

19

Frequency Measurement

- 19.1 Overview of Frequency Measurements and Calibration
- 19.2 The Specifications: Frequency Uncertainty and Stability
Frequency Uncertainty • Stability
- 19.3 Frequency Standards
Quartz Oscillators • Atomic Oscillators
- 19.4 Transfer Standards
WWVB • LORAN-C • Global Positioning System (GPS)
- 19.5 Calibration Methods
- 19.6 Future Developments

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 $f = 1/T$. The International System of Units (SI) states that the period should always be expressed in units of seconds (s), and the frequency should always be expressed in hertz (Hz). The frequency of electric signals often is measured in units of kilohertz (kHz) or megahertz (MHz), where 1 kHz equals 1000 (10^3) cycles per second, and 1 MHz equals 1 million (10^6) cycles per second.

Average frequency over a time interval can be measured very precisely. Time interval is one of the four basic standards of measurement (the others are length, mass, and temperature). Of these four basic standards, time interval (and frequency) can be measured with the most resolution and the least uncertainty. In some fields of metrology, 1 part per million (1×10^{-6}) is considered quite an accomplishment. In frequency metrology, measurements of 1 part per billion (1×10^{-9}) are routine, and even 1 part per trillion (1×10^{-12}) is commonplace.

Devices that produce a known frequency are called *frequency standards*. These devices must be calibrated so that they remain within the tolerance required by the user's application. The discussion begins with an overview of frequency calibrations.

19.1 Overview of Frequency Measurements and Calibration

Frequency calibrations measure the performance of frequency standards. The frequency standard being calibrated is called the *device under test* (DUT). In most cases, the DUT is a *quartz*, *rubidium*, or *cesium* oscillator. In order to perform the calibration, the DUT must be compared to a *standard* or *reference*. The standard should outperform the DUT by a specified ratio in order for the calibration to be valid. This ratio is called the *test uncertainty ratio* (TUR). A TUR of 10:1 is preferred, but not always possible. If a smaller TUR is used (5:1, for example), then the calibration will take longer to perform.

Once the calibration is completed, the metrologist should be able to state how close the DUT's output is to its *nameplate frequency*. The nameplate frequency is labeled on the oscillator output. For example, a DUT with an output labeled "5 MHz" is supposed to produce a 5 MHz frequency. The calibration measures the difference between the actual frequency and the nameplate frequency. This difference is called the *frequency offset*. There is a high probability that the frequency offset will stay within a certain range of values, called the *frequency uncertainty*. The user specifies a frequency offset and an associated uncertainty requirement that the DUT must meet or exceed. In many cases, users base their requirements on the specifications published by the manufacturer. In other cases, they may "relax" the requirements and use a less demanding specification. Once the DUT meets specifications, it has been successfully calibrated. If the DUT cannot meet specifications, it fails calibration and is repaired or removed from service.

The reference used for the calibration must be *traceable*. The International Organization for Standardization (ISO) definition for *traceability* is:

The property of the result of a measurement or the value of a standard whereby it can be related to stated references, usually national or international standards, through an unbroken chain of comparisons all having stated uncertainties [1].

In the United States, the "unbroken chain of comparisons" should trace back to the National Institute of Standards and Technology (NIST). In some fields of calibration, traceability is established by sending the standard to NIST (or to a NIST-traceable laboratory) for calibration, or by sending a set of reference materials (such as a set of artifact standards used for mass calibrations) to the user. Neither method is practical when making frequency calibrations. Oscillators are sensitive to changing environmental conditions and especially to being turned on and off. If an oscillator is calibrated and then turned off, the calibration could be invalid when the oscillator is turned back on. In addition, the vibrations and temperature changes encountered during shipment can also change the results. For these reasons, laboratories should make their calibrations on-site.

Fortunately, one can use *transfer standards* to deliver a frequency reference from the national standard to the calibration laboratory. Transfer standards are devices that receive and process radio signals that provide frequency traceable to NIST. The radio signal is a link back to the national standard. Several different types of signals are available, including NIST radio stations WWV and WWVB, and radionavigation signals from LORAN-C and GPS. Each signal delivers NIST traceability at a known level of uncertainty. The ability to use transfer standards is a tremendous advantage. It allows traceable calibrations to be made simultaneously at a number of sites as long as each site is equipped with a radio receiver. It also eliminates the difficult and undesirable practice of moving frequency standards from one place to another.

Once a traceable transfer standard is in place, the next step is developing the technical procedure used to make the calibration. This procedure is called the *calibration method*. The method should be defined and documented by the laboratory, and ideally a measurement system that automates the procedure should be built. ISO/IEC Guide 25, *General Requirements for the Competence of Calibration and Testing Laboratories*, states:

The laboratory shall use appropriate methods and procedures for all calibrations and tests and related activities within its responsibility (including sampling, handling, transport and storage, preparation of items, estimation of uncertainty of measurement, and analysis of calibration and/or test data). They shall be consistent with the accuracy required, and with any standard specifications relevant to the calibrations or test concerned.

In addition, Guide 25 states:

The laboratory shall, wherever possible, select methods that have been published in international or national standards, those published by reputable technical organizations or in relevant scientific texts or journals [2,3].

Calibration laboratories, therefore, should automate the frequency calibration process using a well-documented and established method. This helps guarantee that each calibration will be of consistently high quality, and is essential if the laboratory is seeking ISO registration or laboratory accreditation.

Having provided an overview of frequency calibrations, we can take a more detailed look at the topics introduced. This chapter begins by looking at the specifications used to describe a frequency calibration, followed by a discussion of the various types of frequency standards and transfer standards. A discussion of some established calibration methods concludes this chapter.

19.2 The Specifications: Frequency Uncertainty and Stability

This section looks at the two main specifications of a frequency calibration: uncertainty and stability.

Frequency Uncertainty

As noted earlier, a frequency calibration measures whether a DUT meets or exceeds its uncertainty requirement. According to ISO, uncertainty is defined as:

Parameter, associated with the result of a measurement, that characterizes the dispersion of values that could reasonably be attributed to the measurand [1].

When we make a frequency calibration, the measurand is a DUT that is supposed to produce a specific frequency. For example, a DUT with an output labeled 5 MHz is supposed to produce a 5 MHz frequency. Of course, the DUT will actually produce a frequency that is not exactly 5 MHz. After calibration of the DUT, we can state its frequency offset and the associated uncertainty.

Measuring the frequency offset requires comparing the DUT to a reference. This is normally done by making a *phase comparison* between the frequency produced by the DUT and the frequency produced by the reference. There are several calibration methods (described later) that allow this determination. Once the amount of phase deviation and the *measurement period* are known, we can estimate the frequency offset of the DUT. The measurement period is the length of time over which phase comparisons are made. Frequency offset is estimated as follows, where Δt is the amount of phase deviation and T is the measurement period:

$$f(\text{offset}) = \frac{-\Delta t}{T} \quad (19.1)$$

If we measure $+1 \mu\text{s}$ of phase deviation over a measurement period of 24 h, the equation becomes:

$$f(\text{offset}) = \frac{-\Delta t}{T} = \frac{1 \mu\text{s}}{86,400,000,000 \mu\text{s}} = -1.16 \times 10^{-11} \quad (19.2)$$

The smaller the frequency offset, the closer the DUT is to producing the same frequency as the reference. An oscillator that accumulates $+1 \mu\text{s}$ of phase deviation per day has a frequency offset of about -1×10^{-11} with respect to the reference. Table 19.1 lists the approximate offset values for some standard units of phase deviation and some standard measurement periods.

