

Frequency Measurement

42.1	Frequency Accuracy and Stability	42-2
	Frequency Accuracy • Frequency Stability	
42.2	Frequency Standards.....	42-8
	Quartz Oscillators • Atomic Oscillators • Rubidium Oscillators • Cesium Oscillators • Hydrogen Masers • Disciplined Oscillators • Choosing a Frequency Standard for a Metrology Laboratory	
42.3	Calibration and Measurement Methods.....	42-19
42.4	Future Developments.....	42-25
	References.....	42-25

Michael A.
Lombardi
*National Institute of
Standards and Technology*

Frequency is the rate of occurrence of a repetitive event. If T is the period of a repetitive event, then the frequency is its reciprocal, $f = 1/T$. The International System of Units (SI) states that the period should be expressed in seconds (s), and the frequency should be expressed in hertz (Hz). The frequency of electrical signals is often stated in units of kilohertz (kHz), megahertz (MHz), or gigahertz (GHz), where 1 kHz equals one thousand (10^3) events per second, 1 MHz equals one million (10^6) events per second, and 1 GHz equals one billion (10^9) events per second.

Frequency metrology first became a topic of interest around 1920, when the burgeoning radio industry began erecting radio transmitters all over the world. These transmitters had to stay near their assigned frequencies to avoid interference with signals from other stations. In addition, the many millions of radio receivers that were soon manufactured had to be able to tune to a desired frequency so that the selected station could be heard. The original requirements for transmitter accuracy were low, about one part per thousand (1×10^{-3}), but at the time, they posed a challenging metrology problem. Within a few years, however, the development of the quartz oscillator soon made that type of accuracy trivial. By the 1930s, commercial quartz oscillators accurate to about one part per million (1×10^{-6}) were widely available, and by the late 1960s, quartz technology with about the same accuracy found its way into low-cost wristwatches and clocks. The production of commercial atomic oscillators also began in earnest in the 1960s, resulting in huge improvements in accuracy and resulting in many new technologies. For example, the infrastructure that we now take for granted, including telecommunication networks and the electric power grid, requires frequency accurate to about 1×10^{-11} to be simultaneously generated around the world during every hour of every day. In calibration and metrology laboratories, frequency measurements accurate to one part per ten trillion (1×10^{-13}) have now become routine, and the Global Positioning System (GPS) depends upon oscillators that are stable to parts in 10^{14} for multiple hours. This level of performance separates frequency metrology from most of the other fields of metrology, where one part per billion (1×10^{-9}) is often either unattainable or considered a major accomplishment.

This chapter provides an overview of frequency measurements. It focuses on the measurement of the electrical signals produced by oscillators. For our purposes, an *oscillator* is a device that

produces electrical signals at a specific frequency, typically in the form of either a sine or a square wave. [Section 42.1](#) begins by discussing the concepts of accuracy and stability, which are essential to understanding oscillator specifications. [Section 42.2](#) then describes the various types of oscillators used as frequency standards, including quartz and atomic oscillators, and oscillators disciplined to agree with an external reference signal. [Section 42.3](#) describes the methods and techniques used to measure frequency and calibrate oscillators. [Section 42.4](#) provides a brief look at the likely future of frequency metrology.

42.1 Frequency Accuracy and Stability

This section looks at the two main specifications used to characterize an oscillator: accuracy and stability. A good understanding of the basic concepts introduced in this section is necessary when evaluating equipment or performing measurements.

42.1.1 Frequency Accuracy

The accuracy of an oscillator is the difference between its actual frequency, as determined by measurement, and its *nominal frequency*. The nominal frequency is labeled on the oscillator output and refers to an ideal frequency with zero uncertainty. For example, an oscillator with an output labeled “10 MHz” would ideally produce perfect 10 MHz signals, but its actual signals will differ from its nominal frequency by some amount. The difference between the actual frequency and the nominal frequency is called the *frequency offset* and determines the accuracy of an oscillator at a given point in time or over a specified interval.

Frequency offset is measured by comparing a test oscillator to a more accurate reference oscillator. There are several established measurement methods (described later in [Section 42.3](#)) that can provide this comparison in either the *frequency domain* or the *time domain*. The standard equation for estimating frequency offset in the frequency domain is

$$f_{\text{off}} = \frac{f_{\text{meas}} - f_{\text{nom}}}{f_{\text{nom}}}, \quad (42.1)$$

where

f_o is the frequency offset

f_{meas} is the actual frequency in hertz reported by the measurement

f_{nom} is the nominal frequency in hertz that the oscillator would ideally produce

Note that in practice, f_{meas} has an associated measurement uncertainty, but f_{nom} is always an ideal value with no uncertainty. Note also that the nominal frequency is included in both the numerator and the denominator. Thus, the resulting value for f_o is dimensionless, and not in units of hertz.

This equation is often simplified in the literature as

$$f_{\text{off}} = \frac{\Delta f}{f}, \quad (42.2)$$

where

f_o is the dimensionless frequency offset

Δf is the difference between the measured and nominal frequency in hertz

f is the nominal frequency in hertz

For example, if an oscillator labeled as 10 MHz (10^7 Hz) produces a frequency that is higher than nominal by 1 Hz, the equation becomes

$$f_{\text{off}} = \frac{1}{10^7} = 1 \times 10^{-7} \tag{42.3}$$

In many cases, the frequency offset of an oscillator is obtained in the *time domain* by measuring *time interval*. This works because frequency is the reciprocal of period, which is expressed as a time interval. A mathematical definition of frequency is

$$f = \frac{1}{T} \tag{42.4}$$

where

T is the period of the signal in seconds

f is the frequency in hertz

it can also be expressed as

$$f = T^{-1} \tag{42.5}$$

If we perform mathematical differentiation on the frequency expression with respect to time and substitute in the result, we can show that the average dimensionless difference in frequency is equivalent to the average dimensionless difference in time or that $\Delta f/f$ is equivalent to $\Delta t/T$. For example,

$$\Delta f = -T^{-2} \Delta t = -\frac{\Delta t}{T^2} = -\frac{\Delta t}{T} f \tag{42.6}$$

therefore,

$$f_{\text{off}} = \frac{\Delta f}{f} = -\frac{\Delta t}{T} \tag{42.7}$$

where

Δt is the difference between two time interval measurements

T is the elapsed time between the two measurements

To keep the sign correct, note that the first reading must be subtracted from the second, therefore,

$$f_{\text{off}} = \frac{TI_2 - TI_1}{T} = -\frac{\Delta t}{T} \tag{42.8}$$

To illustrate this, consider a simple example where a time interval (TI_1) is measured, followed by another time interval measurement (TI_2) one second (10^9 ns) later. If $TI_2 - TI_1 = 100$ ns, this produces the same value for f_o that was previously shown in Equation 42.3:

$$f_{\text{off}} = \frac{100}{10^9} = 1 \times 10^{-7} \tag{42.9}$$

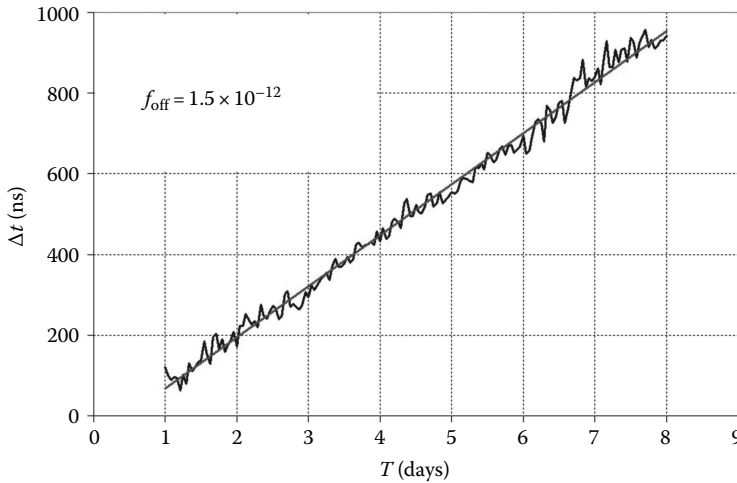


FIGURE 42.1 A sample phase graph used to estimate frequency accuracy.

In practice, more than two interval measurements are recorded, and it is common to graph the results of continuous time interval measurements recorded over multiple hours or days. These graphs are known as *phase* or *time difference* graphs. Phase graphs (Figure 42.1) use the standard Cartesian x/y format. The x -coordinate indicates elapsed time. This is the quantity T shown earlier in Equation 42.8.

The values plotted as the y -coordinate represent the change in phase between the two electrical signals that are being compared to each other. However, since the phase changes are usually measured with instruments that express their results in units of time and not in radians or degrees, the y -coordinate is labeled to show the change in time or the Δt quantity in Equation 42.8.

Average frequency accuracy over a given interval can be estimated from the slope of a phase graph. In practice, a linear least squares line is often fitted to the phase data, and the slope of the least squares line is used to estimate Δt . In many cases, the slope of the least squares line is nearly identical to the slope of the actual data, because of the nearly constant frequency offset between the two oscillators being compared. In fact, to get a good estimate of frequency accuracy, the measurement period must be long enough to show this linear slope and detect a trend. A “clean” phase plot ensures that Δt is really a measure of the test oscillator’s performance and indicates that neither the measurement system nor the reference oscillator has degraded the results by introducing excessive noise.

To illustrate this, Figure 42.1 shows a sample phase graph of an oscillator that was compared to a reference for a period of 7 days. During this period, the total accumulated time difference, Δt , was nearly 1000 ns, as indicated by both the actual data and the least squares line that was fitted to the data. From the slope of the least squares line, we can estimate that $f_o = 1.5 \times 10^{-12}$. The actual data are noisier than the least squares line, because some noise is contributed by the measurement system and the oscillators involved in the comparison. Even so, a strong linear trend is easily detected, and we can be comfortable that this is a good estimate of frequency accuracy.

