlation length of the delay variability can be

as short as tens of kilometres. Even without applying these correction factors, the current LORAN system can usually provide time with an uncertainty on the order of microseconds and frequency with an uncertainty of less than 1×10^{-11} when the measurements are averaged over 1 day. A frequency uncertainty of 1×10^{-12} can be realized in many situations. These capabilities are adequate for many commercial applications, including acting as a primary reference frequency for telecommunications networks^{12,13}.

6. Summary and conclusions

I have discussed the three basic methods that are commonly used to transmit time and frequency information: one-way, two-way and (several versions of) common view. I have presented the advantages and limitations of each method. A comparison of these methods as they are used by the BIPM to construct International Atomic Time is in [24, 25].

The one-way method is most useful for three types of applications: (1) where the path delay can be determined from ancillary data. For example, the one-way method is used when the path delay can be determined with an accuracy that is adequate for the application by the use of only the parameters in the navigation message transmitted by GPS satellites. (2) Where the delay is small enough that it can be ignored. For example, the delay is ignored when the low-frequency signals from the NIST radio station WWVB¹⁴ are used to keep a wall clock or wristwatch on time. (3) Where the primary application is for frequency distribution and the one-way delay is sufficiently stable that its fluctuations do not make an appreciable contribution to the uncertainty budget of the measurement. Many low-frequency radio stations have a delay that has a large diurnal variation but is stable when it is averaged over 24 h. Radio station WWVB was initially widely used as a frequency reference in this way. (Calibrating the frequency of a device by averaging for 24 h implicitly assumes that the frequency of the device is sufficiently stable over the averaging time that the calibration is meaningful. Inexpensive quartz-crystal oscillators usually do not satisfy this criterion.)

The two-way method uses active stations at both ends of the path. It depends on the fact that the path delay is the same in both directions, so that the one-way delay can be estimated as one-half of the measured round-trip value. It is widely used in digital time services such as ACTS and NTP and in comparing the time scales and primary frequency standards of national laboratories and timing centres by means of two-way satellite time transfer [26].

The common-view method depends on the equality of the path delays between a single transmitter and multiple receivers.

⁹ The technical specifications for the LORAN system are at <http://www.navcen.uscg.gov/pubs/loran>.

¹⁰ The information will be transmitted in a new '9th pulse' as described in http://www.navcen.uscg.gov/eloran/ldc.v1-3_mod1_20061020.pdf. This description is provisional, and the details are still being developed.

¹¹ The design of e-LORAN is described at <http://www.navcen.uscg.gov/eloran/overview.htm>. This web site also contains links to other pages that provide additional technical information about the additional signals that will be added to the current LORAN system and the format of the information that will be transmitted.

¹² Synchronization interface standards for digital networks, American National Standards Institute (ANSI) standard T1.101-1987.

¹³ Performance estimates are in the presentation at <http://www.loran.org/library/2004NIST.pdf>.

¹⁴ The station is located at Fort Collins, Colorado. See tf.nist.gov/stations/wwvb.htm. Similar low-frequency stations are operated by other timing laboratories. For example, a similar station, DCF77, is operated by the Physikalisch-Technische Bundesanstalt (PTB). It is located at Mainflingen, Germany, approximately 25 km from Frankfurt and is described at http://www.ptb.de/en/org/4/44/442/DCF77_1_e.htm.

Both the characteristics of the path and the stability of the common transmitter are not important when the paths are exactly equal and are attenuated in first order when the path difference is small and when fluctuations in the delays are coherent. The method can be realized with a single physical transmitter or with several transmitters that are synchronized to a common reference clock. Although the method can operate in real time when a single physical transmitter is used as the source of the signal, the connection between multiple physical transmitters and the common reference clock is often not available in real time with adequate accuracy, so that the 'all-in-view' version of the method typically depends on post-processing to realize the full capability of the technique.

A particularly powerful version of the common-view method uses an ensemble of reference stations to characterize the orbits and clocks of GPS satellites very accurately; the receiver is then characterized with respect to this ensemble in a post-processed analysis. This PPP method usually uses the phase of the carrier transmitted by the GPS satellites as well as the code. If the method is used to compare clocks at widely-separated locations (Europe and the United States, for example), then the method benefits if data from additional stations are added to the solution. If the additional stations are near the sites whose clocks are being compared then data from these stations are particularly useful in evaluating short-wavelength variations in the contributions of the troposphere and the ionosphere to the path delays [27].

All of these methods have residual uncertainties, and these uncertainties will become more important (and more limiting) as station clocks improve and as the applications that depend on time transfer start to require better time and frequency synchronization. Evaluating these uncertainties in a quantitative manner is difficult; one way of estimating them is to compare two clocks by means of two different techniques whose uncertainties are as uncorrelated as possible. At present, the two methods that have shown the smallest error budgets are two-way satellite time transfer and GPS carrier phase analysis. When these two methods are used to compare the frequencies of the clocks at widely-separated timing laboratories (PTB in Germany and NIST in Boulder), they differ by about 4×10^{-16} after the time differences have been averaged for about 40 days [28]. There is no way of partitioning this uncertainty between the two techniques.

The performance of clocks is improving much faster than our ability to compare them when they are geographically separated by a large distance, and improvements in methods of time transfer will be necessary if the full capabilities of these new devices are to be realized. The uncertainty quoted in the previous paragraph is comparable to the total uncertainty budget for many current primary frequency standards, so that these techniques are already only marginally able to compare the current generation of frequency standards at widely-separated locations without degrading the performance of the standards by the noise in the comparison process. The techniques will not be adequate for the next generation of devices, which are already in the laboratory-prototype stage.

