Time and Frequency

Michael A. Lombardi
National Institute of Standards and Technology

GLOSSARY

Accuracy Degree of conformity of a measured or calculated value to its definition.

Allan deviation $\sigma_{\gamma}, (\tau)$ Statistic used to estimate frequency stability.

Coordinated Universal Time (UTC) International atomic time scale used by all major countries.

Nominal frequency Ideal frequency, with zero uncertainty relative to its definition.

Q Quality factor of an oscillator, estimated by dividing the resonance frequency by the resonance width.

Resonance frequency Natural frequency of an oscillator, based on the repetition of a periodic event.

Second Duration of 9,192,631,770 periods of the radiation corresponding to the transition between two hyperfine levels of the ground state of the cesium-133 atom.

Stability Statistical estimate of the frequency or time fluctuations of a signal over a given time interval.

Synchronization Process of setting two or more clocks to the same time.

Syntonization Process of setting two or more oscillators to the same frequency.

Time scale Agreed-upon system for keeping time.

THIS article is an overview of time and frequency technology. It introduces basic time and frequency concepts and describes the devices that produce time and frequency signals and information. It explains how these devices work and how they are measured. Section I introduces the basic concepts of time and frequency and provides some historical background. Section II discusses time and frequency measurements and the specifications used to state the measurement results. Section III discusses time and frequency standards. These devices are grouped into two categories: quartz and atomic oscillators. Section IV discusses time and frequency transfer, or the process of using a clock or frequency standard to measure or set a device at another location.

I. CONCEPTS AND HISTORY

Few topics in science and technology are as relevant as time and frequency. Time and frequency standards and measurements are involved in nearly every aspect of daily life and are a fundamental part of many technologies.

Time and frequency standards supply three basic types of information. The first type, *date* and *time-of-day*,

Encyclopedia of Physical Science and Technology, Third Edition, Volume 16

records when an event happened. Date and time-of-day can also be used to ensure that events are *synchronized*, or happen at the same time. Everyday life is filled with examples of the use of date and time-of-day information. Date information supplied by calendars records when birthdays, anniversaries, and other holidays are scheduled to occur. The time-of-day information supplied by watches and clocks helps keep our lives on schedule. Meeting a friend for dinner at 6 P.M. is a simple example of synchronization. If our watches agree, we should both arrive at about the same time.

Date and time-of-day information has other, more sophisticated uses. Airplanes flying in a formation require synchronized clocks. If one airplane banks or turns at the wrong time, it could result in a collision and loss of life. When a television station broadcasts a network program, it must start broadcasting the network feed at the instant it arrives. If the station and network clocks are not synchronized, part of the program is skipped. Stock market transactions require synchronized clocks so that the buyer and seller can agree on the same price at the same time. A time error of a few seconds could cost the buyer or seller many thousands of dollars. Electric power companies also use synchronization. They synchronize the clocks in their power grids, so they can instantly transfer power to the parts of the grid where it is needed most, and to avoid electrical overload.

The second type of information, *time interval*, is the duration or elapsed time between two events. Our age is simply the time interval since our birth. Most workers are paid for the time interval during which they work, usually measured in hours, weeks, or months. We pay for time interval as well—30 min on a parking meter, a 20-min cab ride, a 5-min long-distance phone call, or a 30-sec radio advertising spot.

The standard unit of time interval is the *second*. However, many applications in science and technology require the measurement of time intervals much shorter than 1 sec, such as *milliseconds* (10^{-3} sec), *microseconds* (10^{-6} sec), *nanoseconds* (10^{-9} sec), and *picoseconds* (10^{-12} sec).

The third type of information, *frequency*, is the rate of a repetitive event. If T is the period of a repetitive event, then the frequency f is its reciprocal, $1/T$. The International System of Units (SI) states that the period should be expressed as seconds (sec), and the frequency should be expressed as hertz (Hz). The frequency of electrical signals is measured in units of kilohertz (kHz), megahertz (MHz), or gigahertz (GHz), where 1 kHz equals 1000 (10^3) events per second, 1 MHz equals 1 million (10^6) events per second, and 1 GHz equals 1 billion (10^9) events per second. Many frequencies are encountered in everyday life. For example, a quartz wristwatch works by counting the oscillations of a crystal whose frequency is 32,768 Hz. When the crystal has oscillated 32,768 times, the watch records that 1 sec has elapsed. A television tuned to channel 7 receives a video signal at a frequency of 175.25 MHz. The station transmits this frequency as closely as possible, to avoid interference with signals from other stations. A computer that processes instructions at a frequency of 1 GHz might connect to the Internet using a T1 line that sends data at a frequency of 1.544 MHz.

Accurate frequency is critical to communications networks. The highest-capacity networks run at the highest frequencies. Networks use groups of oscillators that produce nearly the same frequency, so they can send data at the fastest possible rates. The process of setting multiple oscillators to the same frequency is called *syntonization*.

Of course, the three types of time and frequency information are closely related. As mentioned, the standard unit of time interval is the second. By counting seconds, we can determine the date and the time-of-day. And by counting the events per second, we can measure the frequency.

A. The Evolution of Time and Frequency Standards

All time and frequency standards are based on a *periodic event* that repeats at a constant rate. The device that produces this event is called a *resonator*. In the simple case of a pendulum clock, the pendulum is the resonator. Of course, a resonator needs an energy source before it can move back and forth. Taken together, the energy source and resonator form an *oscillator*. The oscillator runs at a rate called the *resonance frequency*. For example, a clock's pendulum can be set to swing back and forth at a rate of once per second. Counting one complete swing of the pendulum produces a time interval of 1 sec. Counting the total number of swings creates a *time scale* that establishes longer time intervals, such as minutes, hours, and days. The device that does the counting and displays or records the results is called a *clock*. The frequency uncertainty of a clock's resonator relates directly to the timing uncertainty of the clock as shown in Table I.

Throughout history, clock designers have searched for stable resonators. As early as 3500 B.C., time was kept by observing the movement of an object's shadow between sunrise and sunset. This simple clock is called a *sundial*, and the resonance frequency is based on the apparent motion of the sun. Later, water clocks, hourglasses, and calibrated candles allowed dividing the day into smaller units of time. Mechanical clocks first appeared in the early 14th century. Early models used a verge and foliet mechanism for a resonator and had an uncertainty of about 15 min/day ($\cong 1 \times 10^{-2}$).

A timekeeping breakthrough occurred with the invention of the *pendulum clock*, a technology that dominated

TABLE I Relationship of Frequency Uncertainty to Time Uncertainty

Frequency uncertainty	Measurement period	Time uncertainty
$\pm 1.00 \times 10^{-3}$	1 sec	±1 msec
$\pm 1.00 \times 10^{-6}$	1 sec	±1 μsec
$\pm 1.00 \times 10^{-9}$	1 sec	±1 nsec
$\pm 2.78 \times 10^{-7}$	1 hr	±1 msec
$\pm 2.78 \times 10^{-10}$	1 hr	±1 μsec
$\pm 2.78 \times 10^{-13}$	1 hr	±1 nsec
$\pm 1.16 \times 10^{-8}$	1 day	±1 msec
$\pm 1.16 \times 10^{-11}$	1 day	±1 μsec
$\pm 1.16 \times 10^{-14}$	1 day	±1 nsec

timekeeping for several hundred years. Prior to the invention of the pendulum, clocks could not count minutes reliably, but pendulum clocks could count seconds. In the early 1580s, Galileo Galilei observed that a given pendulum took the same amount of time to swing completely through a wide arc as it did a small arc. Galileo wanted to apply this natural periodicity to time measurement and began work on a mechanism to keep the pendulum in motion in 1641, the year before he died. In 1656, the Dutch scientist Christiaan Huygens invented an escapement that kept the pendulum swinging. The uncertainty of Huygens's clock was less than 1 min/day ($\cong 7 \times 10^{-4}$) and later was reduced to about 10 sec/day ($\cong 1 \times 10^{-4}$). The first pendulum clocks were weight-driven, but later versions were powered by springs. In fact, Huygens is often credited with inventing the spring-and-balance wheel assembly still found in some of today's mechanical wristwatches.

Huge advances in accuracy were made by John Harrison, who built and designed a series of clocks in the 1720s that kept time to within fractions of a second per day (parts in 10^6). This performance was not improved upon until the 20th century. Harrison dedicated most of his life to solving the British navy's problem of determining longitude, by attempting to duplicate the accuracy of his land clocks at sea. He built a series of clocks (now known as H1 through H5) in the period from 1730 to about 1770. He achieved his goal with the construction of H4, a clock much smaller than its predecessors, about the size of a large pocket watch. H4 used a spring and balance wheel escapement and kept time within fractions of a second per day during several sea voyages in the 1760s.

The practical performance limit of pendulum clocks was reached in 1921, when W. H. Shortt demonstrated a clock with two pendulums, one a slave and the other a master. The slave pendulum moved the clock's hands and freed the master pendulum of tasks that would disturb its regularity. The pendulums used a battery as their power supply. The Shortt clock kept time to within a few seconds per year ($\cong 1 \times 10^{-7}$) and was used as a primary standard in the United States.

Joseph W. Horton and Warren A. Marrison of Bell Laboratories built the first clock based on a quartz crystal oscillator in 1927. By the 1940s, quartz clocks had replaced pendulums as primary laboratory standards. Quartz crystals resonate at a nearly constant frequency when an electric current is applied. Uncertainties of <100 μsec/day ($\cong 1 \times 10^{-9}$) are possible, and low-cost quartz oscillators are found in electronic circuits and inside nearly every wristwatch and wall clock.

Quartz oscillators still have shortcomings since their resonance frequency depends on the size and shape of the crystal. No two crystals can be precisely alike or produce exactly the same frequency. Quartz oscillators are also sensitive to temperature, humidity, pressure, and vibration. These limitations made them unsuitable for some high-level applications and led to the development of atomic oscillators.