The frequency offset values in Table 19.1 can be converted to units of frequency (Hz) if the nameplate frequency is known. To illustrate this, consider an oscillator with a nameplate frequency of 5 MHz that is high in frequency by 1.16×10^{-11} . To find the frequency offset in hertz, multiply the nameplate frequency by the dimensionless offset value:

$$\left(5 \times 10^6\right) \left(+1.16 \times 10^{-11}\right) = 5.80 \times 10^{-5} = +0.0000580 \text{ Hz} \quad (19.3)$$

TABLE 19.1 Frequency Offset Values for Given Amounts of Phase Deviation

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

The nameplate frequency is 5 MHz, or 5,000,000 Hz. Therefore, the actual frequency being produced by the frequency standard is:

$$5,000,000 \text{ Hz} + 0.0000580 \text{ Hz} = 5,000,000.0000580 \text{ Hz} \quad (19.4)$$

To do a complete uncertainty analysis, the measurement period must be long enough to ensure that one is measuring the frequency offset of the DUT, and that other sources are not contributing a significant amount of uncertainty to the measurement. In other words, must be sure that Δt is really a measure of the DUT's phase deviation from the reference and is not being contaminated by noise from the reference or the measurement system. This is why a 10:1 TUR is desirable. If a 10:1 TUR is maintained, many frequency calibration systems are capable of measuring a 1×10^{-10} frequency offset in 1 s [4].

Of course, a 10:1 TUR is not always possible, and the simple equation given for frequency offset (Equation 19.1) is often too simple. When using transfer standards such as LORAN-C or GPS receivers (discussed later), radio path noise contributes to the phase deviation. For this reason, a measurement period of at least 24 h is normally used when calibrating frequency standards using a transfer standard. This period is selected because changes in path delay between the source and receiver often have a cyclical variation that averages out over 24 h. In addition to averaging, curve-fitting algorithms and other statistical methods are often used to improve the uncertainty estimate and to show the *confidence level* of the measurement [5]. [Figure 19.1](#) shows two simple graphs of phase comparisons between a DUT and a reference frequency. Frequency offset is a linear function, and the top graph shows no discernible variation in the linearity of the phase. This indicates that a TUR of 10:1 or better is being maintained. The bottom graph shows a small amount of variation in the linearity, which could mean that a TUR of less than 10:1 exists and that some uncertainty is being contributed by the reference.

To summarize, frequency offset indicates how accurately the DUT produces its nameplate frequency. Notice that the term *accuracy* (or *frequency accuracy*) often appears on oscillator specification sheets instead of the term frequency offset, since frequency accuracy and frequency offset are nearly equivalent terms. Frequency uncertainty indicates the limits (upper and lower) of the measured frequency offset. Typically, a 2σ uncertainty test is used. This indicates that there is a 95.4% probability that the frequency offset will remain within the stated limits during the measurement period. Think of frequency offset as the result of a measurement made at a given point in time, and frequency uncertainty as the possible dispersion of values over a given measurement period.

Stability

Before beginning a discussion of *stability*, it is important to mention a distinction between frequency offset and stability. Frequency offset is a measure of how well an oscillator produces its nameplate frequency, or how well an oscillator is adjusted. It does not tell us about the inherent quality of an oscillator. For example, a high-quality oscillator that needs adjustment could produce a frequency with




FIGURE 19.1 Simple phase comparison graphs.




FIGURE 19.2 Comparison of unstable and stable frequencies.

a large offset. A low-quality oscillator may be well adjusted and produce (temporarily at least) a frequency very close to its nameplate value.

Stability, on the other hand, indicates how well an oscillator can produce the same frequency over a given period of time. It does not indicate whether the frequency is “right” or “wrong,” only whether it stays the same. An oscillator with a large frequency offset could still be very stable. Even if one adjusts the oscillator and moves it closer to the correct frequency, the stability usually does not change. Figure 19.2 illustrates this by displaying two oscillating signals that are of the same frequency between t_1 and t_2 . However, it is clear that signal 1 is unstable and is fluctuating in frequency between t_2 and t_3 .

Stability is defined as the statistical estimate of the frequency fluctuations of a signal over a given time interval. *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. A typical oscillator specification sheet might list stability estimates for intervals of 1, 10, 100, and 1000 s [6, 7].

Stability estimates can be made in the *frequency domain* or *time domain*, but time domain estimates are more widely used. To estimate stability in the time domain, one must start with a set of frequency offset measurements y_i that consists of individual measurements, y_1, y_2, y_3 , etc. Once this dataset is obtained, one needs to determine the dispersion or scatter of the y_i as a measure of oscillator noise. The larger the dispersion, or scatter, of the y_i , the greater the instability of the output signal of the oscillator.

Normally, classical statistics such as *standard deviation* (or *variance*, the square of the standard deviation) are used to measure dispersion. Variance is a measure of the numerical spread of a dataset with respect to the average or mean value of the data set. 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 is usually nonstationary, because it contains time-dependent noise contributed by the frequency offset. 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 changes each time a new measurement is added.

For these reasons, a nonclassical statistic is used to estimate stability in the time domain. This statistic is often called the *Allan variance*, but because it is actually the square root of the variance, its proper name is the *Allan deviation*. By recommendation of the Institute of Electrical and Electronics Engineers (IEEE), the Allan deviation is used by manufacturers of frequency standards as a standard specification for stability. The equation for the Allan deviation is:

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(M-1)} \sum_{i=1}^{M-1} (\bar{y}_{i+1} - \bar{y}_i)^2} \quad (19.5)$$

where M is the number of values in the y_i series, and the data are equally spaced in segments τ seconds long. Note that while classical deviation subtracts the mean from each measurement before squaring their summation, the Allan deviation 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 dataset, but instead imply that the individual measurements in the dataset are obtained by averaging.

Table 19.2 shows how stability is estimated. The first column is a series of phase measurements recorded at 1 s intervals. Each measurement is larger than the previous measurement. This indicates that the DUT is offset in frequency from the reference and this offset causes a phase deviation. By differencing the raw phase measurements, the phase deviations shown in the second column are obtained. The third column divides the phase deviation (Δt) by the 1 s measurement period to get the frequency offset. Since the phase deviation is about 4 ns^{-1} , it indicates that the DUT has a frequency offset of about 4×10^{-9} . The frequency offset values in the third column form the y_i data series. The last two columns show the first differences of the y_i and the squares of the first differences. Since the sum of the squares equals 2.2×10^{-21} , the equation (where $\tau = 1 \text{ s}$) becomes:

$$\sigma_y(\tau) = \sqrt{\frac{2.2 \times 10^{-21}}{2(9-1)}} = 1.17 \times 10^{-11} \quad (19.6)$$

TABLE 19.2 Using Phase Measurements to Estimate Stability (data recorded at 1 s intervals)

Phase Measurements (ns)	Phase Deviation (ns), Δt	Frequency Offset $\Delta f/1 \text{ s } (y_i)$	First Differences $(y_{i+1} - y_i)$	First Difference Squared $(y_{i+1} - y_i)^2$
3321.44	—	—	—	—
3325.51	4.07	4.07×10^{-9}	—	—
3329.55	4.04	4.04×10^{-9}	-3×10^{-11}	9×10^{-22}
3333.60	4.05	4.05×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3337.65	4.05	4.06×10^{-9}	$+2 \times 10^{-11}$	4×10^{-22}
3341.69	4.04	4.04×10^{-9}	-2×10^{-11}	4×10^{-22}
3345.74	4.05	4.05×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3349.80	4.06	4.06×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3353.85	4.05	4.05×10^{-9}	-1×10^{-11}	1×10^{-22}
3357.89	4.04	4.04×10^{-9}	-1×10^{-11}	1×10^{-22}