A number of groups have demonstrated frequency comparisons by means of the two-way method with signals

transmitted on optical fibres [29], but it is not clear that this technique can be applied over continental distances. The limitations are both technical and financial. It is not easy to maintain the symmetry of the path delay in very long fibres, and the cost of fibres dedicated to time and frequency distribution may be too great to make the method practical.

There is no fundamentally new method of time comparison on the horizon, so that improvements will have to come from a better understanding of the limitations I have described in the methods that we are currently using.

References

- [1] Mills D L 2006 *Computer Network Time Synchronization: The Network Time Protocol* (Boca Raton, FL: CRC Press) See also the many references in this book
- [2] Misra P and Enge P 2006 *Global Positioning System: Signals, Measurements and Performance* (Lincoln, MA: Ganga-Jamuna Press)
- [3] Kaplan E and Hegarty C 2006 *Understanding GPS: Principles and Applications* (Boston: Artech House)
- [4] Smith E K and Weintraub S 1953 The constants in the equation for the atmospheric refractive index at radio frequencies *Proc. IRE* **41** 1035–7
- [5] Owens J C 1967 Optical refractive index of air: dependence on pressure, temperature and composition *Appl. Opt.* **6** 51–8
- [6] Svensson K-F 1959 Measurements of the dispersion of air for wavelengths from 2302 to 6907 Angstroms *Ark. Fys.* **16** 361–84
- [7] Hopfield H S 1969 Two-quartic tropospheric refractivity profile for correcting satellite data *J. Geophys. Res.* **74** 4487–99
- [8] Niell A E 1996 Global mapping functions for the atmosphere delay at radio wavelengths *J. Geophys. Res.* **101** 3227–46
- [9] Choi K, Bilich A, Larson K M and Axelrad P 2004 Modified sidereal filtering: implications for high-rate GPS positioning *Geophys. Res. Lett.* **31** L22608
- [10] Allan D W and Thomas C 2004 Technical directives for standardization of GPS time receiver software to be implemented for improving the accuracy of GPS common-view time transfer *Metrologia* **31** 69–79
- [11] Levine J, Weiss M, Davis D D, Allan D W and Sullivan D B 1989 The NIST automated computer time service *J. Res. Natl. Inst. Stand. Technol.* **94** 311–21
- [12] IEEE TC9 2002 *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurements and Control Systems, IEEE Standard 1588-2002* (New York: Institute of Electrical and Electronic Engineers)
- [13] Levine J 2008 Improvements to the NIST network time protocol servers *Metrologia* **45** S12–22
- [14] Kirchner D 1999 Two-way satellite time and frequency transfer (TWSTFT): principle, implementation, and current performance *Review of Radio Sciences* vol 1996–1999 (Oxford: Oxford University Press) pp 27–44
- [15] Lewandowski W 2006 International atomic time scales and related time transfer *J. Metrol. Soc. India* **4** 223–7
- [16] Petit G and Jiang Z 2008 GPS all in view time transfer for TAI computation *Metrologia* **45** 35–45
- [17] Lewandowski W, Foks A, Jiang Z, Nawrocki J and Nogas P 2005 Recent progress in GLONASS time transfer *Proc. Joint IEEE Frequency Control Symp. and Precise Time and Time Interval Planning and Applications Meeting (Vancouver, Canada)* pp 728–34
- [18] Foks A, Lewandowski W and Nawrocki J 2005 Frequency biases calibration of GLONASS P-code time receivers *Proc. 19th European Time and Frequency Forum (Besançon, France)* (on CDROM)

- [18] Petit G and Jiang Z 2008 Precise point positioning for TAI computation *Int. J. Navig. Observation* **2008** 562878
- [19] Defraigne P and Petit G 2003 Time transfer to TAI using geodetic receivers *Metrologia* **40** 184–8
- [20] Petit G 2004 Stability and accuracy of geodetic GPS time links compared to two way time transfer *Proc. ATF 2004 (Seoul, Korea)* pp 1–7
- [21] Petit G and Jiang Z 2004 Stability of geodetic GPS time links and their comparison to two way time transfer *Proc. 36th Precise Time and Time Interval Planning and Applications Meeting* pp 31–9
- [22] Dach R, Hugentobler U, Fridez P and Meindl M 2007 *Bernese GPS Software Version 5.0* (Bern, Switzerland: Astronomical Institute, University of Bern)
- [23] Ray J and Senior K 2005 Geodetic Techniques for time and frequency comparisons using GPS phase and code measurements *Metrologia* **42** 215–32
- [24] Lewandowski W 2004 International atomic time and time transfer *Artif. Satell.* **40** 71–81
- [25] Arias E F, Jiang Z, Lewandowski W and Petit G 2005 BIPM comparison of time transfer techniques *Proc. Joint IEEE Frequency Control Symposium and Precise Time and Time Interval Planning and Applications Meeting (Vancouver, Canada)* pp 312–5
- [26] Bauch A *et al* 2006 Comparison between frequency standards in Europe and the USA at the 10^{-15} accuracy level *Metrologia* **43** 109–20
- [27] Hackman C and Levine J 2006 Towards sub 10^{-16} transcontinental GPS carrier phase frequency transfer: a simulation study *Proc. IEEE Int. Frequency Control Symp. (Miami, FL)* pp 779–87
- [28] Hackman C, Levine J and Parker T E 2006 A long-term comparison of GPS carrier phase frequency transfer and two-way satellite time transfer *Proc. 38th Precise Time and Time Interval (PTTI), Meeting (Reston, VA)* pp 485–96
- [29] Foreman S M, Ludlow A D, de Miranda M H G, Stalnaker J E, Diddams S A and Ye J 2007 Coherent optical phase transfer over a 32 km fiber with 1s instability at 10^{-17} *Phys. Rev. Lett.* **99** 153601