In the 1930s, I. I. Rabi and his colleagues at Columbia University introduced the idea of using an atomic resonance as a frequency. The first atomic oscillator, based on the ammonia molecule, was developed at the National Bureau of Standards (now the National Institute of Standards and Technology) in 1949. A Nobel Prize was awarded in 1989 to Norman Ramsey, Hans Dehmelt, and Wolfgang Paul for their work in atomic oscillator development, and many other scientists have made significant contributions to the technology. Atomic oscillators use the quantized energy levels in atoms and molecules as the source of their resonance frequency. The laws of quantum mechanics dictate that the energies of a bound system, such as an atom, have certain discrete values. An electromagnetic field at a particular frequency can boost an atom from one energy level to a higher one. Or, an atom at a high energy level can drop to a lower level by emitting energy. The resonance frequency (f) of an atomic oscillator is the difference between the two energy levels divided by Planck's constant (h):

$$f = \frac{E_2 - E_1}{h}.$$

The principle underlying the atomic oscillator is that since all atoms of a specific element are identical, they should produce the exact same frequency when they absorb or release energy. In theory, the atom is a perfect pendulum whose oscillations are counted to measure the time interval. Quartz and the three main types of atomic oscillators (rubidium, hydrogen, and cesium) are described in detail in Section III.

Table II summarizes the evolution of time and frequency standards. The uncertainties listed for modern standards

TABLE II The Evolution of Time and Frequency Standards

Standard	Resonator	Date of origin	Timing uncertainty (24 hr)	Frequency uncertainty (24 hr)
Sundial	Apparent motion of sun	3500 B.C.	NA	NA
Verge escapement	Verge and foliet mechanism	14th century	15 min	1×10^{-2}
Pendulum	Pendulum	1656	10 sec	7×10^{-4}
Harrison chronometer (H4)	Pendulum	1759	300 msec	3×10^{-6}
Shortt pendulum	Two pendulums, slave and master	1921	10 msec	1×10^{-7}
Quartz crystal	Quartz crystal	1927	10 μsec	1×10^{-10}
Rubidium gas cell	^{87}Rb resonance (6,834,682,608 Hz)	1958	100 nsec	1×10^{-12}
Cesium beam	^{133}Cs resonance (9,192,631,770 Hz)	1952	1 nsec	1×10^{-14}
Hydrogen maser	Hydrogen resonance (1,420,405,752 Hz)	1960	1 nsec	1×10^{-14}
Cesium fountain	^{133}Cs resonance (9,192,631,770 Hz)	1991	100 psec	1×10^{-15}

represent current (year 2000) devices, and not the original prototypes. Note that the performance of time and frequency standards has improved by 13 orders of magnitude in the past 700 years and by about 9 orders of magnitude in the past 100 years.

B. Time Scales and the International Definition of the Second

The second is one of seven base units in the International System of Units (SI). The base units are used to derive other units of physical quantities. Use of the SI means that physical quantities such as the second and hertz are defined and measured in the same way throughout the world.

There have been several definitions of the SI second. Until 1956, the definition was based on the *mean solar day*, or one revolution of the earth on its axis. The *mean solar second* was defined as 1/86,400 of the mean solar day and provided the basis for several astronomical time scales known as Universal Time (UT).

UT0: The original mean solar time scale, based on the rotation of the earth on its axis. UT0 was first kept with pendulum clocks. When quartz clocks became available, astronomers noticed errors in UT0 due to polar motion and developed the UT1 time scale.

UT1: The most widely used astronomical time scale, UT1 improves upon UT0 by correcting for longitudinal shifts of the observing station due to polar motion. Since the earth's rotational rate is not uniform the uncertainty of UT1 is about 2 to 3 msec per day.

UT2: Mostly of historical interest, UT2 is a smoothed version of UT1 that corrects for deviations in the period of the earth's rotation caused by angular momenta of the earth's core, mantle, oceans and atmosphere.

The *ephemeris second* served as the SI second from 1956 to 1967. The ephemeris second was a fraction of the tropical year, or the interval between the annual vernal equinoxes, which occur on or about March 21. The tropical year was defined as 31,556,925.9747 ephemeris sec. Determining the precise instant of the equinox is difficult, and this limited the uncertainty of Ephemeris Time (ET) to ±50 msec over a 9-year interval. ET was used mainly by astronomers and was replaced by *Terrestial Time* (TT) in 1984, equal to International Atomic Time (TAI) + 32.184 sec. The uncertainty of TT is ±10 μsec.

The era of atomic time keeping formally began in 1967, when the SI second was redefined based on the resonance frequency of the cesium atom:

The duration of 9,192,631,770 periods of the radiation corresponding to the transition between two hyperfine levels of the ground state of the cesium-133 atom.

Due to the atomic second, time interval and frequency can now be measured with less uncertainty and more resolution than any other physical quantity. Today, the best time and frequency standards can realize the SI second with uncertainties of $\cong 1 \times 10^{-15}$. Physical realizations of the other base SI units have much larger uncertainties (Table III).

International Atomic Time (TAI) is an atomic time scale that attempts to realize the SI second as closely as possible. TAI is maintained by the Bureau International des Poids et Measures (BIPM) in Sevres, France. The BIPM averages data collected from more than 200 atomic time and frequency standards located at more than 40 laboratories, including the National Institute of Standards and Technology (NIST).

Coordinated Universal Time (UTC) runs at the same rate as TAI. However, it differs from TAI by an integral number of seconds. This difference increases when *leap*

TABLE III Uncertainties of Physical Realizations of the Base SI Units

SI base unit	Physical quantity	Uncertainty
Candela	Luminous intensity	10^{-4}
Mole	Amount of substance	10^{-7}
Kelvin	Thermodynamic temperature	10^{-7}
Ampere	Electric current	10^{-8}
Kilogram	Mass	10^{-8}
Meter	Length	10^{-12}
Second	Time interval	10^{-15}

seconds occur. When necessary, leap seconds are added to UTC on either June 30 or December 31. The purpose of adding leap seconds is to keep atomic time (UTC) within ±0.9 sec of astronomical time (UT1). Some time codes contain a UT1 correction that can be applied to UTC to obtain UT1.

Leap seconds have been added to UTC at a rate of slightly less than once per year, beginning in 1972. UT1 is currently losing about 700 to 800 msec per year with respect to UTC. This means that atomic seconds are shorter than astronomical seconds and that UTC runs faster than UT1. There are two reasons for this. The first involves the definition of the atomic second, which made it slightly shorter than the astronomical second to begin with. The second reason is that the earth's rotational rate is gradually slowing down and the astronomical second is gradually getting longer. When a positive leap second is added to UTC, the sequence of events is as follows.

23 hr 59 min 59 sec
23 hr 59 min 60 sec
0 hr 0 min 0 sec

The insertion of the leap second creates a minute that is 61 sec long. This "stops" UTC for 1 sec, so that UTI can catch up.

II. TIME AND FREQUENCY MEASUREMENT

Time and frequency measurements follow the conventions used in other areas of metrology. The frequency standard or clock being measured is called the *device under test* (DUT). The measurement compares the DUT to a *standard* or *reference*. The standard should outperform the DUT by a specified ratio, ideally by 10:1. The higher the ratio, the less averaging is required to get valid measurement results.

The test signal for time measurements is usually a pulse that occurs once per second (1 pps). The pulse width and polarity vary from device to device, but TTL levels are commonly used. The test signal for frequency measurements is usually a frequency of 1 MHz or higher, with 5 or 10 MHz being common. Frequency signals are usually sine waves but can be pulses or square waves.

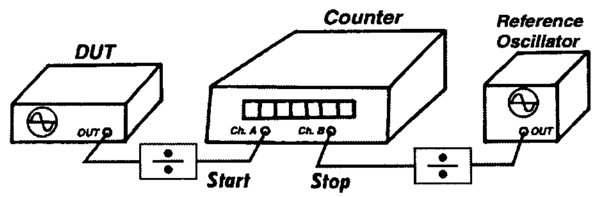

FIGURE 1 Measurement using a time interval counter.

This section examines the two main specifications of time and frequency measurements—*accuracy* and *stability*. It also discusses some instruments used to measure time and frequency.

A. Accuracy

Accuracy is the degree of conformity of a measured or calculated value to its definition. Accuracy is related to the offset from an ideal value. For example, *time offset* is the difference between a measured on-time pulse and an ideal on-time pulse that coincides exactly with UTC. *Frequency offset* is the difference between a measured frequency and an ideal frequency with zero uncertainty. This ideal frequency is called the *nominal frequency*.

Time offset is usually measured with a *time interval counter* (TIC) as shown in Fig. 1. A TIC has inputs for two signals. One signal starts the counter and the other signal stops it. The time interval between the start and the stop signals is measured by counting cycles from the time base oscillator. The resolution of low-cost TICs is limited to the period of their time base. For example, a TIC with a 10-MHz time base oscillator would have a resolution of 100 nsec. More elaborate TICs use interpolation schemes to detect parts of a time base cycle and have a much higher resolution—1-nsec resolution is commonplace, and even 10-psec resolution is available.

Frequency offset can be measured in either the *frequency domain* or the *time domain*. A simple frequency domain measurement involves directly counting and displaying the frequency output of the DUT with a *frequency counter*. The reference for this measurement is either the counter's internal time base oscillator, or an external time base (Fig. 2). The counter's resolution, or the number of digits it can display, limits its ability to measure frequency offset. The frequency offset is determined as

$$f(\text{offset}) = \frac{f_{\text{measured}} - f_{\text{nominal}}}{f_{\text{nominal}}},$$

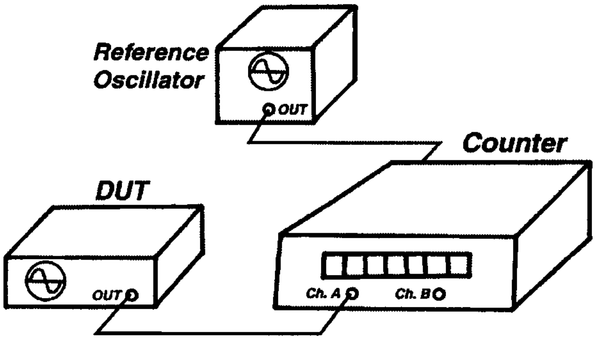

FIGURE 2 Measurement using a frequency counter.

where f_{measured} is the reading from the frequency counter, and f_{nominal} is the frequency labeled on the oscillator's nameplate.