Using the same data, the Allan deviation for $\tau = 2 \text{ s}$ can be computed by averaging pairs of adjacent values and using these new averages as data values. For $\tau = 4 \text{ s}$, take the average of each set of four adjacent values and use these new averages as data values. More data must be acquired for longer averaging times. Keep in mind that the confidence level of a stability estimate improves as the averaging time increases. In the above example, we have eight samples for our $\tau = 1 \text{ s}$ estimate. However, one would have only two samples for an estimate of $\tau = 4 \text{ s}$. The confidence level of the estimate (1σ) can be roughly estimated as:

$$\frac{1}{\sqrt{M}} \times 100\% \quad (19.7)$$

In the example, where $M = 9$, the error margin is 33%. With just two samples, the estimate could be in error by as much as 70%. With 10^4 samples, the error margin is reduced to 1%.

A sample Allan deviation graph is shown in [Figure 19.3](#). It shows the stability improving as the measurement period gets longer. Part of this improvement is because measurement system noise becomes less of a factor as the measurement period gets longer. At some point, however, the oscillator will reach its *flicker floor*, and from a practical point of view, no further gains will be made by averaging. The flicker floor is the point where the white noise processes begin to be dominated by nonstationary processes such as frequency drift. Most quartz and rubidium oscillators reach their flicker floor at a measurement period of 10^3 s or less, but cesium oscillators might not reach their flicker floor for 10^5 s (more than 1 day). [Figure 19.3](#) shows a sample Allan deviation graph of a quartz oscillator that is stable to about 5×10^{-12} at 100 s and is approaching its flicker floor [8-10].

Do not confuse stability with frequency offset when reading a specifications sheet. For example, a DUT with a frequency offset of 1×10^{-8} might still reach a stability of 1×10^{-12} in 1000 s. This means that the output frequency of the DUT is stable, even though it is not particularly close to its nameplate frequency. To help clarify this point, [Figure 19.4](#) is a graphical representation of the relationship between frequency offset (accuracy) and stability.

19.3 Frequency Standards

Frequency standards all have an internal device that produces a periodic, repetitive event. This device is called the *resonator*. Of course, the resonator must be driven by an energy source. Taken together, the energy source and the resonator form an *oscillator*. There are two main types of oscillators used as frequency standards: quartz oscillators and atomic oscillators.




FIGURE 19.3 A sample Allan deviation graph.




FIGURE 19.4 The relationship between frequency uncertainty (accuracy) and stability.

Quartz Oscillators

Quartz crystal oscillators first appeared during the 1920s and quickly replaced pendulum devices as laboratory standards for time and frequency [11]. Today, more than 10^9 quartz oscillators are manufactured annually for applications ranging from inexpensive wristwatches and clocks to communications networks and space tracking systems [12]. However, calibration and standards laboratories usually calibrate only the more expensive varieties of quartz oscillators, such as those found inside electronic instruments (like frequency counters) or those designed as standalone units. The cost of a high-quality quartz oscillator ranges from a few hundred to a few thousand dollars.

The quartz crystal inside the oscillator can be made of natural or synthetic quartz, but all modern devices are made of synthetic material. The crystal serves as a mechanical resonator that creates an oscillating voltage due to the *piezoelectric effect*. This effect causes the crystal to expand or contract as voltages are applied. The crystal has a *resonance frequency* that is determined by its physical dimensions and the type of crystal used. 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 frequency or a multiple




FIGURE 19.5 Block diagram of quartz oscillator.

of that frequency. Figure 19.5 is a simplified circuit diagram that shows the basic elements of a quartz oscillator. The amplifier provides the energy needed to sustain oscillation.

Quartz oscillators are sensitive to environmental parameters such as temperature, humidity, pressure, and vibration [12, 13]. When environmental parameters change, the fundamental resonance frequency also changes. There are several types of quartz oscillator designs that attempt to reduce the environmental problems. The *oven-controlled crystal oscillator* (OCXO) encloses the crystal in a temperature-controlled chamber called an oven. When an OCXO is first 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 alternative solution to the temperature problem is the *temperature-compensated crystal oscillator* (TCXO). In a TCXO, the output signal from a special temperature sensor (called a thermistor) generates a correction voltage that is applied to a voltage-variable reactance (called a 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 much less expensive. Therefore, TCXOs are normally used in small, usually portable units when high performance over a wide temperature range is not required. A third type of quartz oscillator is the *microcomputer-compensated crystal oscillator* (MCXO). The MCXO uses a microprocessor and compensates for temperature using digital techniques. The MCXO falls between a TCXO and an OCXO in both price and performance.

All quartz oscillators are subject to *aging*, which is defined as “a systematic change in frequency with time due to internal changes in the oscillator.” Aging is usually observed as a nearly linear change in the resonance frequency. Aging can be positive or negative and, occasionally, a reversal in aging direction is observed. Often, the resonance frequency decreases, which may indicate that the crystal is getting larger. Aging has many possible causes, including contamination of the crystal due to deposits of foreign material, changes in the oscillator circuitry, or changes in the quartz material or crystal structure. The vibrating motion of the crystal can also contribute to aging. High-quality quartz oscillators age at a rate of 5×10^{-9} per year or less.

In spite of the temperature and aging problems, the best OCXOs can achieve frequency offsets as small as 1×10^{-10} . Less expensive oscillators produce less impressive results. Small ovenized oscillators (such as those used as timebases in frequency counters) typically are offset in frequency by $\pm 1 \times 10^{-9}$ and cost just a few hundred dollars. The lowest priced quartz oscillators, such as those found in wristwatches and electronic circuits, cost less than \$1. However, because they lack temperature control, these oscillators have a frequency offset of about $\pm 1 \times 10^{-6}$ in the best case and may be offset as much as $\pm 1 \times 10^{-4}$.

Since the frequency offset of a quartz oscillator changes substantially over long periods of time, regular adjustments might be needed to keep the frequency within tolerance. For example, even the best quartz

oscillators need regular adjustments to maintain frequency within $\pm 1 \times 10^{-10}$. On the other hand, quartz oscillators have excellent short-term stability. An OCXO might be stable to 1×10^{-12} at 1 s. The limitations in short-term stability are mainly due to noise from electronic components in the oscillator circuits.

Atomic Oscillators

Atomic oscillators use the quantized energy levels in atoms and molecules as the source of their resonance frequency. The laws of quantum mechanics dictate that the energies of a bound system, such as an atom, have certain discrete values. An electromagnetic field 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) [14]:

$$f = \frac{E_2 - E_1}{h} \quad (19.8)$$

All atomic oscillators are *intrinsic standards*, since their frequency is inherently derived from a fundamental natural phenomenon. There are three main types of atomic oscillators: 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. Locking the quartz oscillator to the atomic frequency is advantageous. Most of the factors that degrade the long-term performance of a quartz oscillator disappear, because the atomic resonance frequency is much less sensitive to environmental conditions than the quartz resonance frequency. As a result, the long-term stability and uncertainty of an atomic oscillator are much better than those of a quartz oscillator, but the short-term stability is unchanged.