If necessary, it is easy to convert a dimensionless frequency offset estimate to units of frequency (Hz) if the nominal frequency is known. To illustrate this, consider an oscillator with a nominal frequency of 10 MHz that is high in frequency by 1×10^{-11} . To find the frequency offset in hertz, multiply the nominal frequency by the dimensionless offset:

$$f_{\text{nom}} \times f_{\text{off}} = (1 \times 10^7) (+1 \times 10^{-11}) = 1 \times 10^{-4} = +0.0001 \text{ Hz} = +0.1 \text{ mHz} \quad (42.10)$$

The actual frequency in this case is 10,000,000.0001 Hz, which is obtained by simply adding the offset frequency to the nominal frequency.

To summarize, frequency accuracy and frequency offset are equivalent terms that indicate how closely an oscillator produces its nominal frequency at a given point in time or over a given interval. Frequency accuracy can be estimated in either the frequency domain or the time domain. The accuracy of an oscillator can usually be at least temporarily improved by adjusting it to agree with a more accurate reference [1,2].

42.1.2 Frequency Stability

Frequency stability indicates how well an oscillator can produce the same frequency offset over a given time interval. Any frequency that “stays the same” is a stable frequency, regardless of whether the frequency is “right” or “wrong” with respect to its nominal value. To understand the difference between stability and accuracy, consider that an oscillator in need of adjustment might produce a stable frequency with a large offset. Or, an unstable oscillator that was recently adjusted might temporarily produce an accurate frequency near its nominal value. The stability of an oscillator cannot be changed by adjustment, so unlike accuracy, it tells us something about the inherent quality of an oscillator. In fact, the accuracy of an oscillator over a given interval can never be better than its stability. Figure 42.2 shows the relationship between stability and accuracy.

Frequency stability is normally estimated with statistics that quantify the frequency fluctuations of an oscillator’s output over a given time interval. The fluctuations are measured with respect to a mean frequency offset, and the larger the dispersion of the fluctuations, the greater the instability of the oscillator. *Short-term stability* usually refers to fluctuations over intervals less than 100 s but is commonly used to discuss an oscillator’s stability at an interval of 1 s. *Long-term stability* can refer to any measurement interval greater than 100 s, but commonly is used to discuss stability over intervals of 1 day or longer.

Normally, metrologists rely on classical statistics such as *standard deviation* (or *variance*, the square of the standard deviation) to estimate dispersion. Variance is an estimate of the numerical spread of a dataset with respect to its average or mean value. However, variance works only with stationary data, where the results must be time independent. This assumes the noise is *white*, meaning that its power is evenly distributed across the frequency band of the measurement. Oscillator data are usually nonstationary. For stationary data, the mean and variance will converge to particular values as the number of measurements increases. With nonstationary data, the mean and variance never converge to any particular values. Instead, there is a moving mean that might change each time a new measurement is added [3].

For these reasons, frequency metrologists generally rely on nonclassical statistics to estimate and specify the frequency stability of oscillators [4]. The most common statistic employed for stability estimates is often called the *Allan variance*, but because it is actually the square root of the variance, its proper name is the *Allan deviation* (ADEV). Similar to the standard deviation, ADEV is better suited for




FIGURE 42.2 The relationship between frequency accuracy and stability.

frequency metrology because it has the advantage of being convergent for most types of oscillator noise. The equation for ADEV using frequency measurements and nonoverlapping samples is

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(M-1)} \sum_{i=1}^{M-1} (\bar{y}_{i+1} - \bar{y}_i)^2} \quad (42.11)$$

where \bar{y}_i is the i th in a series of M dimensionless frequency measurements averaged over a measurement or sampling interval designated as τ . Note that while classical deviation subtracts the mean from each measurement before squaring their summation, ADEV subtracts the previous data point. Since stability is a measure of frequency fluctuations and not of frequency offset, the differencing of successive data points is done to remove the time-dependent noise contributed by the frequency offset. Also, note that the \bar{y} values in the equation do not refer to the average or mean of the entire dataset, but instead imply that the individual measurements in the dataset can be obtained by averaging.

The equation for ADEV using phase measurements and nonoverlapping samples is

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(N-2)\tau^2} \sum_{i=1}^{N-2} [x_{i+2} - 2x_{i+1} + x_i]^2}, \quad (42.12)$$

where x_i is the i th in a set of N phase measurements spaced by the measurement interval τ .

To improve the confidence of a stability estimate, ADEV is normally used with overlapping samples that allow estimating stability with all possible combinations of the dataset. The equation for ADEV using phase measurements and overlapping samples is

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(N-2m)\tau^2} \sum_{i=1}^{N-2m} [x_{i+2m} - 2x_{i+m} + x_i]^2}, \quad (42.13)$$

where the averaging factor, m , has been added to Equation 42.12. To understand the averaging factor, consider that τ_0 is the basic measurement interval or the shortest interval at which data are taken. To obtain stability estimates for longer intervals, τ_0 is simply multiplied by m ; thus, $\tau = m\tau_0$. Even though the overlapping samples are not statistically independent, the number of degrees of freedom still increases, thus improving the confidence in the stability estimate [5].

One important advantage of ADEV over classical statistics is its ability to estimate stability over different intervals from the same dataset. Most ADEV graphs found in the literature use the *octave* method, where each successive value of τ is twice as long as the previous value. This method saved computational time, but as computers have become faster, it has become more common to estimate ADEV for all possible values of τ . A typical ADEV graph plots $\log \tau$ on the x -coordinate to indicate the averaging period and $\log \sigma_y(\tau)$ on the y -coordinate to indicate dimensionless frequency stability. These graphs are often referred to colloquially as “sigma- τ ” graphs. ADEV graphs generally show the stability improving as the averaging period increases, until the point where the oscillator reaches its noise floor, or *flicker floor*, when no further gains will be made by averaging additional measurements. Figure 42.3 shows a sample ADEV graph that shows stability estimates using the octave method for intervals of τ ranging from 1 s to more than 2 h. This device was stable to about 3×10^{-12} at $\tau = 1$ s and reached a noise floor near 5×10^{-13} at $\tau = 512$ s. When τ exceeded 1000 s, the oscillator had begun to change frequency. Thus, further averaging degraded, rather than improved, the results.

In addition to estimating stability, ADEV can help identify the types of oscillator noise. Five noise types are commonly discussed in the time and frequency literature: *white phase*, *flicker phase*, *white frequency*, *flicker frequency*, and *random walk frequency*. A brief description of each noise type is provided




FIGURE 42.3 A sample ADEV graph used to estimate frequency stability.

in Table 42.1. The noise type can be identified from the slope of the line on an ADEV graph (Figure 42.4). Note that ADEV cannot distinguish between white phase and flicker phase noise. However, the modified ADEV, $\text{Mod } \sigma_y(\tau)$, can make this distinction, and numerous other variations of ADEV exist for specific applications, such as the improved identification of oscillator noise or improved estimates of long-term stability [4,5].

ADEV and similar statistics have proven to be very useful. However, they appear so often in the literature that confusing stability with accuracy has become a common mistake. It is important to know

TABLE 42.1 Oscillator Noise Types

Noise Type	Description	ADEV Slope
White phase	Fluctuations in the phase of a signal that have the same power at all frequencies across a given bandwidth. The stability is improving at a rate proportional to the averaging period. Mod ADEV can distinguish between white and flicker phase noise (ADEV cannot) and identifies white phase noise as having a slope of $\tau^{-3/2}$.	τ^{-1}
Flicker phase	Also known as $1/f$ phase noise. As the frequency goes up, the intensity of the noise goes down. For example, if the frequency doubles, the power of the noise is cut in half. Unlike white phase noise, flicker phase noise is not evenly distributed across the frequency band.	τ^{-1}
White frequency	Fluctuations in the frequency of a signal that have the same power at all frequencies across a given bandwidth. The stability is still improving, but at a rate proportional to the square root of the averaging period.	$\tau^{-1/2}$
Flicker frequency	Also known as $1/f$ frequency noise. The oscillator has reached a noise floor that shows its best possible stability (often called the “flicker” floor). There is nothing to be gained by more averaging.	τ^0
Random walk	Successive random steps in frequency. The difference between two steps is nearly constant, but the direction of the steps is random. Even so, it is clear that the oscillator frequency is now changing. A slope of τ^1 is sometimes used to distinguish frequency drift from random walk.	$\tau^{1/2}$




FIGURE 42.4 Identification of oscillator noise types.

that the two specifications mean different things and can have very different values for the same oscillator. For example, an oscillator accurate to only 1×10^{-8} might still be stable to 1×10^{-13} at $\tau = 1$ s. This means that even though the frequency of the oscillator is changing by only a small amount during short intervals, it is not particularly close to its nominal value.

42.2 Frequency Standards

As noted previously, stability measurements can tell us something about the inherent quality of an oscillator, and the stability of an oscillator is closely related to its *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. Obviously, a high resonance frequency and a narrow resonance width are both advantages when seeking a high Q . Stability and Q are generally correlated, because a high Q means that an oscillator has to stay close to its natural resonance frequency.

This section discusses the various types of oscillators used as frequency standards. It begins by discussing quartz oscillators, which achieve the highest Q of any mechanical-type device. It then discusses oscillators with higher Q factors, based on the atomic resonance of rubidium, hydrogen, and cesium.

This is followed by a discussion of disciplined oscillators. These devices can be either quartz oscillators or atomic oscillators, but their frequency is automatically adjusted to agree with an external reference.

This section concludes with a discussion of which type of frequency standard is best suited for use in a metrology laboratory. Table 42.2 provides a summary [2,6–8].

