Time transfer using a satellite navigation system

Several satellite navigation systems are currently in operation or are scheduled for deployment in the near future. These include the United States Global Positioning System (GPS), the Russian Global Navigation Satellite System (GLONASS) satellites, and the European Galileo satellites. In addition to determining the position and velocity of a receiver, signals from these satellites can be used to distribute time and frequency information and to compare clocks at remote stations. A number of different techniques are commonly used for these purposes, and we will describe these techniques in this article. Although the different satellite systems differ in technical details, the methods we will describe can be used with any one of them with only minor modifications.

In the following discussion, we will assume that the receiving stations are at known positions in the coordinate system used by the navigation satellites and are not moving in an Earth-fixed coordinate system. In addition, the uncertainties in the coordinate values should be consistent with the requirements of the timing application. The details will vary with the geometry, but time transfer with an uncertainty of a few nanoseconds generally requires a position uncertainty of less than 1 m. In practice, the vertical position is usually the one with the largest uncertainty because there is usually a significant correlation between the vertical position of the station and the time offset of its clock. Therefore, the vertical position uncertainty often limits the accuracy of the time-difference measurement. This uncertainty can be minimized by the use of data from satellites that are uniformly distributed in elevation with respect to the receiver.

Signals transmitted by the satellites. The satellites transmit two signals that are important for time transfer. The first is a pseudorandom code—a series of binary 1s and 0s that is generated by a deterministic algorithm driven by the atomic clock onboard. Although the sequence is fully deterministic, it satisfies many of the statistical tests for randomness. It is not feasible to compute the next bit in the series given the previous values nor is it feasible to invert the series to determine the parameters of the generating algorithm. The second signal is a data stream that includes an ephemeris message, a group of parameters that characterize the orbit of the satellite and that can be used to determine the position of the satellite at any time. The ephemeris message also contains an estimate of the time and frequency offsets of its clock relative to the system time scale (to be defined in the following) and Coordinated Universal Time (UTC), the international time scale computed by the International Bureau of Weights and Measures (BIPM, in French). In the case of the GPS satellites, the system time scale is called GPS time, and the ephemeris message includes an estimate of the difference between this timescale and UTC as maintained at the U.S. Naval Observatory. These data are transmitted using a carrier frequency of about 1.5 GHz. A second pseudorandom code is generally transmitted on a second frequency, but access to this code is often restricted to authorized users.

Pseudorange. All of the methods we will discuss start from the pseudorange—a measurement of the apparent time difference between the clock in the satellite and the clock at the receiver. To measure this quantity, the receiver generates a copy of the pseudorandom code transmitted by the satellite and varies the offset time of the code generator to maximize the correlation between the local and received copies of the code. The period of the pseudorandom code used by the GPS satellites is approximately 1 μs, and the cross-correlation hardware can determine the time offset that maximizes the correlation to a few percent of the period, so that the noise in the correlation process is on the order of tens of nanoseconds. The pseudorandom code implicitly defines the satellite time, and the receiver combines this measurement with the data in the navigation message to compute the pseudorange. The time offset of the station clock from the clock in the satellite is computed by applying various corrections to the pseudorange as described in the next sections. The clock-parameter data in the ephemeris message can then be used to refer this computed time difference back to the system time and ultimately to UTC.

Receiver outputs. Timing receivers use the corrected pseudorange computation in several different ways. Some receivers generate an output pulse that is derived from the average of the times extracted from the pseudorange calculations for all of the satellites in view. In addition to the composite output pulse, these receivers often produce a separate data stream containing the contribution of each satellite in view to the computed average time. (This is useful for detecting a bad satellite, whose computed time difference is very different from the computations using the others.) In many cases, the output pulse is offset from the time extracted from the pseudorange data, and this offset is also included in the data stream. Other timing receivers accept an input pulse and measure the time difference between this pulse and the satellite time scale, again using all satellites in view.

