An Introduction to Frequency Calibration
Part I

Michael Lombardi
National Institute of Standards and Technology
Time and Frequency Division
Boulder, Colorado, USA

This article is part I of a two-part series on frequency calibrations. Part I introduces the topic of frequency calibrations, discusses the specifications involved, and describes the different types of oscillators. Part II will discuss how you can obtain traceability to NIST through the use of transfer standards. It will also describe how frequency calibrations are made and present an automated frequency calibration system developed at NIST.

Introduction

Frequency calibrations share some similarities with the many other types of calibrations routinely performed by metrologists. According to International Organization for Standardization (ISO) guidelines, a calibration is:

> *A set of operations that establish, under specified conditions, the relationship between values of quantities indicated by a measuring instrument or measuring system, and the corresponding standard or known values derived from the standard.*

In other words, a calibration measures the performance of the *device under test* (DUT). In the field of frequency calibrations, the DUT is a device that produces frequency. In most cases, this device is based on 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 the frequency labeled on the oscillator output. For example, a DUT with an output labeled "5 MHz" is supposed to produce a 5 MHz frequency. The actual output frequency is the *measurand* or the quantity to be measured. The DUT is calibrated by determining how close the measurand is to the nameplate frequency.

The difference between the nameplate frequency and the actual frequency produced by the DUT is called *the frequency offset* and is a measure of *the frequency uncertainty* of the DUT. Calibration laboratories specify an uncertainty requirement that the DUT must meet or exceed. In many cases the lab bases this requirement on the specifications published by the manufacturer. In other cases they may relax the requirements a bit 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 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.*

In the United States, the "unbroken chain of comparisons" should trace back to NIST. In some fields of calibration, traceability is established by sending the standard to NIST or another site for calibration or by sending a set of reference materials (like a set of weights used for mass calibrations) to the user. Neither method is practical when making frequency calibrations. Oscillators are sensitive to being turned on and off. If an oscillator is calibrated and then turned off, the calibration may 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 the calibrations on-site.

Fortunately, there is an easy way to prove that a frequency standard is traceable to NIST. You can receive a frequency signal by radio that is referenced to NIST. This radio link serves as the "unbroken chain" back to the national standard. Several different types of signals are available, including NIST radio stations WWV and WWVB, and radio navigation signals from LORAN-C and GPS. Each signal delivers NIST traceability at a known level of uncertainty. The signal used depends upon the level of uncertainty required.

The ability to use radio signals is a tremendous advantage. The radio signal serves as a transfer standard that delivers a frequency reference from the national standard to the calibration laboratory. A transfer standard 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 oscillators 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 which 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.*

Calibration laboratories, therefore, should automate the calibration process using a well documented and established method. One way to accomplish this is to subscribe to a calibration service such as the NIST Frequency Measurement and Analysis Service (FMAS), which will be described in Part II. Using a proven calibration method helps guarantee that each calibration will be of consistently high quality. This is essential if the laboratory is seeking ISO registration or laboratory accreditation.

The Specifications: Frequency Uncertainty and Stability

Frequency and time interval can be measured with greater precision than all other physical quantities. In some fields of calibration, one part per million (1×10^{-6}) is considered quite an accomplishment. In the world of frequency calibrations, measurements of one part per billion (1×10^{-9}) are routine, and even one part per trillion (1×10^{-12}) is commonplace.

Frequency Uncertainty

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

> *A parameter, associated with the result of a measurement, that characterizes the dispersion of values that could reasonably be attributed to the measurand.*

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 5 MHz frequency. Of course, the DUT will produce a frequency that isn’t exactly 5 MHz. After we calibrate the DUT, we can state its frequency uncertainty, or the amount that the DUT is in error.

Measuring the frequency uncertainty of a DUT requires comparing it to a standard. This is normally done by making a *phase comparison* between the frequency produced by the DUT and the frequency produced by the standard. A phase comparison is used to measure the difference between two frequencies. 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 will change, and by measuring the rate of change, we 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 phase change (called *a phase step).* The purpose of a phase comparison is to measure 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 1 shows a phase comparison between two frequencies that are represented as sine waves. You can think of the top sine wave as a signal from the DUT and the bottom sine wave as a signal from the reference. Vertical lines have been drawn through the point where each sine wave begins and ends. The bottom of Figure 1 shows the spacing between these two lines. The spacing represents the phase difference between the two signals.

