Time and Frequency Standards

Here is a thorough review of the types of frequency standards, comparing their accuracy, stability and cost

By Richard Percival
Quartzlock Limited

As we move towards a world increasingly dominated by technology, the need for accurate and reliable frequency and time standards has risen. The market consists of a variety of devices varying in accuracy, stability and price. This article will briefly review some of the medium performance frequency standards and discuss in depth the atomic frequency standards and their indirect derivatives — GPS, Rubidium Gas Cell, Cesium Beam and Hydrogen Masers. The article will then outline the principle behind the Primary Reference Clock currently under development, with an examination of what the future holds for high performance frequency and time standards.

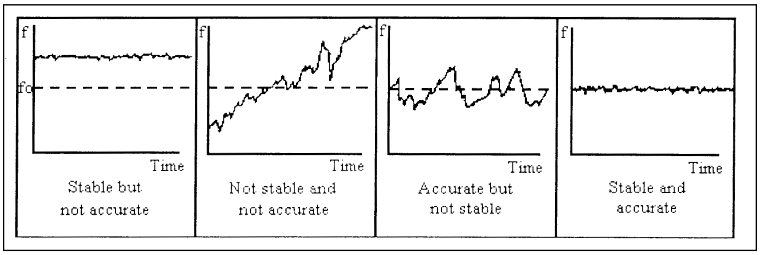

▲ **Figure 1. Comparisons of the basic performance definitions for frequency sources: stability and accuracy.**

The first question: what is a clock? In basic terms a clock may be thought of as a resonator or frequency source with a counter to keep track of the number of oscillations [Allan 2]. On the atomic scale the oscillator would be the Cs^{133} or Rb^{87} atom and the counter would be a fast electronic one. The oscillator is generally referred to as the clock's frequency standard, the oscillations of which are determined by the laws of physics. An example of a simple oscillator would be the quartz crystal found in wrist watches. Quartz crystal oscillators are very stable over short averaging times, but their frequencies can drift quite rapidly. Drift is due to both internal changes (aging) and external environmental changes [NIST 1]. This is why your watch needs periodic calibration against a more accurate clock. Based on the price, however, the accuracy of the quartz crystal is remarkably good.

When a higher degree of accuracy is required for scientific purposes, atomic clocks easily outperform any physical oscillator. The frequency standard is obtained by the oscillation of an electromagnetic signal associated with a quantum transition between two energy levels in an atom. The advantage of using atomic resonant transitions as the reference frequency is that they are determined by fundamental constants, which result from the basic interactions between elementary particles.

To understand how time and frequency standards are compared against one another, it is necessary to define some technical terms which will be used throughout this article.

The most important term used within the time and frequency community is *stability*. It is defined as "the statistical estimate of the frequency fluctuations of a signal over a given period of time." [Lombardi 1]. While this does not indicate the accuracy of a signal, it does indicate whether it changes. It is the stability which indicates the quality of the oscillator.

The other term very often quoted is the *accuracy* or the *fractional frequency offset*. This indicates to the user how accurately the supplied frequency standard agrees with the frequency specified. It is really an indication of how well-adjusted the oscillator is. It indicates coarsely whether the oscillator is subject to drift or aging to any significant degree. It does not, however, indicate the inherent quality of an oscillator.

LF tracking receivers

LF tracking receivers are intended solely for indoor use and will operate at both Low Frequency (LF) and Very Low Frequency (VLF) frequencies. The range will often exceed 2500 km. These receivers are compact lightweight instruments that provide outputs that are phase locked to a standard frequency transmission at a certain frequency. In the US, this will usually be WWVB transmission at 60 kHz. Many such transmitters are traceable to well-recognized national frequency standards via published data “post facto.” This series of standards yields a choice of price/performance tradeoffs and is able to suit many medium to high frequency and (relative) time applications without the expensive acquisition and maintenance of atomic standards. Within reasonable bounds, the results are traceable to a primary reference source. In the UK the primary service is radiated from Droitwich with fill-in coverage from transmitters at Aberdeen and Westerglen on 198 kHz [Quartzlock 1]. The France Inter transmission on 162 kHz provides a corrected cesium reference.

