Appendix A
An Introduction to Frequency Calibrations

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 electrical signals often is measured in units of kilohertz (kHz) or megahertz (MHz), where 1 kHz equals one thousand (10^3) cycles per second and 1 MHz equals one 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, one part per million (1×10^{-6}) is considered quite an accomplishment. In frequency metrology, measurements of one part per billion (1×10^{-9}) are routine, and even one 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. Let's begin our discussion with an overview of frequency calibrations.

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*. Often called the nominal

frequency, the nameplate frequency is labeled on the oscillator's 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 an uncertainty requirement for the frequency offset 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 always make their calibrations on-site.

Fortunately, we 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 signals are available, including NIST radio stations WWV, WWVH, 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 17025, *General Requirements for the Competence of Testing and Calibration Laboratories*, states:

> *The laboratory shall use appropriate methods and procedures for all tests and/or calibrations within its scope. These include sampling, handling, transport, storage and preparation of items to be tested and/or calibrated, and, where appropriate, an estimatation of the measurement uncertainty as well as statistical techniques for analysis of test and/or calibration data.*

In addition, Guide 17025 states:

> *The laboratory shall use test and/or calibration methods, including methods for sampling, which meet the needs of the client and which are appropriate for the test and/or calibrations it undertakes .... When the client does not specify the method to be used, the laboratory shall select appropriate methods that have been published either in international, regional, or national standards, or by reputable technical organizations, or in relevant scientific texts or journals, or as specified by the manufacturer of the equipment* [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.

Now that we've provided an overview of frequency calibrations, we'll take a more detailed look at the topics introduced. We'll begin by looking at the specifications used to describe a frequency calibration. Then, we'll discuss the various types of frequency standards and transfer standards. We'll conclude with a discussion of how the NIST Frequency Measurement and Analysis Service provides a complete solution to the frequency calibration problem.

The Specifications: Frequency Offset and Stability

In this section, we'll look at the two main specifications of a frequency calibration, frequency offset and stability. We'll define frequency offset and stability and show how they are measured. Keep in mind during this discussion that frequency offset is often referred to simply as *accuracy* (or *frequency accuracy*), and that stability is nearly the same thing as frequency uncertainty.

Frequency Offset

When we make a frequency calibration, our 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 signal at a frequency of 5 MHz. Of course, the DUT will actually produce a frequency that isn't exactly 5 MHz. After we calibrate the DUT, we can state its frequency offset and the associated uncertainty.

Measuring the frequency offset of a DUT requires comparing it 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 us to do this. Once we know the amount of *phase deviation* and the *measurement period,* 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:

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

To illustrate, let's say that we measure +1 µs (microsecond) of phase deviation over a measurement period of 24 hours (h). The unit used for measurement period (h) must be converted to the unit used for phase deviation (s). The equation then becomes:

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

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

Table A.1. Frequency offset values for given amounts of phase deviation.

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

The frequency offset values in Table A.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:

$$(5 \times 10^{6}) \, (+1.16 \times 10^{-11}) = 5.80 \times 10^{-5} = +0.000 \ 058 \ 0 \text{ Hz}$$

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

$$5 \ 000 \ 000 \text{ Hz} + 0.000 \ 058 \ 0 \text{ Hz} = 5 \ 000 \ 000.000 \ 058 \ 0 \text{ Hz}$$

To do a complete uncertainty analysis, the measurement period must be long enough to insure that we are 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, we must be sure that Δt is really a measure of only 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 we gave for frequency offset is often too simple. When transfer standards such as LORAN-C or GPS receivers are used (discussed later), radio path noise contributes to the phase deviation. To get around this problem, a measurement period of at least 24 hours 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 hours. 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].
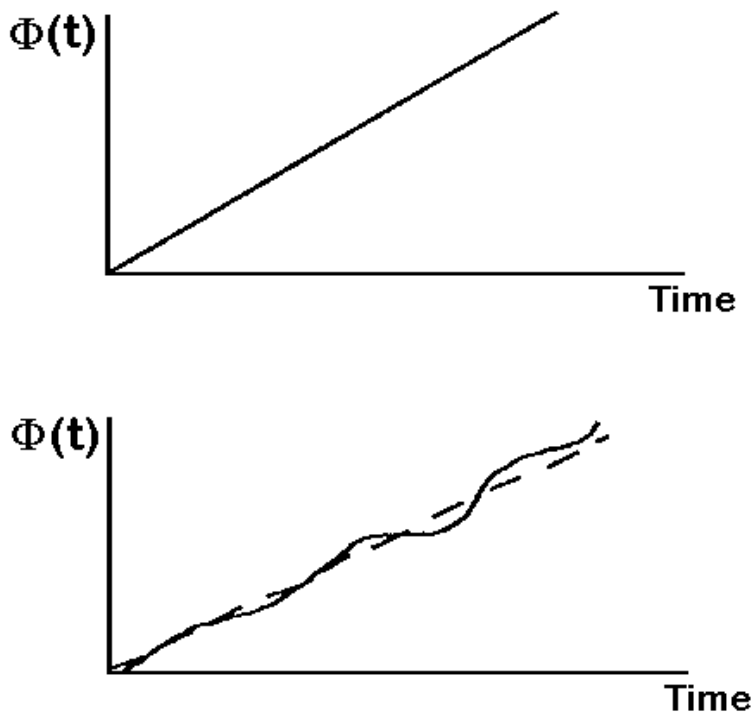

Figure A.1. Simple phase comparison graphs.

Figure A.1 shows two graphs of phase comparisons between a DUT and a reference. The top graph shows no discernible phase noise. This indicates that a TUR of 10:1 or better is being maintained. The bottom graph shows a small amount of phase noise, 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 is the quantity of greatest interest to a calibration laboratory because it tells us how close a DUT is to its nameplate frequency.

You will probably notice that the term *frequency accuracy* (or just *accuracy*) often appears on oscillator specification sheets instead of the term frequency offset. Frequency accuracy and frequency offset are equivalent terms that refer to the result of a measurement at a given time. Frequency uncertainty indicates the limits (upper and lower) of the frequency offset. ISO defines uncertainty as a:

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

In other words, the frequency uncertainty shows us the possible range of values (or limits) for the frequency offset. It is now standard practice to use a 2σ uncertainty test. This means that there is a 95.4 % probability that the frequency offset will stay within the stated range during the measurement period. The range of values is obtained by both adding the frequency uncertainty to and subtracting it from the average (or mean) frequency offset. Therefore, frequency uncertainty is sometimes stated with a "plus or minus" sign ($\pm 1 \times 10^{-12}$) to show the upper and lower limits of the offset. If the "$\pm$" symbol is omitted, it is still implied. The largest contributor to the frequency uncertainty is usually the stability of the device under test. Stability is the topic of the next section.