42.2.1 Quartz Oscillators

Billions (10^9) of quartz crystal oscillators are manufactured annually. Most are miniature, inexpensive devices that are embedded inside wristwatches, clocks, computers, cellular phones, and nearly every type of electronic circuit. However, only the larger, more expensive varieties of quartz oscillators are used as frequency standards. These devices are sometimes sold as stand-alone instruments but are more typically found inside test and measurement equipment, such as counters, signal generators, and oscilloscopes.

A quartz crystal inside the oscillator serves as resonator. The crystal strains (expands or contracts) when a voltage is applied. Reversing the polarity of the applied voltage will reverse the strain and force

TABLE 42.2 Summary of Oscillator Types

Oscillator Type	Quartz (OCXO)	Rubidium	Cesium	Active Hydrogen Maser	GPSDO
Primary standard	No	No	Yes	No	No
Resonance frequency	Mechanical (varies)	6.834682610904 GHz	9.19263177 GHz	1.420405751768 GHz	NA
Quality factor, Q	$\sim 10^6$	$\sim 10^7$	$\sim 10^8$	$\sim 10^9$	NA
Frequency accuracy (1 day average)	1×10^{-6} to 1×10^{-10}	5×10^{-9} to 5×10^{-12}	1×10^{-12} to 1×10^{-14}	$\sim 1 \times 10^{-13}$	1×10^{-12} to 5×10^{-14}
Stability, $\sigma_y(\tau)$, $\tau = 1$ s	1×10^{-11} to 1×10^{-13}	5×10^{-11} to 5×10^{-12}	1×10^{-11} to 5×10^{-12}	$\sim 2 \times 10^{-13}$	1×10^{-10} to 1×10^{-12}
Stability, $\sigma_y(\tau)$, $\tau = 1$ day	1×10^{-10}	5×10^{-12}	8×10^{-14} to 2×10^{-14}	$\sim 2 \times 10^{-16}$	1×10^{-12} to 5×10^{-14}
Aging/year	5×10^{-9}	2×10^{-10}	None	$\sim 1 \times 10^{-13}$	None
Phase noise (dbc/Hz, 10 Hz from carrier)	-125 to -140	-90 to -130	-130 to -136	-130 to -142	-90 to -140
Life expectancy	Indefinite	>15 years	5-20 years	>15 years	>15 years
Cost (USD)	\$500-\$5,000	\$2,000-\$10,000	\$30,000-\$80,000	\sim \$200,000	\$1,000 to \$20,000

the crystal to mechanically oscillate. This is known as the *piezoelectric effect*. The energy needed to sustain oscillation is obtained by taking a voltage signal from the resonator, amplifying it, and feeding it back to the resonator. Figure 42.5 is a simplified circuit diagram that shows the basic elements of a quartz crystal oscillator.

The rate of expansion and contraction is the resonance frequency and is determined by the cut and size of the crystal. No two crystals can be exactly alike or produce exactly the same frequency. 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 the 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 10^3 for a wristwatch-type oscillator to higher than 10^6 for the most stable devices. The maximum Q for a high-stability quartz oscillator can be roughly estimated as $Q = 16 \text{ million}/f$, where f is the resonance frequency in megahertz.

Environmental changes can change the resonance frequency of a quartz crystal. Temperature changes are the largest problem, but other parameters such as humidity, pressure, and vibration can also change the frequency. There are several types of design packages that reduce these environmental problems.

The most stable type of quartz oscillator is the *oven-controlled crystal oscillator (OCXO)*, which encloses

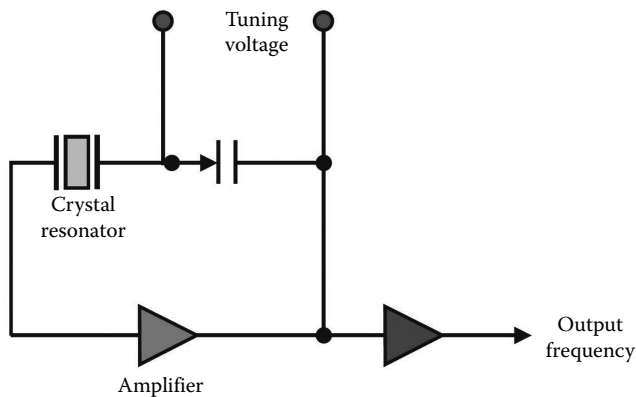


FIGURE 42.5 Block diagram of quartz oscillator.

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 period, 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 alternative solution to the temperature problem is the *temperature-compensated crystal oscillator* (TCXO). These devices include a temperature sensor that 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 it generally costs less. Therefore, TCXOs are used when high stability over a wide temperature range is not required.

The best quartz oscillators have excellent short-term stability. A high-quality OCXO might be stable to 1×10^{-13} at $\tau = 1$ s. However, quartz oscillators are not stable over long intervals. Their long-term stability is limited by *aging*, which causes their frequency to change over time due to internal changes in the oscillator. Aging usually results in a nearly linear change in the resonance frequency that can be either positive or negative. A reversal in the direction of the aging occasionally occurs, and ironically, the aging rate of a quartz oscillator sometimes decreases as the device gets older. 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 less than 5×10^{-9} per year, while a TCXO might age 100 times faster.

The simple design of quartz oscillators makes them very reliable, and many devices have run continuously for decades without failing. However, their accuracy can change rapidly due to aging and environmental factors, and even a high-quality OCXO will need regular adjustments to maintain frequency accurate to within 1×10^{-9} . They are also subject to large frequency shifts when they are turned on after a power outage. For these reasons, quartz oscillators are usually a poor choice as a frequency standard, unless the measurement requirements of a laboratory are very low [8,9].

42.2.2 Atomic Oscillators

Atomic oscillators derive their resonance frequency from the quantized energy levels in atoms. The laws of quantum mechanics dictate that the energies of a bound system, such as an atom, have certain discrete values. An electromagnetic field 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 electromagnetic energy. The resonance frequency (f) of an atomic oscillator is the difference between the two energy levels divided by Planck's constant (h):

$$f = \frac{E_2 - E_1}{h} \quad (42.14)$$

All atomic oscillators are *intrinsic standards*, because their frequency is inherently derived from a fundamental natural phenomenon. There are currently (2011) three types of atomic oscillators sold commercially: rubidium standards, cesium standards, and hydrogen masers (discussed individually in the following sections). All three types contain an internal quartz oscillator that is locked to a resonance frequency generated by the atom of interest. This method causes the factors that degrade the long-term stability of a quartz oscillator to disappear. As a result, the long-term stability of an atomic oscillator is at least several orders of magnitude better than that of a quartz oscillator, but the short-term stability is unchanged [2,8,10].

42.2.3 Rubidium Oscillators

Rubidium oscillators are outperformed by the other types of atomic oscillators, but they have the advantage of being much smaller and less expensive. Because of their low cost and small size, rubidium




FIGURE 42.6 Block diagram of rubidium oscillator.

oscillators are often found in telecommunication networks and as time base oscillators in test and measurement equipment. They operate at 6,834,682,610.904 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 photocell 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 (Figure 42.6).

The Q of a rubidium oscillator is about 10^7 . The shifts in the resonance frequency are mainly caused by collisions between the rubidium atoms and other gas molecules. These frequency shifts limit the long-term stability. Stability at $\tau = 1$ s is typically less than 1×10^{-11} and near 1×10^{-12} at $\tau = 1$ day. There is generally no guaranteed specification for accuracy, but after a warm-up period of a few minutes, a rubidium oscillator will typically be accurate to within parts in 10^{10} or less, and some devices might be as accurate as 5×10^{-12} . However, if an application has an accuracy requirement of parts in 10^9 or smaller, a rubidium oscillator will need to be regularly measured and adjusted, because accuracy better than about 5×10^{-9} cannot be assumed. With regular frequency adjustments, a rubidium can maintain average frequency to within a few parts in 10^{11} or 10^{12} over periods of months or years. The adjustments are made to compensate for the aging and frequency drift that changes the rubidium frequency slowly over time. Manufacturers typically specify the aging rate as less than 5×10^{-11} per month, but this is sometimes conservative, as the frequency of a well-behaved rubidium standard might change by less than 1×10^{-11} over the course of a month. Even so, the frequency change can exceed 1×10^{-10} if left unadjusted for a year, which is unacceptable for some applications [2,6–8].

42.2.4 Cesium Oscillators

Cesium oscillators are primary frequency standards because 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 have inherent accuracy and stability and be close to its nominal frequency without adjustment.

Cesium is a complicated atom with $F = 3$ and $F = 4$ ground states (Figure 42.7). Each atomic state is characterized not only by the quantum number F but also by a second quantum number, m_F , which can have integer values between $-F$ and $+F$. There are 16 possible magnetic states of cesium, but the




FIGURE 42.7 Cesium atomic structure.

transition between the $|4,0\rangle$ and $|3,0\rangle$ states is insensitive to magnetic fields. Thus, the frequency of this transition was chosen to define the SI second.

Figure 42.8 provides a simplified schematic of a cesium beam frequency standard. The design details of a cesium beam standard can vary significantly from model to model, but their basic design principles are similar. As shown on the left side of the figure, ^{133}Cs atoms are heated to a gaseous state in an oven. A beam of atoms emerges from the oven at a temperature near 100°C and travels through a magnetic field, where the beam is split into two beams of atoms with different magnetic states. One beam is absorbed by the getter and is of no further interest. The other beam is deflected into the microwave interrogation cavity (commonly known as the Ramsey cavity).

