17

Fundamentals of Time and Frequency

17.1 Introduction
Coordinated Universal Time (UTC)
17.2 Time and Frequency Measurement
Accuracy • Stability
17.3 Time and Frequency Standards
Quartz Oscillators • Rubidium Oscillators • Cesium Oscillators
17.4 Time and Frequency Transfer
Fundamentals of Time and Frequency Transfer • Radio Time and Frequency Transfer Signals
17.5 Closing

Michael A. Lombardi
National Institute of Standards and Technology

17.1 Introduction

Time and frequency standards supply three basic types of information: *time-of-day, time interval*, and *frequency*. Time-of-day information is provided in hours, minutes, and seconds, but often also includes the *date* (month, day, and year). A device that displays or records time-of-day information is called a *clock*. If a clock is used to label when an event happened, this label is sometimes called a *time tag* or *time stamp*. Date and time-of-day can also be used to ensure that events are *synchronized*, or happen at the same time.

Time interval is the duration or elapsed time between two events. The standard unit of time interval is the second(s). However, many engineering applications require the measurement of shorter time intervals, such as milliseconds (1 ms = 10^{-3} s), microseconds (1 μs = 10^{-6} s), nanoseconds (1 ns = 10^{-9} s), and picoseconds (1 ps = 10^{-12} s). Time is one of the seven base physical quantities, and the second is one of seven base units defined in the International System of Units (SI). The definitions of many other physical quantities rely upon the definition of the second. The second was once defined based on the earth's rotational rate or as a fraction of the tropical year. That changed in 1967 when the era of atomic time keeping formally began. The current definition of the SI second is:

> 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.

Frequency is the rate of a repetitive event. If T is the period of a repetitive event, then the frequency f is its reciprocal, $1/T$. Conversely, the period is the reciprocal of the frequency, $T = 1/f$. Since the period is a time interval expressed in seconds (s), it is easy to see the close relationship between time interval and frequency. The standard unit for frequency is the hertz (Hz), defined as events or cycles per second. The frequency of electrical signals is often measured in multiples of hertz, including kilohertz (kHz), megahertz (MHz), or gigahertz (GHz), where 1 kHz equals one thousand (10^3) events per second, 1 MHz

TABLE 17.1 Uncertainties of Physical Realizations of the Base SI Units

SI Base Unit	Physical Quantity	Uncertainty
Candela	Luminous intensity	1×10^{-4}
Kelvin	Temperature	3×10^{-7}
Mole	Amount of substance	8×10^{-8}
Ampere	Electric current	4×10^{-8}
Kilogram	Mass	1×10^{-8}
Meter	Length	1×10^{-12}
Second	Time interval	1×10^{-15}

equals one million (10^6) events per second, and 1 GHz equals one billion (10^9) events per second. A device that produces frequency is called an *oscillator*. 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 events or cycles per second, we can measure frequency.

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, as shown in Table 17.1 [1–5].

Coordinated Universal Time (UTC)

The world's major metrology laboratories routinely measure their time and frequency standards and send the measurement data to 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). As a result of this averaging, the BIPM generates two time scales, International Atomic Time (TAI), and Coordinated Universal Time (UTC). These time scales realize the SI second as closely as possible.

UTC runs at the same frequency as TAI. However, it differs from TAI by an integral number of seconds. This difference increases when *leap 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 s of an older time scale called UT1, which is based on the rotational rate of the earth. Leap seconds have been added to UTC at a rate of slightly less than once per year, beginning in 1972 [3,5].

Keep in mind that the BIPM maintains TAI and UTC as "paper" time scales. The major metrology laboratories use the published data from the BIPM to steer their clocks and oscillators and generate real-time versions of UTC. Many of these laboratories distribute their versions of UTC via radio signals, which are discussed in section 17.4.

You can think of UTC as the ultimate standard for time-of-day, time interval, and frequency. Clocks synchronized to UTC display the same hour, minute, and second all over the world (and remain within one second of UT1). Oscillators syntonized to UTC generate signals that serve as reference standards for time interval and frequency.

17.2 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*). A measurement compares the DUT to a *standard* or *reference*. The standard should outperform the DUT by a specified ratio, called the *test uncertainty ratio* (*TUR*). Ideally, the TUR should be 10:1 or higher. The higher the ratio, the less averaging is required to get valid measurement results.