Frequency offset measurements in the time domain involve a *phase comparison* between the DUT and the reference. A simple phase comparison can be made with an oscilloscope (Fig. 3). The oscilloscope will display two sine waves (Fig. 4). The top sine wave represents a signal from the DUT, and the bottom sine wave represents a signal from the reference. If the two frequencies were exactly the same, their phase relationship would not change and both would appear to be stationary on the oscilloscope display. Since the two frequencies are not exactly the same, the reference appears to be stationary and the DUT signal moves. By determining the rate of motion of the DUT signal, we can determine its frequency offset. Vertical lines have been drawn through the points where each sine wave passes through zero. The bottom of the figure shows bars whose width represents the phase difference between the signals. This difference increases or decreases to indicate whether the DUT frequency is high or low with respect to the reference.

Measuring high-accuracy signals with an oscilloscope is impractical, since the phase relationship between signals changes very slowly. More precise phase comparisons can be made with a time interval counter, using a setup similar to Fig. 1. Since frequencies like 5 or 10 MHz are usually involved, *frequency dividers* (shown in Fig. 1) or *frequency mixers* are used to convert the test frequency to a lower frequency. Measurements are made from the TIC, but instead of using these measurements directly, we determine the rate of change from reading to reading. This rate of change is called the phase deviation. We can estimate frequency offset as follows, where Δt is the amount of phase deviation, and T is the measurement period:

$$f(\text{offset}) = \frac{-\Delta t}{T}.$$

To illustrate, consider a measurement of $+1\ \mu$sec of phase deviation over a measurement period of 24 hr. The unit used for measurement period (hr) must be converted to the unit used for phase deviation (μsec). The equation becomes

$$f(\text{offset}) = \frac{-\Delta t}{T} = \frac{1\ \mu\text{sec}}{86{,}400{,}000{,}000\ \mu\text{sec}} = -1.16 \times 10^{-11}.$$

As shown, a device that accumulates 1 μsec of phase deviation/day has a frequency offset of about -1.16×10^{-11} with respect to the reference.

Dimensionless frequency offset values can be converted to units of frequency (Hz) if the nominal frequency is known. To illustrate this, consider an oscillator with a nominal frequency of 5 MHz and a frequency offset of $+1.16 \times 10^{-11}$. To find the frequency offset in hertz, multiply the nominal frequency by the offset:

$$(5 \times 10^{6})(+1.16 \times 10^{-11}) = 5.80 \times 10^{-5} = +0.0000580\ \text{Hz}.$$

Then add the offset to the nominal frequency to get the actual frequency:

$$5{,}000{,}000\ \text{Hz} + 0.0000580\ \text{Hz} = 5{,}000{,}000.0000580\ \text{Hz}.$$

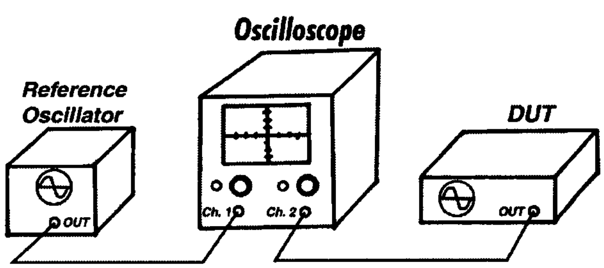

FIGURE 3 Phase comparison using an oscilloscope.

B. Stability

Stability indicates how well an oscillator can produce the same time or frequency offset over a given period of time. It does not indicate whether the time or frequency is "right" or "wrong" but only whether it stays the same. In contrast, accuracy indicates how well an oscillator has been set on time or set on frequency. To understand this difference, consider that a stable oscillator that needs adjustment might produce a frequency with a large offset. Or an unstable oscillator that was just adjusted might temporarily produce a frequency near its nominal value. Figure 5 shows the relationship between accuracy and stability.

Stability is defined as the statistical estimate of the frequency or time fluctuations of a signal over a given time

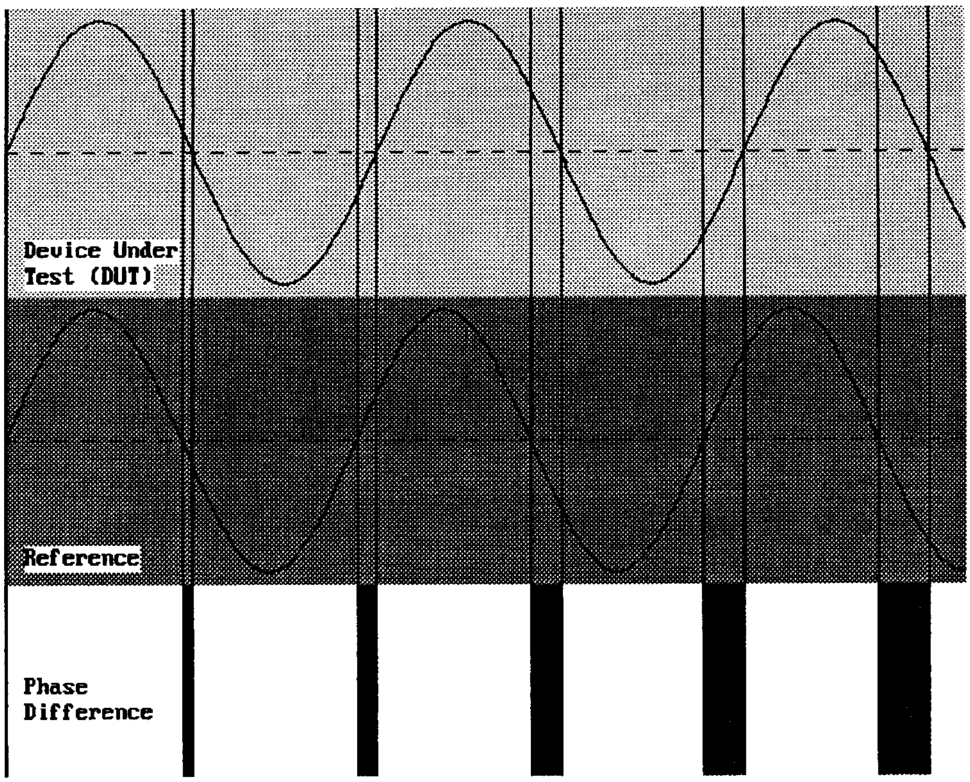

FIGURE 4 Two sine waves with a changing phase relationship.

interval. These fluctuations are measured with respect to a mean frequency or time offset. *Short-term* stability usually refers to fluctuations over intervals less than 100 sec. *Long-term* stability can refer to measurement intervals greater than 100 sec but usually refers to periods longer than 1 day.

Stability estimates can be made in either the frequency domain or the time domain, but time domain estimates are more common and are discussed in this section. To estimate frequency stability in the time domain, we can start with a series of phase measurements. The phase measurements are nonstationary, since they contain a trend contributed by the frequency offset. With nonstationary data, the mean and variance never converge to any particular values. Instead, there is a moving mean that changes each time we add a measurement.

For these reasons, a nonclassical statistic is often used to estimate stability in the time domain. This statistic is sometimes called the *Allan variance*, but since it is the square root of the variance, its proper name is the *Allan deviation*. The equation for the Allan deviation is

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(M-1)} \sum_{i=1}^{M-1} (y_{i+1} - y_i)^2},$$

where y_i is a set of frequency offset measurements that consists of individual measurements, y_1, y_2, y_3, and so on,

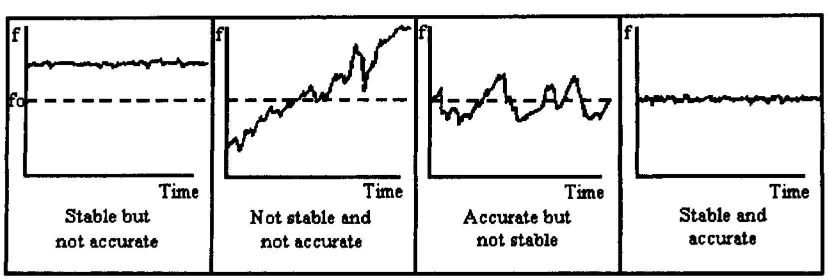

FIGURE 5 The relationship between accuracy and stability.

TABLE IV Using Phase Measurements to Estimate Stability

Phase measurement (nsec), x_i	Phase deviation (nsec), Δt	Frequency offset $\Delta t/\tau(y_i)$	First difference $(y_{i+1} - y_i)$	First difference squared $(y_{i+1} - y_i)^2$
3321.44	(—)	(—)	(—)	(—)
3325.51	4.07	4.07×10^{-9}	(—)	(—)
3329.55	4.04	4.04×10^{-9}	-3×10^{-11}	9×10^{-22}
3333.60	4.05	4.05×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3337.65	4.05	4.06×10^{-9}	$+2 \times 10^{-11}$	4×10^{-22}
3341.69	4.04	4.04×10^{-9}	-2×10^{-11}	4×10^{-22}
3345.74	4.05	4.05×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3349.80	4.06	4.06×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3353.85	4.05	4.05×10^{-9}	-1×10^{-11}	1×10^{-22}
3357.89	4.04	4.04×10^{-9}	-1×10^{-11}	1×10^{-22}

M is the number of values in the y_i series, and the data are equally spaced in segments τ seconds long. Or

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(N-2)\tau^2}\sum_{i=1}^{N-2}[x_{i+2} - 2x_{i+1} + x_i]^2},$$

where x_i is a set of phase measurements in time units that consists of individual measurements, x_1, x_2, x_3, and so on, N is the number of values in the x_i series, and the data are equally spaced in segments τ seconds long.