While inside the Ramsey cavity, the cesium beam is exposed to a microwave frequency from a frequency synthesizer driven by a quartz oscillator. If this frequency is tuned to precisely match cesium resonance, some of the atoms will change their magnetic state. After leaving the Ramsey cavity, the atoms pass through a second magnetic field. These magnets direct only the atoms that changed state to the detector; the other atoms are directed to a getter and absorbed. In essence, the magnets located on

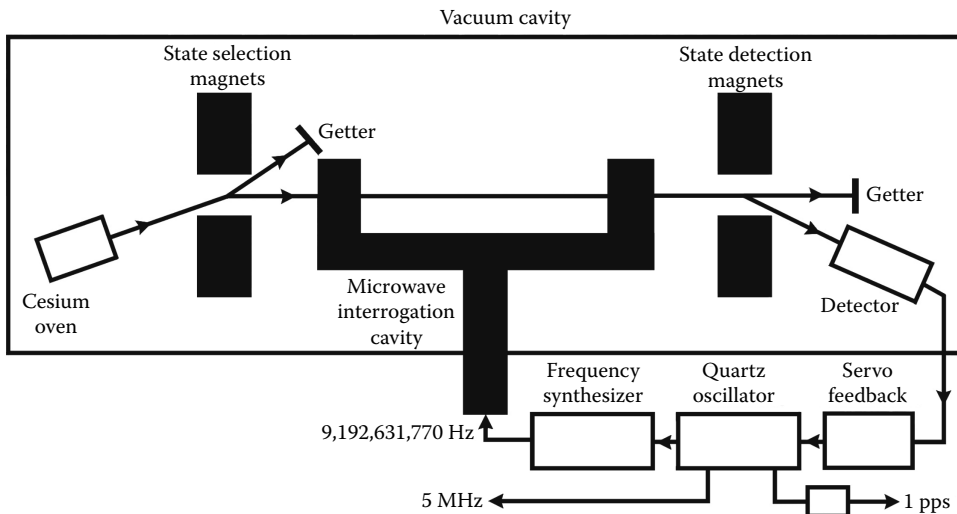


FIGURE 42.8 Block diagram of cesium oscillator.

both sides of the Ramsey cavity serve as a “gate” that allows only those atoms that undergo the desired $|4,0\rangle \leftrightarrow |3,0\rangle$ energy transition to pass through and reach the detector. The detector sends a feedback signal to a servo circuit that continually tunes the quartz oscillator so that the maximum number of atoms reaches the detector, thereby increasing the signal strength. This process is analogous to carefully tuning a radio dial until the loudest and clearest signal is heard and keeps the quartz oscillator frequency locked as tightly as possible to cesium resonance. Output frequencies, such as 1 Hz, 5 and 10 MHz, are then derived from the locked quartz oscillator.

The Q of a commercial cesium standard is about 10^8 . The beam tube is typically less than 0.5 m in length, and the atoms travel at velocities of greater than 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} , normally reaching parts in 10^{14} at $\tau = 1$ day. The frequency offset is typically near 1×10^{-13} after a warm-up period of 30 min. Because the second is defined based on cesium resonance, there should be no change in frequency due to aging. However, in practice a cesium oscillator will slowly change its frequency by a very small amount, typically by parts in 10^{17} per day.

Cesium standards have a limited life expectancy and a high cost. The major component of a cesium oscillator, called the *beam tube*, typically lasts for about 5–10 years, and replacing the beam tube can cost nearly as much as replacing the entire device. When the beam tube fails, a cesium standard will no longer be locked to cesium resonance and will become a free-running quartz oscillator. For this reason, cesium frequency standards should be regularly monitored or checked to ensure that they are working properly [8,10–12].

42.2.5 Hydrogen Masers

The *hydrogen maser* is the most expensive commercially available frequency standard and is therefore found in only a small number of metrology laboratories. The word *maser* is an acronym that stands for microwave amplification by stimulated emission of radiation. The resonance frequency of the hydrogen atom is 1,420,405,751.768 Hz.

There are two types of hydrogen masers. The first type, called an *active maser*, has a microwave cavity that oscillates spontaneously, and a quartz oscillator is phase locked to this active oscillation. The second type, called a *passive maser*, frequency locks a quartz oscillator to the atomic reference in much the same fashion as a rubidium or cesium oscillator. Because active masers derive their output frequency more directly from the atomic resonance, they are more stable than passive masers in both the short and long term. Both types of maser are more stable in the short term than cesium oscillators and are well suited for applications where optimal frequency stability is required. Over long intervals, however, hydrogen masers are less accurate than cesium oscillators. This is due to several factors: their accuracy depends upon a more complex set of conditions, the resonance frequency of their microwave cavity can change over time, and also because the definition of the SI second is based on cesium resonance [2,8,10].

An active hydrogen maser works by sending hydrogen gas through a magnetic gate that only allows 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 (Figure 42.9).

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 about one order of magnitude better than that of a commercial cesium standard. As noted, a hydrogen maser is more stable than

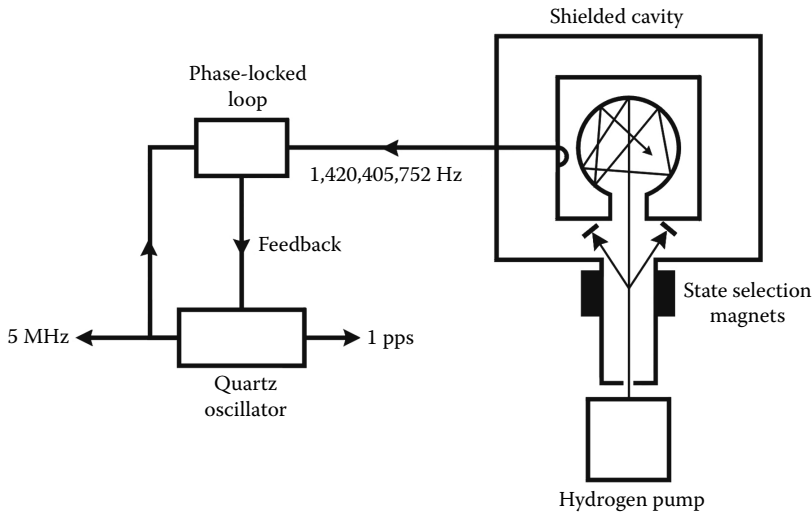


FIGURE 42.9 Block diagram of an active hydrogen maser.

a cesium oscillator for periods ranging from 1 s to perhaps weeks or months. The stability of an active maser will typically reach a few parts in 10^{13} at $\tau = 1$ s and a few parts in 10^{16} at $\tau = 1$ day.

42.2.6 Disciplined Oscillators

Oscillators whose frequency is controlled by an external reference signal are known as *disciplined oscillators*. Unlike free-running oscillators that need to be periodically adjusted to stay within specification, disciplined oscillators are locked to a reference signal and never require manual adjustment. The best disciplined oscillators can generate local signals with nearly the same accuracy and stability as the remote reference. Various types of radio signals have been used to discipline oscillators, but the vast majority of disciplined oscillators in use today (2011) employ signals from the GPS satellites as their reference source. For this reason, this section will focus entirely on GPS-disciplined oscillators (GPSDOs).

Unlike the other types of frequency standards described earlier, a GPSDO requires a small antenna to be mounted on a rooftop location with a clear view of the sky. A GPSDO will normally begin surveying its antenna position as soon as it is turned on. The survey is a one-time process that typically lasts for several hours. When the antenna survey is complete, the GPSDO is ready to use as a frequency standard and will typically produce sine wave signals of 5 and/or 10 MHz.

The basic function of a GPSDO is to receive signals from the GPS satellites and to use the information contained in these signals to control the frequency of a local quartz or rubidium oscillator. GPS signals are kept in agreement with the Coordinated Universal Time scale maintained by the United States Naval Observatory (UTC[USNO]). Nearly all GPSDOs use the coarse acquisition (C/A) code on the L1 carrier frequency (1575.42 MHz) as their incoming reference signal. The satellite signals can be trusted as a reference for two reasons: (1) they originate from atomic oscillators and (2) *they must be accurate and stable* to within parts in 10^{14} over a 12 h averaging period in order for GPS to meet its specifications as a positioning and navigation system. The best GPSDOs transfer as much of the inherent accuracy and stability of the satellite signals as possible to the signals generated by the local quartz or rubidium oscillator.

Many of the methods used to discipline oscillators are proprietary, and GPSDO manufacturers seldom disclose exactly how their products work. However, there are a few basic concepts that apply to most designs. Generally, the local oscillator is controlled with one or more servo loops, with each loop having a fixed or variable time constant. For example, one type of servo loop is a *phase-locked loop*, or PLL. In a GPSDO, the reference input signal to the PLL comes from a GPS receiver that tracks

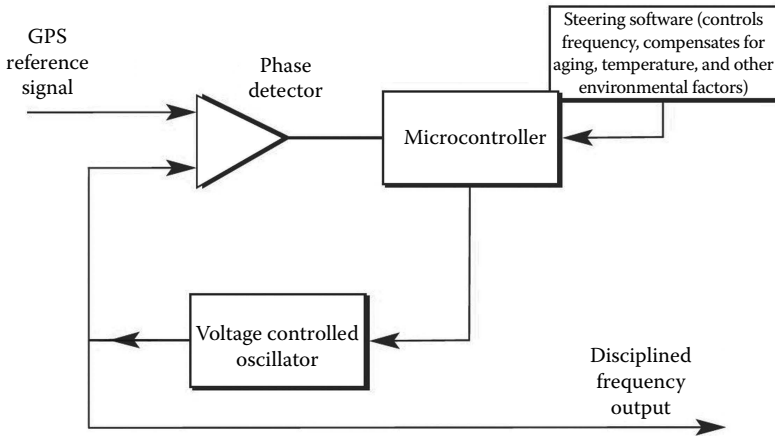


FIGURE 42.10 Block diagram of a GPSDO that steers its local oscillator.

multiple satellites and outputs a 1 pulse per second (pps) signal. A phase detector measures the difference between the 1 pps signal from the GPS receiver and a signal from a voltage-controlled oscillator (VCO). The VCO typically has a nominal frequency of 10 MHz, so its signal is divided to a lower frequency (often all the way down to 1 pps) prior to this phase comparison. A microcontroller reads the output of the phase detector and monitors the phase difference. When the phase difference changes, the software changes the control voltage sent to the VCO, so that the phase difference is held within a given range (Figure 42.10). The GPSDO is locked when the phase of the VCO has a constant offset relative to the phase of the GPS signals. Ideally, the servo loop must be loose enough to ignore the short-term fluctuations of the GPS signals, reducing the amount of phase noise and allowing the VCO to provide reasonably good short-term stability. However, the loop must be tight enough to track GPS closely and to allow the GPS signals to control the VCO frequency in the longer term. The microcontroller software often compensates not only for the phase and frequency changes of the local oscillator but also for the effects of aging, temperature, and other environmental parameters.