System timescale. In positioning or navigation applications, the pseudorange locates the receiver on a spherical surface centered on the satellite. To compute a unique position solution, the receiver uses at least four pseudorange observations to solve for the four unknowns: the three coordinates of its location and the offset of its clock. (There is generally a correlation between the solutions for the position and for the clock, so that timing receivers generally use a fixed position that is treated as a known quantity with no uncertainty.) Since each pseudorange is determined by the difference between the receiver clock and the clock in the satellite being observed, the times of the various satellite clocks must be linked together so that the receiver needs to solve for only one clock offset. This linkage of the satellite clocks is computed on the ground as a weighted sum of the clocks in the system, and the predicted offset of the time and frequency of each clock from this average "system time" is uploaded to the satellite periodically and is then broadcast as part of the ephemeris message. For the GPS satellites, the satellites also broadcast a prediction of the time difference between the system time scale and UTC as maintained at the U.S. Naval Observatory. Other satellite systems will transmit similar information. The existence of a system timescale is required for position and timing applications, but the linkage to UTC is needed only for timing applications, and the details of how the connection to UTC is realized will vary from one satellite system to another.

One-way time transfer. This is the simplest situation. It is used to synchronize the local clock to system time and also is the basis for more sophisticated techniques.

The first step in the analysis is to correct the pseudorange for the time of flight of the signal from the satellite to the receiver. The geometrical delay is about 65 ms (0.065 s) and is estimated using the ephemeris transmitted by the satellite and the known position of the receiver. The computation is usually performed in an Earth-centered-inertial (ECI) frame that corresponds to the Earth-centered-Earth-fixed (ECEF) coordinate system at the instant that the signal is received. This coordinate system simplifies the computations when signals from several satellites are received simultaneously. The calculation must be done by iteration, since the satellite, which is moving with an orbital speed of about 4 km/s, has moved almost 300 m during the time of flight. Depending on the latitude of the receiver, it has also moved by several tens of meters during the time of flight. Therefore, using the position of the receiver at the instant of reception as the origin of the coordinate system, we imagine that the coordinate system is rotated "backward" to compute the satellite position in the ECI frame at the time of transmission. The first-order time of flight is calculated using the geometrical distance between the position of the satellite at the instant of transmission, and the position of the receiver at the instant of reception, keeping in mind that both the receiver and the transmitter are moving during the time of flight. The uncertainty in the time of flight is much smaller than the magnitude of delay itself and is determined by the uncertainty in the position of the receiver and in any errors in the broadcast ephemeris. An uncertainty of 1 m in position translates into an error of about 3 ns in time, so that it is relatively easy to reduce the impact of an uncertainty in the position to tens of nanoseconds.

The second step is to remove the additional time delay due to the path through the ionosphere. The magnitude of this contribution varies, but is typically of order 65 ns. The magnitude of this delay is proportional to the density of charged particles in the ionosphere and varies as the inverse of the square of the carrier frequency, so that it is possible to estimate the delay by measuring its dispersion—the apparent difference in the transit times of signals at two different frequencies. All of the satellite systems transmit a pseudorandom code on two frequencies, but not all receivers can process the second transmission. Simpler, single-frequency receivers can use the estimate of the contribution of the ionosphere to the time of flight using a model broadcast by the satellite in the ephemeris message. This is better than nothing, but is usually not as accurate as the two-frequency technique. Although the magnitude of the correction due to the ionosphere is much smaller than the time of flight, the uncertainty in this correction can be quite large. For example, a single-frequency receiver that does not use any correction for the ionosphere can have a timing error that has a roughly diurnal periodicity with an amplitude that may reach 100 ns. A single-frequency receiver that uses the broadcast estimate of the ionosphere may also have a diurnal timing error, but the amplitude will be smaller, perhaps of order 50 ns. These values depend on the location and on the state of the ionosphere, which is driven by many factors including sunspot activity. A full dual-frequency receiver that measures the dispersion due to the ionosphere will have a much smaller residual uncertainty, as little as a few nanoseconds if the receiver is well calibrated.

The final step is to correct for local effects—the additional delay caused by the refractivity of the troposphere (typically about 6 ns at the zenith, increasing as the reciprocal of the cosine of the zenith angle for satellites at lower elevations due to the increase in the slant path), the delay through the receiving hardware (typically tens of nanoseconds within the receiver itself and about 5 ns per meter of cable between the receiver and the antenna), and small changes in the position of the station due to Earth tides, polar motion, and similar effects. These effects are much smaller than the effects discussed above, but they are more difficult to estimate, so that uncertainties in the magnitudes of these effects make an appreciable contribution to the overall error budget. They are typically of order a few nanoseconds. The effects of multipath reflections, which are discussed in the next section, must also be considered.