Rubidium Oscillators

Rubidium oscillators are the lowest priced members of the *atomic* oscillator group. They offer perhaps the best price/performance ratio of any oscillator. They perform much better than a quartz oscillator and cost much less than a cesium oscillator.

A rubidium oscillator operates at the resonance frequency of the rubidium atom (^{87}Rb), which is 6,834,682,608 Hz. This frequency is synthesized from a lower quartz frequency (typically 5 MHz) and the quartz frequency is steered by the rubidium resonance. The result is a very stable frequency source with the short-term stability of quartz, but much better long-term stability. Since rubidium oscillators are more stable than quartz oscillators, they can be kept within tolerance with fewer adjustments. The initial price of a rubidium oscillator (typically from \$3000 to \$8000) is higher than that of a quartz oscillator, but because fewer adjustments are needed, labor costs are reduced. As a result, a rubidium oscillator might be less expensive to own than a quartz oscillator in the long run.

The frequency offset of a rubidium oscillator ranges from 5×10^{-10} to 5×10^{-12} after its warm-up period. Maintaining frequency within $\pm 1 \times 10^{-11}$ can be done routinely with a rubidium oscillator but is impractical with even the best quartz oscillators. The performance of a well-maintained rubidium oscillator can approach the performance of a cesium oscillator, and a rubidium oscillator is much smaller, more reliable, and less expensive.

Cesium Oscillators

Cesium oscillators are *primary frequency standards* because the SI second is based on the resonance frequency of the cesium atom (^{133}Cs), which is 9,192,631,770 Hz. This means that a cesium oscillator that is working properly should be very close to its nameplate frequency without any adjustment, and there should be no change in frequency due to aging. The time scale followed by all major countries, Coordinated Universal Time (UTC), is derived primarily from averaging the performance of a large ensemble of cesium oscillators.

Cesium oscillators are the workhorses in most modern time and frequency distribution systems. The primary frequency standard for the United States is a cesium oscillator named NIST-7 with a frequency uncertainty of about $\pm 5 \times 10^{-15}$. Commercially available cesium oscillators differ in quality, but their frequency uncertainty should still be $\pm 5 \times 10^{-12}$ or less.

The two major drawbacks of cesium oscillators involve reliability and cost. Reliability is a major issue. The major component of a cesium oscillator, called the *beam tube*, has a life expectancy of about 3 years to 10 years. The beam tube is needed to produce the resonance frequency of the cesium atom, and this frequency is then used to discipline a quartz oscillator. When the beam tube fails, the cesium oscillator performs like an undisciplined quartz oscillator. For this reason, a cesium oscillator needs to be constantly monitored to make sure that it is still delivering a cesium-derived frequency. Cost is also a major issue. The initial purchase price of a cesium oscillator ranges from \$30,000 to \$80,000 and maintenance costs are high. The cost of a replacement beam tube is a substantial fraction of the cost of the entire oscillator. Laboratories that use cesium oscillators need to budget not only for their initial purchase, but for the cost of maintaining them afterward.

Hydrogen Masers

The *hydrogen maser* is the most elaborate and most expensive commercially available frequency standard. Few masers are built and most are owned by observatories or national standards laboratories. The word “maser” is an acronym that stands for Microwave Amplification by Stimulated Emission of Radiation. Masers derive their frequency from the resonance frequency of the hydrogen atom, which is 1,420,405,752 Hz.

There are two types of hydrogen masers. The first type, called an *active maser*, 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. The “passive” technique is also used by rubidium and cesium oscillators. Since active masers derive the output frequency more directly from the atomic resonance, they have better short-term stability than passive masers. However, both types of maser have better short-term stability than a cesium oscillator. Even so, the frequency uncertainty of a maser is still greater than that of a cesium oscillator because its operation is more critically dependent on a complex set of environmental conditions. Although the performance of a hydrogen maser is excellent, its cost is still very high, typically \$200,000 or more [15, 16].

Table 19.3 summarizes the characteristics of the different types of oscillators.

19.4 Transfer Standards

To briefly review, a frequency calibration compares the *device under test* (DUT) to a reference. The DUT is usually a quartz, rubidium, or cesium oscillator. The reference is an oscillator of higher performance than the DUT or a *transfer standard* that receives a radio signal. All transfer standards receive a signal that has a cesium oscillator at its source, and this signal delivers a cesium-derived frequency to the user. This benefits many users, because cesium oscillators are expensive both to buy and maintain, and not all calibration laboratories can afford them. Even if a laboratory already has a cesium oscillator, it still needs to check its performance. The only practical way to do this is by comparing the cesium to a transfer standard.

Transfer standards also provide *traceability*. Most transfer standards receive signals traceable to the national frequency standard maintained by NIST. Some signals, such as those transmitted by HF (high frequency) radio stations WWV and WWVH and the LF (low frequency) station WWVB, are traceable because they are directly controlled by NIST. Other signals, like the LORAN-C and Global Positioning System (GPS) satellite broadcasts, are traceable because their reference is regularly compared to NIST. Some signals broadcast from outside the United States are also considered to be traceable. This is because NIST compares its frequency standard to the standards maintained in other countries.

Some compromises are made when using a transfer standard. Even if the radio signal is referenced to a nearly perfect frequency, its performance is degraded as it travels along the radio path between the transmitter and receiver. To illustrate, consider a laboratory that has a rack-mounted frequency standard

TABLE 19.3 Summary of Oscillator Types

Oscillator type	Quartz (TCXO)	Quartz (MCXO)	Quartz (OCXO)	Rubidium	Cesium	Hydrogen maser
Primary standard	No	No	No	No	Yes	No
Intrinsic standard	No	No	No	Yes	Yes	Yes
Resonance frequency	Mechanical (varies)	Mechanical (varies)	Mechanical (varies)	6.834682608 GHz	9.19263177 GHz	1.420405752 GHz
Leading cause of failure	None	None	None	Rubidium lamp (15 years)	Cesium beam tube (3 to 10 years)	Hydrogen depletion (7 years +)
Power (W)	0.05	0.04	0.6	20	30	>100
Stability, $\sigma_y(\tau)$, $\tau = 1$ s	1×10^{-9}	3×10^{-10}	1×10^{-12}	5×10^{-11} – 5×10^{-12}	5×10^{-11} – 5×10^{-12}	$\leq 1 \times 10^{-12}$
Aging/year	5×10^{-7}	3×10^{-8}	5×10^{-9}	2×10^{-10}	None	$\leq 1 \times 10^{-13}$
Frequency Offset after Warm-up	1×10^{-6}	1×10^{-8}	1×10^{-8} – 1×10^{-10}	5×10^{-10} – 5×10^{-12}	5×10^{-12} – 1×10^{-14}	1×10^{-12} – 1×10^{-13}
Warm-up time	<10 s– 1×10^{-6}	<10 s– 1×10^{-8}	<5 min– 1×10^{-8}	<5 min– 5×10^{-10}	30 min– 5×10^{-12}	24 h– 1×10^{-12}
Cost	\$100	\$1000	\$2000	\$3000–\$8000	\$30,000–\$80,000	\$200,000–\$300,000

TABLE 19.4 Traceability Levels Provided by Various Transfer Standards

Transfer standard	Frequency uncertainty over 24 h measurement period (with respect to NIST)
HF receiver (WWV and WWVH)	$\pm 1 \times 10^{-7}$
LF receiver (LORAN-C and WWVB)	$\pm 1 \times 10^{-12}$
Global Positioning System (GPS) Receiver	$\pm 5 \times 10^{-13}$

that produces a 5 MHz signal. Metrologists need to use this signal on their work bench, so they run a length of coaxial cable from the frequency standard to their bench. The signal is delayed as it travels from the standard to the bench, but because the cable is of fixed length, the delay is constant. Constant delays do not change the frequency. The frequency that goes into one end of the cable is essentially the same frequency that comes out the other end. However, what if the cable length were constantly changing? This would generally cause the frequency to fluctuate. Over long periods of time, these fluctuations will average out, but the short-term frequency stability would still be very poor. This is exactly what happens when one uses a transfer standard. The “cable” is actually a radio path that might be thousands of kilometers in length. The length of the radio path is constantly changing and appears to introduce frequency fluctuations, even though the source of the frequency (a cesium oscillator) is not changing. This makes transfer standards unsuitable as a reference when making short-term stability measurements. However, transfer standards are well suited for measuring frequency uncertainty, because one can minimize and even eliminate the effects of these frequency fluctuations if one averages for a long enough measurement interval. Eventually, the performance of a cesium oscillator will be recovered.