Crystal oscillators

Crystal oscillators are among the most important electronic components in use today and are second only to the atomic sources mentioned below as the most stable frequency devices. Most complex electronic systems rely on a crystal oscillator to provide a stable reference so that other frequencies of the system can be compared to or generated from this reference. More than 1 billion quartz crystal oscillators are produced annually for a variety of applications [Vig].

The link with today’s high-precision oscillators may be traced directly back to the work carried out at Bell Labs in the early 1950s by Warner. Following the development of the transistor by Shockley et al, the first frequency standard based on all-transistor technology was built at Sulzer in 1958 [Norton et al].

The quartz crystal resonator, which is the backbone of the crystal oscillator, uses the piezoelectric effect (by application of an electrical signal, the quartz resonates, and vice versa). As long as this signal is maintained, the crystal will continue to oscillate at a frequency unique to the shape, size and cut of the crystal. During the last 50 years, huge improvements in performance have been made in the design and manufacture of the resonator.

The range of different oscillator types includes simple uncompensated oscillators (XO), temperature compensated (TCXO), microcomputer compensated (MCXO), voltage compensated (VCXO) and oven or double-oven controlled compensated (OCXO). Each has advantages and disadvantages in terms of performance and price. The best oscillator for a given application is determined by a variety of factors, including frequency and accuracy, drift, phase deviation, phase noise, warm-up time and perhaps most important, price. The cost of crystal oscillators varies greatly between the simple XO and the much more accurate and stable OCXO. The historical development of these devices may be found in [Norton et al]. A detailed description of the theory of operation, characteristics and limitations of various compensated oscillators and sources of instability (aging, magnetic field, etc.) is well explained in [Vig].

The Global Positioning System (GPS)

The GPS is owned by the US Department of Defense and is intended to be used mainly as a navigation system. However, due to the method by which precise location is calculated, it is also possible to obtain precise time from the very same system [Kaplan, Kramer and Klische, Ponsonby]. The GPS constellation consists of 24 operational satellites [Davis and Furlong]. Each satellite has on board 2 cesium and 2 rubidium atomic clocks. The GPS system and its Russian counterpart are both extremely complex, using Einstein’s general theory of relativity for the first time in a practical application, thus an explanation is beyond the scope of this article. However, its use for time and frequency dissemination can be described. The GPS carrier signal is operated on two frequencies: L1 is at 1575.42 MHz (10.23 MHz $\times$ 154) and L2 is at 1227.60 MHz (10.23 MHz $\times$ 120). L1 is the only code available to civilians, since it has both the P-code (Precise Code) modulation and the C/A code (Clear/Acquisition) modulation. L2, however, only has the P-Code modulation. The P-code is only available to users trusted by the US DoD. Unfortunately, full accuracy of GPS is denied to the Time Metrology community (and all other “non-authorized” users) by the introduction of pseudo-random noise into the carrier signal [Thomas 1]. As David Allan [Allan 2] said, “If it were turned off, the venetian blinds would go up, and we would be able to see the GPS satellite clocks very clearly.”

At present, a user interested in acquisition of precise

▲ **Figure 2. A receiver can recover accurate frequency information transmitted by GPS satellites.**

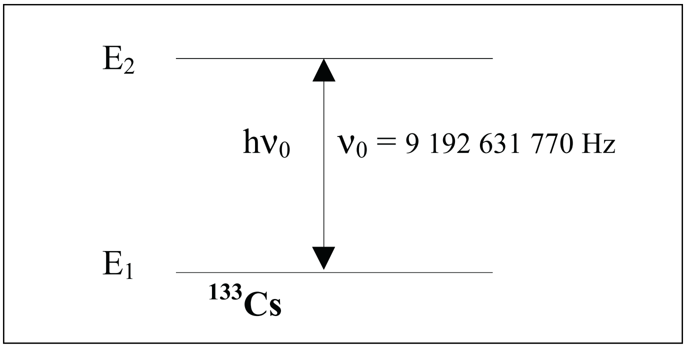

▲ **Figure 3. Energy levels for cesium, hydrogen and rubidium resonances.**

▲ **Figure 4. A commercial rubidium standard.**

time or frequency must accept degradation of stability. A GPS disciplined oscillator (GPS-DO) is an attempt to overcome this short to medium term noise through the slaving of a quartz oscillator or rubidium oscillator to the GPS carrier signal. By combining this stable oscillator with a good disciplining algorithm, the effects of the errors be greatly reduced. It has been shown [Davis and Furlong] that in the short term (<50 s) the stability of a GPS-DO is almost completely determined by the quality of the local slave oscillator and in the long term (>1 day) by the GPS signal itself (generally $<1\times10^{-13}$ at 1 day averaging time).