If the phase relationship between the signals is changing, the “spacing” will either increase or decrease to indicate that the DUT has a frequency offset with respect to

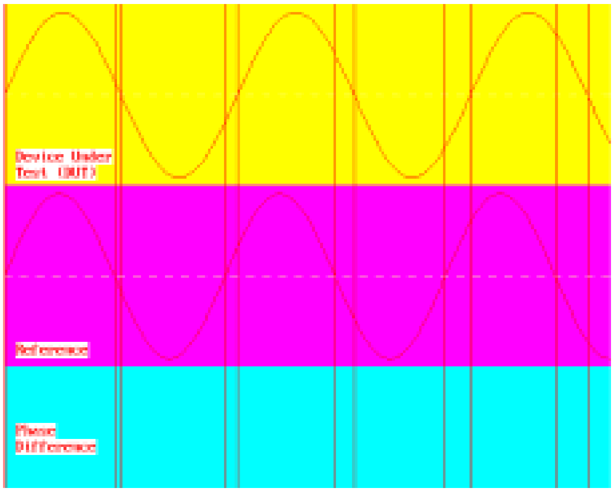

Figure 1. Two sine waves in a changing phase relationship.

the reference. This means that the DUT is producing a different frequency than the reference.

Every frequency calibration system has to be able to measure the amount of phase shift and to convert this quantity to time units (microseconds, for example). One common way of accomplishing this is called the *time interval method.* This involves using a device called a *time interval counter* to measure the time interval between two signals. The time interval method will be explained in more detail in Part II.

Once we know the amount of phase shift (in time units) and the measurement period, we can estimate the frequency uncertainty of the DUT. The measurement period is the length of time over which phase comparisons were made. Frequency uncertainty is estimated as follows:

$$\text{Frequency Uncertainty} = \frac{\text{Phase Shift}}{\text{Measurement Period}}$$

To illustrate, let's say that we measured 1 microsecond of phase shift over a measurement period of 24 hours. The length of the measurement period (hours) must first be converted to the same unit used to measure phase shift (microseconds). When we convert 24 hours to microseconds the equation becomes:

$$\text{Frequency Uncertainty} = \frac{1}{86{,}400{,}000{,}000}$$

As the amount of phase shift gets smaller, the frequency uncertainty gets smaller. The smaller the frequency uncertainty, the closer the DUT is to producing the same frequency as the reference. Since frequency uncertainty values are so small, they are normally converted to scientific notation: 1.16×10^{-11}

Thus, an oscillator that accumulates 1 microsecond of phase shift per day has a frequency uncertainty of about 1×10^{-11} with respect to the reference. Table 1 lists the approximate uncertainty values for some standard units of phase shift and some standard measurement periods.

The simple equation we gave for frequency uncertainty is often too simple. When calibrating one oscillator using another oscillator as a reference, it works fine. This is because both oscillators usually produce clean signals with little noise and nearly all of the phase shift will be caused by the frequency offset of the DUT.

Other quantities (such as measurement system noise) also contribute to the phase shift, but these quantities are so small that they do not significantly change the uncertainty value. However, when we use a transfer standard like LORAN-C or GPS (part II of this article), radio path noise contributes to the amount of phase shift.

To get around this problem, a measurement period of at least 24 hours is normally used when calibrating a high-quality oscillator using a transfer standard. In addition, algorithms that fit a curve to the data are often used and frequency uncertainty is estimated from the slope of the curve.

As you can see from Table 1, the values used to express frequency uncertainty are dimensionless. However, they can easily be converted to a *frequency offset* in Hertz (Hz), if the nameplate frequency is known.

Phase Shift	Measurement Period	Frequency Uncertainty
1 microsecond	1 second	1.0×10^{-6}
1 nanosecond	1 second	1.0×10^{-9}
1 millisecond	1 hour	2.4×10^{-7}
1 microsecond	1 hour	2.4×10^{-10}
1 nanosecond	1 hour	2.4×10^{-13}
1 second	24 hours	1.2×10^{-5}
1 millisecond	24 hours	1.2×10^{-8}
1 microsecond	24 hours	1.2×10^{-11}
1 nanosecond	24 hours	1.2×10^{-14}

Table 1. Approximate uncertainty values for given amounts of phase shift.