Finally, the data in the ephemeris message are used to relate the time transmitted by the clock in the satellite to the system time of the satellite system and then to an international time scale. In the case of the GPS system, the message can be used to relate the time difference to UTC as maintained by the U.S. Naval Observatory (USNO). The values in these messages transmitted by the satellite are predictions calculated on the ground and uploaded into the satellite periodically. The uncertainties in these corrections therefore depend on the stability of the clock in the satellite and on the time that has elapsed since the last upload. These uncertainties generally do not exceed 25 ns and are often much smaller than this value.

The accuracy of the time difference is limited by two sets of effects: the random fluctuations in the cross-correlation of the pseudorandom code, and the systematic errors that result from errors in the position of the satellite or in the magnitudes of the various corrections discussed above. The contribution of the random errors is attenuated by averaging, and the improvement that can be realized by averaging implicitly assumes that the clock in the receiver is sufficiently stable so that the random fluctuations in its time and frequency can be neglected relative to the measurement noise during the averaging time. These considerations can be used to define the range of optimum averaging times—long enough to attenuate the random measurement errors and short enough so that the fluctuations in the time and frequency of the clock do not make an appreciable contribution to the variance of the data. The techniques discussed in the following sections are designed to attenuate the systematic errors.

Multipath reflections. The antennas used with satellite receivers cannot have a strong directional sensitivity, since the satellites are moving with respect to the receiving station and also because a robust position solution depends on observing multiple satellites that are uniformly distributed in elevation and azimuth with respect to the receiver. The antennas are therefore also sensitive to multipath reflections—signals that reach the antenna after reflection from a nearby object. These signals, which combine with the direct signals in the receiver, always travel a longer path than the direct ones and therefore bias both the position and timing solutions.

The amplitude of the multipath signal and its time variation depend in a complicated way on the position of the satellite with respect to the ground station antenna and nearby reflectors. Amplitudes of 5–10 ns and variations on the order of minutes are not unusual, so that the effect of multipath reflections is often the largest unmodeled systematic error for a timing receiver.

One solution to mitigate this effect is to mount the antenna as high as possible above any local reflectors, to use a ground plane, which blocks signals arriving from below the antenna, and choke rings around the antenna, which attenuate signals arriving at very low elevation angles.

The geometrical configuration of the satellite, the ground-station antenna and the reflectors repeats with a nearly sidereal period (about 23 h 56 m), so that the multipath reflection from any satellite also has this periodicity. The BIPM tracking schedule (discussed below) exploits this periodicity by advancing the tracking times used by timing laboratories by 4 min every day. This has the advantage that the multipath effect is the same (in first order) every day, but the disadvantage is that the variation is converted to a systematic offset that is hard to estimate and remove.

Common-view time comparisons. In the common-view method, two (or more) receivers observe the same satellite at the same time. Each receiver computes the one-way time difference as described in the previous section, and these measured time differences are then subtracted to compute the time difference of the clocks at the two stations. The method is illustrated in **Fig. 1**. Receivers 1 and 2 receive a signal transmitted by the satellite at time S. The signals reach the receivers after a time δ, which is the same for both paths. The receivers measure the time difference between the local clock, t, and the received signal, and they then compute the time difference between the local clocks by subtracting these values as shown. Since the path delays are equal, both the path delay and the time of the source cancel in the difference.

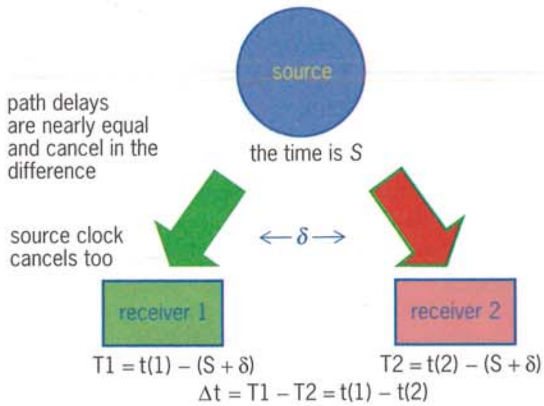

Fig. 1. The common-view method. Receivers 1 and 2 receive a signal transmitted by the satellite at time S. The signals reach the receivers after a time δ, which is the same for both paths. The receivers measure the time difference between the local clock, t, and the received signal, and they then compute the time difference between the local clocks by subtracting these values as shown. Since the path delays are equal, both the path delay and the time of the source cancel in the difference.