Some radio signals have path variations that are so pronounced that they are not well suited for high-level frequency calibrations. To illustrate this, consider the signal broadcast from WWV, located in Fort Collins, CO. WWV is an HF radio station (often called a *shortwave* station) that transmits on 2.5, 5, 10, 15, and 20 MHz. WWV is referenced to the national frequency standard at NIST, but by the time the signal gets to the receiver, much of its potential performance has been lost. Most shortwave users receive the *skywave*, or the part of the signal that travels up to the ionosphere and is reflected back to Earth. Since the height of the ionosphere constantly changes, the path delay constantly changes, often by as much as 500 μs to 1000 μs . Because there is so much variability in the path, averaging leads to only limited improvement. Therefore, even though WWV is traceable to NIST, its frequency uncertainty is limited to $\pm 1 \times 10^{-7}$ when averaged for 1 day.

Other radio signals have more stable paths and much lower uncertainty values. Low-frequency (LF) radio stations (like NIST radio station WWVB and LORAN-C) can provide traceability to NIST with a frequency uncertainty of $\pm 1 \times 10^{-12}$ per day. An LF path is much more stable than an HF path, but still experiences a path delay change when the height of the ionosphere changes at sunrise and sunset. Currently, the most widely used signals originate from the Global Positioning System (GPS) satellites. GPS signals have the advantage of an unobstructed path between the transmitter and receiver. The frequency uncertainty of GPS is about $\pm 5 \times 10^{-13}$ /day. WWVB, LORAN-C, and GPS are described in the next three sections.

Table 19.4 shows some of the transfer standards available and the level of NIST traceability they provide when averaged for a measurement period of at least 24 h [17].

WWVB

Many countries broadcast time and frequency signals in the LF band from 30 kHz to 300 kHz, as well as in the VLF (very low frequency) band from 3 kHz to 30 kHz. Because part of the LF signal is *groundwave* and follows the curvature of the Earth, the path stability of these signals is acceptable for many applications. One such station is WWVB, which is operated by NIST. WWVB transmits on 60 kHz from the same site as WWV in Fort Collins, CO. The signal currently covers most of North America, and a power increase (6 dB, scheduled for 1998) will increase the coverage area and improve the signal-to-noise ratio within the United States.

Although far more stable than an HF path, the WWVB path length is influenced by environmental conditions along the path and by daily and seasonal changes. Path length is important because part of the signal travels along the ground (*groundwave*) and another part is reflected from the ionosphere (*skywave*). The groundwave path is far more stable than the skywave path. If the path is relatively short (less than 1000 km), the receiver might continuously track the groundwave signal because it always arrives first. For longer paths, a mixture of groundwave and skywave is received. And over a very long path, the groundwave might become so weak that it will only be possible to receive the skywave. In this instance, the path becomes much less stable.

The characteristics of an LF path are different at different times of day. During the daylight and nighttime hours, the receiver might be able to distinguish between groundwave and skywave, and path stability might vary by only a few hundred nanoseconds. However, if some skywave is being received, *diurnal* phase shifts will occur at sunrise and sunset. For example, as the path changes from all darkness to all daylight, the ionosphere lowers and the path gets shorter. The path length then stabilizes until either the transmitter or receiver enters darkness. At this point, the ionosphere rises and the path gets longer.

WWVB receivers have several advantages when used as a transfer standard. They are low cost and easy to use, and the received signals are directly traceable to NIST. With a good receiver and antenna system, one can achieve a frequency uncertainty of $\pm 1 \times 10^{-12}$ by averaging for 1 day [18].

LORAN-C

LORAN-C is a radionavigation system that operates in the LF band. Most of the system is operated by the U.S. Department of Transportation (DOT), but some stations are operated by other governments. The system consists of groups of stations called *chains*. Each chain has one master station and from two to five secondary stations. The stations operate at high power, typically 275 kW to 800 kW, and broadcast on a carrier frequency of 100 kHz using the 90 kHz to 110 kHz band.

Since all LORAN-C chains use the same carrier frequency, the chains transmit pulses so that individual stations can be identified. Each chain transmits a pulse group that includes pulses from all of the individual stations. The pulse group is sent at a unique Group Repetition Interval (GRI). For example, the 7980 chain transmits pulses every 79.8 ms. When the pulses leave the transmitter, they radiate in all directions. The *groundwave* travels parallel to the surface of the Earth. The *skywave* travels upward and is reflected off the ionosphere. The pulse shape was designed so that the receiver can distinguish between groundwave and skywave and lock onto the more stable groundwave signal. Most receivers stay locked to the groundwave by tracking the third cycle of the pulse. The third cycle was chosen for two reasons. First, it arrives early in the pulse so we know that it is groundwave. Second, it has more amplitude than the first and second cycles in the pulse, which makes it easier for the receiver to stay locked. Generally, a receiver within 1500 km of the transmitter can track the same groundwave cycle indefinitely and avoid skywave reception. The variations in groundwave path delay are typically quite small (<500 ns per day). However, if the path length exceeds 1500 km, the receiver might lose lock, and jump to another cycle of the carrier. Each cycle jump introduces a 10 μ s phase step, equal to the period of 100 kHz.

Commercially available LORAN-C receivers designed as transfer standards range in price from about \$3000 to \$10,000. These receivers typically provide several different frequency outputs (usually 10 MHz, 1 Hz, and the GRI pulse). LORAN-C navigation receivers are mass produced and inexpensive but do not provide a frequency output. However, it is sometimes possible to find and amplify the GRI pulse on the circuit board and use it as a reference frequency.