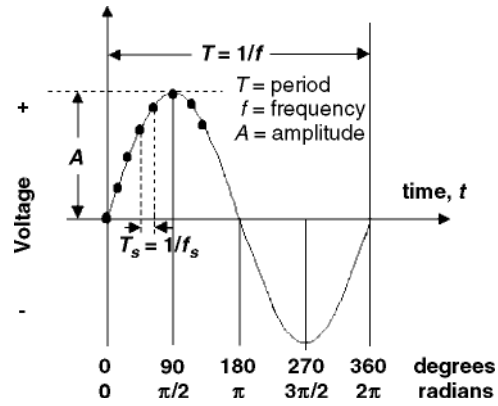

FIGURE 17.1 An oscillating sine wave.

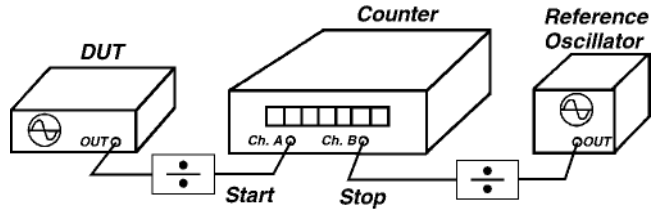

FIGURE 17.2 Measurement using a time interval counter.

The test signal for time measurements is usually a pulse that occurs once per second (1 pps). The pulse width and polarity varies from device to device, but TTL levels are commonly used. The test signal for frequency measurements is usually at a frequency of 1 MHz or higher, with 5 or 10 MHz being common. Frequency signals are usually sine waves, but can also be pulses or square waves. If the frequency signal is an oscillating sine wave, it might look like the one shown in Fig. 17.1. This signal produces one cycle (360∞ or 2π radians of phase) in one period. The signal amplitude is expressed in volts, and must be compatible with the measuring instrument. If the amplitude is too small, it might not be able to drive the measuring instrument. If the amplitude is too large, the signal must be attenuated to prevent overdriving the measuring instrument.

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.

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. 17.2. 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 stop signals is measured by counting cycles from the time base oscillator. The resolution of a low cost TIC is limited to the period of its time base. For example, a TIC with a 10-MHz time base oscillator would have a resolution of 100 ns. More elaborate TICs use interpolation schemes to detect parts of a time base cycle and have much higher resolution—1 ns resolution is commonplace, and 20 ps resolution is available.

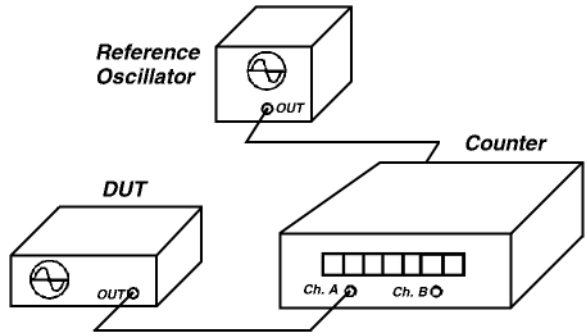

FIGURE 17.3 Measurement using a frequency counter.

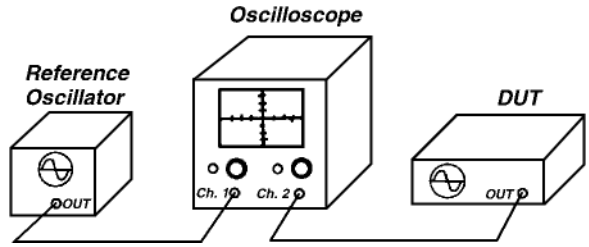

FIGURE 17.4 Phase comparison using an oscilloscope.

Frequency offset can be measured in either the *frequency domain* or *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. 17.3). The counter's resolution, or the number of digits it can display, limits its ability to measure frequency offset. For example, a 9-digit frequency counter can detect a frequency offset no smaller than 0.1 Hz at 10 MHz (1×10^{-8}). The frequency offset is determined as

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

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

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. 17.4). The oscilloscope will display two sine waves (Fig. 17.5). 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 measuring 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. In this case the phase difference is increasing, indicating that the DUT is lower in frequency than the reference.