If the stations are not too far apart, then many of the one-way corrections discussed in the previous section are almost the same for both stations and are therefore attenuated in the subtraction. Therefore, the common-view method is much less sensitive to errors in the satellite ephemeris, in the accuracy of the satellite clock, and in the refractivity of the ionosphere. The common-view method is less effective in attenuating the effect of the troposphere, since it is usually not so well correlated between stations and it has no effect on the local station-dependent effects discussed in the previous section. The method is also not as useful when the geometrical ranges to the receivers are very different. If the stations are not too far apart, a common-view time comparison can have an uncertainty of less than 25 ns. Timing laboratories with well-calibrated receivers located at well-known positions can realize uncertainties of 1–2 ns using this method.

All-in-view time comparisons. The common-view method depends on the fact that a single physical signal is observable at all of the participating receivers. This requirement limits the maximum distance between the receivers, and there comes a point where the method fails because the satellite is not simultaneously visible at the receivers. For example, common view cannot be used between locations in Australia and most parts of the United States.

The physical common view can be replaced with a logical common view in this case. In this method, each station measures the difference between its clock and the satellite system time—not the physical time of the satellite clock. Since the ephemeris messages transmitted by all of the satellites link the

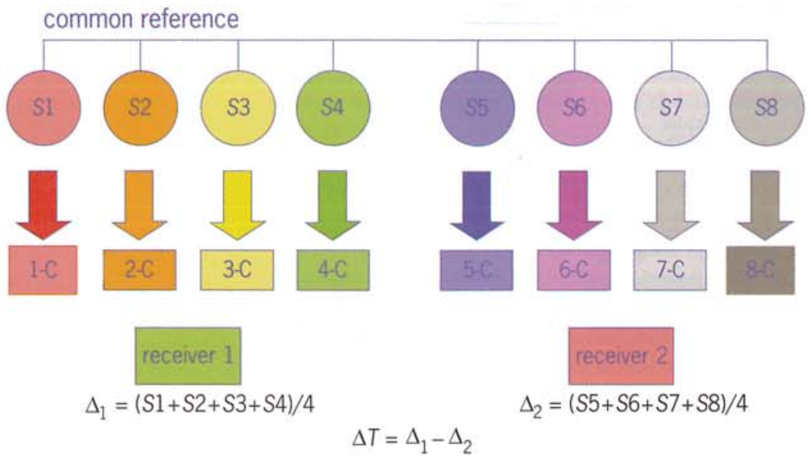

Fig. 2. The all-in-view melting pot method. Receiver 1 observes satellites *S1* through *S4* and uses the data from these satellites to compute the difference between the local clock and the system time. The receiver combines the time offset of the local clock with respect to the satellite clock computed from the pseudorange with the offset of the satellite clock from the system time that is broadcast as part of the ephemeris message. Receiver 2 does the same computation using satellites *S5* through *S8*. The time difference is then the difference between these two computations. The result is a logical common-view using the common system time as the reference rather than the physical signal from a single satellite as in Fig. 1.

physical satellite clock to the system time scale, it is possible to compute the common-view time difference between each station and the system time, even though this is a time scale whose time is not realized by any physical clock as discussed above. The method is shown in **Fig. 2**. Receiver 1 observes satellites *S*1 through *S*4 and uses the data from these satellites to compute the difference between the local clock and the system time. The receiver combines the time offset of the local clock with respect to the satellite clock computed from the pseudorange with the offset of the satellite clock from the system time that is broadcast as part of the ephemeris message. Receiver 2 does the same computation using satellites S5 through S8. The time difference is then the difference between these two computations. The result is a logical common view using the common system time as the reference rather than the physical signal from a single satellite as in the previous discussion.