To illustrate, consider an oscillator with a nameplate frequency of 5 MHz that is high in frequency by 1.16×10^{-11}. What is the frequency offset of the oscillator? To find out, first multiply the nameplate frequency by the uncertainty to get the frequency offset in Hz:

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

The nameplate frequency is 5 MHz, or 5,000,000 Hz. Therefore, the actual frequency is obtained by adding the offset to this value:

5,000,000 Hz + 0.0000580 Hz
= 5,000,000.0000580 Hz

As this discussion has shown, frequency uncertainty tells us how close a DUT is to producing its nameplate frequency. Therefore, frequency uncertainty is the quantity of the greatest interest to a calibration laboratory.

You have probably noticed that the term *accuracy* (or *frequency accuracy)* often appears on oscillator specification sheets instead of the term frequency uncertainty. However, by international agreement, frequency uncertainty is the correct term to use when stating the performance of an oscillator. Accuracy refers to the result of a measurement at a fixed point in time, whereas uncertainty refers to the *dispersion of values* over a given measurement period. We should point out that this section has presented a simplified explanation of uncertainty. A true uncertainty analysis requires using statistical techniques to show the *confidence level* of the measurement.

Stability

There is an important distinction between frequency uncertainty and stability. Frequency uncertainty is a measure of how well an oscillator produces its nameplate frequency or how well an oscillator is adjusted. It doesn't tell us much about the inherent quality of an oscillator. For example, a high-quality oscillator that needs adjustment may produce a frequency with a large offset. A low-quality oscillator may be well adjusted and produce (temporarily at least) a frequency very close to the nameplate value.

Stability, on the other hand, indicates how well an oscillator can reproduce the same frequency over a given period of time. It doesn't indicate whether the frequency output of an oscillator is "right" or "wrong", only if it stays the same. Also, the stability doesn't change when its frequency offset changes. An oscillator with a large frequency offset may still be very stable. If we adjust the oscillator and move it closer to the correct frequency, the stability won't change. We can improve or degrade the frequency uncertainty without affecting the stability at all.

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 seconds. Long-term stability can refer to measurement intervals greater than 100 seconds, but usually refers to periods longer than 1 day. A typical oscillator specification sheet may quote stability specifications for 1, 10, 100, and 1000 second intervals.

A common statistical test used to estimate stability is the *Allan Variance (AVAR)*, also called the *two-sample* or *pair variance*. You may find it interesting that the first thing a stability estimate does is to remove the frequency uncertainty. This is because stability is a measure of frequency fluctuations and not frequency uncertainty. When we estimate stability we are interested in how much the frequency changes and not how far the oscillator is from the nameplate frequency. Once the frequency uncertainty is removed, the *variance* in the data tells us how stable an oscillator is. This variance is caused by noise that is present in the oscillator or in the measurement system.

To better illustrate how stability is estimated, Table 2 lists a series of 10 phase measurements recorded during a frequency calibration. Each of the numbers in the left column represents a l-second average. The unit is nanoseconds. Since there are 10 data points (n=10), the measurement period is 9 seconds. Each number in the series is larger than the previous number. This indicates that DUT is offset in frequency from the reference and this offset is causing a phase shift.

The middle column shows the amount of phase shift per second. Since this quantity is about 4 nano-

Phase measurements (1 second average, unit is nanoseconds)	Phase Shift due to frequency offset (nanoseconds)	Phase Shift due to frequency fluctuations (nanoseconds)
3321.44	(-----)	(-----)
3325.51	4.07	(-----)
3329.55	4.04	-0.03
3333.59	4.04	0.00
3337.65	4.06	+0.02
3341.69	4.04	-0.02
3345.74	4.05	+0.01
3349.80	4.06	+0.01
3353.85	4.05	-0.01
3357.89	4.04	-0.01

Table 2. Ten phase measurements recorded during a frequency calibration.

seconds per second, we can estimate that the DUT has a frequency uncertainty of about 4 x 10^{-9}. The values in the middle column are obtained by taking the difference between each pair of measurements. For example, nl – n2, n2 – n3, and so on. Since there are 10 measurements, there are 9 pairs.