Measuring high accuracy signals with an oscilloscope is impractical, since the phase relationship between signals changes very slowly and the resolution of the oscilloscope display is limited. More precise phase comparisons can be made with a TIC, using a setup similar to Fig. 17.2. If the two input signals have the same frequency, the time interval will not change. If the two signals have different frequencies,

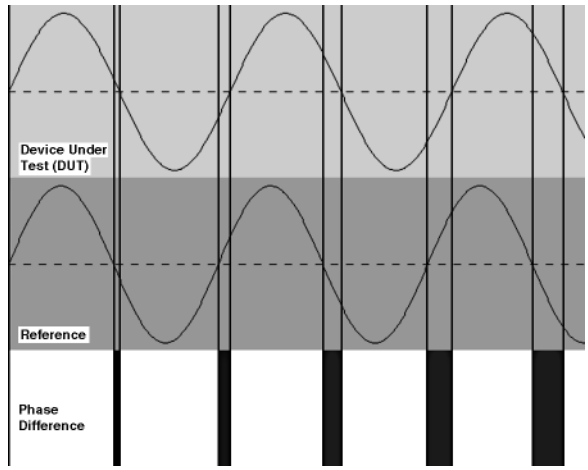

FIGURE 17.5 Two sine waves with a changing phase relationship.

the time interval will change, and the rate of change is the frequency offset. The resolution of a TIC determines the smallest frequency change that it can detect without averaging. For example, a low cost TIC with a single-shot resolution of 100 ns can detect frequency changes of 1×10^{-7} in 1 s. The current limit for TIC resolution is about 20 ps, which means that a frequency change of 2×10^{-11} can be detected in 1 s. Averaging over longer intervals can improve the resolution to <1 ps in some units [6].

Since standard frequencies like 5 or 10 MHz are not practical to measure with a TIC, *frequency dividers* (shown in Fig. 17.2) or *frequency mixers* are used to convert the test frequency to a lower frequency. Divider systems are simpler and more versatile, since they can be easily built or programmed to accommodate different frequencies. Mixer systems are more expensive, require more hardware including an additional reference oscillator, and can often measure only one input frequency (e.g., 10 MHz), but they have a higher signal-to-noise ratio than divider systems.

If dividers are used, 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:

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

where Δt is the amount of phase deviation, and T is the measurement period.

To illustrate, consider a measurement of +1 μs of phase deviation over a measurement period of 24 h. The unit used for measurement period (h) must be converted to the unit used for phase deviation (μs). The equation becomes

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

As shown, a device that accumulates 1 μs of phase deviation/day has a frequency offset of -1.16×10^{-11} with respect to the reference. This simple example requires only two time interval readings to be made, and Δt is simply the difference between the two readings. Often, multiple readings are taken and the frequency offset is estimated by using least squares linear regression on the data set, and obtaining Δt from the slope of the least squares line. This information is usually presented as a phase plot, as shown in Fig. 17.6. The device under test is high in frequency by exactly 1×10^{-9}, as indicated by a phase deviation of 1 ns/s [2,7,8].

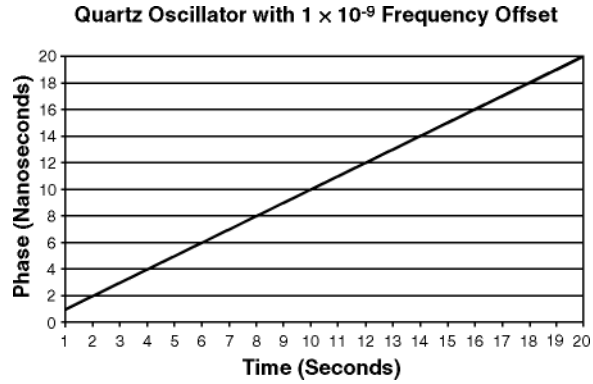

FIGURE 17.6 A sample phase plot.

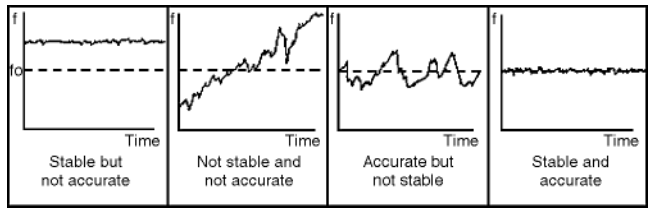

FIGURE 17.7 The relationship between accuracy and stability.