The use of GPS-DO's as stable and accurate frequency and time standards has significant advantages over free-running oscillators (such as rubidium or cesium), despite the obvious disadvantages mentioned above. The GPS system is referenced to the United States Naval Observatory (USNO), thus, a GPS-DO will not require periodic calibration. The devices are light and easily transportable, enabling them to be used anywhere in the world, which is unimaginable for a high quality cesium beam or hydrogen haser. A simple quartz-based GPS-DO can be bought for under $5,000.

Atomic frequency standards: 1. Rubidium

The first realized rubidium frequency standard arose out of the work of Carpenter [Carpenter et al] and Arditi [Arditi]. The first commercial devices came onto the market primarily due to the work of Packard and Schwartz. Unlike much of the research done on frequency standards at that time, the rubidium maser was high on the research agenda. It was understood that such a device would have extremely good short-term stability relative to size and price. In 1964, Davidovits brought such research to fruition, with the first operational rubidium frequency standard [Vanier and Audoin]. Zepler et al at Plessey, UK did much work on making the rubidium into a small compact unit [Zepler et al, Bennett et al].

The Rubidium Frequency Standard, like its more expensive cousin, the Hydrogen Maser, may be operated either as a passive or as an active device. The passive rubidium frequency standard has proved the most useful, as it may be reduced to the smallest size while retaining excellent frequency stability. The applications for such a device abound in the communication, space and navigation fields.

The rubidium standard may be thought of as consisting of a cell containing the rubidium in its vapor state, placed into a microwave cavity resonant at the hyperfine frequency of the ground state. Optical pumping ensures state selection. The cell contains a buffer gas primarily to inhibit wall relaxation and Doppler broadening. The rubidium frequency standard essentially consists of a voltage-controlled crystal oscillator, which is locked to a highly stable atomic transition in the ground state of the Rb^{87} atom. [Zepler et al, Bennett et al].

Rubidium accuracy is comparable with that of the standard cesium with an operating life approximately five times that of cesium. Furthermore, the cost of a replacement physics package ranges from free to about $50. Moreover, the stability of the rubidium frequency standard over short time-scales (hundreds of seconds) is better than cesium (cesium standards are more stable over longer time periods, from hours to years). The phase noise of a good quality compact rubidium standard is –145 dBc/Hz at 10 kHz from the carrier, identical to the popular HP5071 cesium beam atomic clock [Vanier and Audoin].

There are, however, a few drawbacks to the use of rubidium as a frequency standard. In the past, these included the limited life of the rubidium lamp (since improved to >10 years), although the cesium beam is affected to a greater degree than this, while the H-maser operates differently and is not affected. The thermal stability of rubidium is inferior to that of either cesium or H-masers, and previously required access to a primary reference signal or synchronization source (GPS or Glonass) to maintain long-term cesium-level accuracy [Vanier and Audoin].

The cost of a rubidium frequency standard instrument (around $5000) is significantly cheaper than a cesium standard. Due to its small size, low weight and environmental tolerance, the rubidium frequency stan-