Stability

Before beginning our discussion of *stability*, we should mention an important distinction between frequency offset and stability. Frequency offset is a measure of how closely an oscillator produces its nameplate frequency, or how well an oscillator is adjusted. It doesn't tell us about the quality of an oscillator.
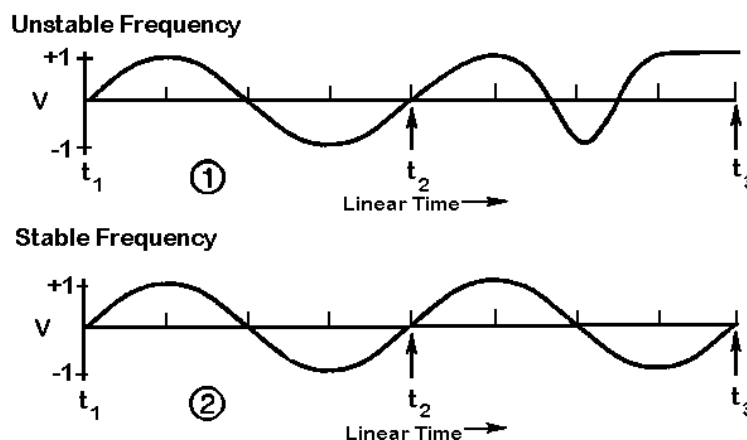

Figure A.2. Comparison of unstable and stable frequencies.

For example, a stable oscillator that needs adjustment might produce a frequency with a large offset. An unstable oscillator that is well adjusted might temporarily produce a frequency very close to its nameplate value.

Stability indicates how well an oscillator can produce the same frequency over a given period of time. It doesn't indicate whether the frequency is "right" or "wrong," only whether it stays the same. Also, the stability doesn't necessarily change when the frequency offset changes. You can adjust an oscillator and move its frequency either further away from or closer to its nameplate frequency without changing its stability at all. Figure A.2 illustrates this by displaying two oscillating signals that are of the same frequency between t1 and t2. However, it's clear that signal 1 is unstable and is fluctuating in frequency between t2 and t3.

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*, and statistical tools exist to easily convert from one domain to the other. Time domain estimates are more widely used, since time interval counters are often used to measure frequency. To estimate stability in the time domain, we must start with a set of frequency offset measurements y_i that consists of individual measurements, y_1, y_2, y_3, and so on. Once this data set is obtained, we need 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 data set 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, since 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, we have a moving mean that changes each time we add a new measurement.

For these reasons, a non classical statistic is used to estimate stability in the time domain. This statistic is often called the *Allan variance*, but since 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:

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

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, successive data points are differenced to remove the time-dependent noise contributed by the frequency offset.

Table A.2. Using phase measurements to estimate stability (1 s data intervals).

Phase Measurements (ns)	Phase Deviation (ns), Δt	Frequency Offset $\Delta t / 1$ 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}

Table A.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, we obtain the phase deviations shown in the second column. 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/s, 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 $\tau = 1$ s) becomes:

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

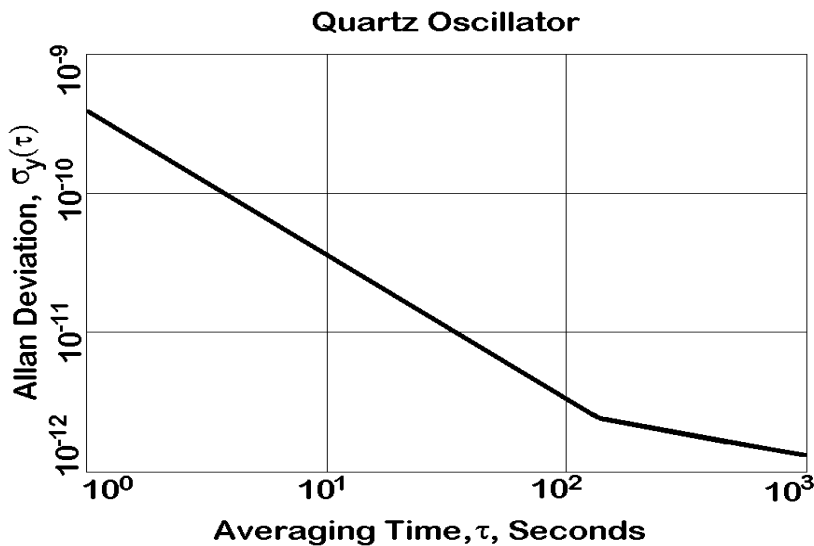Using the same data, the Allan deviation for $\tau = 2$ s can be computed by averaging pairs of adjacent values and using these new averages as data values. For $\tau = 4$ s, take the average of each set of four adjacent values and use these new averages as data values. More data must be acquired for longer averaging times. Keep in mind that the confidence level of a stabililty estimate

Figure A.3. A sample Allan deviation graph.

improves as the averaging time increases. In the above example, we have eight samples for our $\tau = 1$ s estimate. However, we have only two samples for an estimate of $\tau = 4$ s.

The confidence level of our estimate (1σ) can be roughly estimated as:

$$\frac{1}{\sqrt{M}} \text{ x } 100\%$$

In our example, where $M = 9$, our error margin is 33 %. With just 2 samples, our estimate may 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 A.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 may not reach their flicker floor for 10^5 s (more than 1 day). Figure A.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, 9, 10].

Be sure not to confuse stability with frequency offset when you read 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 A.4 is a graphical representation of the relationship between frequency offset (accuracy) and stability.
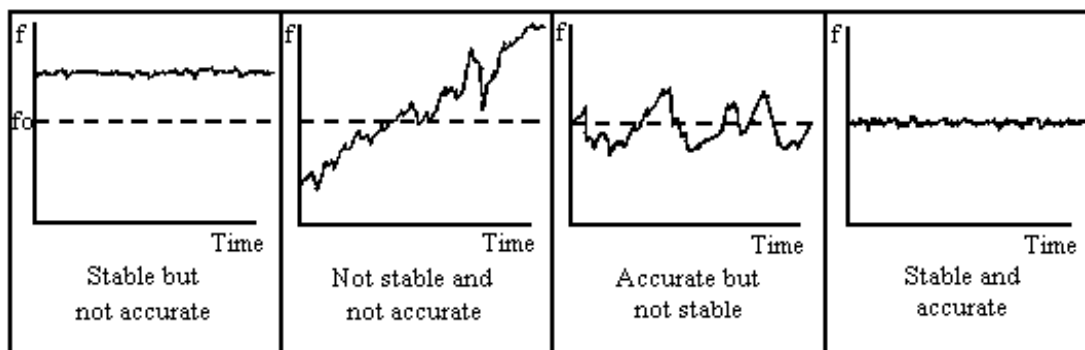

Figure A.4. The relationship between frequency
offset (accuracy) and stability.