Another type of GPSDO design does not correct the frequency of the local oscillator. Instead, the output of a free-running local oscillator is sent to a frequency synthesizer. Steering corrections are then applied to the output of the synthesizer (Figure 42.11). Modern direct digital synthesizers (DDSs) have

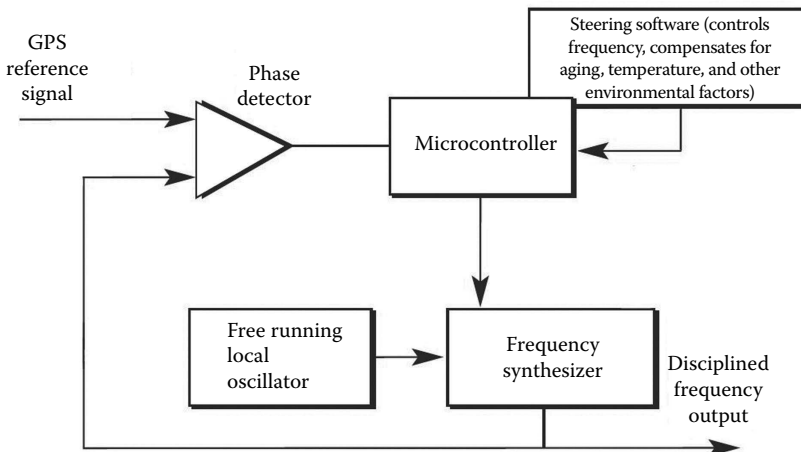


FIGURE 42.11 Block diagram of a GPSDO that adjusts a frequency synthesizer.

excellent resolution and allow very small frequency corrections to be made. For example, 1 μHz resolution at 10 MHz allows instantaneous frequency corrections of 1×10^{-13} . In addition, allowing the local oscillator to free run often results in better performance than the VCO method, where unexpected shifts in the control voltage can produce unwanted adjustments in the output frequency.

A reasonably good metric to use when evaluating GPSDO performance is its frequency stability at $\tau = 1$ day, as estimated with ADEV. Stability of 1×10^{-13} or less at $\tau = 1$ day normally indicates a device of high quality, and many (but certainly not all) GPSDOs can reach this specification. To demonstrate how the performance of GPSDOs can vary significantly, a test was conducted between two rubidium-based GPSDOs at National Institute of Standards and Technology (NIST). Both devices had the same type of rubidium local oscillator and cost approximately the same amount. During the test, both GPSDO devices were connected to the same GPS antenna with an antenna splitter. The antenna's position had previously been surveyed with an uncertainty of less than 1 m, and these precise coordinates were keyed into both units. The 10 MHz outputs of both devices were then simultaneously compared to the US national frequency standard at NIST, UTC(NIST), for a period of 80 days. The results are shown in the phase graph in Figure 42.12.

The results show that the frequency of Device A was tightly controlled. The peak-to-peak phase variation over the entire 80 day period was just 38 ns, with most of this variation due to the difference between UTC(USNO), the reference for GPS, and UTC(NIST) during the measurement interval. The frequency accuracy, as estimated from the slope of the phase, was about 1×10^{-15} . In sharp contrast, the frequency of Device B was very loosely controlled, and the phase plot shows a very large peak-to-peak phase variation of 588 ns, much larger than the dispersion of the GPS timing signals. Figure 42.13 shows the long-term frequency stability of both devices as estimated with ADEV, for values of τ ranging from 1 h to about 3 weeks. Device A is more stable than Device B at all averaging periods by roughly a factor of 10. Stability at $\tau = 1$ day, the metric discussed earlier, is about 6×10^{-14} for Device A, comparable to the performance of a cesium standard. However, the stability of Device B is worse by more than a factor of 10, about 70×10^{-14} .

In spite of the wide performance disparity between Device A and Device B, these examples still illustrate that a GPSDO that is even loosely locked to the satellite signals should be inherently accurate (parts in 10^{13} or better) and inherently stable in the long term. This is because the signals broadcast by the GPS satellites are continuously steered to agree with UTC and GPSDOs that simply "follow" the satellites will produce frequency that closely agrees with UTC over long intervals.

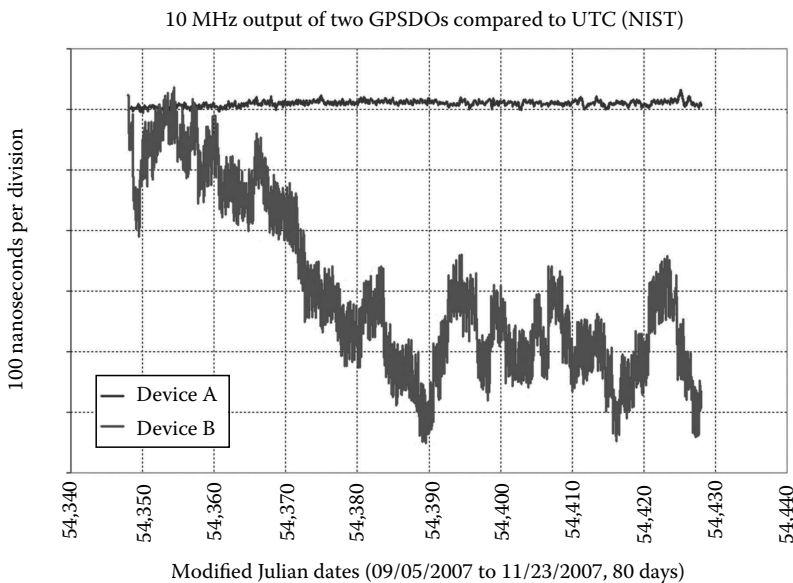


FIGURE 42.12 Phase comparison of two GPSDOs.

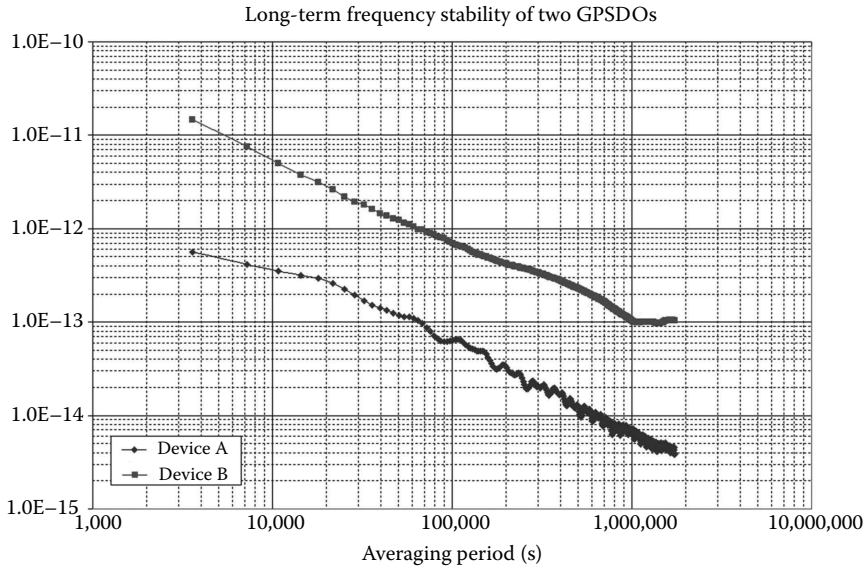


FIGURE 42.13 Stability comparison of two GPSDOs.

GPSDOs are generally reliable, but users should be aware that they will eventually fail if the GPS signal is unavailable. The most likely cause of failure is probably RF interference, because GPS signals are highly susceptible to intentional or unintentional interference due to their low power levels. A GPSDO can stop tracking satellites if there are interfering signal only a few orders of magnitude more powerful than the minimum received GPS signal strength, which is -160 dBW on Earth for the L1 carrier, equivalent to 10^{-16} W. When the GPS signal is unavailable, a GPSDO continues to produce frequency but begins relying on its *holdover* capability. In many cases, a GPSDO will simply become a free-running local oscillator while in holdover mode, in which case, its frequency accuracy will probably be several orders of magnitude worse than normal, perhaps parts in 10^9 or 10^{10} for a device with a rubidium local oscillator after a few hours without GPS and parts in 10^7 or 10^8 for a quartz based device under the same conditions. Some devices implement holdover algorithms that continue to steer the local oscillator without GPS, but the performance of these algorithms will degrade over time. As is the case with a cesium standard, it is important to verify that a GPSDO is working properly [13].

42.2.7 Choosing a Frequency Standard for a Metrology Laboratory

As noted earlier, a quartz oscillator is generally a poor choice as a frequency standard, and only those laboratories with the most demanding frequency stability requirements will be able to justify the expense of a hydrogen maser. Thus, when a metrology laboratory decides which frequency standard to buy, it will likely be choosing between a rubidium oscillator, a cesium oscillator, and a GPSDO. The specifications of the various oscillator types are summarized in Table 42.2. The specifications listed in the table were obtained from manufacturer’s specification sheets and from the results of measurements performed by NIST.