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:

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

Stability

Stability indicates how well an oscillator can produce the same time or frequency offset over a given time interval. It doesn't 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 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 17.7 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 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 s. *Long-term* stability can refer to measurement intervals greater than 100 s, but usually refers to periods longer than 1 day.

Stability estimates can be made in either the frequency domain or time domain, and can be calculated from a set of either frequency offset or time interval measurements. In some fields of measurement, stability is estimated by taking the standard deviation of the data set. However, standard deviation only

works with stationary data, where the results are time independent, and the noise is *white*, meaning that it is evenly distributed across the frequency band of the measurement. Oscillator data is usually nonstationary, since it contains time dependent noise contributed by the frequency offset. With stationary data, the mean and standard deviation will converge to particular values as more measurements are made. With nonstationary data, the mean and standard deviation never converge to any particular values. Instead, there is a moving mean that changes each time we add a measurement.

For these reasons, a non-classical 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 ($\sigma_y(\tau)$) 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 containing y_1, y_2, y_3, and so on, 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 containing 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. Note that while standard deviation subtracts the mean from each measurement before squaring their summation, the Allan deviation subtracts the previous data point. This differencing of successive data points removes the time dependent noise contributed by the frequency offset.

An Allan deviation graph is shown in Fig. 17.8. 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 such as flicker noise or random walk. The device measured in Fig. 17.8 has a noise floor of ~5×10^{-11} at τ = 100 s.

Practically speaking, a frequency stability graph also tells us how long we need to average to get rid of the noise contributed by the reference and the measurement system. The noise floor provides some indication of the amount of averaging required to obtain a TUR high enough to show us the true frequency

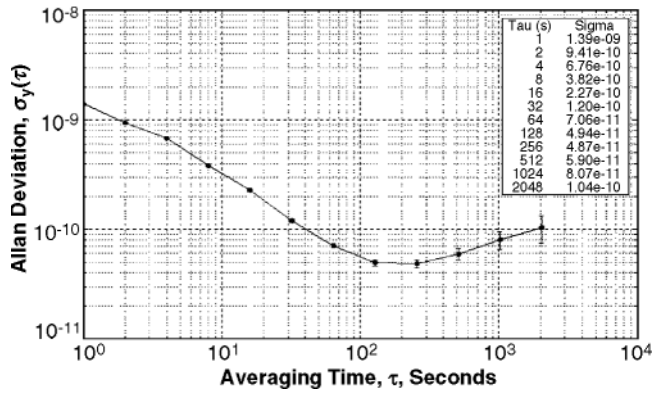

FIGURE 17.8 A frequency stability graph.

TABLE 17.2 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.

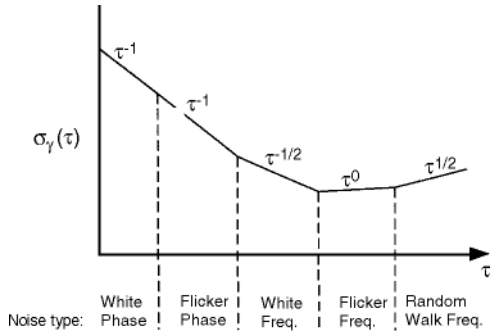

FIGURE 17.9 Using a frequency stability graph to identify noise types.

offset of the DUT. If the DUT is an atomic oscillator (section 17.4) and the reference is a radio controlled transfer standard (section 17.5) we might have to average for 24 h or longer to have confidence in the measurement result.

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 help identify the amount of averaging needed to remove these noise types (Fig. 17.9). The first type of noise to be removed by averaging is phase noise, or the rapid, random fluctuations in the phase of the signal. Ideally, only the device under test would contribute phase noise to the measurement, but in practice, some phase noise from the measurement system and reference needs to be removed through averaging. Note that the Allan deviation does not distinguish between white phase noise and flicker phase noise. Table 17.2 shows several other statistics used to estimate stability and identify noise types for various applications.

Identifying and eliminating sources of oscillator noise can be a complex subject, but plotting the first order differences of a set of time domain measurements can provide a basic understanding of how noise is removed by averaging. Figure 17.10 was made using a segment of the data from the stability graph in Fig. 17.8. It shows phase plots dominated by white phase noise (1 s averaging), white frequency noise (64 s averages), flicker frequency noise (256 s averages), and random walk frequency (1024 s averages). Note that the white phase noise plot has a 2 ns scale, and the other plots use a 100 ps scale [8–12].