LORAN-C Performance

The frequency uncertainty of LORAN-C is degraded by variations in the radio path. The size of these variations depends on the signal strength, the distance between the receiver and the transmitter, the weather and atmospheric conditions, and the quality of the receiver and antenna. [Figure 19.6](#) shows the type of path variations one can typically expect. It shows a 100 s phase comparison between the LORAN-C

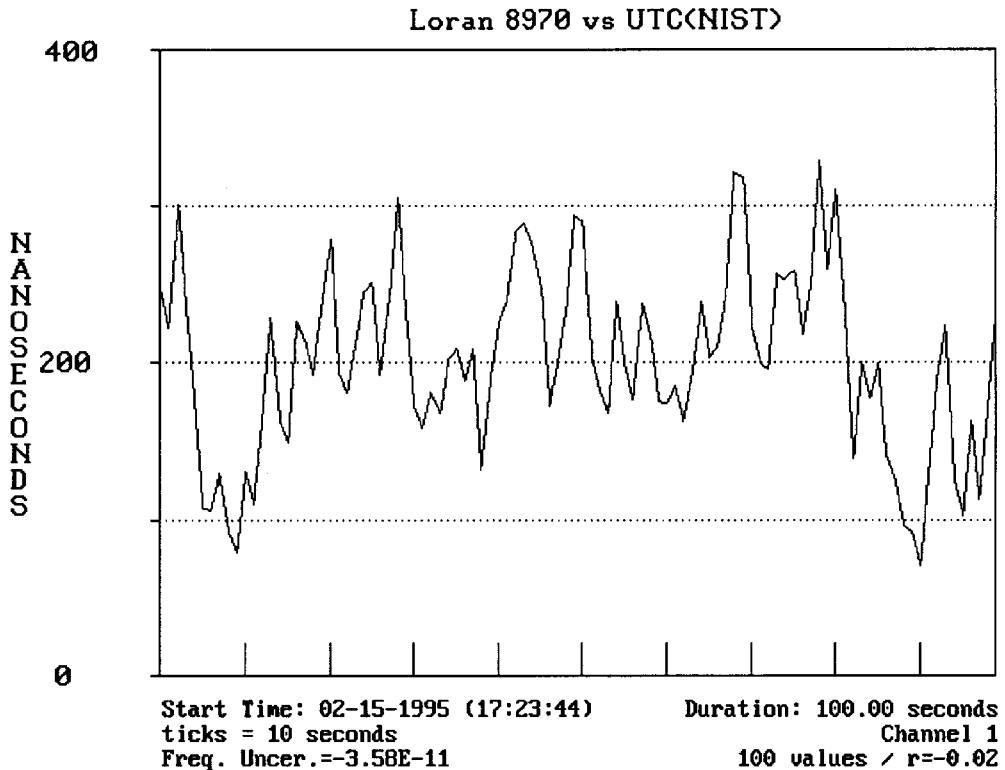


FIGURE 19.6 LORAN-C path variations during 100 s measurement period.

8970 chain as received in Boulder, CO, and the NIST national frequency standard. The 8970 signal is broadcast from the master station in Dana, IN, a site about 1512 km from Boulder. The frequency offset between the NIST frequency standard and the cesium standard at the LORAN-C transmitter would produce a phase shift of <1 ns per 100 s. Therefore, the graph simply shows the LORAN-C path noise or the difference in the arrival times of the GRI pulses due to path variations. The range of this path noise is about 200 ns.

The path variations cause the short-term stability of LORAN-C to be poor. However, because the path variations average out over time, the long-term stability is much better. This means that we can use LORAN-C to calibrate nearly any frequency standard if we average for a long enough interval. The better the frequency standard, the longer it takes to obtain a TUR approaching 10:1. For example, a quartz oscillator with a frequency uncertainty of 1×10^{-7} could be calibrated in 100 s. The quartz would accumulate 10 μ s of phase shift in 100 s, and this would completely hide the path noise shown in Figure 19.6. However, a rubidium oscillator with a frequency uncertainty of 1×10^{-11} would accumulate only 1 ns of phase shift in 100 s, or about 1 μ s per day. For this reason, a measurement period of at least 24 h is recommended when using LORAN-C to calibrate atomic oscillators. Figure 19.7 shows the results of a 96 h calibration of a cesium oscillator using LORAN-C. The thick line is a least squares estimate of the frequency offset. Although the path noise is clearly visible, the slope of the line indicates that the cesium oscillator is low in frequency by $\approx 3.4 \times 10^{-12}$.

Global Positioning System (GPS)

The Global Positioning System (GPS) is a radionavigation system developed and operated by the U.S. Department of Defense (DOD). The system consists of a constellation of 24 Earth orbiting satellites (21 primary satellites and 3 in-orbit spares). Each satellite carries four atomic frequency standards (two




FIGURE 19.7 LORAN-C compared to cesium oscillator over 96 h interval.

rubidiums and two cesiums) that are referenced to the United States Naval Observatory (USNO) and traceable to NIST. The 24 satellites orbit the Earth in six fixed planes inclined 55° from the equator. Each satellite is 20,200 km above the Earth and has an 11 h, 58 min orbital period, which means a satellite will pass over the same place on Earth 4 min earlier each day. Since the satellites continually orbit the Earth, GPS should be usable anywhere on the Earth's surface.

Each GPS satellite broadcasts 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* code (called the C/A code) with a chip rate of 1023 chips per millisecond. The second is a *precision* code (called the P code) with a chip rate of 10230 chips per millisecond. The C/A code repeats every millisecond. The P code repeats only every 267 days, but for practical reasons is reset every week. The C/A code is broadcast on L1, and the P code is broadcast on both L1 and L2 [19, 20].

For national security reasons, the DOD started the *Selective Availability* (SA) program in 1990. SA intentionally increases the positioning and timing uncertainty of GPS by adding about 300 ns of noise to both the C/A code and the P code. The resulting signal is distributed through the *Standard Positioning Service* (SPS). The SPS is intended for worldwide use, and can be used free of charge by anyone with a GPS receiver. The *Precise Positioning Service* (PPS) is available only to users authorized by the United States military. PPS users require a special receiver that employs cryptographic logic to remove the effects of SA. Because PPS use is restricted, nearly all civilian GPS users use the SPS.

GPS Receiving Equipment

At this writing (1996), GPS receivers have become common-place in the consumer electronics market, and some models cost \$200 or less. However, receivers suitable for use as a transfer standard are much

less common, and more expensive. Receivers range in price from less than \$100 for an OEM timing board, to \$20,000 or more for the most sophisticated models. The price often depends on the quality of the receiver's timebase oscillator. Lower priced models have a low-quality timebase that must be constantly steered to follow the GPS signal. Higher priced receivers have better timebases (some have internal rubidium oscillators) and can ignore many of the GPS path variations because their oscillator allows them to coast for longer intervals [21].

Most GPS receivers provide a 1 *pulse per second* (pps) output. Some receivers also provide a 1 kHz output (derived from the C/A code) and at least one standard frequency output (1, 5, or 10 MHz). To use these receivers, one simply mounts the antenna, connects the antenna to the receiver, and turns the receiver on. The antenna is often a small cone or disk (normally about 15 cm in diameter) that must be mounted outdoors where it has a clear, unobstructed view of the sky. Once the receiver is turned on, it performs a sky search to find out which satellites are currently above the horizon and visible from the antenna site. It then computes its three-dimensional coordinate (latitude, longitude, and altitude as long as four satellites are in view) and begins producing a frequency signal. The simplest GPS receivers have just one channel and look at multiple satellites using a sequencing scheme that rapidly switches between satellites. More sophisticated models have parallel tracking capability and can assign a separate channel to each satellite in view. These receivers typically track from 5 to 12 satellites at once (although more than 8 will only be available in rare instances). By averaging data from multiple satellites, a receiver can remove some of the effects of SA and reduce the frequency uncertainty.

GPS Performance

GPS has many technical advantages over LORAN-C. The signals are usually easier to receive, the equipment is often less expensive, and the coverage area is much larger. In addition, the performance of GPS is superior to that of LORAN-C. However, like all transfer standards, the short-term stability of a GPS receiver is poor (made worse by SA), and it lengthens the time required to make a calibration. As with LORAN-C, a measurement period of at least 24 h is recommended when calibrating atomic frequency standards using GPS.