As Table 42.2 indicates, a GPSDO will have better long-term frequency accuracy and stability than a stand-alone rubidium oscillator, and the GPSDO will never require adjustment. A GPSDO will normally cost more than a stand-alone rubidium standard, but in most cases, the performance and convenience of the GPSDO will easily justify the higher cost. Therefore, for most metrology laboratories, a GPSDO is probably a better choice.

Choosing between a cesium standard and a GPSDO is more difficult. As noted earlier, the SI second is defined based on energy transitions of the cesium atom; thus, cesium oscillators are often the preferred choice of frequency standard for laboratories with the best measurement capabilities and the most

demanding performance requirements. However, as also noted earlier, cesium oscillators are expensive and have a limited life expectancy, and not all laboratories can afford them. Assuming that a calibration laboratory can afford a cesium standard, should they still economize by choosing a GPSDO as their frequency standard? There are several pros and cons related to GPSDOs that should be considered before answering this question. First the pros:

- A GPSDO costs much less than a cesium standard to initially purchase, sometimes as much as 90% less. It also costs less to own, because there is no cesium beam tube to replace. This means that a calibration laboratory could buy multiple GPSDOs for less than the cost of a cesium standard and use the additional standards for crosschecks and redundancy.
- Unlike a cesium standard, a GPSDO can recover time by itself (time-of-day and an on-time pulse synchronized to UTC). This is important if a laboratory needs time synchronization capability.
- Cesium standards seldom require adjustment, but a GPSDO will never require adjustment, because its frequency is controlled by the signals from the GPS satellites.

Now, the cons:

- GPSDOs generally have poorer short-term stability and higher phase noise than cesium standards.
- GPSDOs require an outdoor antenna that must be located in an area with access to the roof. A cesium standard can be operated anywhere where electric power is available.
- Cesium standards are autonomous and independent sources of frequency, which means they can operate without input from another source. GPSDO can operate properly only where signals from the GPS satellites are available and are not suitable for applications that need an autonomous frequency source.

Based on these criteria, it seems likely that many laboratories that can afford a cesium standard will undoubtedly choose a GPSDO as a lower-cost alternative that meets all of their requirements. However, a certain percentage of laboratories do require a cesium standard, and some laboratories will operate both types of standards. Even if a laboratory already owns a cesium standard, it might be wise to acquire a GPSDO as a secondary standard so that the two devices can be compared to each other to ensure that both are working properly. Table 42.3 lists suppliers of the various types of frequency standards.

TABLE 42.3 Suppliers of Frequency Standards

Company	Website	Rubidium Oscillators	Cesium Oscillators	Hydrogen Masers	GPS Disciplined Oscillators
Accubeat	www.accubeat.com	X			X
Arbiter	www.arbiter.com				X
Brandywine	www.brandywinecomm.com	X			X
EndRun	www.endruntechnologies.com				X
FEI-Zyfer	www.fei-zyfer.com				X
Fluke	www.fluke.com	X			X
Frequency electronics	www.frequelec.com	X	X		
Meinberg	www.meinberg.de				X
Oscilloquartz	www.oscilloquartz.com	X	X		X
Precision test systems	www.ptsyst.com	X			X
Precise time and frequency	www.ptnc.com	X		X	X
Spectracom	www.spectracomcorp.com	X			X
Stanford research	www.thinksrs.com	X			
Symmetricom	www.Symmetricom.com	X	X	X	X
Trak	www.trak.com				X
Trimble	www.trimble.com				X

42.3 Calibration and Measurement Methods

The objective of a frequency calibration is to measure the accuracy and/or stability of the *device under test* (DUT), which will be one of the oscillator types described in Section 42.2. During the calibration, the DUT is compared to a *standard* or *reference*. In most cases, both the DUT and the reference produce oscillating sine wave signals, as illustrated in Figure 42.14. The sine wave signals produce one cycle (2π radians of phase) in one period. The period is measured in units of time, and the amplitude is measured in units of voltage. The nominal frequency of the DUT is normally 1 MHz or higher, with 5 or 10 MHz being common.

In order for the calibration to be valid, the reference must outperform the DUT. The ratio by which the reference outperforms the DUT is called the *test uncertainty ratio* (TUR). A TUR of 10:1 is preferred, but not always possible. If a smaller TUR is used (4:1, for example), then the calibration will take longer to perform, because more measurements and more averaging will be required.

To further validate a measurement result, it is necessary to establish the *traceability* of a frequency measurement to the International System (SI) of units. Because frequency is the reciprocal of time interval, establishing traceability in frequency metrology means establishing traceability to the SI second through an unbroken and documented chain of calibrations. The SI second is a virtual and not a physical standard, so the chain of calibrations typically extends back to a national standard maintained by a laboratory that contributes to the UTC, such as NIST in the United States. The traceability chain will have only one link if a DUT is calibrated by NIST or an equivalent laboratory. However, there are often several links (calibrations) involved when establishing a traceability chain back to the SI.

Once a suitable reference and a plan for traceability have been chosen, the next step is to select a calibration or measurement method. The remainder of this section discusses several established techniques that are commonly used to measure frequency. With the appropriate hardware and software, measurement systems designed around these techniques can produce data that can be used to estimate either frequency accuracy or stability.

The comparison of the sine wave signal produced by the DUT to the sine wave signal produced by the reference can be made in either the time domain or in the frequency domain, but time domain

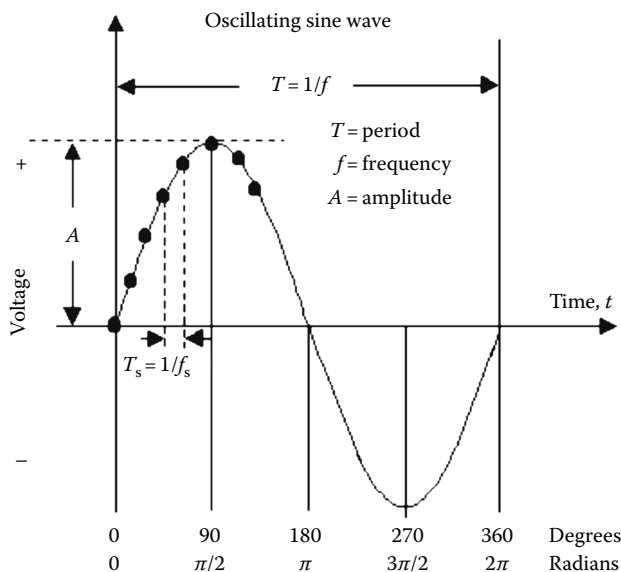


FIGURE 42.14 An oscillating sine wave.




FIGURE 42.15 Phase comparison with an oscilloscope.

comparisons are more common. A simple time domain comparison that nicely illustrates the concept is the *oscilloscope pattern drift method*. This method requires a two-channel oscilloscope. In this example, both the DUT and reference signals are 10 MHz sine waves. The reference signal is produced by a cesium oscillator and the DUT signal is produced by a quartz oscillator. The scope is triggered with the reference signal on channel 2, and the DUT signal is connected to channel 1. The amplitude and position of both waveforms are adjusted on the oscilloscope display so that they overlap (Figure 42.15).

If the two frequencies were exactly the same, both sine waves would appear to be stationary on the oscilloscope display. However, because the two frequencies are not the same, their phase relationship will be continuously changing. The reference sine wave will appear to be stationary and the DUT sine wave will move. The direction of the sine wave motion will determine whether the DUT frequency is low or high with respect to the reference. A stopwatch can be started when the two signals are exactly in phase and stopped after one complete *cycle slip* has occurred and they are exactly in phase again. Frequency accuracy can be estimated as $\Delta t/T$ (see Equation 42.8), where $\Delta t = 100$ ns (the period of 10 MHz) and T is the elapsed time indicated on the stopwatch. If one cycle slip occurs per second, the frequency accuracy of the DUT is 1×10^{-7} [1].

As noted, this method is very useful for demonstrating the concept of a phase comparison, and oscilloscopes are indispensable for simple measurements and for viewing waveforms. However, the shortcomings of the pattern drift method quickly become obvious: If DUT were accurate to 1×10^{-12} , a single cycle slip would take more than a day to occur. For this reason and others, you will probably seldom use an oscilloscope to measure frequency, especially if you have a universal counter available.

Universal counters can be configured as either a *frequency counter (FC)* or a *time interval counter (TIC)*. FCs are the most common instrument used for frequency measurements, and they are especially handy, because they can quickly measure a DUT's frequency and instantly display the results. The reference for this type of measurement is the FC's time base oscillator, which is usually a quartz oscillator of unknown accuracy (typically no better than 1×10^{-8}). For this reason, the laboratory's best oscillator




FIGURE 42.16 TIC measurement system.

should be connected to the FC's external time base input. Frequency accuracy can be estimated by using Equation 42.1, where f_{meas} is the reading taken from the FC display. The smallest frequency offset that an FC can detect with a single reading will be determined by its *resolution*, which is limited by the number of digits on the counter's display. For example, a 10-digit FC will be unable to detect a frequency change smaller than 1×10^{-9} without averaging when measuring a 10 MHz signal, but a 12-digit counter can reduce this value by two orders of magnitude to 1×10^{-11} .