The right column is obtained by taking the difference between each pair of values in the middle column. The numbers in the right column are very small, since the frequency offset has now been entirely removed. These numbers represent the frequency fluctuations of the DUT. The variance in these numbers is used to estimate the stability of the DUT.

A typical AVAR plot is shown in Figure 2. It shows the stability improving as the measurement period gets longer. Part of this improvement is because measurement system errors become less of a factor as the measurement period gets longer. At some point, however, the oscillator reaches its noise floor and the stability estimates stop improving. Most oscillators reach their noise floor at a measurement period of 1000 seconds or less. The values along the x-axis represent the length of the measurement period in seconds.

Each division represents a measurement period 10 times longer than the previous division. For example, a value of 1 represents a measurement period of 10 seconds (10^1). A value of 2 represents a measurement period of 100 seconds (10^2).

The values along the y-axis show the results of the stability estimate. A value of -10 means that the oscillator has a stability of 1 x 10^{-10}. Be sure not to confuse the stability with the frequency uncertainty when you read the specifications of an oscillator. For example, an oscillator with a frequency uncertainty of 1 x 10^{-8} may still reach a stability of 1 x 10^{-12} in 1000 seconds. This means that the oscillator is not particularly close to its nameplate frequency, but it does produce a stable frequency that changes very little over time.

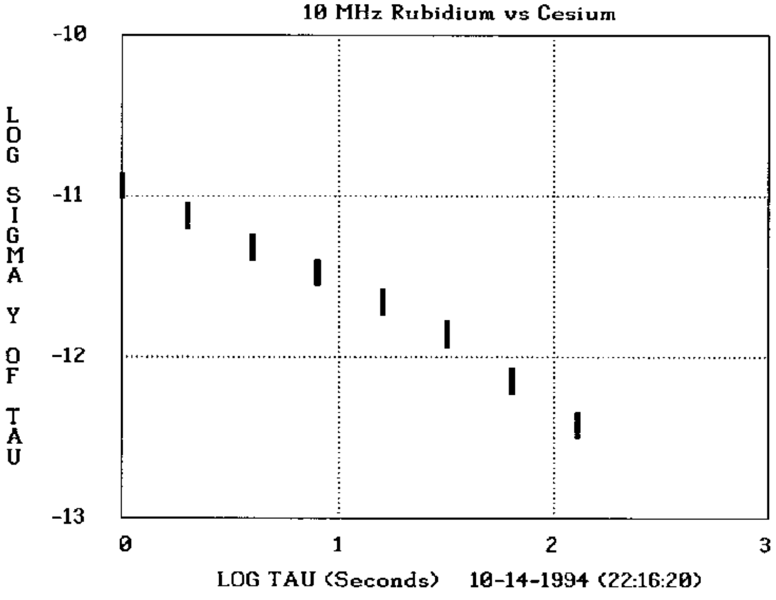

Figure 2. Stability graph (AVAR) of rubidium oscillator.

Oscillators

As stated earlier, the DUT in a frequency calibration is usually a quartz, rubidium, or cesium oscillator. A few high level laboratories may even have an advanced type of oscillator known as a *hydrogen maser. In this section we'll take a brief look at the various types of oscillators available.*

Quartz Oscillators

Quartz oscillators (also called crystal oscillators) are easily the most common type of oscillator. Simple quartz oscillators are found in wristwatches and in many types of electronic circuits. However, calibration laboratories usually only calibrate the more expensive varieties of quartz oscillators, those found inside electronic instruments (like frequency counters) or those designed as standalone units. The cost of a 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. It serves as a mechanical resonator due to the *piezoelectric effect.* This effect causes the crystal to expand or contract when a voltage is applied. The frequency produced by the crystal, called the *resonance frequency, is* determined by the physical dimensions of the crystal and the type of crystal used. The output frequency of a quartz oscillator is either the fundamental resonance frequency or a multiple.