To illustrate this, [Figure 19.8](#) shows a 100 s comparison between GPS and a cesium oscillator. The cesium oscillator has a frequency offset of $\approx 1 \times 10^{-13}$, and its accumulated phase shift during the 100 s measurement period is < 1 ns. Therefore, the noise on the graph can be attributed to GPS path variations and SA.

[Figure 19.9](#) shows a 1-week comparison between a GPS receiver and the same cesium oscillator used in [Figure 19.8](#). The range of the data is 550 ns. The thick line is a least squares estimate of the frequency offset. Although the GPS path noise is still clearly visible, one can see the linear trend contributed by the frequency offset of the cesium; this trend implies that the cesium oscillator is low in frequency by $\approx 5 \times 10^{-13}$.

19.5 Calibration Methods

To review, frequency standards are normally calibrated by comparing them to a traceable reference frequency. In this section, a discussion of how this comparison is made, is presented. To begin, look at the electric signal produced by a frequency standard. This signal can take several forms, as illustrated in [Figure 19.10](#). The dashed lines represent the supply voltage inputs (ranging from 5 V to 15 V for CMOS), and the bold solid lines represent the output voltage.

If the output frequency is an oscillating sine wave, it might look like the one shown in [Figure 19.11](#). This signal produces one cycle (2π radians of phase) in one period. Frequency calibration systems compare a signal like the one shown in [Figure 19.11](#) to a higher quality reference signal. The system then measures and records the change in phase between the two signals. If the two frequencies were exactly the same, their phase relationship would not change. Because the two frequencies are not exactly the same, their phase relationship will change; and by measuring the rate of change, one can determine the frequency offset of the DUT. Under normal circumstances, the phase changes in an orderly, predictable fashion. However, external factors like power outages, component failures, or human errors can cause a sudden




FIGURE 19.8 GPS compared to cesium oscillator over 100 s interval.

phase change, or *phase step*. A calibration system measures the total amount of phase shift (caused either by the frequency offset of the DUT or a phase step) over a given measurement period.

Figure 19.12 shows a phase comparison between two sinusoidal frequencies. The top sine wave represents a signal from the DUT, and the bottom sine wave represents a signal from the reference. Vertical lines have been drawn through the points where each sine passes through zero. The bottom of the figure shows “bars” that indicate the size of the phase difference between the two signals. If the phase relationship between the signals is changing, the phase difference will either increase or decrease to indicate that the DUT has a frequency offset (high or low) with respect to the reference. Equation 19.1 estimates the frequency offset. In Figure 19.12, both Δt and T are increasing at the same rate and the phase difference “bars” are getting wider. This indicates that the DUT has a constant frequency offset.

There are several types of calibration systems that can be used to make a phase comparison. The simplest type of system involves directly counting and displaying the frequency output of the DUT with a device called a *frequency counter*. This method has many applications but is unsuitable for measuring high-performance devices. The DUT is compared to the counter’s timebase (typically a TCXO), and the uncertainty of the system is limited by the performance of the timebase, typically $\approx 1 \times 10^{-8}$. Some counters allow use of an external timebase, which can improve the results. The biggest limitation is that frequency counters display frequency in hertz and have a limited amount of resolution. Detecting small changes in frequency could take many days or weeks, which makes it difficult or impossible to use this method to adjust a precision oscillator or to measure stability [22]. For this reason, most high-performance calibration systems collect time series data that can be used to estimate both frequency uncertainty and stability. A discussion follows on how phase comparisons are made using two different methods, the *time interval method* and the *dual mixer time difference method* [9, 10].

The time interval method uses a device called a *time interval counter* (TIC) to measure the time interval between two signals. A TIC has inputs for two electrical signals. One signal starts the counter and the




FIGURE 19.9 GPS compared to cesium oscillator over 1 week interval.




FIGURE 19.10 Oscillator outputs.

other signal stops it. If the two signals have the same frequency, the time interval will not change. If the two signals have different frequencies, the time interval will change, although usually very slowly. By looking at the rate of change, one can calibrate the device. It is exactly as if there were two clocks. By reading them each day, one can determine the amount of time one clock gained or lost relative to the other clock. It takes two time interval measurements to produce any useful information. By subtracting the first measurement from the second, one can tell whether time was gained or lost.

TICs differ in specification and design details, but they all contain several basic parts known as the *timebase*, the *main gate*, and the *counting assembly*. The timebase provides evenly spaced pulses used to measure time interval. The timebase is usually an internal quartz oscillator that can often be phase locked to an external reference. It must be stable because timebase errors will directly affect the measurements.




FIGURE 19.11 An oscillating signal.




FIGURE 19.12 Two signals with a changing phase relationship.

The main gate controls the time at which the count begins and ends. Pulses that pass through the gate are routed to the counting assembly, where they are displayed on the TIC's front panel or read by computer. The counter can then be reset (or *armed*) to begin another measurement. The stop and start inputs are usually provided with level controls that set the amplitude limit (or *trigger level*) at which the counter responds to input signals. If the trigger levels are set improperly, a TIC might stop or start when it detects noise or other unwanted signals and produce invalid measurements.




FIGURE 19.13 Measuring time interval.

Figure 19.13 illustrates how a TIC measures the interval between two signals. Input A is the start pulse and Input B is the stop pulse. The TIC begins measuring a time interval when the start signal reaches its trigger level and stops measuring when the stop signal reaches its trigger level. The time interval between the start and stop signals is measured by counting cycles from the timebase. The measurements produced by a TIC are in time units: milliseconds, microseconds, nanoseconds, etc. These measurements assign a value to the phase difference between the reference and the DUT.

The most important specification of a TIC is *resolution*; that is, the degree to which a measurement can be determined. For example, if a TIC has a resolution of 10 ns, it can produce a reading of 3340 ns or 3350 ns, but not a reading of 3345 ns. Any finer measurement would require more resolution. In simple TIC designs, the resolution is limited to the period of the TIC's timebase frequency. For example, a TIC with a 10 MHz timebase is limited to a resolution of 100 ns if the unit cannot resolve time intervals smaller than the period of one cycle. To improve this situation, some TIC designers have multiplied the timebase frequency to get more cycles and thus more resolution. For example, multiplying the timebase frequency to 100 MHz makes 10 ns resolution possible, and 1 ns counters have even been built using a 1 GHz timebase. However, a more common way to increase resolution is to detect parts of a timebase cycle through interpolation and not be limited by the number of whole cycles. Interpolation has made 1 ns TICs commonplace, and 10 ps TICs are available [23].

A time interval system is shown in Figure 19.14. This system uses a TIC to measure and record the difference between two signals. Low-frequency start and stop signals must be used (typically 1 Hz). Since oscillators typically produce frequencies like 1, 5, and 10 MHz, the solution is to use a *frequency divider* to convert them to a lower frequency. A frequency divider could be a stand-alone instrument, a small circuit, or just a single chip. Most divider chips divide by multiples of 10, so it is common to find circuits that divide by a thousand, a million, etc. For example, dividing a 1 MHz signal by 10^6 produces a 1 Hz signal. Using low-frequency signals reduces the problem of counter overflows and underflows and helps prevent errors that could be made if the start and stop signals are too close together. For example, a TIC might make errors when attempting to measure a time interval of <100 ns.

The time interval method is probably the most common method in use today. It has many advantages, including low cost, simple design, and excellent performance when measuring long-term frequency uncertainty or stability. There are, however, two problems that can limit its performance when measuring short-term stability. The first problem is that after a TIC completes a measurement, the data must be processed before the counter can reset and prepare for the next measurement. During this delay, called *dead time*, some information is lost. Dead time has become less of a problem with modern TICs where




FIGURE 19.14 Time interval system.