Table IV shows how the Allan deviation is calculated. The left column contains a series of phase measurements recorded once per second ($\tau = 1$ sec) in units of nanoseconds. These measurements have a trend; note that each value in the series is larger than the previous value. By subtracting pairs of values, we remove the trend and obtain the phase deviations (Δt) shown in the second column. The third column divides the phase deviation (Δt) by τ to get the frequency offset values, or the y_i data series. The last two columns show the first differences of the y_i and the squares of the first differences.

Since the sum of the squares equals 2.2×10^{-21}, the frequency stability using the first equation (at $\tau = 1$ sec) is

$$\sigma_y(\tau) = \sqrt{\frac{2.2 \times 10^{-21}}{2(9-1)}} = 1.17 \times 10^{-11}.$$

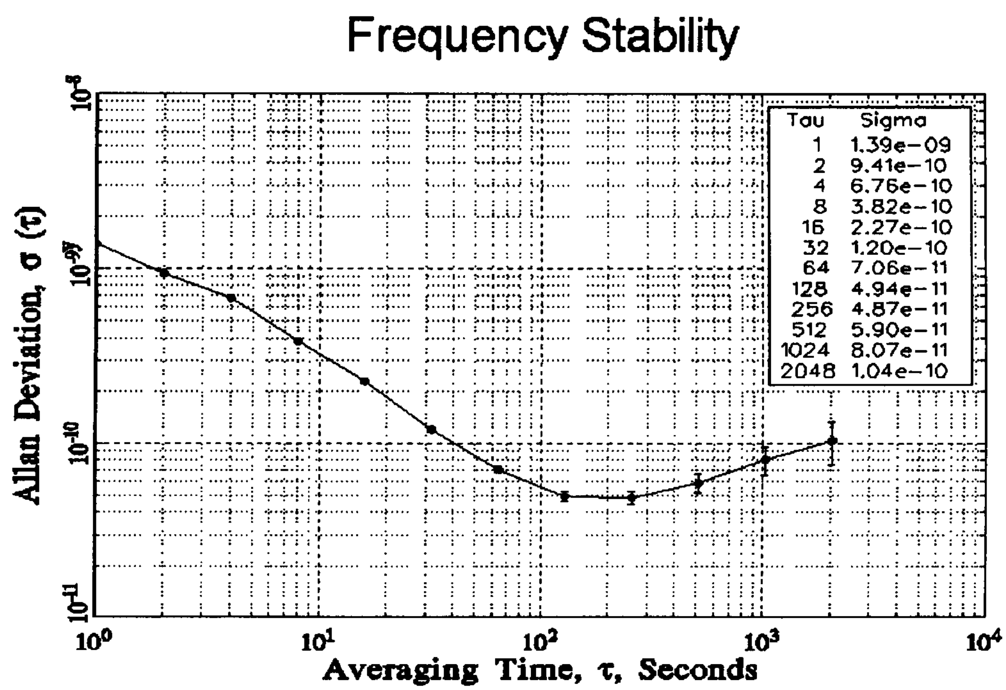

FIGURE 6 A graph of frequency stability.

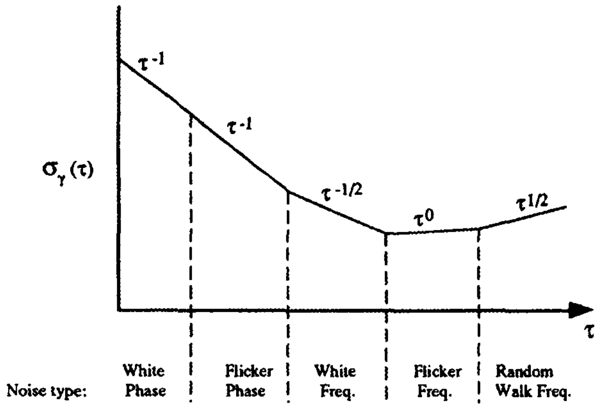

FIGURE 7 Using a frequency stability graph to identify noise types.

A graph of the Allan deviation is shown in Fig. 6. It shows the stability of the device improving as the averaging period (τ) gets longer, since some noise types can be removed by averaging. At some point, however, more averaging no longer improves the results. This point is called the *noise floor*, or the point where the remaining noise consists of nonstationary processes like aging or random walk. The device measured in Fig. 6 is has a noise floor of $\cong 5 \times 10^{-11}$ at $\tau = 100$ sec.

Five noise types are commonly discussed in the time and frequency literature: *white phase*, *flicker phase*, *white frequency*, *flicker frequency*, and *random walk frequency*. The slope of the Allan deviation line can identify the amount of averaging needed to remove these noise types (Fig. 7). Note that the Allan deviation does not distinguish between white phase noise and flicker phase noise. Several other statistics are used to estimate stability and identify noise types for various applications (Table V).

C. Uncertainty Analysis

Time and frequency metrologists must often perform an uncertainty analysis when they calibrate or measure a device. The uncertainty analysis states the measurement error with respect to a national or international standard, such as UTC(NIST) or UTC. Two simple ways to estimate measurement uncertainty are discussed here. Both use the concepts of accuracy and stability discussed above.

One common type of uncertainty analysis involves making multiple measurements and showing that a single measurement will probably fall within a stated range of values. The standard deviation (or an equivalent statistic) is usually both added and subtracted from the mean to form the upper and lower bounds of the range. The stated probability that a given measurement will fall within this range is usually 1σ (68.3%) or 2σ (95.4%).

In time and frequency metrology, the mean value is usually the accuracy (mean time or mean frequency offset), and the deviation in the mean is usually calculated using one of the statistics listed in Table IV. For example, if a device has a frequency offset of 2×10^{-9} and a 2σ stability of 2×10^{-10}, there is a 95.4% probability that the frequency offset will be between 1.8 and 2.2 parts in 10^9.

The second type of uncertainty analysis involves adding the *systematic* and *statistical* uncertainties to find the combined uncertainty. For example, consider a time signal received from a radio station where the mean path delay is measured as 9 msec (time offset), and the deviation in the path delay is measured as 0.5 msec (stability). In this example, 9 msec is the systematic uncertainty and 0.5 msec is the statistical uncertainty. For some applications, it is convenient to simply add the two numbers together and state the combined uncertainty as <10 msec.

III. TIME AND FREQUENCY STANDARDS

The stability of time and frequency standards is closely related to their quality factor, or Q. The Q of an oscillator is its resonance frequency divided by its resonance width. The resonance frequency is the natural frequency of the oscillator. The resonance width is the range of possible values where the oscillator will run. A high-Q resonator will not oscillate at all unless it is near its resonance frequency. Obviously a high resonance frequency and a narrow resonance width are both advantages when seeking a high Q.

TABLE V Statistics Used to Estimate Time and Frequency Stability and Noise Types

Name	Mathematical notation	Description
Allan deviation	$\sigma_y(\tau)$	Estimates frequency stability. Particularly suited for intermediate to long-term measurements.
Modified Allan deviation	MOD $\sigma_y(\tau)$	Estimates frequency stability. Unlike the normal Allan deviation, it can distinguish between white and flicker phase noise, which makes it more suitable for short-term stability estimates.
Time deviation	$\sigma_x(\tau)$	Used to measure time stability. Clearly identifies both white and flicker phase noise, the noise types of most interest when measuring time or phase.
Total deviation	$\sigma_{y,\mathrm{TOTAL}}(\tau)$	Estimates frequency stability. Particularly suited for long-term estimates where τ exceeds 10% of the total data sample.

TABLE VI Summary of Oscillator Types

Oscillator type	Quartz TCXO	Quartz OCXO	Rubidium	Commercial cesium beam	Hydrogen maser
Q	10^4 to 10^6	3.2×10^6 (5 MHz)	10^7	10^8	10^9
Resonance frequency	Various	Various	6.834682608 GHz	9.192631770 GHz	1.420405752 GHz
Leading cause of failure	None	None	Rubidium lamp (15 years +)	Cesium beam tube (3 to 25 years)	Hydrogen depletion (7 years +)
Stability, $\sigma_y(\tau)$, $\tau = 1$ sec	1×10^{-8} to 1×10^{-9}	1×10^{-12}	5×10^{-11} to 5×10^{-12}	5×10^{-11} to 5×10^{-12}	1×10^{-12}
Noise floor, $\sigma_y(\tau)$	1×10^{-9} ($\tau = 1$ to 10^2 sec)	1×10^{-12} ($\tau = 1$ to 10^2 sec)	1×10^{-12} ($\tau = 10^3$ to 10^5 sec)	1×10^{-14} ($\tau = 10^5$ to 10^7 sec)	1×10^{-15} ($\tau = 10^3$ to 10^5 sec)
Aging/year	5×10^{-7}	5×10^{-9}	1×10^{-10}	None	$\cong 1 \times 10^{-13}$
Frequency offset after warm-up	1×10^{-6}	1×10^{-8} to 1×10^{-10}	5×10^{-10} to 5×10^{-12}	5×10^{-12} to 1×10^{-14}	1×10^{-12} to 1×10^{-13}
Warm-up period	<10 sec to 1×10^{-6}	<5 min to 1×10^{-8}	<5 min to 5×10^{-10}	30 min to 5×10^{-12}	24 hr to 1×10^{-12}

Generally speaking, the higher the Q, the more stable the oscillator, since a high Q means that an oscillator will stay close to its natural resonance frequency.

This section discusses quartz oscillators, which achieve the highest Q of any mechanical-type device. It then discusses oscillators with much higher Q factors, based on the atomic resonance of rubidium, hydrogen, and cesium. The performance of each type of oscillator is summarized in Table VI.