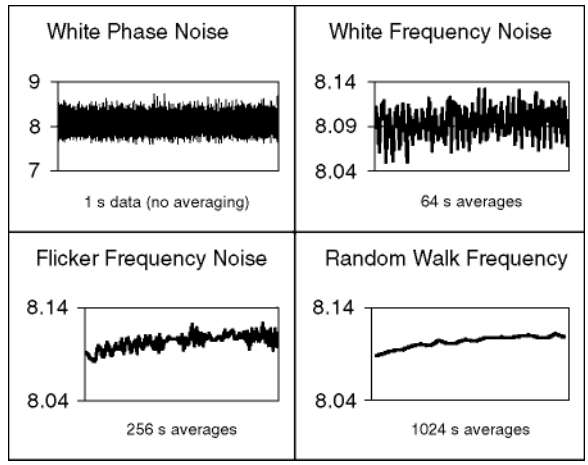

FIGURE 17.10 Phase plots of four noise types.

17.3 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 s. 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*. Table 17.3 shows how the frequency uncertainty of a clock's resonator corresponds to the timing uncertainty of a clock.

Throughout history, clock designers have searched for more stable resonators, and the evolution of time and frequency standards is summarized in Table 17.4. The uncertainties listed for modern standards represent current (year 2001) 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 nine orders of magnitude in the past 100 years.

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 frequencies where the oscillator will oscillate. 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*. 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 begins by discussing 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 and cesium. Atomic oscillators use the quantized energy levels in atoms and molecules as the source of their resonance. 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

Table 17.3 Relationship of Frequency Uncertainty to Time Uncertainty

Frequency Uncertainty	Measurement Period	Time Uncertainty
$\pm 1.00 \times 10^{-3}$	1 s	± 1 ms
$\pm 1.00 \times 10^{-6}$	1 s	± 1 μs
$\pm 1.00 \times 10^{-9}$	1 s	± 1 ns
$\pm 2.78 \times 10^{-7}$	1 h	± 1 ms
$\pm 2.78 \times 10^{-10}$	1 h	± 1 μs
$\pm 2.78 \times 10^{-13}$	1 h	± 1 ns
$\pm 1.16 \times 10^{-8}$	1 day	± 1 ms
$\pm 1.16 \times 10^{-11}$	1 day	± 1 μs
$\pm 1.16 \times 10^{-14}$	1 day	± 1 ns

TABLE 17.4 The Evolution of Time and Frequency Standards

Standard	Resonator	Date of Origin	Timing Uncertainty (24 h)	Frequency Uncertainty (24 h)
Sundial	Apparent motion of the sun	3500 B.C.	NA	NA
Verge escapement	Verge and foliet mechanism	14th century	15 min	1×10^{-2}
Pendulum	Pendulum	1656	10 s	1×10^{-4}
Harrison chronometer (H4)	Spring and balance wheel	1759	350 ms	4×10^{-6}
Shortt pendulum	Two pendulums, slave and master	1921	10 ms	1×10^{-7}
Quartz crystal	Quartz crystal	1927	10 μs	1×10^{-10}
Rubidium gas cell	^{87}Rb resonance (6,834,682,608 Hz)	1958	100 ns	1×10^{-12}
Cesium beam	^{133}Cs resonance (9,192,631,770 Hz)	1952	1 ns	1×10^{-14}
Hydrogen maser	Hydrogen resonance (1,420,405,752 Hz)	1960	1 ns	1×10^{-14}
Cesium fountain	^{133}Cs resonance (9,192,631,770 Hz)	1991	100 ps	1×10^{-15}

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 exactly the same frequency when they absorb or release energy. In theory, the atom is a perfect "pendulum" whose oscillations are counted to measure time interval. The discussion of atomic oscillators is limited to devices that are commercially available, and excludes the primary and experimental standards found in laboratories such as NIST. Table 17.5 provides a summary [1,4,8].

Quartz Oscillators

Quartz crystal oscillators are by far the most common time and frequency standards. An estimated two 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.