FIGURE 19.15 Dual mixer time difference system.

faster hardware and software has reduced the processing time. A more significant problem when measuring short-term stability is that the commonly available frequency dividers are sensitive to temperature changes and are often unstable to ± 1 ns. Using more stable dividers adds to the system cost [10, 22].

A more complex system, better suited for measuring short-term stability, uses the Dual Mixer Time Difference (DMTD) method, shown in Figure 19.15. This method uses a frequency synthesizer and two mixers in parallel. The frequency synthesizer is locked to the reference oscillator and produces a frequency lower than the frequency output of the DUT. The synthesized frequency is then heterodyned (or mixed) both with the reference frequency and the output of the DUT to produce two beat frequencies. The beat frequencies are out of phase by an amount proportional to the time difference between the DUT and reference. A *zero crossing detector* is used to determine the zero crossing of each beat frequency cycle. An event counter (or *scaler*) accumulates the coarse phase difference between the oscillators by counting the number of whole cycles. A TIC provides the fine-grain resolution needed to count fractional cycles. The resulting system is much better suited for measuring short-term stability than a time interval system. The oscillator, synthesizer, and mixer combination provides the same function as a divider, but the noise is reduced, often to ± 5 ps or less.

Other variations of these calibration methods are used and might be equally valid for a particular application [9, 10, 21, 22]. Keep in mind that a well-defined and documented calibration method is mandatory for a laboratory seeking ISO registration or compliance with a laboratory accreditation program.

19.6 Future Developments

In future years, the performance of frequency standards will continue to improve. One promising development is the *cesium-fountain* standard. This device works by laser cooling cesium atoms and then lofting them vertically through a microwave cavity. The resonance frequency is detected as the atoms rise and fall under the influence of gravity. Many laboratories are working on this concept, which should reduce the frequency uncertainty realized with existing atomic-beam cesium standards [24]. Eventually, a *trapped-ion* standard could lead to improvements of several orders of magnitude. This standard derives its resonance frequency from the systematic energy shifts in transitions in certain ions that are held motionless in an electromagnetic trap. The frequency uncertainty of such a device could be as small as $\pm 1 \times 10^{-18}$ [25]. Also, new statistical tools could improve the ability to estimate oscillator stability, particularly at long term [26].

The future of transfer standards should involve more and more reliance on satellite receivers, and their performance should continue to improve. Ground-based systems such as LORAN-C are expected to be phased out [27]. The frequency uncertainty of GPS receivers will improve when the Selective Availability (SA) program is discontinued in the early part of the next century [28].

References

1. International Organization for Standardization (ISO), *International Vocabulary of Basic and General Terms in Metrology (VIM)*, Geneva, Switzerland, 1993.
2. ISO/IEC Guide 25, *General Requirements for the Competence of Calibration and Testing Laboratories*, International Organization for Standardization (ISO), 1990.
3. ANSI/NCSL Z540-1-1994, *Calibration Laboratories and Measuring and Test Equipment — General Requirements*, American National Standards Institute, 1994.
4. M. A. Lombardi, An introduction to frequency calibration. Part I, *Cal Lab Int. J. Metrol.*, January-February, 17-28, 1996.
5. B. N. Taylor and C. E. Kuyatt, Guidelines for evaluating and expressing the uncertainty of NIST measurement results, *Natl. Inst. of Stan. and Technol. Tech. Note 1297*, 1994.
6. IEEE, IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology, *IEEE 1139*, Piscataway, NJ, 1988.
7. D. W. Allan, H. Hellwig, P. Kartaschoff, J. Vanier, J. Vig, G. M. R. Winkler, and N. F. Yannoni, Standard Terminology for Fundamental Frequency and Time Metrology, *Characterization of Clocks and Oscillators — Natl. Inst. of Stan. and Technol. Tech. Note 1337*, 1990, 139-145.
8. J. Jespersen, Introduction to the time domain characterization of frequency standards, *Proc. 25th Annu. Precise Time and Time Interval (PTTI) Meeting*, Pasadena, CA, December 1991, 83-102.
9. S. R. Stein, Frequency and time — their measurement and characterization, *Precision Frequency Control*, Vol. 2, E. A. Gerber and A. Ballato, Eds., Academic Press, New York, 1985, 191-232.
10. D. A. Howe, D. W. Allan, and J. A. Barnes, Properties of signal sources and measurement methods, *Characterization of Clocks and Oscillators*, D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, Eds., *Natl. Inst. of Stan. Technol. Tech. Note 1337*, 1990, 14-60.
11. W. A. Marrison, The evolution of the quartz crystal clock, *Bell Systems Tech.*, 27(3), 510-588, 1948.
12. J. R. Vig, Introduction to Quartz Frequency Standards, *Army Research and Development Technical Report*, SLCET-TR-92-1, October 1992.
13. F. L. Walls and J. Gagnepain, Environmental sensitivities of quartz oscillators, *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 39, 241-249, March 1992.

14. W. M. Itano and N. F. Ramsey, Accurate measurement of time, *Sci. Am.*, 269(1), 56-65, 1993.
15. L. Lewis, An introduction to frequency standards, *Proc. IEEE*, 79(7), 927-935, 1991.
16. J. Vanier and C. Audoin, *The Quantum Physics of Atomic Frequency Standards*, Adam Hilger, Bristol, England, 2 Vols., 1989.
17. M. A. Lombardi, An introduction to frequency calibration. Part II, *Cal Lab Int. J. Metrology*, March-April, 28-34, 1996.
18. R. Beehler and M. A. Lombardi, NIST time and frequency services, *Natl. Inst. of Stan. and Technol. Special Publ. 432*, 1991.
19. B. Hoffmann-Wellenhof, H. Lichtenegger, and J. Collins, *GPS: Theory and Practice*, 3rd ed., Springer-Verlag, New York, 1994.
20. ARINC Researc Corporation, NAVSTAR Global Positioning System: User's Overview, NAVSTAR GPS Joint Program Office, Los Angeles, YEE-82-009D, March 1991.
21. T. N. Osterdock and J. A. Kusters, Using a new GPS frequency reference in frequency calibration operations, *IEEE Int. Freq. Control Symp.*, 1993, 33-39.
22. R. J. Hesselberth, Precise frequency measurement and calibration, *Proc. Natl. Conf. Stan. Lab.*, 1988, 47-1-47-7.
23. V. S. Zhang, D. D. Davis, and M. A. Lombardi, High resolution time interval counter, *Precise Time and Time Interval Conf. (PTTI)*, 1994, 191-200.
24. A. Clairon, P. Laurent, G. Santarelli, S. Ghezali, S. N. Lea, and M. Bahoura, A cesium fountain frequency standard: preliminary results, *IEEE Trans. Instrum. Meas.*, 44(2), 128-131, 1995.
25. W. M. Itano, Atomic ion frequency standards, *Proc. IEEE*, 79(7), 936-942, 1991.
26. D. A. Howe, An extension of the Allan variance with increased confidence at long term, *IEEE Int. Freq. Control Symp.*, 1995, 321-330.
27. U.S. Dept. of Transportation, Federal Radionavigation Plan, *DOT-VNTSC-RSPA-95-1*, DOD-4650.5, 1994 (new version published every 2 years).
28. G. Gibbons, A national GPS policy, *GPS World*, 7(5), 48-50, 1996.