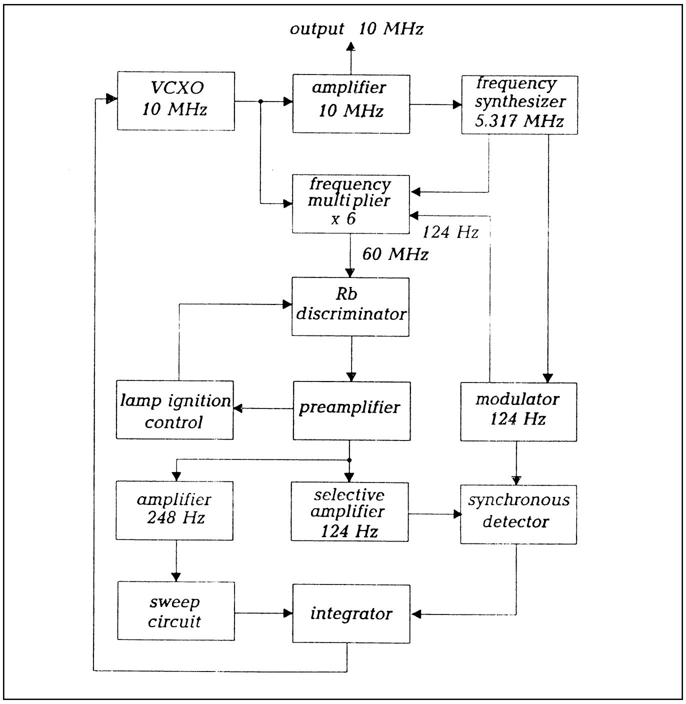

▲ **Figure 5. Block diagram of a rubidium standard using a VCXO locked to the rubidium atomic transitions.**

dard is ideal for mobile applications. Rubidium atomic clocks are beginning to be implemented into the new generation of GPS satellites. Rubidium is also extremely quick to reach operational performance, reaching five parts in 10^{10} within five minutes of turn-on.

GPS disciplined oscillator, (GPS-DO), slaving a rubidium oscillator to a GPC carrier signal, is challenging the dominance once held by cesium within timing centers. A major advantage, if the disciplining algorithm is correctly implemented, is that superior rubidium short-term stability may be obtained, while the characteristic frequency aging is removed. This gives the GPS-DO cesium-derived accuracy at 1/5 the price. It is also able to give Coordinated Universal Time. The future for such devices, if they are accepted as traceable to the national time scale [Davis and Furlong], is very promising.

2. Cesium beam

The heart of a cesium beam clock is the cesium beam resonator. The operation is based on the population difference between the two hyperfine levels of the ground state of the Cs^{133} atom on which the definition of the second is based. It works by passing a beam of state-selected cesium atoms through an excited microwave cavity. On exciting the cavity further, state selection is used to select atoms that have made a microwave transition, and eventually obtain a signal that is maximum when the microwave excitation frequency equals the resonance frequency of the atoms. Some form of frequency or phase modulation of the microwave excitation is used to allow precise determination of the line center.

The major advantage of the cesium beam is its very high accuracy and reproductability. However, the short-term stability of commercial cesium beam standards is rather poor in comparison to hydrogen masers. There exist two types of cesium beam devices — standard and high performance. They differ significantly in accuracy, stability and hence price.

The initial work into developing a Cesium Beam Frequency Standard was carried out by NIST (then NBS) with the first unit built in 1955 [Essen and Parry]. After 1958, the first commercial units became available. This led to the present definition of the SI second which is “the the duration of 9,191,631,770 periods of radiation corresponding to the transition between the hyperfine levels of the ground state of the cesium atom” [BIPM].

Laboratory cesium beam frequency standards are actually built for the purpose of realizing the SI definition of the second with as much accuracy as possible. They are true primary frequency standards because their accuracy is “the normalized uncertainty of the measured or estimated frequency difference between the realized value of the hyperfine transition frequency and the unperturbed transition frequency” [ITU p18]. Neither rubidium nor hydrogen masers are capable of this, regardless of their superior performance in other areas. At present, the best accuracy of a primary Cesium is about 1 to 2×10^{-14} [Bauch]. It is also important to realize that there is a great difference between high performance \$1 million primary cesium standards and those available commercially.

Cesium beam frequency standards are important where high accuracy, reproducibility and negligible drift are required. High end commercial units are capable of the following performance: accuracy $\sim1\times10^{-12}$, drift $<1\times10^{-15}$/day, flicker floor stability level $<1\times10^{-14}$, short-term stability $<8\times10^{-12} \times t^{-1/2}$, and temperature coefficient $<1\times10^{-15}$/ °C, summarized in Figure 6.