TABLE 17.5 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 (life expectancy >15 years)	Cesium beam tube (life expectancy of 3 to 25 years)	Hydrogen depletion (life expectancy >7 years)
Stability, $\sigma_y(\tau)$, $\tau = 1$ s	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 s)	1×10^{-12} ($\tau = 1$ to 10^2 s)	1×10^{-12} ($\tau = 10^3$ to 10^5 s)	1×10^{-14} ($\tau = 10^5$ to 10^7 s)	1×10^{-15} ($\tau = 10^3$ to 10^5 s)
Aging/year	5×10^{-7}	5×10^{-9}	1×10^{-10}	None	$\sim 1 \times 10^{-13}$
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 s to 1×10^{-6}	<5 min to 1×10^{-8}	<5 min to 5×10^{-10}	30 min to 5×10^{-12}	24 h to 1×10^{-12}

A quartz crystal inside the oscillator is the resonator. It can be made of either 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 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 $Q = 1.6 \times 10^7/f$, where *f* is the resonance frequency in megahertz.

Environmental changes due to temperature, humidity, pressure, and vibration can change the resonance frequency of a quartz crystal, but there are several designs that reduce these environmental effects. 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 is used to generate 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 $\tau = 1$ s) 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 direction of aging occurs. Aging has many possible causes including a build-up of foreign material on the crystal, changes in the oscillator circuitry,

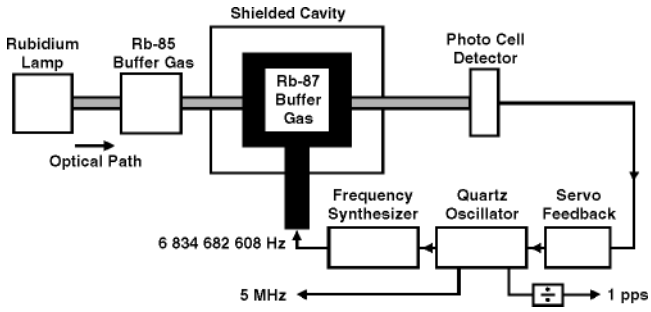

FIGURE 17.11 Rubidium oscillator.

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 better long-term accuracy and stability [4,13,14].

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 its output is used to tune 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 from the quartz oscillator and provided as outputs (Fig. 17.11).

Rubidium oscillators continue to get smaller and less expensive, and offer perhaps the best price-to-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 due mainly to collisions of the rubidium atoms with other gas molecules. These shifts limit the long-term stability. Stability ($\sigma_y(\tau)$, at $\tau = 1$ s) 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 or hours, so they meet the accuracy requirements of most applications without adjustment.

Cesium Oscillators

Cesium oscillators are *primary frequency standards* since the SI second is defined from 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

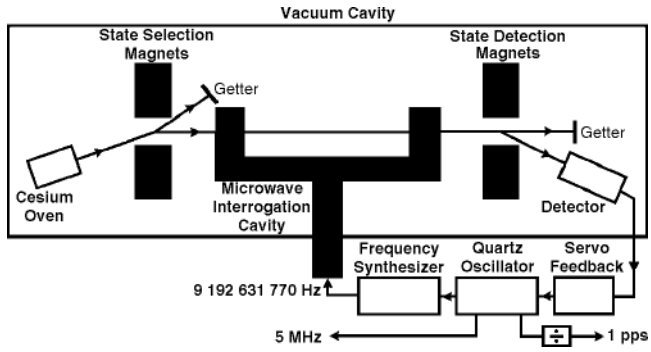

FIGURE 17.12 Cesium beam oscillator.

frequency derived from a quartz oscillator. If the microwave frequency matches the resonance frequency of cesium, the cesium atoms 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. 17.12).

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/s inside the tube. This limits the observation time to a few milliseconds, and the resonance width to a few hundred hertz. Stability ($\sigma_y(\tau)$, at $\tau = 1$ s) 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.

17.4 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 radio spectrum), telephone lines, and the Internet. Synchronization requires both an on-time pulse and a time code. Syntonization requires extracting a stable frequency from the broadcast. The frequency can come from the carrier itself, or from a time code or other information modulated onto the carrier.

This section discusses both the fundamentals of time and frequency transfer and the radio signals used as calibration references. Table 17.6 provides a summary.

Fundamentals of Time and Frequency Transfer

Signals used for time and frequency transfer are generally referenced to atomic oscillators that are steered to agree as closely as possible with UTC. Information is sent from a transmitter (A) to a receiver (B) and is delayed by τ_{ab}, commonly called the *path delay* (Fig. 17.13).