The two biggest factors that influence quartz oscillator performance are temperature and aging. Both factors change the fundamental resonance frequency. Temperature changes can cause a slight change in the elastic properties of the crystal. To get the best performance, the crystal is often enclosed in a temperature-controlled chamber called an *oven.* When a quartz oscillator is first turned on, it goes through a warm-up period while the temperature of the crystal resonator and its oven stabilizes. The warm-up period may last several days or more. During this time, the performance of the oscillator continuously improves until it reaches its normal operating temperature. An alternate solution to the temperature problem is the *temperature compensated crystal oscillator* (TCXO). A TCXO is normally less stable than a crystal with good oven control. Therefore, TCXO's are normally used in small, usually portable units when high performance over a wide temperature range is not required.

Aging is a common trait of all quartz oscillators. It is a nearly linear change in the resonance frequency over time. Often, the resonance frequency decreases, which may indicate that the crystal is getting larger. Aging has many possible causes: crystal contamination due to deposits of foreign material, reforming of loose surface material, or changes in the internal crystal structure. The vibrating motion of the crystal contributes to all of these causes. High quality quartz oscillators age at a rate of 1×10^{-11} per day or less.

In spite of temperature and aging problems, the best quartz oscillators may still achieve uncertainties as small as 1×10^{-11} when properly adjusted. These oscillators typically include an oven for temperature control. The more expensive models may produce several output frequencies which are obtained by dividing and multiplying the resonance frequency.

Less expensive oscillators produce less impressive results. Small ovenized oscillators (like those used as timebases in frequency counters) typically have uncertainties ranging from about 1×10^{-7} to 1×10^{-9}, but cost just a few hundred dollars. At the other end of the spectrum, the tiny quartz oscillators found in wristwatches and electronic circuits may cost less than $1 in single quantity. These oscillators only achieve uncertainties of about 1×10^{-7} in the best case and may be off by as much as 1×10^{-4} in the worst case. Since they are not ovenized, they are quite sensitive to temperature changes.

Since the frequency uncertainty of a quartz oscillator changes substantially over time, adjustments should be made regularly if they are required to perform at their highest level. For example, a quartz oscillator may be capable of 1×10^{-11} performance, but may need regular adjustments to stay at that level. On the other hand, quartz oscillators have excellent short-term stability. A high quality quartz oscillator may be stable to 1×10^{-13} for a measurement period of 1 second. The limitations in stability are mainly due to noise from electronic components in the oscillator circuits.

Atomic Oscillators

Rubidium, cesium, and hydrogen maser oscillators all belong to the category of *atomic oscillators.* These oscillators all work basically the same way. They all contain an internal *voltage-controlled crystal oscillator,* known as a VCXO. The VCXO is locked to a resonance frequency that is generated by the atom of interest. Locking the VCXO to the atomic frequency provides two advantages. First, since the quartz oscillator is controlled by a superior frequency reference, its long-term stability and uncertainty improve. Second, most of the factors that degrade the performance of a quartz oscillator disappear, since the atomic resonance frequency is much less sensitive to environmental conditions than the quartz resonance frequency.

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 contains a quartz oscillator whose frequency is locked to the resonance frequency of the rubidium atom, which is 6,834,682,608 Hz. The result is a very stable frequency source that changes frequency much more slowly than a quartz oscillator without rubidium control. Since rubidium oscillators are more stable, they give better results with fewer adjustments than quartz oscillators. They cost more, but their costs have fallen in recent years. The typical price range for rubidium oscillators is from $3,000 to $8,000. Also, a rubidium oscillator may be less expensive than a quartz oscillator in the long run, because fewer labor costs are involved in keeping the rubidium oscillator adjusted.

The frequency uncertainty of a rubidium oscillator ranges from 5×10^{-10} to 5×10^{-12}. Maintaining frequency to within 1×10^{-11} can be done routinely with a rubidium oscillator and is very difficult to do over long periods of time with even the best quartz oscillators. A well maintained rubidium oscillator can even approach the performance of a cesium oscillator, and a rubidium oscillator is much smaller, more reliable, and less expensive. And although the short term stability of a rubidium oscillator is not as good as a quartz oscillator, its long term stability is much better.

Cesium Oscillators

Cesium oscillators are an *intrinsic frequency standard,* which makes them significant for several reasons. First, the internationally agreed upon definition of the second is based on the resonance frequency of the cesium atom, which is 9,192,631,770 Hz. Second, 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. And finally, since the second is defined based on cesium, a cesium oscillator is assumed to be correct. This means that a cesium oscillator that is working properly should be very close to its correct frequency without any adjustment. There is no change in frequency due to aging.