Several frequency offsets afflict current cesium beam standards. These are the existence of additional levels very close to the levels of interest, which lead to frequency offsets; residual first and second order Doppler effects due to atomic motion; very small distortion of the atomic wave function; state selection difficulties; coupling of the interrogation magnetic field; and sensitivity of the measured resonance frequency to the quality of the frequency modulated signal used to make the transition. However, primary cesium standards use a complex method of evaluating such shifts and compensating for them [Guinot and Azoubib].

In addition, cesium is much more expensive than rubidium, with a far greater weight (>25 kg). Cesium reliability also depends upon the life-time of the beam tube, which has proven to be poor. Improvements in the beam tube performance have led to reduced tube life. This is a big issue for cesium users, with a new tube being needed every 3 to 7 years at a cost of >$35,000.

Despite these drawbacks, cesium plays a major role in the frequency and time community, with over 80 percent of clocks used in the estimation of TAI/UTC. This is mainly due to their lower long-term drift than hydrogen masers. However, their usage is limited for applications like Very Long Baseline Interferometry (VLBI), where short-term stability requirements are most important.

Unit	Standard	High Performance	Digital Controlled	Primary Laboratory
Accuracy	7E-12	7E-12	1E-12	1E-14
$\sigma_y(\tau=100s)$	3E-12	<1E-12	<1E-12	<1E-13
Flicker Floor (σ_y^{min})	1E-13	3E-14	5E-15	3E-15
$\sigma_y(\tau=1$ month)	1E-13	3E-14	5E-15	N/A
Temperature (per K)	1E-13	1E-13	<1E-15	1E-15
Magnetic Field (per 10^{-4} t)	1E-12	1E-13	<1E-14	N/A

▲ **Figure 6. Performance summary for cesium atomic clocks.**

Recent developments in cesium beam technology have led to standards that operate by optical pumping. This leads to better state selection and atom detection. The most high profile frequency standard developed using this technology is the new primary cesium frequency standard, NIST 7, in operation at NIST. Several of the deleterious frequency shifts affecting current cesium standards are reduced, consequently, performance is very much improved. Such improvements are critical if cesium beam standards are to compete with rubidium-based GPS receivers.

3. Hydrogen masers

The first hydrogen maser was the brainchild Ramsey who, at Harvard University in 1960, succeeded in making the first operational model [Goldenburg et al, Ramsey]. He completed the design after the discovery of the maser effect in 1955 by Townes, Basov and Prokhorov [ITU p. 11]. Kleppner realized that the hydrogen maser was able to deliver an extremely stable frequency reference signal. This led to the consideration of using it as a primary frequency standard. The frequency of the hyperfine transition of the hydrogen atom has been measured with accuracy of the order of 1×10^{-12}.

A major difference among hydrogen masers, cesium beam, and rubidium cell frequency standards is that there is no direct access to the population difference change in hydrogen masers. This is because there is, at present, no efficient means of detecting hydrogen atoms. The basic principle of operation is that the strong coupling between the atomic medium and the microwave field in the resonant cavity makes it very easy to see the necessary hyperfine transition via amplification of the microwave field by stimulating the emission of radiation. If the amplification is large enough, the oscillation may be sustained. The hydrogen maser may be operated actively, as an oscillator or passively, as an amplifier.

The principle of operation for the active hydrogen maser is based upon a 5 MHz quartz crystal oscillator, phase locked to the hyperfine transition of the hydrogen atom. The atomic hydrogen signal generated within the physics package is then picked up by an antenna in the cavity and coupled to the synchronization unit, which synchronizes the 5 MHz crystal oscillator signal phase to the hydrogen signal phase producing the 5 MHz and 100 MHz spectrally pure sine-wave signal output (Figure 7).

The Active Hydrogen Maser (AHM) provides the best known frequency stability for a commercially-available standard. It excels in the domain of 1 second to 1 day. At a 1-hour averaging time, it exceeds the stability of the best known cesium oscillators by a factor of up to 100, with an Allan deviation of $\sim 2 \times 10^{-15}$. The AHM owes its high stability to the following factors [Vanier & Audoin]:

1. The resonant line is narrow due to the long storage time spent in the storage volume (1 s).
2. If the amplifying elements are isolated atoms the noise level of the maser is very low.
3. The storage of atoms at low pressure leading to free and unperturbed movement during radiation.
4. The exposure of atoms to a standing wave leads to removal of the first order Doppler shift. Also the average velocity for stored atoms is very low.