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*. Two main types of oscillators are used as frequency standards: quartz oscillators and atomic oscillators.

Quartz Oscillators

Quartz crystal oscillators first appeared during the 1920s and quickly replaced pendulum devices as laboratory standards for time and frequency [11]. More than 10^9 quartz oscillators are produced annually for applications ranging from 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 (such as frequency counters) or those designed as stand-alone units. The cost of a high-quality quartz oscillator ranges from a few hundred to a few thousand dollars.
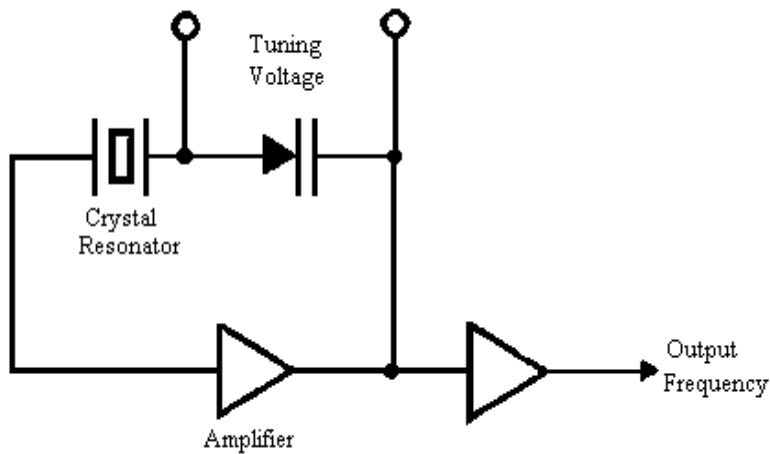
The quartz crystal inside the oscillator can be made of natural or synthetic quartz, but all modern devices are made of synthetic material. The crystal serves as a mechanical resonator which creates an oscillating voltage due to the *piezoelectric effect*. This effect causes the crystal to expand or 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

Figure A.5. Block diagram of quartz oscillator.

multiple of that frequency. Figure A.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 alternate 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 over time 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 might 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 no more than 5×10^{-9} per year.

The best quartz oscillators have excellent short-term stability. An OCXO might be stable to 1×10^{-13} at 1 s. The limiting factor in the short-term stability is often noise from electronic components in the oscillator circuits. However, due to aging and other factors, the long term stability of a quartz oscillator is poor. Even the best OCXOs

require periodic adjustments to stay within 1×10^{-10} of their nameplate frequency. The TCXOs found in test equipment like counters and signal generators typically have frequency offsets ranging from 1×10^{-7} to 1×10^{-9}. Quartz oscillators without temperature control (such as those found in wristwatches, computers, radios, and so on) typically have frequency offsets measured in parts per million.

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 energy 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]:

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

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, since 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 since fewer adjustments are needed, labor costs are reduced. As a result, a rubidium oscillator might actually be less expensive to own than a quartz oscillator when used as a frequency standard.

The typical frequency offset of a rubidium oscillator ranges from 5×10^{-10} to 5×10^{-12}. Stability is typically about 1×10^{-12} at one day. Maintaining frequency within 1×10^{-11} can be done routinely with a rubidium oscillator but is nearly impossible 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, although some hydrogen masers also contribute to UTC.

Cesium oscillators are the workhorses in most modern time and frequency distribution systems. The primary frequency standard for the United States is a large cesium fountain oscillator named NIST-F1 that has a frequency uncertainty of about 1.7×10^{-15}. Commercially available cesium oscillators use cesium beam technology, and are small enough to fit in a standard equipment rack. They differ in quality, but their frequency offset should still be no greater than $\pm 5 \times 10^{-12}$ after a brief warmup period, and is typically parts in 10^{13}. Stability generally reaches parts in 10^{14} at one day, and it might take days or weeks before a cesium oscillator reaches its noise floor.

Reliability and cost are both issues to consider when purchasing a cesium oscillator. The main component of a cesium oscillator, called the *beam tube*, has a life expectancy of about 3 to 25 years, depending upon the type of tube and the amount of beam current used. 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 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 afterwards.

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. "Maser" is an acronym that stands for "Microwave Amplification by Stimulated Emission of Radiation."   Masers use 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. Both types of maser have better short-term stability than a cesium oscillator. However, since the maser's performance depends upon a complex set of environmental conditions, its frequency uncertainty is greater than that of a cesium oscillator. And due to their complexity and low volume of production, masers typically cost $200,000 or more [15,16].

Table A.3 summarizes the characteristics of the different types of oscillators.

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, since not all calibration laboratories can afford to buy and maintain a cesium. Even if a laboratory has a cesium oscillator, it still needs to check its performance, and the only practical way to do this is by using a transfer standard.

Transfer standards provide *traceability*. Most transfer standards receive signals traceable to the national frequency standard maintained by NIST. Some signals, such as

Table A.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	1.42040575 GHz
Leading Cause of Failure	None	None	None	Rubidium Lamp (15 years or more)	Cesium Beam Tube (3 to 25 years)	Hydrogen Depletion (7 years or more)
Stability, $\sigma_y(\tau)$, $\tau =1s$	1×10^{-9}	1×10^{-10}	1×10^{-12}	5×10^{-11} to 5×10^{-12}	5×10^{-11} to 5×10^{-12}	1×10^{-12}
Noise Floor, $\sigma_y(\tau)$	1×10^{-9} ($\tau = 1$ to 10^2 s)	1×10^{-10} ($\tau = 1$ to 10^2 s)	1×10^{-12} ($\tau = 1$ to 10^2 s)	1×10^{-12} ($\tau = 10^3$ to 10^5 s)	1×10^{-14} ($\tau = 10^5$ to 10^7 s)	1×10^{-15} ($\tau = 10^3$ to 10^5 s)
Aging/year	5×10^{-7}	5×10^{-8}	5×10^{-9}	2×10^{-10}	None	$\sim 1 \times 10^{-13}$
Frequency Offset after warm up	1×10^{-6}	1×10^{-7} to 1×10^{-8}	1×10^{-8} to 1×10^{-10}	5×10^{-10} to 5×10^{-12}	5×10^{-12} to 1×10^{-14}	1×10^{-12} to 1×10^{-13}
Warm-Up Time	< 10 s to 1×10^{-6}	< 10 s to 1×10^{-8}	< 5 min to 1×10^{-8}	< 5 min to 5×10^{-10}	30 min to 5×10^{-12}	24 hours to 1×10^{-12}
Cost	$100	$1000	$2000	$3000 to $8000	$30,000 to $80,000	$200,000 to $300,000