A. Quartz Oscillators

Quartz crystal oscillators are by far the most common time and frequency standard. An estimated 2 billion (2×10^9) quartz oscillators are manufactured annually. Most are small devices built for wristwatches, clocks, and electronic circuits. However, they are also found inside test and measurement equipment, such as counters, signal generators, and oscilloscopes, and interestingly enough, inside every atomic oscillator.

A quartz crystal inside the oscillator is the resonator. It can be made of natural or synthetic quartz, but all modern devices use synthetic quartz. The crystal strains (expands or contracts) when a voltage is applied. When the voltage is reversed, the strain is reversed. This is known as the *piezoelectric effect.* Oscillation is sustained by taking a voltage signal from the resonator, amplifying it, and feeding it back to the resonator. The rate of expansion and contraction is the resonance frequency and is determined by the cut and size of the crystal. The output frequency of a quartz oscillator is either the fundamental resonance or a multiple of the resonance, called an *overtone frequency.* Most high-stability units use either the third or the fifth overtone to achieve a high Q. Overtones higher than fifth are rarely used because they make it harder to tune the device to the desired frequency. A typical Q for a quartz oscillator ranges from 10^4 to 10^6. The maximum Q for a high-stability quartz oscillator can be estimated as $Q = 16$ million$/f$, where f is the resonance frequency in megahertz.

Environmental changes such as temperature, humidity, pressure, and vibration can change the resonance frequency of a quartz crystal, and there are several designs that reduce the environmental problems. The *oven-controlled crystal oscillator* (OCXO) encloses the crystal in a temperature-controlled chamber called an oven. When an OCXO is turned on, it goes through a "warm-up" period while the temperatures of the crystal resonator and its oven stabilize. During this time, the performance of the oscillator continuously changes until it reaches its normal operating temperature. The temperature within the oven, then remains constant, even when the outside temperature varies. An alternate solution to the temperature problem is the *temperature-compensated crystal oscillator* (TCXO). In a TCXO, the signal from a temperature sensor generates a correction voltage that is applied to a voltage-variable reactance, or varactor. The varactor then produces a frequency change equal and opposite to the frequency change produced by temperature. This technique does not work as well as oven control but is less expensive. Therefore, TCXOs are used when high stability over a wide temperature range is not required.

Quartz oscillators have excellent short-term stability. An OCXO might be stable ($\sigma_y \tau$, at $\tau = 1$ sec) to 1×10^{-12}. The limitations in short-term stability are due mainly to noise from electronic components in the oscillator circuits. Long-term stability is limited by *aging*, or a change in frequency with time due to internal changes in the oscillator. Aging is usually a nearly linear change in the resonance frequency that can be either positive or negative, and occasionally, a reversal in aging direction occurs. Aging has many possible causes including a buildup of foreign material on the crystal, changes in the oscillator circuitry, or changes in the quartz material or crystal structure. A high-quality OCXO might age at a rate of $<5 \times 10^{-9}$ per year, while a TCXO might age 100 times faster.

Due to aging and environmental factors such as temperature and vibration, it is hard to keep even the best quartz oscillators within 1×10^{-10} of their nominal frequency without constant adjustment. For this reason, atomic oscillators are used for applications that require higher long-term accuracy and stability.

B. Rubidium Oscillators

Rubidium oscillators are the lowest priced members of the atomic oscillator family. They operate at 6,834,682,608 Hz, the resonance frequency of the rubidium atom (^{87}Rb), and use the rubidium frequency to control the frequency of a quartz oscillator. A microwave signal derived from the crystal oscillator is applied to the ^{87}Rb vapor within a cell, forcing the atoms into a particular energy state. An optical beam is then pumped into the cell and is absorbed by the atoms as it forces them into a separate energy state. A photo cell detector measures how much of the beam is absorbed and tunes a quartz oscillator to a frequency that maximizes the amount of light absorption. The quartz oscillator is then locked to the resonance frequency of rubidium, and standard frequencies are derived and provided as outputs (Fig. 8).

Rubidium oscillators continue to get smaller and less expensive, and offer perhaps the best price/performance ratio of any oscillator. Their long-term stability is much better than that of a quartz oscillator and they are also smaller, more reliable, and less expensive than cesium oscillators.

The Q of a rubidium oscillator is about 10^7. The shifts in the resonance frequency are caused mainly by collisions of the rubidium atoms with other gas molecules. These shifts limit the long-term stability. Stability ($\sigma_y\tau$, at $\tau = 1$ sec) is typically 1×10^{-11}, and about 1×10^{-12} at 1 day. The frequency offset of a rubidium oscillator ranges from 5×10^{-10} to 5×10^{-12} after a warm-up period of a few minutes, so they meet the accuracy requirements of most applications without adjustment.

C. Cesium Oscillators

Cesium oscillators are *primary frequency standards* since the SI second is defined using the resonance frequency of the cesium atom (^{133}Cs), which is 9,192,631,770 Hz. A properly working cesium oscillator should be close to its nominal frequency without adjustment, and there should be no change in frequency due to aging.

Commercially available oscillators use *cesium beam* technology. Inside a cesium oscillator, ^{133}Cs atoms are heated to a gas in an oven. Atoms from the gas leave the oven in a high-velocity beam that travels through a vacuum tube toward a pair of magnets. The magnets serve as a gate that allows only atoms of a particular magnetic energy state to pass into a microwave cavity, where they are exposed to a microwave frequency derived from a quartz oscillator. If the microwave frequency matches the resonance frequency of cesium, the cesium atoms will change their magnetic energy state.

The atomic beam then passes through another magnetic gate near the end of the tube. Those atoms that changed their energy state while passing through the microwave cavity are allowed to proceed to a detector at the end of the tube. Atoms that did not change state are deflected away from the detector. The detector produces a feedback signal that continually tunes the quartz oscillator in a way that maximizes the number of state changes so that the greatest number of atoms reaches the detector. Standard output frequencies are derived from the locked quartz oscillator (Fig. 9).

The Q of a commercial cesium standard is a few parts in 10^8. The beam tube is typically <0.5 m in length, and the atoms travel at velocities of >100 m per second inside the tube. This limits the observation time to a few

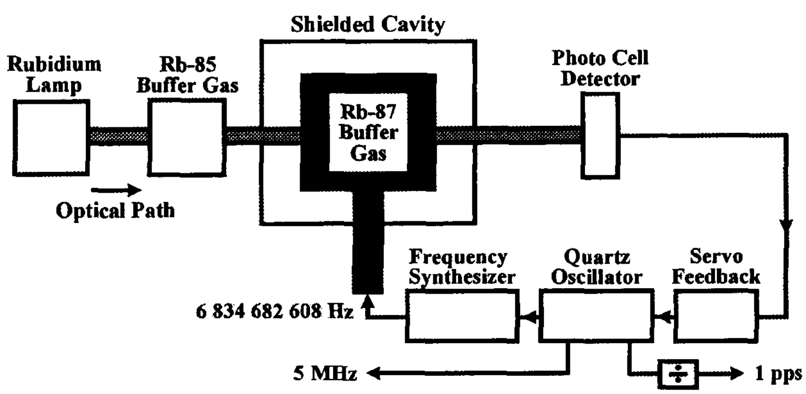

FIGURE 8 Rubidium oscillator.

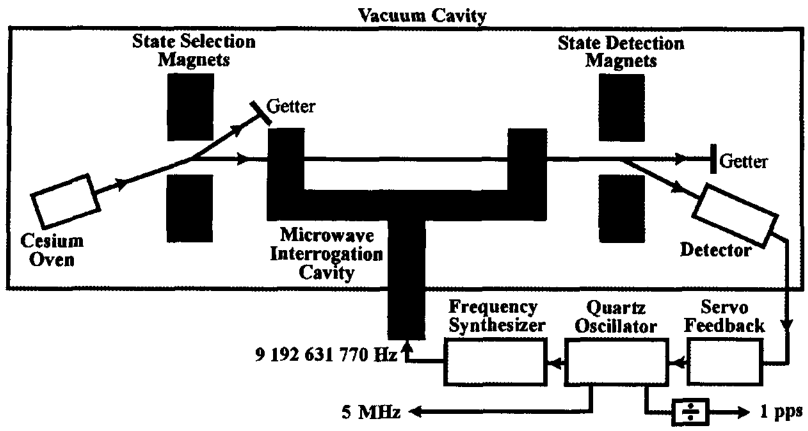

FIGURE 9 Cesium beam oscillator.

milliseconds, and the resonance width to a few hundred hertz. Stability ($\sigma_y\tau$, at $\tau = 1$ sec) is typically 5×10^{-12} and reaches a noise floor near 1×10^{-14} at about 1 day, extending out to weeks or months. The frequency offset is typically near 1×10^{-12} after a warm-up period of 30 min.

The current state-of-the-art in cesium technology is the *cesium fountain* oscillator, named after its fountain-like movement of cesium atoms. A cesium fountain named NIST-F1 serves as the primary standard of time and frequency for the United States.

A cesium fountain works by releasing a gas of cesium atoms into a vacuum chamber. Six infrared laser beams are directed at right angles to each other at the center of the chamber. The lasers gently push the cesium atoms together into a ball. In the process of creating this ball, the lasers slow down the movement of the atoms and cool them to temperatures a few thousandths of a degree above absolute zero. This reduces their thermal velocity to a few centimeters per second.

Two vertical lasers gently toss the ball upward and then all of the lasers are turned off. This little push is just enough to loft the ball about a meter high through a microwave-filled cavity. Under the influence of gravity, the ball then falls back down through the microwave cavity. The round trip up and down through the microwave cavity lasts for about 1 sec and is limited only by the force of gravity pulling the atoms to the ground. During the trip, the atomic states of the atoms might or might not be altered as they interact with the microwave signal. When their trip is finished, another laser is pointed at the atoms. Those atoms whose states were altered by the microwave signal emit photons (a state known as *fluorescence*) that are counted by a detector. This process is repeated many times while the microwave signal in the cavity is tuned to different frequencies. Eventually, a microwave frequency is found that alters the states of most of the cesium atoms and maximizes their fluorescence. This frequency is the cesium resonance (Fig. 10).