At present, the PHM's stability outperforms the best available cesium by a factor of 10. One of the big advantages of the PHM is that it is not constrained by the 3- to 7-year life cycle of the cesium tube. Figure 8 shows a commercial PHM unit.

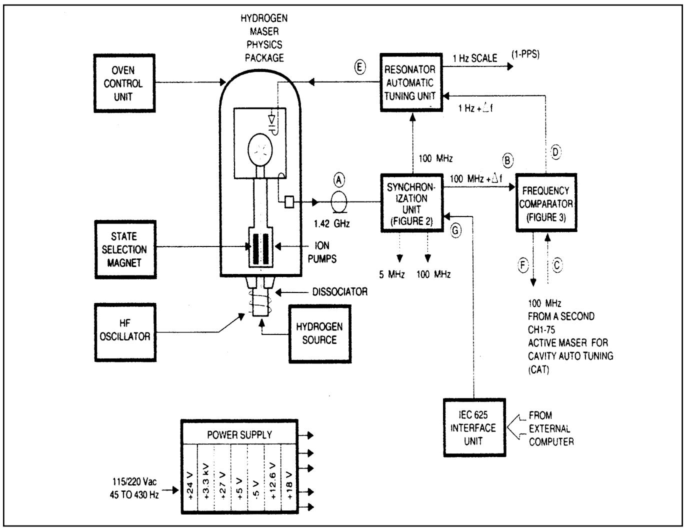

▲ **Figure 7. Hydrogen maser block diagram.**

Long term frequency stability of H-masers depends on whether the Cavity Pulling effect is eliminated by the cavity auto-tuning system (CAT) [Koshelyaevsky and Pushkin]. In CAT systems, long term frequency stability similar to laboratory primary cesium beam frequency standards has been observed over several years. The accuracy of the H-masers is on par with the best commercially available cesium. Reproducibility is the degree to which the frequency standard reproduces its normal output frequency without the need for calibration against another frequency standards. The H-maser has a reproducibility of 10^{-14}, an order of magnitude better than cesium.

There are many present and future applications for hydrogen masers, such as:

1. In radio astronomy for very long baseline interferometry (VLBI) timing for data recording.
2. Laboratory standard/in-house reference. Two active hydrogen masers coupled together by auto cavity tuning form the UK time scale.
3. The ultimate Stratum 1 primary reference clock standard for telephone networks.
4. Satellite ground station clock/frequency standard.
5. Test equipment reference for measuring the quality of GPS disciplined oscillators.
6. Deep space missions clocks.
7. Scientific research, e.g. testing Einstein's theory of special relativity [Wolf and Petit].

Relatively little research has been carried out in the West concerning the development of the H-maser. In contrast, work in the former Soviet Union was far more advanced, with many hundreds of devices having been manufactured and sold [Demidov and Uljanov]. Until now, the market has been rather limited for H-masers, but with the increased emphasis on high frequency stability (e.g. VLBI) rather than high accuracy, the demand is growing rapidly. One example of how this technology is being applied in the west is through a teaming relationship between Quartzlock (UK) Ltd. and KVARZ Institute of Electronic Measurements in Russia. These units have been built in substantial quantities for the last 30 years and form the major time-scales in Germany, Brazil, Spain, Japan, Russia, China, Belgium, Taiwan, S. Korea and the UK. The instruments are produced in Russia and are shipped to Quartzlock, where they are prepared for delivery. It is also possible for these instruments to be tested by NIST in the USA [Allan and Weiss, Weiss and Walls], PTB in Germany, BIPM in France [Azoubib and Thomas] or at NPL in the UK.

H-Masers will continue to be useful for applications where very high stability is required for intervals between 1 second and 105 seconds. However, the medium term frequency stability is limited by the cavity thermal noise and goes as $t^{-1/2}$ [Demidov 1]. Similarly, the short term frequency stability depends on the cavity thermal noise and on the electronic noise in the first stage of amplification. A reduction in the operating temperature of both the cavity and system electronics should lead to an improvement in the frequency stability. The development of a cryogenic H-maser has stimulated a lot of research interest. Early results show that it can become a field operable atomic frequency standard with fractional frequency stability in the 10^{-18} range [Vessot].