The advantage of this method is that all of the satellites in view can be used to compute the time difference between the two stations even when no satellite is simultaneously in view at both sites. There will be multiple pseudorange measurements at each epoch, and we would expect that the measurement noise would be attenuated. Conversely, the accuracy of the relationships between the physical satellite clocks and the system time scale is a requirement of the all-in-view method but is not important for common view; the accuracy of the satellite ephemerides is also more important for the all-in-view method, since the errors in the ephemeris are not attenuated as they are in the common-view differences. Orbital and clock parameters broadcast by the each satellite in the ephemeris message are often not sufficiently accurate for all-in-view measurements, and postprocessed values are often used.

At very short baselines the two types of common view are effectively equivalent, since the same satellites and physical signals are in view at both stations. The all-in-view method is the only choice at very long baselines. The comparison at baselines of intermediate length depends on the magnitudes of the different error contributions discussed above. The full advantage of all-in-view methods can be realized only with post-processed ephemerides, which are available only after some delay, so that they are not suited to real-time clock comparisons. Both methods are limited by local effects, especially those that cannot be attenuated by averaging, such as the effect of the troposphere, uncertainties in the calibration of the receiver delay, and similar problems. Even the advantages of averaging are useful only if the local clock is sufficiently stable to support longer averaging times.

To facilitate common-view and all-in-view analyses, the BIPM has defined a standard format for publishing the data, and all timing laboratories use this format for data interchange. The format specification uses a 13-min average of the time difference between the satellite and receiver clocks, and this average is computed by the use of 52 1-s time-difference measurements. The averaging algorithm and the number of points used to compute the average were based on considerations that were appropriate for the first generation of satellite receivers, and are no longer relevant for newer receivers that can observe multiple satellites at the same time. A shorter averaging time would be appropriate in many situations. For example, many geodetic receivers, which measure the phase of the carrier (to be discussed in the next section), use an averaging time of 30 s.

Using the phase of the carrier. All of the methods described above depend on the pseudorange value, which is the time difference measured using the transmitted code. The transmitted carrier is derived from the same clock as the code, so that it has the same stability. Therefore, the pseudorange could also be estimated using the phase difference between the received carrier and the local receiver clock. Since there is no way of distinguishing one cycle of the carrier from another one, the resulting time-difference measurement would be ambiguous modulo the period of the carrier, which is somewhat less than 1 ns—about 1000 times smaller than the period of the civilian pseudorandom code that is used by most receivers. Since the measurement noise in a time difference is typically some fraction of the period of the signal that is being observed, carrier-phase measurements can have more resolution than code-based estimates, assuming that the integer cycle count can be determined and that cycle jumps can be detected and removed. However, there is usually no way of determining the correct integer cycle count using the carrier phase data alone, and most analyses use the code-based time difference to assist in identifying the correct integer cycle count. Thus, while carrier-phase measurements have greater resolution than those based on the code, the accuracy is often no better. The difficulty of detecting cycle slips in the carrier-phase data depends on the stability of the clock in the receiver, since a more stable clock makes it easier to detect time steps in the data that are due to cycle slips. Comparing measurements using multiple satellites is also used to detect a cycle slip in the data from one of them.

The increase in resolution that can be realized by the use of the phase of the carrier requires a corresponding increase in the accuracy of the various corrections that we discussed above in connection with the pseudorange. In practice, the carrier-phase methods are used with post-processed ephemerides, since the parameters broadcast by the satellites are not accurate enough to take full advantage of the increased resolution that is possible using the phase of the carrier. The post-processed ephemerides are available from the International Geophysical Service (IGS). The IGS has a number of different products with varying delays from real time. Under favorable conditions, the carrier-phase method has a statistical uncertainty of about 50–100 ps using an averaging time on the order of minutes. The accuracy of the data is generally somewhat poorer than this value, since the integer cycle must be determined from the code data, and the long-term accuracy is generally no better than what can be achieved using a code-based analysis.

The time delay in the availability of post-processed ephemerides has been slowly decreasing as more sophisticated computational methods are developed at the analysis centers. There are now "ultra-rapid" ephemerides with delays of hours.