To illustrate path delay, consider a radio signal broadcast over a path 1000 km long. Since radio signals travel at the speed of light (~3.3 μs/km), we can calibrate the path by applying a 3.3-ms correction to

TABLE 17.6 Summary of Time and Frequency Transfer Signals and Methods

Signal or Link	Receiving Equipment	Time Uncertainty (24 h)	Frequency Uncertainty (24 h)
Dial-Up Computer Time Service	Computer, client software, modem, and phone line	<15 ms	Not recommended for frequency measurements
Internet Time Service	Computer, client software, and Internet connection	<1 s	Not recommended for frequency measurements
HF Radio (3 to 30 MHz)	HF receiver and antenna	1 to 20 ms	10^{-6} to 10^{-9}
LF Radio (30 to 300 kHz)	LF receiver and antenna	1 to 100 μs	10^{-10} to 10^{-12}
Global Positioning System (GPS)	GPS receiver antenna	<20 ns	$<2 \times 10^{-13}$

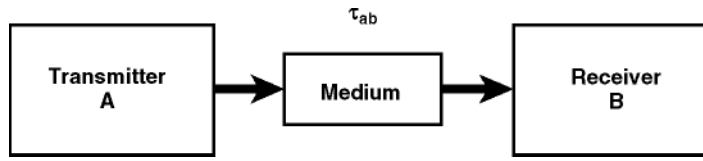

FIGURE 17.13 One-way time and frequency transfer.

our measurement. 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 s of UTC, there is no need to worry about a 100-ms path delay through a network. And, of course, 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.

More sophisticated transfer systems estimate and remove all or part of the path delay. This is usually 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 ms for all users, the time can be sent 20 ms early. This advancement of the timing signal removes at least some of the delay for all users.

A better technique is to compute τ_{ab} and to apply a correction to the received signal. A correction for τ_{ab} can be computed if the position of both the transmitter and 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 (GPS) provides the best of both worlds—each GPS satellite broadcasts its position and the receiver can use coordinates from multiple satellites to compute its own position.

The transmitted information often includes 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; the month, day, and year; and advance warning of daylight saving time and leap seconds.

Radio Time and Frequency Transfer Signals

There are many types of radio receivers designed to receive time and frequency signals. Some are designed primarily to produce time-of-day information or an on-time pulse, others are designed to output standard frequencies, and some can be used for both time and frequency transfer. The following sections look at three types of time and frequency radio signals that distribute UTC—high frequency (HF), low frequency (LF), and GPS satellite signals.

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 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 broadcast the same program, and at least one frequency should be usable at all times. The stations can also be heard by telephone; dial (303) 499-7111 for WWV or (808) 335-4363 for WWVH.

WWV and WWVH signals 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 ms 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 then reflected back to earth. Since the height of the ionosphere changes throughout the day, 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 manual synchronization of analog and digital clocks, simple frequency calibrations, and calibrations of stop watches and timers. However, LF and GPS signals are better choices for more demanding applications [2,7,15].

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 two major advantages—they can often be received indoors without an external antenna and several stations broadcast a time code. 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 17.7). 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 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 those obtained with HF reception. Short paths make it possible to continuously track the ground wave. 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 the carrier. For example, a cycle slip on a 60 kHz carrier introduces a 16.67 μs phase step. However, a strong ground wave signal can produce very good results. An LF receiver that

TABLE 17.7 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

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 h 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 daylight saving time, leap years, and leap seconds. The time code is received and displayed by wristwatches, alarm clocks, wall clocks, and other consumer electronic products [2,7,15].

Global Positioning System (GPS)

The GPS is a navigation system developed and operated by the U.S. Department of Defense (DoD) that is usable nearly anywhere on the 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 h 58 m, which means that each 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 *pseudo random 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 chipping rate of 1023 chips per millisecond. The second is a *precision* (*P*) code with a chipping rate of 10230 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 receiving antenna must have a clear view of the sky [16].

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 maintained within 100 ns of UTC (NIST). The oscillators provide the reference for both the carrier and the code broadcasts.

GPS signals now dominate the world of high performance time and frequency transfer, since they provide reliable reception and exceptional results with minimal effort. A GPS receiver can automatically compute its latitude, longitude, and altitude from 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 GPS receiver calibrated for equipment delays has a timing uncertainty of <20 ns relative to UTC (NIST), and the frequency uncertainty is often $< 2 \times 10^{-13}$ when averaged for 1 day. Figure 17.14 shows an Allan deviation plot of the output of a low cost GPS receiver. The stability is near 1×10^{-13} after about 1 day of averaging.