hose 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, such as the LORAN-C and Global Positioning System (GPS) broadcasts, are traceable because their received signals are regularly compared to the NIST standard. Some signals broadcast from outside the United States are also 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 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 since the cable is of fixed length the delay is constant. Constant delays don't 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 you use 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 long term measurements, since you can minimize these frequency fluctuations if you average for a long enough measurement interval. Eventually, you will recover the performance of a cesium oscillator.

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, Colorado. WWV is a 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 your 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 to 1000 μs. Since there is so much variability in the path, averaging leads to only limited improvement. Therefore, although WWV is traceable to NIST, its frequency uncertainty is limited to parts in 10^9 when averaged for one day.

Other radio signals have more stable paths and much lower uncertainty values. Low frequency (LF) radio stations (such as WWVB and LORAN-C) can provide traceability to NIST with a frequency uncertainty of 1×10^{-12}/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

Table A.4. Typical frequency uncertainties of various transfer standards.

Transfer Standard	Frequency Uncertainty over 24 h measurement period (with respect to NIST)
HF receiver (WWV and WWVH)	$\pm 5 \times 10^{-9}$
LF receiver (LORAN-C and WWVB)	$\pm 1 \times 10^{-12}$
Global Positioning System receiver (GPS)	$\pm 2 \times 10^{-13}$

frequency uncertainty of GPS is about 2×10^{-13}/day. WWVB, LORAN-C and GPS are described in the next three sections.

Table A.4 shows some of the transfer standards available, as well as their frequency uncertainty with respect to NIST when averaged for at least 24 h [17].

WWVB

Many countries broadcast time and frequency signals in the LF band from 30 to 300 kHz, as well as in the VLF (very low frequency) band from 3 to 30 kHz. Since part of the LF signal is *groundwave* and follows the curvature of the earth, the path stability of these signals can be quite good. One such station is NIST's WWVB, which transmits on 60 kHz from the same site as WWV in Fort Collins, Colorado and provides coverage to most of North America.

Although far more stable than an HF path, the WWVB path length is still 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 since

it always arrives first. For longer paths, a mixture of groundwave and skywave is received. And over a very long path, the groundwave can become so weak that it will be possible to receive only 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 occur at sunrise and sunset. For instance, 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, you can achieve a frequency uncertainty of 1×10^{-12} by averaging for one 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 to 800 kW, and broadcast on a carrier frequency of 100 kHz using the 90 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 of the ionosphere. The pulse shape was designed so that the receiver can distinguish between groundwave and skywave and lock on to 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/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 µs phase step, equal to the period of 100 kHz.

LORAN-C Performance

        The frequency uncertainty of LORAN-C is degraded by variations in the radio path. The size of these variations depends upon the signal strength, your distance from the transmitter, the weather and atmospheric conditions, and the quality of your receiver and antenna. Path variations cause the short-term stability of LORAN-C to be poor. However, since the path variations average out over time, the long-term stability is quite good. This means that we can use LORAN-C to calibrate nearly any frequency standard if we can average for a long enough interval. For this reason, a measurement period of at least 24 h is recommended when using LORAN-C to calibrate atomic oscillators.   Figure A.6 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 3.4×10^{-12}.
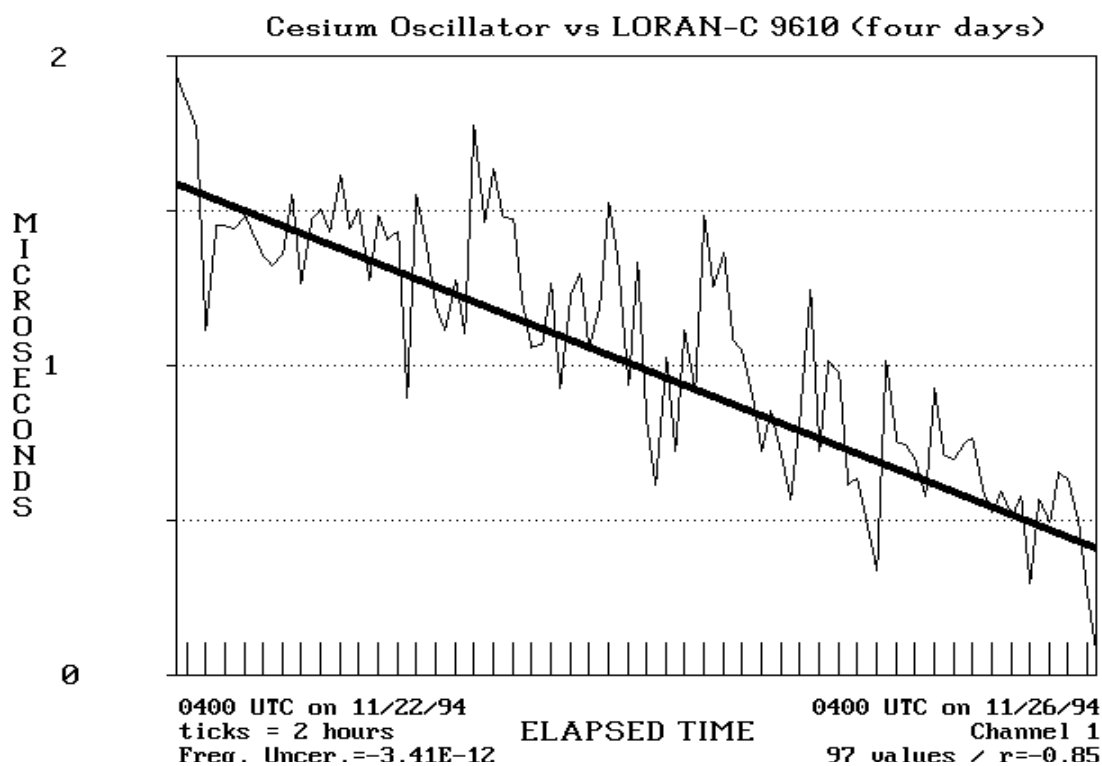

Figure A.6. LORAN-C compared to cesium oscillator over 96 h interval.