Frequency transfer. Frequency comparisons are generally performed by observing the evolution of the time difference between two clocks over some averaging time. The frequency difference is then expressed as seconds per second, which is a dimensionless parameter. For example, a clock that gained 1 μs per day with respect to a reference device would have a frequency of 1 μs/86,400 s $= 1.16 \times 10^{-11}$. Using the same ideas, the accuracy of a frequency comparison is determined by the residual noise of the time-difference process divided by the averaging time between observations. For example, if the residual noise in the time-difference measurements is 50 ns, then the frequency transfer noise using a 1-h (3600-s) average will be about 1.39×10^{-11}. For a given uncertainty in the measured time differences, the best frequency transfer will be determined by the maximum averaging time, which is the time over which the parameters of the local clock are constant. This maximum averaging time ranges from seconds for quartz-crystal oscillators to days for hydrogen masers. In practice, the stochastic fluctuations in the systematic corrections we have discussed limit the maximum averaging time to about a month even under ideal conditions (which cannot be realized routinely), and this limits the minimum uncertainty of frequency comparisons to about 3×10^{-16}. This limit is much better than is needed for almost all current applications, but it will not be small enough to compare the next generation of primary frequency standards, which will have stabilities that are smaller than this limit.

Applications. Satellite time signals are used in numerous applications—far too many to enumerate in this article. The signals are widely used to synchronize the time of computer systems and network elements. A number of commercially-available devices can provide time signals in the Network Time Protocol (NTP) format (an Internet standard that defines messages used to transmit time over digital networks) that are synchronized using signals from a satellite constellation.

More demanding applications include providing a frequency reference for the telecommunications network, which requires a frequency accuracy of 1×10^{-11}, and the NASA deep-space network, which requires time synchronization on the order of nanoseconds. Radio astronomy observatories, such as Arecibo in Puerto Rico, have similar requirements.

The data are used as one of the primary techniques for comparing the times of National Metrology Institutes and timing laboratories such as the National Institute for Standards and Technology and the Naval Observatory in the United States and similar laboratories in other countries. These comparisons depend on timing accuracies on the order of 1 ns and frequency comparisons with an uncertainty of less than 5×10^{-15}. These time comparisons form the basis for the computation of TAI and UTC. These comparisons are important for the international interoperability of timing, communications, and navigation systems and for evaluating the next generation of primary frequency standards, which are being developed in many laboratories in different countries. *See* OPTICAL CLOCKS AND RELATIVITY.

The future. There will be a significant increase in the number of satellites that are available for time and frequency comparisons in the next few years. In addition to the U.S. Global Positioning System, the Russian GLONASS system is currently operational, the European Galileo system will be operational in a few years, the Chinese are planning a system, and a number of other countries are planning either regional or global satellite systems that will be useful for time and frequency distribution using the methods that have been described. An important aspect will be interoperability—the ability to combine the data from the different satellite systems in a single time or frequency comparison. The accuracy of the time and frequency data from all of these systems will be limited by the considerations have been discussed, so that improving the accuracy of time and frequency comparisons will depend on how well these limitations are addressed by all of the satellite operators. Improving the accuracy of the broadcast ephemerides and the stability of the clocks on the satellites will facilitate real-time clock comparisons, and more sophisticated receivers that can reject multipath signals and have greater immunity to local environmental perturbations will also be very useful.

The jamming of satellite signals, either inadvertently or intentionally, is also likely to become important in the future, and receivers that have greater immunity to extraneous signals will become more important.

For background information *see* ATOMIC CLOCK; ATOMIC TIME; EPHEMERIS; FREQUENCY MEASUREMENT; RADIO ASTRONOMY; SATELLITE NAVIGATION SYSTEMS; SPACECRAFT GROUND INSTRUMENTATION; TIME in the McGraw-Hill Encyclopedia of Science & Technology. Judah Levine

Bibliography. C. Hackman and D. B. Sullivan (eds.), *Time and Frequency Measurement*, American Association of Physics Teachers, College Park, MD, 1996; J. Levine, Introduction to time and frequency metrology, *Rev. Sci. Instrum.*, 70:2567–2596, 1999, DOI:10.1063/1.1149844.