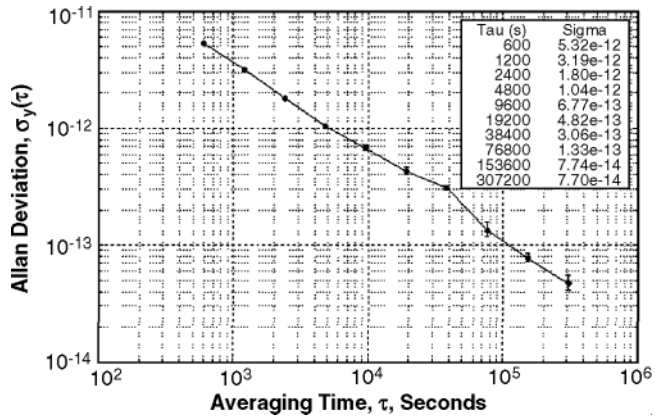

FIGURE 17.14 Frequency stability of GPS receiver.

17.5 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 advance has made many new products and technologies possible. While it is impossible to predict what the future holds, we can be certain that oscillator Qs will continue to increase, measurement uncertainties will continue to decrease, and new technologies will continue to emerge.

References

1. Jespersen, J., and Fitz-Randolph, J., *From Sundials to Atomic Clocks: Understanding Time and Frequency,* 2nd ed., Dover, Mineola, New York, 1999.
2. Kamas, G., and Lombardi, M. A., *Time and Frequency Users Manual,* NIST Special Publication 559, U.S. Government Printing Office, Washington, DC, 1990.
3. Levine, J., Introduction to time and frequency metrology, *Rev. Sci. Instrum.*, 70, 2567, 1999.
4. Hackman, C., and Sullivan, D. B., Eds., *Time and Frequency Measurement,* American Association of Physics Teachers, College Park, Maryland, 1996.
5. ITU Radiocommunication Study Group 7, *Selection and Use of Precise Frequency and Time Systems,* International Telecommunications Union, Geneva, Switzerland, 1997.
6. Novick, A. N., Lombardi, M. A., Zhang, V. S., and Carpentier, A., A high performance multi-channel time interval counter with an integrated GPS receiver, in *Proc. 31st Annu. Precise Time and Time Interval (PTTI) Meeting,* Dana Point, California, p. 561, 1999.
7. Lombardi, M. A., Time measurement and frequency measurement, in *The Measurement, Instrumentation, and Sensors Handbook,* Webster, J. G., Eds., CRC Press, Boca Raton, Florida, 1999, chap. 18–19.
8. Sullivan, D. B., Allan, D. W., Howe, D. A., and Walls, F. L., Eds., *Characterization of Clocks and Oscillators,* NIST Technical Note 1337, U.S. Government Printing Office, Washington, DC, 1990.
9. Jespersen, J., Introduction to the time domain characterization of frequency standards, in *Proc. 23rd Annu. Precise Time and Time Interval (PTTI) Meeting,* Pasadena, California, p. 83, 1991.
10. IEEE Standards Coordinating Committee 27, *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities,* Institute of Electrical and Electronics Engineers, New York, 1999.
11. Walls, F. L., and Ferre-Pikal, E. S., Measurement of frequency, phase noise, and amplitude noise, in *Wiley Encyclopedia of Electrical and Electronics Engineering,* John Wiley and Sons, New York, 1999, 12, 459.

12. Howe, D. A., An extension of the Allan variance with increased confidence at long term, *IEEE Int. Freq. Control Symp.,* 321, 1995.
13. Vig, J. R., Introduction to quartz frequency standards, *Army Research and Development Technical Report,* SLCET-TR-92-1, October 1992.
14. Hewlett-Packard Company, *Fundamentals of Quartz Oscillators,* HP Application Note 200-2, 1997.
15. Carr, J. J., *Elements of Electronic Instrumentation and Measurement,* 3rd ed., Prentice-Hall, NJ, 1996.
16. Hoffmann-Wellenhof, B., Lichtenegger, H., and Collins, J., *GPS: Theory and Practice, 3rd* ed., Springer-Verlag, New York, 1994.