Global Positioning System (GPS)

GPS is a radionavigation system developed and operated by the U.S. Department of Defense (DOD). The system consists of a constellation of at least 24 earth-orbiting satellites (21 primary satellites and 3 in-orbit spares).   Each satellite carries its own atomic frequency standards (cesium and/or rubidium oscillators) 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 55º from the equator. Each satellite is 20 200 km above the earth's surface and has an orbital period of 11 h, 58 min, 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.

The GPS satellites broadcast on two carrier frequencies: L1 at 1575.42 MHz, and L2 at 1227.6 MHz. Each satellite broadcasts a spread-spectrum waveform, called a pseudo-random noise (PRN) code on L1 and L2, and each satellite is identified by the PRN code it transmits. There are two types of PRN codes. The first type is a coarse acquisition (C/A) code with a chip rate of 1023 chips per millisecond. The second type is a precision (P) code with a chip rate of 10230 chips per millisecond. The C/A code is broadcast on L1, and the P code is broadcast on both L1 and L2. 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]. GPS reception is line-of-sight, which means that the antenna must have a clear view of the sky. If a clear sky view is available, the signals can be received nearly anywhere on earth.

Each satellite carries either rubidium or cesium oscillators, or a combination of both. The on-board oscillators provide the reference for both the carrier and code broadcasts. They are steered from USDOD ground stations and are referenced to Coordinated Universal Time (UTC) maintained by the USNO. By mutual agreement UTC(USNO) and UTC(NIST) are maintained within 100 ns of each other, and the frequency offset between the two time scales is always $<1 \times 10^{-13}$.

GPS Receiving Equipment

Most GPS receivers provide a 1 *pulse per second* (pps) output. Some receivers also provide a standard frequency output (1, 5, or 10 MHz). To use these receivers, you simply mount the antenna, connect the antenna to the receiver, and turn 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 be available only in rare instances). By averaging data from multiple satellites, a receiver can reduce the frequency uncertainty [20].

GPS Performance

GPS has many technical advantages over LF radio signals. The signals are usually easier to receive, the equipment is often less expensive, the coverage area is much larger (worldwide), and the performance is better. However, as with all transfer standards, the short-term stability of a GPS receiver is not particularly good and this 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 A.7 shows a 100 s comparison between GPS and a cesium oscillator. The cesium oscillator has a frequency uncertainty of 1×10^{-13}, and its accumulated phase shift during the 100 s measurement period is <1 ns. Therefore, most of the phase noise on the graph can be attributed to GPS path variations.
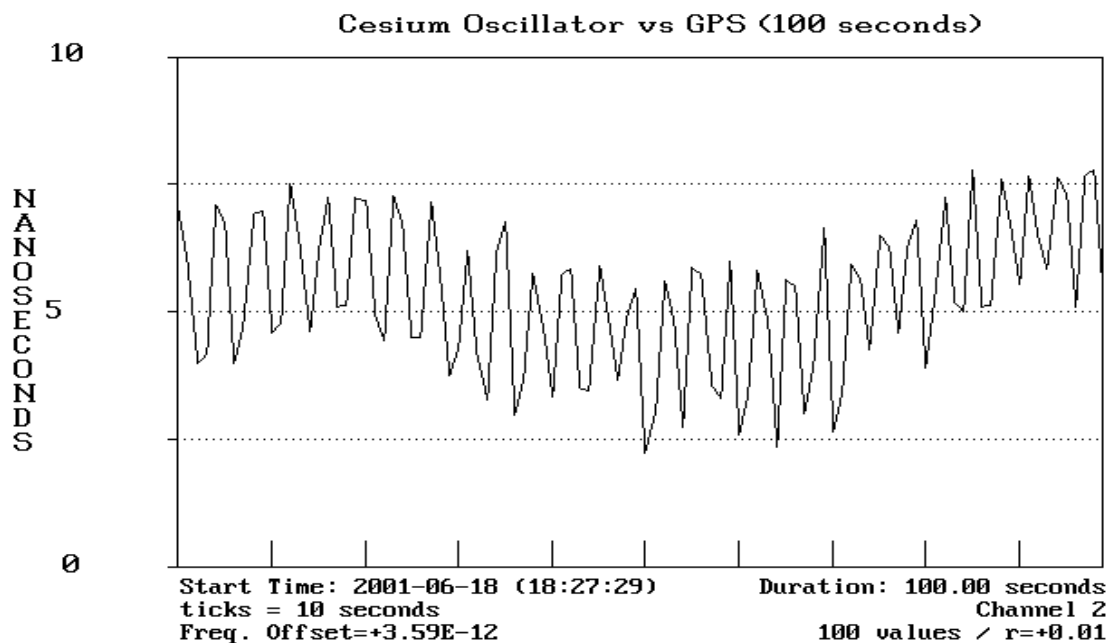

Figure A.7. GPS compared to cesium oscillator over 100 s interval.