TICs are commonly used to measure frequency in the time domain (Figure 42.16). When the *time interval method* is used to measure frequency, it is no longer practical to work directly with 10 MHz signals, because low-frequency input signals must be used to start and stop the counter. The solution is to use a *frequency divider* to convert standard frequency signals to a much lower frequency, typically to 1 Hz. Frequency dividers can be stand-alone instruments, integrated into the oscillator design or, in some cases, integrated into the TIC. Dividing to 1 Hz can be accomplished by blocking $f_{\text{nom}} - 1$ cycles of the frequency from passing through to the counter. The use of low-frequency signals reduces the problem of counter overflows and underflows (cycle ambiguity) and helps prevent errors that can occur when the start and stop signals are too close together. It is also important to make sure that the TIC is triggering at the correct voltage level on the input signal. This involves either carefully adjusting the trigger levels or using a TIC with fixed trigger levels and converting both signals to an identical shape and amplitude prior to the comparison.

A TIC has inputs for two signals. Typically, a signal from the DUT starts the counter and a signal from the reference stops the counter. The time interval reading will change, typically very slowly, to indicate the difference in frequency between the two signals. A single reading from an FC can immediately produce useful information, because it can be differenced from the nominal frequency. However, in the case of a TIC, at least two readings are required to produce useful frequency information. As indicated in Equation 42.8, a quick estimation of the DUT's frequency offset can be made by recording a single TIC reading, waiting for a specified period, and then recording a second TIC reading. The difference between the two readings divided by the measurement period ($\Delta t/T$) provides an estimate of the frequency offset. In practice, data are usually collected continuously, typically every second, and then averaged over longer intervals of time to get a better estimate.

Modern universal counters typically have a single-shot time interval resolution of less than 1 ns. The best devices have a resolution near 10 ps, which enables them to detect frequency changes of 1×10^{-11} in 1 s, which is equivalent to the capability of a 12-digit FC. Dedicated TICs can have a resolution near 1 ps, making it possible in theory to detect a frequency change of 1×10^{-12} in 1 s. However, the ability of a time interval system to detect small frequency changes is limited by factors other than TIC resolution, including trigger errors, time base errors, and noise from frequency dividers, and 1×10^{-11} at 1 s usually represents best case performance [1,14,15].

The time interval method is widely implemented and is an excellent way to measure long-term accuracy and stability. However, systems with higher resolution are needed to measure short-term stability.




FIGURE 42.17 Simple heterodyne measurement system.

These systems continue to use counters but employ heterodyne techniques, rather than frequency dividers, to obtain the necessary low-frequency signals. The reason for this is straightforward: When dividers are used, both the frequency of the DUT and its phase fluctuations (noise) are divided by the same amount, and thus, no resolution is gained. The resolution of the system remains the same as the resolution of the counter. In contrast, heterodyne techniques convert the DUT frequency to a lower frequency, known as the intermediate frequency (IF) or *beat frequency*, without dividing the device's phase fluctuations (noise). Thus, phase information is preserved in the heterodyne process, and when the DUT changes frequency relative to the reference, the beat frequency changes by the same amount.

The resolution of the counter is potentially improved by the downconversion or *heterodyne factor*, which equals $f_{\text{nom}}/f_{\text{beat}}$. Unlike divider systems, which typically produce 1 Hz signals, the beat frequency is usually not lower than 10 Hz, because it must be large enough to account for a wide range of possible frequency differences between the DUT and reference, including situations where a quartz oscillator is involved in the comparison.

Figure 42.17 shows a simple heterodyne system where a low-pass filter (LPF) and double-balanced mixer are used in conjunction with a universal counter in frequency or period mode. Although not shown in the figure, in most designs, the reference oscillator is intentionally offset in frequency, normally by 100 Hz or 1 kHz, by use of a frequency synthesizer. The double-balanced mixer mixes the reference frequency with the DUT signal, producing a beat frequency that is nominally the same as the intentional frequency offset of the reference. If a 10 MHz frequency is downconverted to 1 kHz, the resolution of the counter can be potentially enhanced by a heterodyne factor of 10^4 ($10^7/10^3$). The beat frequency is sent through a LPF to remove carrier frequency harmonics and then measured with either an FC or a period counter. If an FC is used, the frequency offset $f_o = (f_{\text{ic}} - f_{\text{beat}})/f_{\text{nom}}$, where f_{ic} is the reading on the FC, f_{beat} is the nominal beat frequency, and f_{nom} is the nominal frequency of the DUT.

The *dual-mixer time difference* (DMTD) method (Figure 42.18) combines the best features of the time interval and heterodyne methods. The DMTD method uses a transfer oscillator and two double-balanced mixers in parallel. The transfer oscillator is offset from the nominal frequencies by a small amount, typically by 10 Hz. The transfer oscillator does not have to be particularly stable or accurate because the noise it produces is common to both measurement channels and cancels when the phase difference is computed. Even so, the transfer oscillator signal is usually obtained by locking a frequency synthesizer to the reference oscillator. Typically, both the DUT and the reference must have the same nominal frequency, usually 5 or 10 MHz, and DMTD systems are often designed to measure only one nominal frequency. However, the use of frequency synthesizers allows some DMTD systems to measure a number of DUT frequencies within a specified range (e.g., 1–30 MHz) and can even allow the DUT and the reference to have different nominal frequencies.

The transfer oscillator signal is heterodyned with both the reference oscillator and the DUT to produce two beat frequencies. The two beat frequencies are out of phase by an amount proportional to the time difference between the DUT and reference, and this phase difference is measured with a TIC. Before being sent to the TIC, the beat frequencies pass through a LPF that removes harmonics

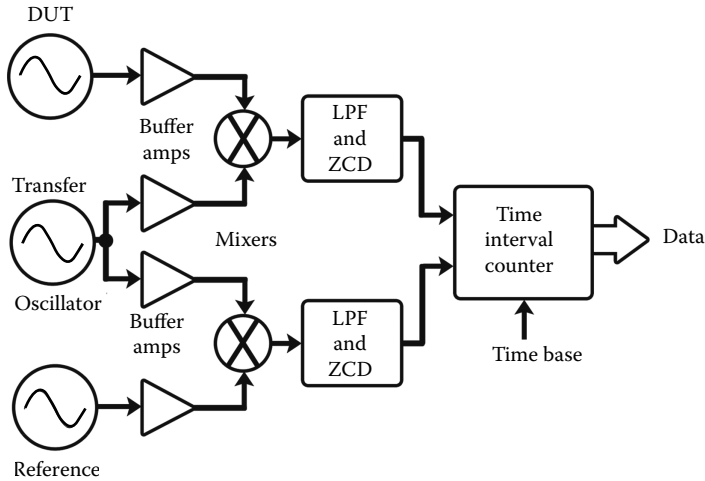


FIGURE 42.18 DMTD measurement system.

and a *zero-crossing detector* (ZCD) that determines the zero crossing of each beat frequency cycle. A DMTD system works best if the zero crossings are coincident or can be interpolated to a common epoch. us, a variable phase shifter is sometimes used to put the DUT signal nominally in phase with the reference before the DUT signal is mixed with the transfer oscillator. A er the mixing process, some systems use an event counter to count the whole beat note cycles to eliminate the ambiguity of the zero crossings; other systems time-tag the zero crossings. e resolution of a DMTD system is determined by the resolution of the TIC or the time-tagging hardware, divided by the heterodyne factor. For example, if the TIC resolution is 100 ns and the heterodyne factor is 10^6 (based on a 10 MHz nominal frequency and a 10 Hz beat frequency), then the DMTD resolution can be estimated at $10^{-7} \text{ s}/10^6$ or 0.1 ps. is high resolution allows some DMTD systems to detect frequency changes of 1×10^{-13} in 1 s or two or three orders of magnitude better than a typical universal counter. It should be noted that all components of heterodyne or DMTD systems must be carefully chosen for optimum stability and not all frequency measurement systems are stable enough to support their theoretical resolution [1,5,14,16,17].

Recently developed frequency measurement systems employ digital signal processing (DSP) techniques to achieve even higher resolutions. ese systems rapidly sample the DUT and reference signals with analog-to-digital converters and then t the digitized signals with a sinusoidal waveform at an IF that is much lower than the carrier frequency. Figure 42.19 shows the relationship between the oscillator signal, the sampling clock, and the resulting IF sine wave. Again, the idea is simply to compare two

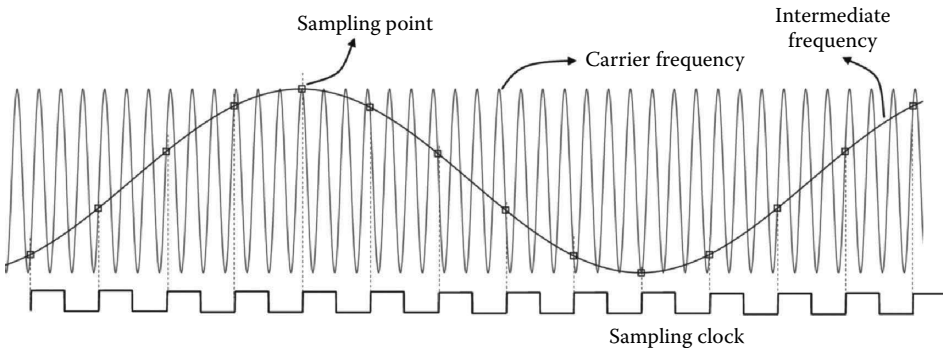


FIGURE 42.19 Digital sampling of the carrier frequency.

low-frequency signals, but note that the DSP hardware and software replaces the analog mixers shown in Figures 42.17 and 42.18 or the dividers shown in Figure 42.16. After two low-frequency signals are obtained for comparison, the subsequent filtering and phase detection is also accomplished with DSP techniques.

At this writing (2011), the available digital sampling frequency measurement systems are relatively expensive but can outperform DMTD systems by a wide margin. The best systems have a resolution of a few femtoseconds (10^{-15} s) and can detect frequency changes of less than 5×10^{-15} at 1 s. These systems can also measure phase noise, can simultaneously estimate time and frequency domain statistics with the same hardware, and can measure oscillators with a wide range of nominal frequencies [18].