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 NIST7 with a frequency uncertainty of about 1×10^{-14}. Commercially available cesium oscillators differ in quality, but their frequency uncertainty should be at least 5×10^{-12}.

The two major drawbacks of cesium oscillators are reliability and cost. Reliability is a major issue. The major component of a cesium oscillator, called the *beam tube,* has a life expectancy of from about 3 to 10 years. The beam tube is needed to produce the resonance frequency of the cesium atom, and this frequency is then used to control a quartz oscillator. When the beam tube fails, a cesium oscillator becomes an undisciplined quartz oscillator. Even though a cesium oscillator may not need to be adjusted, it 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. Maintenance costs are also high. In many cases, replacing a beam tube costs nearly as much as replacing 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 of all commercially available frequency standards. Few masers are built, and most are owned by national standards laboratories. 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.

Description	Ovenized Quartz Oscillator	Rubidium Oscillator	Cesium Oscillator	Hydrogen Maser
Intrinsic Standard	No	No	Yes	No
Fundamental Failure Mechanism	None	Rubidium Lamp Life is about 15 years	Cesium Beam Tube (3 to 10 years)	Hydrogen Depletion (7 years plus)
Weight	1-2 lbs.	1-4 lbs.	25-70 lbs.	>100 lbs.
Power Usage	2-5 watts	10-18 watts	25-35 watts	> 100 watts
Stability at 1 second	1×10^{-12} to 1×10^{-13}	5×10^{-11} to 1×10^{-11}	5×10^{-11} to 5×10^{-12}	about 1×10^{-12}
Stability at 1 day	parts in 10^{-10}	1×10^{-12} to 1×10^{-13}	3×10^{-13} to 3×10^{-14}	1×10^{-14} to 5×10^{-15}
Frequency Uncertainty	1×10^{-7} to 1×10^{-11}	5×10^{-10} to 5×10^{-12}	5×10^{-12} to 1×10^{-14}	1×10^{-12} to 1×10^{-13}
Warm-Up Time	< 60 minutes to 1×10^{-8}	<5 minutes to 5×10^{-10}	30 minutes to 5×10^{-12}	24 hours to 1×10^{-12}
Typical Output Frequencies	10 kHz to 100 MHz	1, 5, 10 MHz 1 Hz	1, 5, 10 MHz 1 Hz	1, 5, 10 MHz 1 Hz
Cost	\$300 to \$3,000	\$3,000 to \$8,000	\$30,000 to \$80,000	\$70,000 to \$250,000

Table 3. Specifications of the different types of oscillators.

There are two types of hydrogen masers. The first type, called an *active maser,* phase locks the VCXO to the resonance frequency. This means that the frequency output of the maser is directly derived from the resonance frequency. The second type, called a *passive maser,* frequency-locks the VCXO to the atomic reference. This is the same technique used by rubidium and cesium oscillators. Active masers have better short-term stability than passive masers, but both types of maser have better short-term stability than a cesium oscillator. However, masers are not intrinsic standards, and their frequency uncertainty is greater than that of a cesium oscillator.

The passive maser lends itself better to size reduction and mass production than the active maser. For this reason, most commercially-available masers will probably be passive in coming years. Although the performance of a maser is excellent, its cost is still very high, ranging from about \$70,000 for the least expensive passive masers, to \$250,000 or more for an active maser.

For more information: Michael Lombardi, National Institute of Standards and Technology, Time and Frequency Division, 325 Broadway, Boulder, CO 80303, voice: 303-497-3212, fax: 303-497-3228, lombardi@bldrdoc.gov

References

1. International Organization for Standardization (ISO), *International Vocabulary of Basic and General Terms in Metrology (VIM),* 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. George Kamas and Michael A. Lombardi, *Time and Frequency User's Manual,* NIST Special Publication 559, 1990.

5. Brian W. Petley, "Time and Frequency in Fundamental Metrology," *Proc. of the IEEE,* vol. 79, no. 7, July 1991.

6. Barry N. Taylor and Chris E. Kuyatt, *Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results,* NIST Technical Note 1297, 1993.