Figure A.8 shows a one-week comparison between a GPS receiver and the same cesium oscillator used in Figure A.7. The range of the data is well under 100 ns. The thick line is the result of a linear least squares fit. Although phase noise caused by GPS path variations is visible, we can clearly see the linear trend contributed by the small frequency offset of the cesium ($<1 \times 10^{-13}$).
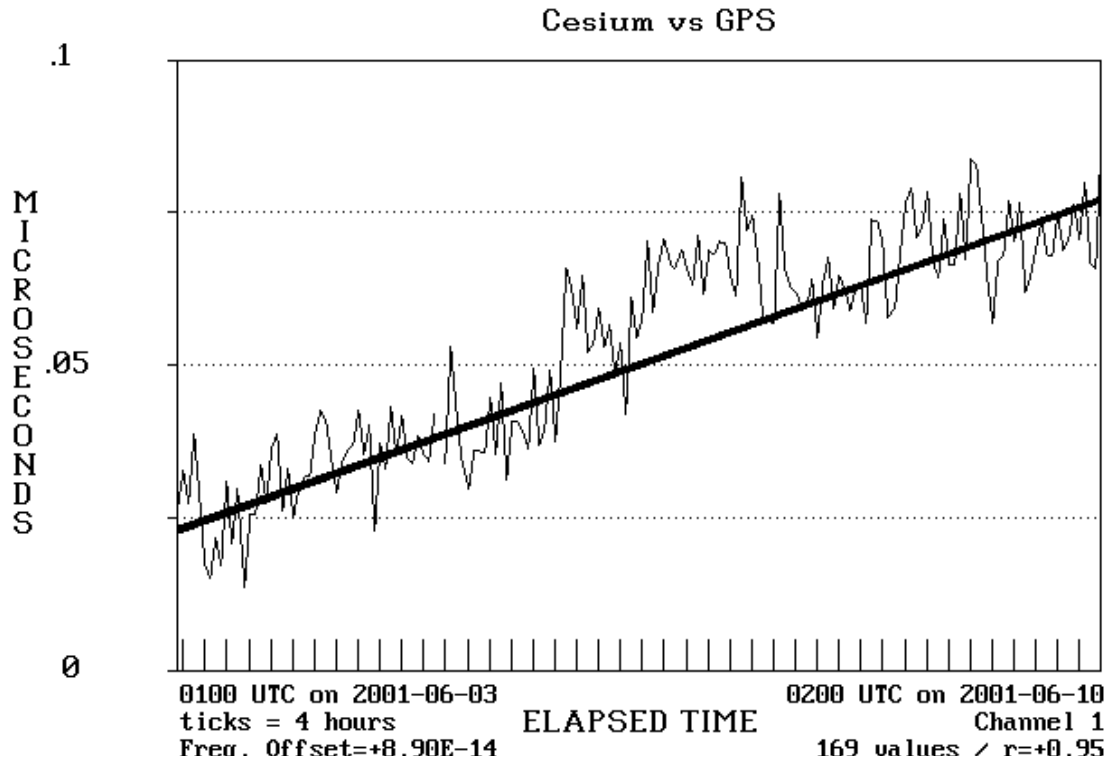

Figure A.8. GPS compared to cesium oscillator over 1 week interval.

Calibration Methods

To review, frequency standards are normally calibrated by comparing them to a traceable reference frequency. In this section, we'll discuss how this comparison is made. To begin with, let's look at the electrical signal produced by a frequency standard. This signal can take several forms, as illustrated in Figure A.9. The dashed lines represent the supply voltage inputs (ranging from 5 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 A.10. This signal produces one cycle (2π radians of phase) in one period. Frequency calibration systems compare a signal like the one shown in Figure

A.10 to a reference signal of higher quality, and measure and record the change in phase between the two signals. If the two frequencies were exactly the same, their phase relationship would not change. Since the two frequencies are not exactly the same, their phase relationship does change, and the rate of change determines the frequency offset of the DUT.
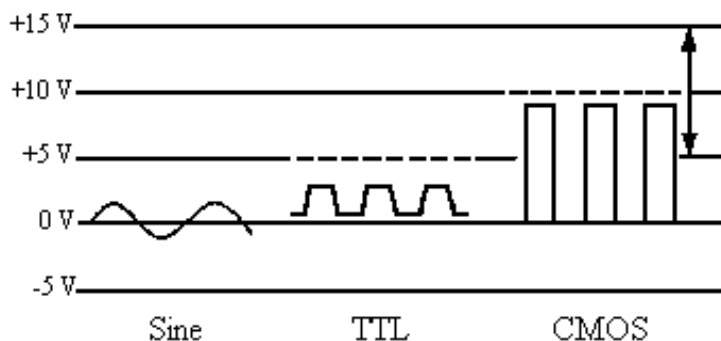

Figure A.9. Oscillator outputs.

Under normal circumstances, the phase changes in an orderly, predictable fashion. However, external factors such as power outages, component failures, or human errors can cause a sudden 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.
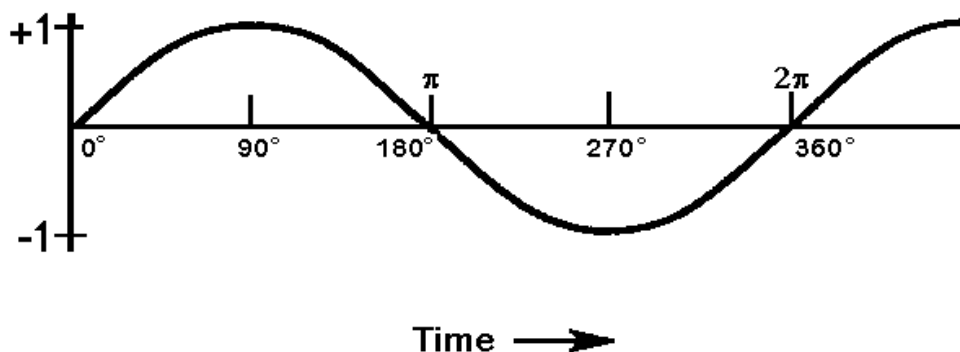

Figure A.10. An oscillating signal.

Figure A.11 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 wave crosses 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. Earlier, we introduced this simple equation to estimate frequency offset:

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

In Figure A.11, both Δt and T are increasing at the same rate and the phase difference "bars" are getting uniformly wider. This indicates that the frequency from the DUT is stable, but is offset with respect to the reference.
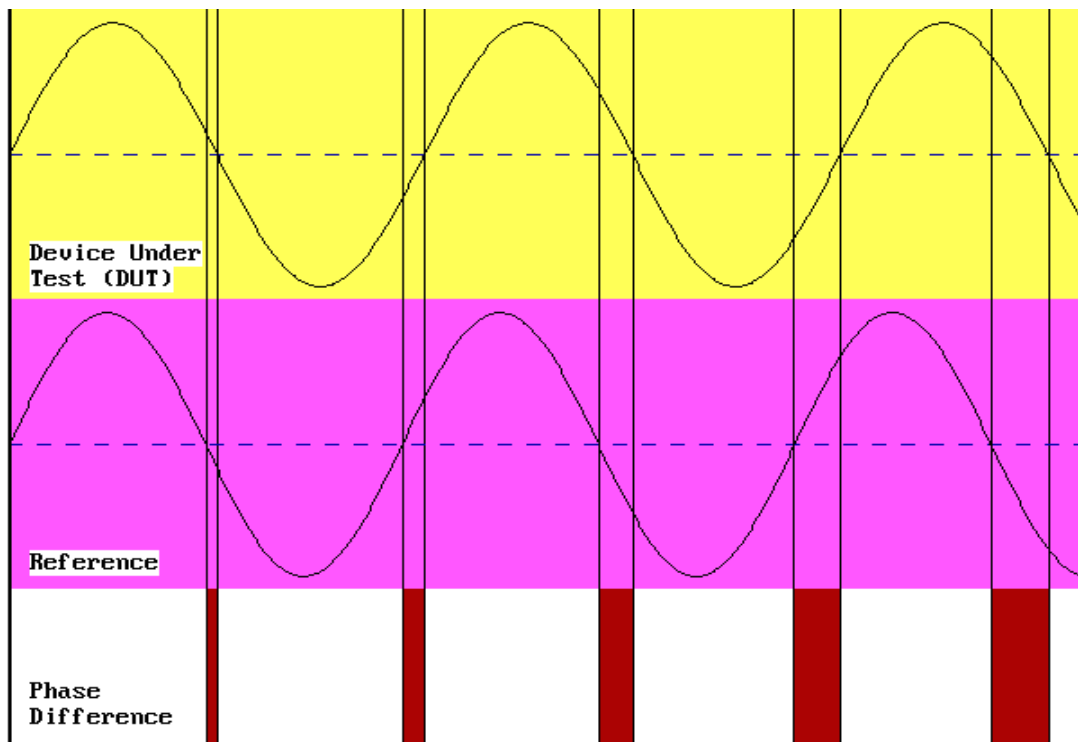