To provide a practical, “real-world” comparison of the most commonly used frequency measurement systems, Figure 42.20 compares the FC, time interval, single-mixer, and DMTD methods by showing an ADEV graph of the same OCXO measured with all four systems. The same universal counter, with an external time base connection to an atomic oscillator, was used for the first three methods. The counter chosen was typical of those found in many metrology laboratories, with a 12-digit display and a time interval resolution of 150 ps. For the FC method, the counter was configured as an FC and used by itself. For the time interval method, the counter was configured as a TIC and used with a low-cost divider circuit. For the mixer method, the counter was used in conjunction with a low-cost double-balanced mixer and a frequency synthesizer that generated a signal offset by 1 kHz from the reference. A commercially available system was used for the DMTD measurement.

Only the DMTD system was able to correctly show that the OCXO reaches its noise floor at an averaging period of 10–20 s, when its stability briefly drops below 3×10^{-12} . Due to a lack of resolution and due to measurement system noise that exceeds the oscillator noise at short intervals, the other three systems incorrectly indicate that the OCXO noise floor is not reached until after several hundred seconds of averaging, when a random walk noise process had already begun. The FC and TIC methods provide similar results at all intervals, but the TIC estimates are smoother and have fewer variations. The single-mixer system was not optimized for best performance, but was somewhat quieter than the FC and TIC systems at short averaging periods. It is interesting to note that when the averaging interval exceeds several hundred seconds, all four methods produce similar results. This demonstrates that the choice of a measurement system is not critical when measuring long-term stability or accuracy, but is critical when measuring short-term stability.

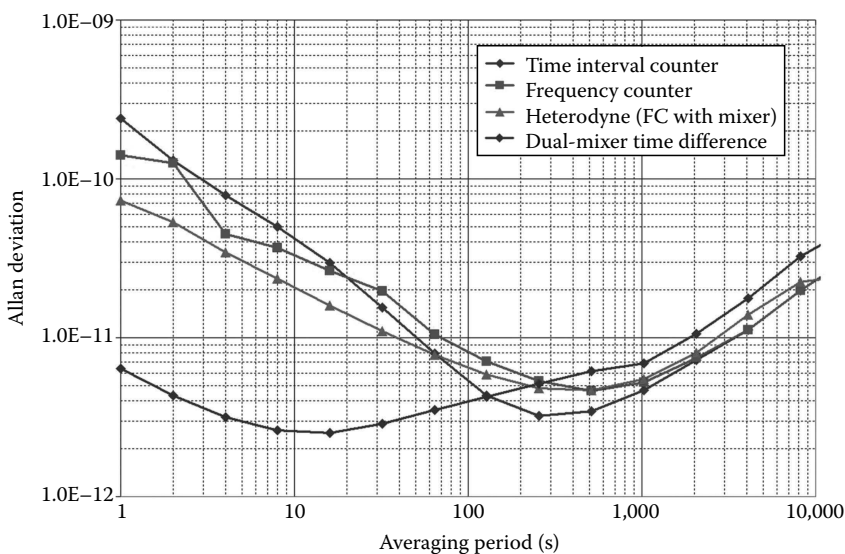


FIGURE 42.20 ADEV graph comparing different measurement methods.

42.4 Future Developments

Cesium fountain standards currently (2011) serve as the national standards of frequency at NIST and a number of other laboratories. These devices implement techniques such as a vertical microwave cavity and laser cooling to increase the interrogation period of the cesium atoms. By doing so, they reduce the resonance width and increase the oscillator Q to about 10^{10} . As a result, their accuracy and stability is about a factor of 100 times better than commercial cesium standards [19]. Perhaps cesium fountains or another type of laser-cooled microwave frequency standard will eventually be sold commercially.

It seems almost certain that optical atomic standards will have a huge impact on the future of frequency metrology. Optical standards operate at much higher resonance frequencies than microwave standards; the stabilized lasers that serve as their resonators typically operate at a frequency near 10^{15} Hz, as opposed to less than 10^{10} Hz for cesium. As a result, these standards promise accuracies and stabilities that are several orders of magnitude better than the best microwave standards. Optical frequency standards have been constructed at NIST utilizing single-ion techniques based on mercury ($^{199}\text{Hg}^+$) and aluminum ($^{27}\text{Al}^+$), as well as neutral atom techniques based on calcium (^{40}Ca), ytterbium (^{174}Yb), and strontium (^{87}Sr). It appears likely that the SI second will eventually be redefined, with the new definition based on one of these optical atomic transitions [20].

The accurate measurement of optical frequencies had historically been difficult, involving large and complex chains of frequency-doubled and frequency-mixed lasers. This changed in the early part of the twenty-first century, with the advent of the self-referenced femtosecond laser frequency combs. These devices, which are now sold commercially, have made the linkage of optical to microwave frequencies, or optical to optical frequencies, relatively simple. A self-referenced femtosecond laser frequency comb generates a series of discrete, equally frequency-spaced modes. The mode spacing, f_{rep} , is given by the repetition rate of a mode-locked laser, which is typically around 1 GHz. The frequency of an individual mode can be expressed as $f(m) = f_0 + m \times f_{\text{rep}}$, where m is an integer. The frequency offset f_0 is measured by a method called self-referencing. An optical standard can be compared to a conventional frequency standard by referencing the comb to a cesium standard or a GPSDO and then measuring the heterodyne beat frequency f_{beat} between the optical standard and the nearest tooth of the comb. This will lead to determination of the frequency of the optical standard in terms of f_0 , f_{rep} , and f_{beat} , if the integer m is known. Frequency combs also allow the ratio of two optical frequencies to be compared without being limited by the accuracy of conventional standards [21]. Frequency comb comparisons of optical standards have successfully measured relative frequency differences of parts in 10^{19} , an astounding 16 orders of magnitude smaller than the requirements that existed less than a century ago for the measurements of radio frequencies.

References

1. G. Kamas and S. L. Howe, eds., *Time and Frequency User's Manual*, National Bureau of Standards Special Publication 559, Washington, DC, 256p, November 1979.
2. C. Audoin and B. Guinot, *The Measurement of Time: Time, Frequency, and the Atomic Clock*, Cambridge University Press, Cambridge, U.K., 335p, 2001.
3. J. Jespersen, Introduction to the time domain characterization of frequency standards, *Proceedings of the 25th Annual Precise Time and Time Interval (PTTI) Meeting*, Pasadena, CA, pp. 83–102, December 1991.
4. IEEE, *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*, IEEE Std. 1139, New York, 1999.
5. W. J. Riley, *Handbook of Frequency Stability Analysis*, National Institute of Standards and Technology Special Publication 1065, Gaithersburg, MD, 136p, July 2008.
6. H. Hellwig, *Frequency Standards and Clocks: A Tutorial Introduction*, National Bureau of Standards Technical Note 616, Washington, DC, 72p, June 1977.
7. L. Lewis, An introduction to frequency standards, *Proc. IEEE*, 79(7), 927–935, 1991.

8. F. G. Major, *The Quantum Beat: The Physical Principles of Atomic Clocks*, Springer-Verlag, New York, 475p, 1998.
9. J. R. Vig and A. Ballato, Frequency control devices, in *Ultrasonic Instruments and Devices*, Chapter 7, Academic Press, New York, pp. 637–701, 1999.
10. J. Vanier and C. Audoin, *The Quantum Physics of Atomic Frequency Standards*, 2 Vols., Adam Hilger, Bristol, England, 1989.
11. L. Cutler, Fifty years of commercial caesium clocks, *Metrologia*, 42, S90–S99, June 2005.
12. J. A. Kusters, L. S. Cutler, E. D. Powers, Long-term experience with cesium beam frequency standards, *Proceedings of the 1999 IEEE Frequency Control Symposium and European Frequency and Time Forum (EFTF)*, Besançon, France, pp. 159–163, April 1999.
13. M. A. Lombardi, The use of GPS disciplined oscillators as primary frequency standards for calibration and metrology laboratories, *Measure: J. Meas. Sci.*, 3(3), 56–65, September 2008.
14. D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, Eds., *Characterization of Clocks and Oscillators*, National Institute of Standards and Technology, Washington, DC, Technical Note 1337, 352p, March 1990.
15. Hewlett-Packard Corporation, *Fundamentals of the Electronic Counters*, Hewlett-Packard Application Note 200, 5965-7660E, Palo Alto, CA, 44p, March 1997.
16. Hewlett-Packard Corporation, *Simplify Frequency Stability Measurements with Built-in Allan Variance Analysis*, Hewlett-Packard Application Note 358-12, 5952-8006, Palo Alto, CA, 32p, January 1990.
17. J. Levine, Introduction to time and frequency metrology, *Rev. Sci. Instrum.*, 70(6), 2567–2596, June 1999.
18. S. R. Stein, J. Grove, J. Hein, R. A. Martin, B. Mesander, M. Phillips, P. Schweiger, J. Shima, and W. Solbrig, Comparison of heterodyne and direct-sampling techniques for phase-difference measurements, *Proceedings of the 2005 NCSLI International Workshop and Symposium*, Washington, DC, 10p, July 2005.
19. M. A. Lombardi, T. P. Heavner, and S. R. Jefferts, NIST primary frequency standards and the realization of the SI second, *Measure: J. Meas. Sci.*, 2(4), 74–89, December 2007.
20. L. Lorini, N. Ashby, A. Brusch, S. A. Diddams, R. E. Drullinger, E. Eason, T. Fortier et al., Recent atomic clock comparisons at NIST, *Eur. Phys. J. Special Topics*, 35, 19–35, 2008.
21. S. Diddams, The evolving optical frequency comb, *J. Opt. Soc. Am. B*, 27(11), B51–B62, November 2010.