The Q of a cesium fountain is about 10^{10}, or about 100 times higher than a traditional cesium beam. Although the resonance frequency is the same, the resonance width is much narrower (<1 Hz), due to the longer observation times made possible by the combination of laser cooling and the fountain design. The combined frequency uncertainty of NIST-F1 is estimated at $<2 \times 10^{-15}$.

D. Hydrogen Masers

The *hydrogen maser* is the most elaborate and expensive commercially available frequency standard. The word *maser* is an acronym that stands for microwave amplification by stimulated emission of radiation. Masers operate at the resonance frequency of the hydrogen atom, which is 1,420,405,752 Hz.

A hydrogen maser works by sending hydrogen gas through a magnetic gate that allows only atoms in certain energy states to pass through. The atoms that make it through the gate enter a storage bulb surrounded by a tuned, resonant cavity. Once inside the bulb, some atoms drop to a lower energy level, releasing photons of microwave frequency. These photons stimulate other atoms to drop their energy level, and they in turn release additional photons. In this manner, a self-sustaining microwave field builds up in the bulb. The tuned cavity around the bulb helps to redirect photons back into the system to keep the oscillation going. The result is a microwave signal that is locked to the resonance frequency of the hydrogen atom and that is continually emitted as long as new atoms are fed into the system. This signal keeps a quartz crystal oscillator in step with the resonance frequency of hydrogen (Fig. 11).

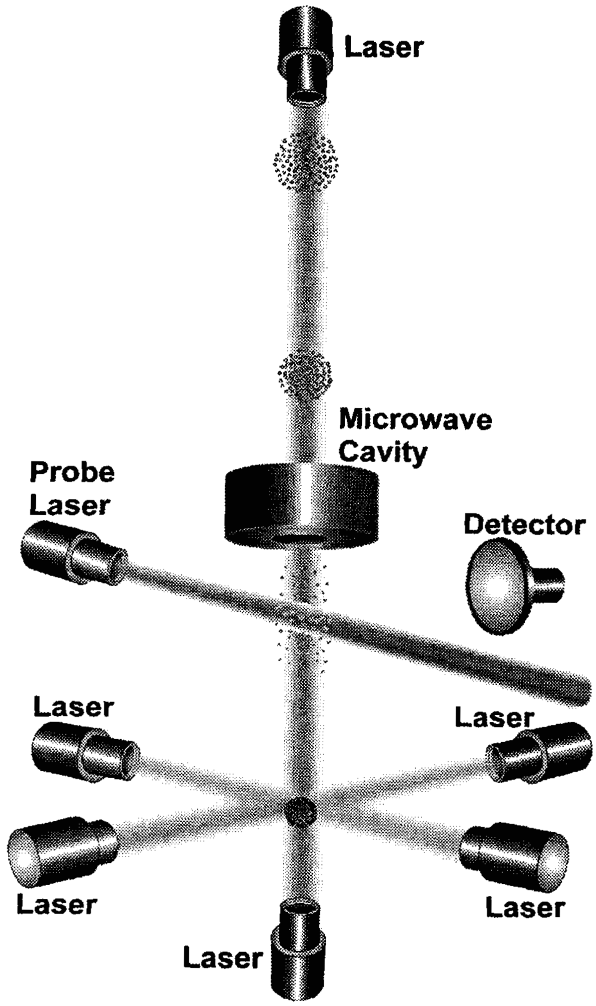

FIGURE 10 Cesium fountain oscillator.

The resonance frequency of hydrogen is much lower than that of cesium, but the resonance width of a hydrogen maser is usually just a few hertz. Therefore, the Q is about 10^9, or at least one order of magnitude better than a commercial cesium standard. As a result, the short-term stability is better than a cesium standard for periods out to a few days—typically $<1 \times 10^{-12}$ ($\sigma_y\tau$, at $\tau = 1$ sec) and reaching a noise floor of $\cong 1 \times 10^{-15}$ after about 1 hr. However, when measured for more than a few days or weeks, a hydrogen maser might fall below a cesium oscillator's performance. The stability decreases because of changes in the cavity's resonance frequency over time.

E. Future Standards

Research conducted at NIST and other laboratories should eventually lead to frequency standards that are far more stable than current devices. Future standards might use

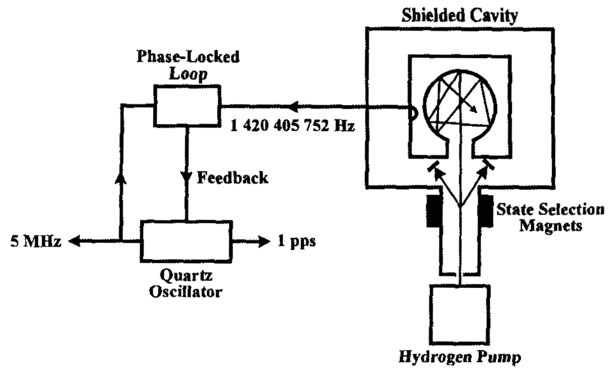

FIGURE 11 Hydrogen maser oscillator.

the resonance frequency of trapped, electrically charged ions. Trapping ions and suspending them in a vacuum allows them to be isolated from disturbing influences and observed for periods of 100 sec or longer. Much of this work has been based on the mercury ion ($^{199}Hg^+$), since its resonance frequency in the microwave realm is about 40.5 GHz, or higher than that of other atoms appropriate for this trapping technique. With a resonance width of 10 mHz or less, the Q of a mercury ion standard can reach 10^{12}.

The most promising application of trapped ions is their use in optical frequency standards. These devices use ion traps that resonate at optical, rather than microwave frequencies. The resonance frequency of these devices is about 10^{15} Hz; for example, the $^{199}Hg^+$ ion has an optical wavelength of just 282 nm. Although long observation times are difficult with this approach, experiments have shown that a resonance width of 1 Hz might eventually be possible. This means that the Q of an optical frequency standard could reach 10^{15}, several orders of magnitude higher than the best microwave experiments.

IV. TIME AND FREQUENCY TRANSFER

Many applications require clocks or oscillators at different locations to be set to the same time (*synchronization*) or the same frequency (*syntonization*). *Time and frequency transfer* techniques are used to compare and adjust clocks and oscillators at different locations. Time and frequency transfer can be as simple as setting your wristwatch to an audio time signal or as complex as controlling the frequency of oscillators in a network to parts in 10^{13}.

Time and frequency transfer can use signals broadcast through many different media, including coaxial cables, optical fiber, radio signals (at numerous places in the spectrum), telephone lines, and the Internet. Synchronization requires both an on-time pulse and a time code.

TABLE VII Summary of Time and Frequency Transfer Signals and Methods

Signal or link	Receiving equipment	Time uncertainty (24 hr)	Frequency uncertainty (24 hr)
Dial-up computer time service	Computer, software, modem, and phone line	<15 msec	NA
Network time service	Computer, software, and Internet connection	<1 sec	NA
HF radio (3 to 30 MHz)	HF receiver	1 to 20 msec	10^{-6} to 10^{-9}
LF radio (30 to 300 kHz)	LF receiver	1 to 100 μsec	10^{-10} to 10^{-12}
GPS one-way	GPS receiver	<50 nsec	$\cong 10^{-13}$
GPS common-view	GPS receiver, tracking schedule (single channel only), data link	<10 nsec	$<1 \times 10^{-13}$
GPS carrier phase	GPS carrier phase tracking receiver, orbital data for postprocessing corrections, data link	<50 nsec	$<1 \times 10^{-14}$
Two-way satellite	Receiving equipment, transmitting equipment, data link	<1 nsec	$<1 \times 10^{-14}$

Syntonization requires extracting a stable frequency from the broadcast, usually from the carrier frequency or time code.

This section discusses both the fundamentals of time and frequency transfer and the radio and network signals used. Table VII provides a summary.

A. Fundamentals of Time and Frequency Transfer

The largest contributor to time transfer uncertainty is *path delay*, or the signal delay between the transmitter and the receiver. For example, consider a radio signal broadcast over a 1000-km path. Since radio signals travel at the speed of light ($\cong$3.3 μsec/km), we can calibrate the path by estimating the path delay as 3.3 msec and applying a 3.3-msec correction to our measurement. The more sophisticated time transfer systems are self-calibrating and automatically correct for path delay.

Path delay is not important to frequency transfer systems, since on-time pulses are not required. Instead, frequency transfer requires only a stable path where the delays remain relatively constant. The three basic types of time and frequency transfer methods are described below.

1. One-Way Method

This is the simplest and most common way to transfer time and frequency information. Information is sent from a transmitter to a receiver and is delayed by the path through the medium (Fig. 12). To get the best results, the user must estimate τ_{ab} and calibrate the path to compensate for the delay. Of course, for many applications the path delay is simply ignored. For example, if our goal is simply to synchronize a computer clock within 1 sec of UTC, there is no need to worry about a 100-msec delay through a network.

More sophisticated one-way transfer systems estimate and remove all or part of the τ_{ab} delay. This is usually

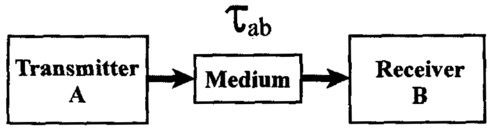

FIGURE 12 One-way transfer.

done in one of two ways. The first way is to estimate τ_{ab} and send the time out early by this amount. For example, if τ_{ab} is at least 20 msec for all users, the time can be sent 20 msec early. This advancement of the timing signal will remove at least some of the delay for all users.

A better technique is to compute τ_{ab} and to apply a correction to the broadcast. A correction for τ_{ab} can be computed if the position of both the transmitter and the receiver are known. If the transmitter is stationary, a constant can be used for the transmitter position. If the transmitter is moving (a satellite, for example), it must broadcast its position in addition to broadcasting time. The Global Positioning System provides the best of both worlds—each satellite broadcasts its position and the receiver can use coordinates from multiple satellites to compute its own position.