Figure A.11. Two signals with a changing phase relationship.

Several types of calibration systems can be used to compare phase. 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 time base (typically a OCXO), and the uncertainty of the system is limited by the uncertainty of the time base, typically $\pm 1 \times 10^{-8}$. Some counters allow use of an external time base, 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 may take many days or weeks, which might make it difficult or impossible to use this method to adjust a

precision oscillator or to measure stability. For this reason, most high performance calibration systems collect time series data that can be used to estimate both frequency uncertainty and stability. We'll discuss how phase comparisons are made using the *time interval 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 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, you can calibrate the device. It is exactly as if you had two clocks. By reading them each day, you could determine the amount of time one clock gained or lost relative to the other clock. It takes two time interval measurements to produce useful information. By subtracting the first measurement from the second, we can tell whether time was gained or lost.
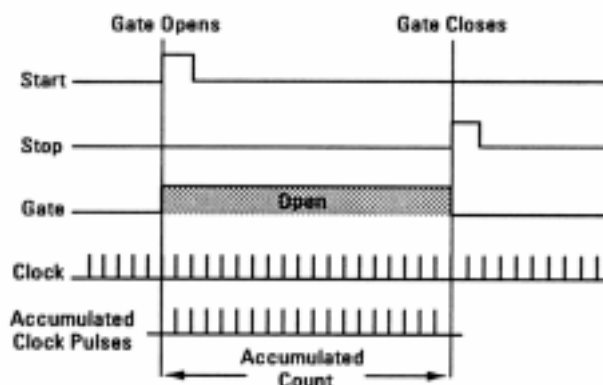

Figure A.12. Measuring time interval.

TIC's differ in specification and design details, but they all contain several basic parts known as the *time base*, the *main gate*, and the *counting assembly*. The time base provides evenly spaced pulses used to measure time interval. The time base is usually an internal quartz oscillator that can often be phase locked to an external reference. It must be stable because time base errors will directly affect the measurements. The main gate controls the time at which the count begins and ends. Pulses passing 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 A.12 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 time base. The measurements produced by a TIC are in time units: milliseconds, microseconds, nanoseconds, and so on.   These measurements assign a value to the phase difference between the reference and DUT.

The most important specification of a TIC is *resolution*, which is the degree to which a measurement can be determined. For example, if a TIC has a resolution of 10 ns, it could produce a reading of 3340 ns or 3350 ns but not a reading of 3345 ns. This is because 10 ns is the smallest significant difference the TIC can measure. Any finer measurement would require more resolution. In traditional TIC designs, the resolution is limited to the period of the TIC's time base frequency. For example, a TIC with a 10 MHz time base would be limited to a resolution of 100 ns. This is because traditional TIC designs count whole time base cycles to measure time interval and cannot resolve time intervals smaller than the period of one cycle. To improve this situation, some TIC designers have multiplied the time base frequency to get more cycles and thus more resolution. For example, multiplying the time base frequency to 100 MHz makes 10 ns resolution possible, and 1 ns counters have even been built using a 1 GHz time base. However, a more common way to increase resolution is to detect parts of a time base cycle through interpolation and not be limited by the number of whole cycles. Interpolation has made 1 ns TICs commonplace, and even 20 ps TICs are available [21, 22].
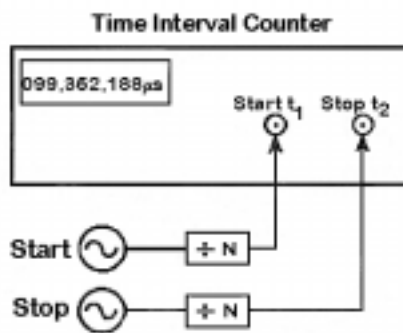

Figure A.13. Time interval system.

A time interval system is shown in Figure A.13.   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 such as 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 circuits divide by multiples of 10, so it is common to find circuits that divide by one thousand, one million, and so on. 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 offset or stability.

The NIST Frequency Measurement and Analysis System

We have now completed our discussion of how frequency calibrations are made. We've seen that a frequency calibration system must include several basic components: an oscillator to calibrate, a reference (preferably a NIST-traceable transfer standard), and a phase comparison device. We'll conclude our tutorial with a brief look at how the NIST Frequency Measurement and Analysis System (FMAS) implements the concepts we've discussed.

A block diagram of the FMAS is shown in Figure A.14. The FMAS fits in a metal equipment rack (75 cm tall), and is computer controlled. Software developed at NIST controls all aspects of the calibration process. It makes measurements and stores and graphs them automatically. Up to five oscillators can be calibrated at once. The system produces calibration graphs that show the performance of each DUT for intervals ranging from 2 s to 150 days.
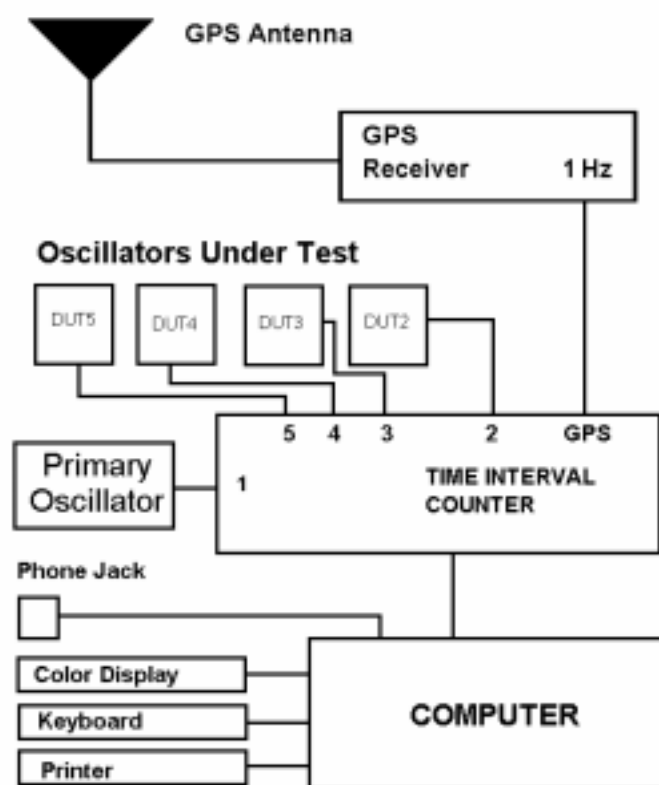

Figure A.14. The NIST Frequency Measurement and Analysis System.

By looking at Figure A.14, you can see all the various components of a frequency calibration system. The FMAS uses a GPS receiver as a transfer standard. The GPS receiver produces a 1 kHz output frequency that provides traceability to NIST with an uncertainty of 2×10^{-13}.

The FMAS makes phase comparisons using the time interval method. It includes a TIC with a single-shot resolution of less than 40 ps. The TIC achieves its high resolution through an interpolation scheme and includes a built-in multiplexer that enables it to switch between five inputs. This allows the FMAS to calibrate up to five oscillators simultaneously. The TIC also includes built-in divider circuitry and can directly accept an input frequency ranging from 1 Hz to 120 MHz on each of the five channels. The TIC is software controlled, so users don't have to set trigger levels. The system is connected to NIST through a telephone modem so that NIST personnel can troubleshoot the system remotely and analyze the calibration data.

As you can see, the FMAS implements many of the concepts discussed in this overview. It provides a calibration laboratory with a well-defined and documented calibration method, and is a major asset to laboratories seeking ISO compliance or laboratory accreditation.

References

[1] International Organization for Standardization (ISO). International Vocabulary of Basic and General Terms in Metrology. Switzerland: ISO; 1993. 59 p.

[2] International Organization for Standardization (ISO). ISO/IEC Guide 17025, General Requirements for the Competence of Testing and Calibration Laboratories. Switzerland: ISO; 1999. 26 p.

[3] American National Standards Institute (ANSI). ANSI/ISO/IEC Guide 17025:2000, General Requirements for the Competence of Testing and Calibration Laboratories. Boulder, Colorado: NCSL International; 2001. 26 p.

[4] Lombardi, Michael A. An Introduction to Frequency Calibration: Part I. Cal Lab Int. J. Metrology: 17-28. 1996 January-February.

[5] Taylor, B. N.; Kuyatt, C. E. Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results. Nat. Inst. Stand. Technol. Tech. Note 1297; 1994 September. 24 p.

[6] IEEE Standards Coordinating Committee 27 on Time and Frequency. IEEE Standard 1139-1999, IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilitites. New York, NY: IEEE, 1999. 31 p.

[7] Sullivan, D. B.; Allan, D. W.; Howe, D. A.; Walls, F. L.; ed. Characterization of Clocks and Oscillators. Nat. Inst. Stand. Technol. Tech Note 1337; 1990 March. 352 p.

[8] Jespersen, James. Introduction to the Time Domain Characterization of Frequency Standards. Proc. 23th Annu. Precise Time and Time Interval (PTTI) Meeting; 1991 December; Pasadena, CA. 83-102.

[9] Stein, S. R. Frequency and Time - Their Measurement and Characterization, chapter 12 in Precision Frequency Control, Vol. 2. Gerber, E. A.; Ballato, A.; ed. New York, NY: Academic Press; 1985. 191-232. Reprinted in [7].

[10] Howe, D. A.; Allan, D. W.; Barnes, J. A. Properties of Signal Sources and Measurement Methods. Proc. 35th Annu. Sym. Freq. Cont; 1981. 47 p. Reprinted in [7].

[11] Marrison, W. A. The Evolution of the Quartz Crystal Clock. Bell Systems Tech. 27(3): 510-588; 1948.

[12] Vig, J. R. Introduction to Quartz Frequency Standards. Army Research and Development Technical Report: SLCET-TR-92-1; 1992. 49 p.

[13] Walls, F. L.; Gagnepain, J. Environmental Sensitivities of Quartz Oscillators. IEEE Trans. Ultrasonics, Ferroelectrics, Frequency Control 39(2): 241-249; 1992 March.

[14] Itano, W. M.; Ramsey, N. F. Accurate Measurement of Time. Sci. Am. 269(1): 56-65; 1993 July.

[15] Lewis, L. An Introduction to Frequency Standards. Proc. IEEE 79(7): 927-935; 1991 July.

[16] Major, F. G. The Quantum Beat: The Physical Principles of Atomic Clocks. New York: Springer Verlag; 1998, 475 p.

[17] Lombardi, Michael A. An Introduction to Frequency Calibration: Part II. Cal Lab Int. J. Metrology: 28-34. 1996 March-April.

[18] Beehler, R.; Lombardi, M. A. NIST Time and Frequency Services. *Nat. Inst. Stand. Technol. Spec. Publ. 432;* 1991 June. 30 p.

[19] Hoffmann-Wellenhof, B.; Lichtenegger, H.; Collins, J. Global Positioning System: Theory and Practice. New York: Springer-Verlag; 1997. 370 p.

[20] Lombardi, Michael A; Nelson, Lisa. M.; Novick, Andrew N.; Zhang, Victor S. Time and Frequency Measurements Using the Global Positioning System. Cal Lab Int. J. Metrology: 26-33. 2001 July-September.

[21] Zhang, Victor S., Davis, Dick D.; Lombardi, Michael A. High Resolution Time Interval Counter. Proc. 26th Annu. Precise Time and Time Interval (PTTI) Meeting; 1994 December; Reston, VA. 191-200.

[22] Novick, Andrew N., Lombardi, Michael A.; Zhang, Victor S.; Carpentier, Anthony. High Performance Multi-Channel Time Interval Counter with Integrated GPS Receiver. Proc. 31st Annu. Precise Time and Time Interval (PTTI) Meeting; 1999 December; Dana Point, CA. 561-568.