One-way time transfer systems often include a *time code* so that a clock can be set to the correct time-of-day. Most time codes contain the UTC hour, minute, and second. Some contain date information, a UT1 correction, and advance warning of daylight savings time and leap seconds.

2. Common-View Method

The common-view method involves a single reference transmitter (R) and two receivers (A and B). The transmitter is in common view of both receivers. Both receivers compare the simultaneously received signal to their local clock and record the data. Receiver A receives the signal over the path τ_{ra} and compares the reference to its local clock (R − Clock A). Receiver B receives the signal over

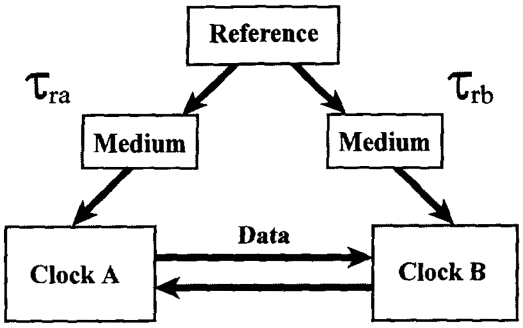

FIGURE 13 Common-view transfer.

the path τ_{rb} and records (R − Clock B). The two receivers then exchange and difference the data (Fig. 13).

Common-view directly compares two time and frequency standards. Errors from the two paths (τ_{ra} and τ_{rb}) that are common to the reference cancel out, and the uncertainty caused by path delay is nearly eliminated. The result of the measurement is (Clock A − Clock B) − ($\tau_{ra} - \tau_{rb}$).

3. Two-Way Method

The two-way method requires two users to both transmit and receive through the same medium at the same time. Sites A and B simultaneously exchange time signals through the same medium and compare the received signals with their own clocks. Site A records A − (B + τ_{ba}) and site B records B − (A + τ_{ab}), where τ_{ba} is the path delay from A to B, and τ_{ab} is the path delay from A to B. The difference between these two sets of readings produces 2(A − B) − ($\tau_{ba} - \tau_{ab}$). Since the path is reciprocal ($\tau_{ab} = \tau_{ba}$), the path delay cancels out of the equation (Fig. 14).

The two-way method is used for international comparisons of time standards using spread spectrum radio signals at C- or Ku-band frequencies, and a geostationary satellite as a transponder. The stability of these comparisons is usually <500 psec ($\sigma_x\tau$, at $\tau = 1$ sec), or $<1 \times 10^{-14}$ for frequency, even when the clocks are separated by thousands of kilometers.

The two-way method is also used in telecommunications networks where transmission of a signal can be done in software. Some network and telephone time signals use a variation of two-way, called the *loop-back* method. Like

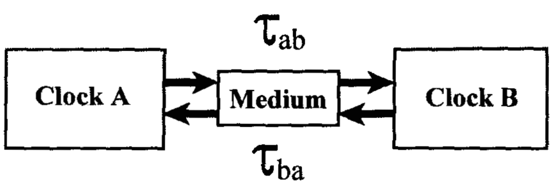

FIGURE 14 Two-way transfer.

the two-way method, the loop-back method requires both users to transmit and receive, but not at the same time. For example, a signal is sent from the transmitter (A) to the receiver (B) over the path τ_{ab}. The receiver (B) then echoes or reflects the signal back to the transmitter (A) over the path τ_{ba}. The transmitter then adds the two path delays ($\tau_{ab} + \tau_{ba}$) to obtain the round-trip delay and divides this number by 2 to estimate the one-way path delay. The transmitter then advances the next time signal by the estimated one-way delay. Since users do not transmit and receive at the same time, the loop-back method has larger uncertainties than the two-way method. A reciprocal path cannot be assumed, since we do not know if the signal from A to B traveled the same path as the signal from B to A.

B. Radio Time and Frequency Transfer Signals

There are many types of radio receivers designed to receive time and frequency information. Radio clocks come in several different forms. Some are tabletop or rack-mount devices with a digital time display and a computer interface. Others are available as cards that plug directly into a computer.

The uncertainty of a radio time transfer system consists of the uncertainty of the received signal, plus delays in the receiving equipment. For example, there is cable delay between the antenna and the receiver. There are equipment delays introduced by hardware, and processing delays introduced by software. These delays must be calibrated to get the best results. When doing frequency transfer, equipment delays can be ignored if they remain relatively constant.

The following sections look at the three types of radio signals most commonly used for time and frequency transfer—high frequency (HF), low frequency (LF), and Global Positioning System (GPS) satellite signals.

1. HF Radio Signals (Including WWV and WWVH)

High-frequency (HF) radio broadcasts occupy the radio spectrum from 3 to 30 MHz. These signals are commonly used for time and frequency transfer at moderate performance levels. Some HF broadcasts provide audio time announcements and digital time codes. Other broadcasts simply provide a carrier frequency for use as a reference.

HF time and frequency stations (Table VIII) include NIST radio stations WWV and WWVH. WWV is located near Fort Collins, Colorado, and WWVH is on the island of Kauai, Hawaii. Both stations broadcast continuous time and frequency signals on 2.5, 5, 10, and 15 MHz, and WWV also broadcasts on 20 MHz. All frequencies carry the same program, and at least one frequency should be usable at all times. The stations can also be heard by

TABLE VIII HF Time and Frequency Broadcast Stations

Call sign	Country	Frequency(ies) (MHz)	Always on?	Language
ATA	India	10	No	English
BPM	China	2.5, 5, 10, 15	No	Chinese
BSF	Taiwan	5, 15	Yes	No voice
CHU	Canada	3.33, 7.335, 14.670	Yes	English/French
DUW21	Philippines	3.65	No	No voice
EBC	Spain	4.998, 15.006	No	No voice
HD2IOA	Ecuador	1.51, 3.81, 5, 7.6	No	Spanish
HLA	Korea	5	No	Korean
LOL1	Argentina	5, 10, 15	No	Spanish
LQB9	Argentina	8.167	No	No voice
LQC28	Argentina	17.551	No	No voice
PLC	Indonesia	11.440	No	No voice
PPEI	Brazil	8.721	No	No voice
PPR	Brazil	4.244, 8.634, 13.105, 17.194	No	No voice
RID	Russia	5.004, 10.004, 15.004	Yes	No voice
RTA	Russia	10, 15	No	No voice
RWM	Russia	4.996, 9.996, 14.996	Yes	No voice
ULW4	Uzbekistan	2.5, 5, 10	No	No voice
VNG	Australia	2.5, 5, 8.638, 12.984, 16	Yes	English
WWV	United States	2.5, 5, 10, 15, 20	Yes	English
WWVH	United States	2.5, 5, 10, 15	Yes	English
XBA	Mexico	6.976, 13.953	No	No voice
XDD	Mexico	13.043	No	No voice
XDP	Mexico	4.8	No	No voice
YVTO	Venezuela	5	Yes	Spanish

telephone; dial (303) 499-7111 for WWV and (808) 335-4363 for WWVH.

WWV and WWVH can be used in one of three modes.

- The audio portion of the broadcast includes seconds pulses or ticks, standard audio frequencies, and voice announcements of the UTC hour and minute. WWV uses a male voice, and WWVH uses a female voice.
- A binary time code is sent on a 100-Hz subcarrier at a rate of 1 bit per second. The time code contains the hour, minute, second, year, day of year, leap second, and Daylight Saving Time (DST) indicators and a UT1 correction. This code can be read and displayed by radio clocks.
- The carrier frequency can be used as a reference for the calibration of oscillators. This is done most often with the 5- and 10-MHz carrier signals, since they match the output frequencies of standard oscillators.

The time broadcast by WWV and WWVH will be late when it arrives at the user's location. The time offset depends upon the receiver's distance from the transmitter but should be <15 msec in the continental United States. A good estimate of the time offset requires knowledge of HF radio propagation. Most users receive a signal that has traveled up to the ionosphere and was reflected back to earth. Since the height of the ionosphere changes, the path delay also changes. Path delay variations limit the received frequency uncertainty to parts in 10^9 when averaged for 1 day.

HF radio stations such as WWV and WWVH are useful for low-level applications, such as the synchronization of analog and digital clocks, simple frequency calibrations, and calibrations of stopwatches and timers. However, LF and satellite signals are better choices for more demanding applications.

2. LF Radio Signals (Including WWVB)

Before the advent of satellites, low-frequency (LF) signals were the method of choice for time and frequency transfer. While the use of LF signals has diminished in the laboratory, they still have a major advantage—they can be received indoors without an external antenna. This makes them ideal for many consumer electronic products that display time-of-day information.

Many time and frequency stations operate in the LF band from 30 to 300 kHz (Table IX). These stations lack

TABLE IX LF Time and Frequency Broadcast Stations

Call sign	Country	Frequency (kHz)	Always on?
DCF77	Germany	77.5	Yes
DGI	Germany	177	Yes
HBG	Switzerland	75	Yes
JG2AS	Japan	40	Yes
MSF	United Kingdom	60	Yes
RBU	Russia	66.666	No
RTZ	Russia	50	Yes
TDF	France	162	Yes
WWVB	United States	60	Yes

the bandwidth needed to provide voice announcements, but they often provide both an on-time pulse and a time code. The performance of the received signal is influenced by the path length and signal strength. Path length is important because the signal is divided into ground wave and sky wave. The ground wave signal is more stable. Since it travels the shortest path between the transmitter and the receiver, it arrives first and its path delay is much easier to estimate. The sky wave is reflected from the ionosphere and produces results similar to HF reception. Short paths make it possible to track the ground wave continuously. Longer paths produce a mixture of sky wave and ground wave. And over very long paths, only sky wave reception is possible.

Signal strength is also important. If the signal is weak, the receiver might search for a new cycle of the carrier to track. Each time the receiver adjusts its tracking point by one cycle, it introduces a phase step equal to the period of a carrier. For example, a cycle slip on a 60-kHz carrier introduces a 16.67-μsec phase step. However, a strong ground wave signal can produce very good results—a LF receiver that continuously tracks the same cycle of a ground wave signal can transfer frequency with an uncertainty of about 1×10^{-12} when averaged for 1 day.

NIST operates LF radio station WWVB from Fort Collins, Colorado, at a transmission frequency of 60 kHz. The station broadcasts 24 hr per day, with an effective radiated output power of 50 kW. The WWVB time code is synchronized with the 60-kHz carrier and contains the year, day of year, hour, minute, second, and flags that indicate the status of DST, leap years, and leap seconds. The time code is received and displayed by wristwatches, alarm clocks, wall clocks, and other consumer electronic products.

3. Global Positioning System (GPS)

The Global Positioning System (GPS) is a navigation system developed and operated by the U.S. Department of Defense (DoD) that is usable nearly anywhere on earth. The system consists of a constellation of at least 24 satellites that orbit the earth at a height of 20,200 km in six fixed planes inclined 55° from the equator. The orbital period is 11 hr 58 min, which means that a satellite will pass over the same place on earth twice per day. By processing signals received from the satellites, a GPS receiver can determine its position with an uncertainty of <10 m.

The satellites broadcast on two carrier frequencies: L1 at 1575.42 MHz and L2 at 1227.6 MHz. Each satellite broadcasts a spread spectrum waveform, called a *pseudorandom noise* (PRN) code, on L1 and L2, and each satellite is identified by the PRN code it transmits. There are two types of PRN codes. The first type is a *coarse acquisition* (C/A) code, with a chip rate of 1023 chips per millisecond. The second is a *precision* (P) code, with a chip rate of 10,230 chips per millisecond. The C/A code is broadcast on L1, and the P code is broadcast on both L1 and L2. GPS reception is line-of-sight, which means that the antenna must have a clear view of the sky.

Each satellite carries either rubidium or cesium oscillators, or a combination of both. These oscillators are steered from DoD ground stations and are referenced to the United States Naval Observatory time scale, UTC(USNO), which by agreement is always within 100 nsec of UTC(NIST). The oscillators provide the reference for both the carrier and the code broadcasts.

a. GPS one-way measurements. GPS one-way measurements provide exceptional results with only a small amount of effort. A GPS receiver can automatically compute its latitude, longitude, and altitude using position data received from the satellites. The receiver can then calibrate the radio path and synchronize its on-time pulse. In addition to the on-time pulse, many receivers provide standard frequencies such as 5 or 10 MHz by steering an OCXO or rubidium oscillator using the satellite signals. GPS receivers also produce time-of-day and date information.

A quality GPS receiver calibrated for equipment delays has a timing uncertainty of about 10 nsec relative to UTC(NIST) and a frequency uncertainty of about 1×10^{-13} when averaged for 1 day.

b. GPS common-view measurements. The *common-view* method synchronizes or compares time standards or time scales at two or more locations. Common-view GPS is the primary method used by the BIPM to collect data from laboratories that contribute to TAI.

There are two types of GPS common-view measurements. *Single-channel common-view* requires a specially designed GPS receiver that can read a tracking schedule. This schedule tells the receiver when to start making measurements and which satellite to track. Another user at another location uses the same schedule and makes simultaneous measurements from the same satellite. The tracking schedule must be designed so that it chooses satellites visible to both users at reasonable elevation angles. *Multichannel common-view* does not use a schedule. The receiver simply records timing measurements from all satellites in view. In both cases, the individual measurements at each site are estimates of (Clock A − GPS) and (Clock B − GPS). If the data are exchanged, and the results are subtracted, the GPS clock drops out and an estimate of Clock A − Clock B remains. This technique allows time and frequency standards to be compared directly even when separated by thousands of kilometers. When averaged for 1 day, the timing uncertainty of GPS common-view is <5 nsec, and the frequency uncertainty is $<1 \times 10^{-13}$.

c. GPS carrier phase measurements. Used primarily for frequency transfer, this technique uses the GPS carrier frequency (1575.42 MHz) instead of the codes transmitted by the satellites. Carrier phase measurements can be one-way or common-view. Since the carrier frequency is more than 1000 times higher than the C/A code frequency, the potential resolution is much higher. However, taking advantage of the increased resolution requires making corrections to the measurements using orbital data and models of the ionosphere and troposphere. It also requires correcting for cycle slips that introduce phase shifts equal to multiples of the carrier period (≅635 psec for <1). Once the measurements are properly processed, the frequency uncertainty of common-view carrier phase measurements is $<1 \times 10^{-14}$ when averaged for 1 day.

C. Internet and Telephone Time Signals

One common use of time transfer is to synchronize computer clocks to the correct date and time-of-day. This is usually done with a time code received through an Internet or telephone connection.

1. Internet Time Signals

Internet time servers use standard timing protocols defined in a series of RFC (Request for Comments) documents. The three most common protocols are the Time Protocol, the Daytime Protocol, and the Network Time Protocol (NTP). An Internet time server waits for timing requests sent using any of these protocols and sends a time code in the correct format when a request is received.

Client software is available for all major operating systems, and most client software is compatible with either the Daytime Protocol or the NTP. Client software that uses the Simple Network Time Protocol (SNTP) makes the same timing request as an NTP client but does less processing and provides less accuracy. Table X summarizes the various protocols and their port assignments, or the port where the server "listens" for a client request.

NIST operates an Internet time service using multiple servers distributed around the United States. A list of IP addresses for the NIST servers and sample client software can be obtained from the NIST Time and Frequency Division web site: http://www.boulder.nist.gov/timefreq. The uncertainty of Internet time signals is usually <100 msec, but results vary with different computers, operating systems, and client software.

2. Telephone Time Signals

Telephone time services allow computers with analog modems to synchronize their clocks using ordinary telephone lines. These services are useful for synchronizing computers that are not on the Internet or that reside behind an Internet firewall. One example of a telephone service is NISTs Automated Computer Time Service (ACTS), (303) 494-4774.

TABLE X Internet Time Protocols

Protocol name	Document	Format	Port assignment(s)
Time protocol	RFC-868	Unformatted 32-bit binary number contains time in UTC seconds since January 1, 1900	Port 37, tcp/ip, udp/ip
Daytime protocol	RFC-867	Exact format not specified in standard. Only requirement is that the time code is sent as ASCII characters	Port 13, tcp/ip, udp/ip
Network time protocol (NTP)	RFC-1305	The server provides a data packet with a 64-bit time stamp containing the time in UTC seconds since January 1, 1900, with a resolution of 200 psec. NTP provides an accuracy of 1 to 50 msec. The client software runs continuously and gets periodic updates from the server.	Port 123, udp/ip
Simple network time protocol (SNTP)	RFC-1769	The data packet sent by the server is the same as NTP, but the client software does less processing and provides less accuracy.	Port 123, udp/ip

ACTS requires a computer, a modem, and client software. When a computer connects to ACTS it receives a time code containing the month, day, year, hour, minute, second, leap second, and DST indicators and a UT1 correction. The last character in the ACTS time code is the on-time marker (OTM). To compensate for the path delay between NIST and the user, the server sends the OTM 45 msec early. If the client returns the OTM, the server can calibrate the path using the *loop-back* method. Each time the OTM is returned, the server measures the round-trip path delay and divides this quantity by 2 to estimate the one-way path delay. This path calibration reduces the uncertainty to <15 msec.

V. CLOSING

As noted earlier, time and frequency standards and measurements have improved by about nine orders of magnitude in the past 100 years. This rapid advancement has made many new products and technologies possible. While it is impossible to predict what the future holds, we can be certain that oscillator Q's will continue to get higher, measurement uncertainties will continue to get lower, and new technologies will continue to emerge.

SEE ALSO THE FOLLOWING ARTICLES

MICROWAVE COMMUNICATIONS • QUANTUM MECHANICS • RADIO SPECTRUM UTILIZATION • REAL-TIME SYSTEMS • SIGNAL PROCESSING • TELECOMMUNICATIONS

BIBLIOGRAPHY

Allan, D. W., Ashby, N., and Hodge, C. C. (1997). "The Science of Timekeeping," Hewlett–Packard Application Note 1289, United States.

Hackman, C., and Sullivan, D. B. (eds.) (1996). "Time and Frequency Measurement," American Association of Physics Teachers, College Park, MD.

IEEE Standards Coordinating Committee 27 (1999). "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities," Institute of Electrical and Electronics Engineers, New York.

ITU Radiocommunication Study Group 7 (1997). "Selection and Use of Precise Frequency and Time Systems," International Telecommunications Union, Geneva, Switzerland.

Jespersen, J., and Fitz-Randolph, J. (1999). "From Sundials to Atomic Clocks: Understanding Time and Frequency," NIST Monograph 155, U.S. Government Printing Office, Washington, DC.

Kamas, G., and Lombardi, M. A. (1990). "Time and Frequency Users Manual," NIST Special Publication 559, U.S. Government Printing Office, Washington, DC.

Levine, J. (1999). "Introduction to time and frequency metrology," *Rev. Sci. Instrum.* **70,** 2567–2596.

Seidelmann, P. K. (ed.) (1992). "Explanatory Supplement to the Astronomical Almanac," University Science Books, Mill Valley, CA.

Sullivan, D. B., Allan, D. W., Howe, D. A., and Walls, F. L. (eds.) (1990). "Characterization of Clocks and Oscillators," NIST Technical Note 1337, U.S. Government Printing Office, Washington, DC.

Vig, J. R. (1992). "Introduction to Quartz Frequency Standards," United States Army Research and Development Technical Report SLCET-TR-92-1.

Walls, F. L., and Ferre-Pikal, E. S. (1999). Frequency standards, characterization. Measurement of frequency, phase noise, and amplitude noise. *In* "Wiley Encyclopedia of Electrical and Electronics Engineering," Vol. 7, pp. 767–774, Vol. 12, pp. 459–473, John Wiley and Sons, New York.