▲ **Figure 8. A commercial Passive Hydrogen Maser.**

Primary reference clocks (software clocks)

The objective of a primary reference clock is to combine a number of diverse frequency standards into a single output to achieve improved performance and higher reliability. The output should not show any sudden phase glitches if the output of one or more of the sources should fail. In addition, it is desirable that the random behavior of the output phase of a source be detectable in comparison with the phase of other sources. A simplified bock diagram is shown in Figure 9.

The sources may be divided into two classes. The first is the free running oscillator type of source, which may show a frequency offset from an internationally defined time scale. The hydrogen maser and the rubidium oscillator are in this class. They may have excellent stability, but can still have an unknown frequency offset.

The second class is the type which supplies a replica of a time scale elsewhere. The GPS standard and the LF tracking receiver are of this type. They provide a link to an ensemble of standards, which are monitored and referred to other internationally maintained time scales. The quality of the link controls the short and medium term stability of the standard. In the long term, the stability of this type will approach that of the host.

The local first class standards are likely to have better short and medium term stability than second class standards. For example, the hydrogen maser will be superior to either the GPS standard or the LF tracking receiver for averaging times up to weeks. The final outputs of the PRC should have short-term quality of the hydrogen maser, but long-term quality of the GPS.

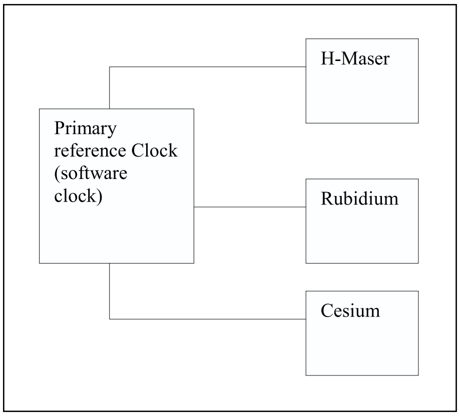

▲ **Figure 9. Three device software clock.**

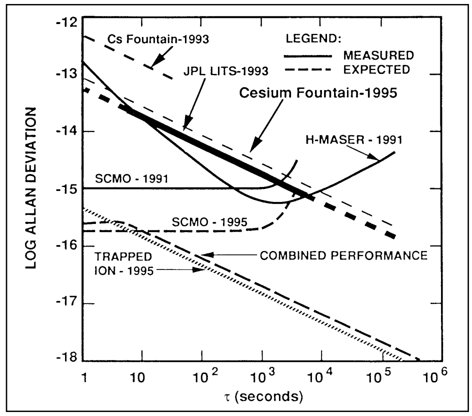

▲ **Figure 10. Peformance of newly-developed standards.**

The redundancy requirement poses several problems. The outputs of the various standards can only be combined if they are all in phase. If we consider a simple system where one source is considered to be the master, then the other sources may be phase aligned to the master. If one source should fail, the amplitude of the output will only drop by the ratio of the number of sources. This simple system would work quite well if all the sources were of the same class and the same degree of frequency (phase) stability.

If we are combining sources of different performance levels, the output should be that of the best standard, in this case the hydrogen maser. However, if that unit should fail, the output may fall to an unacceptably low level. Thus the performance requirement conflicts with the redundancy requirement.

One could improve the short-term stability of the lesser sources by phase comparison with the hydrogen maser. One could then use a higher percentage of the improved (phase adjusted) standard in the final output. If the hydrogen maser did fail, the reference source for the phase adjustment would vanish, and the basic noise characteristics of the lesser sources would reappear. The next best source (e.g., rubidium oscillator) could then be designated the master and used for phase alignment, although it may be difficult to achieve a smooth changeover.

A better solution might be to derive a notional time scale based on phase measurement of all the sources and phase align them to the notional time scale. This way, the need for a master is eliminated. Careful consideration of the weighting of each source contribution to the notional time scale would be required, with each contribution at a different averaging time. To clarify, a source of the second class would have more weighting for the phase averaged over a long period of time as it is referenced to an international time scale.

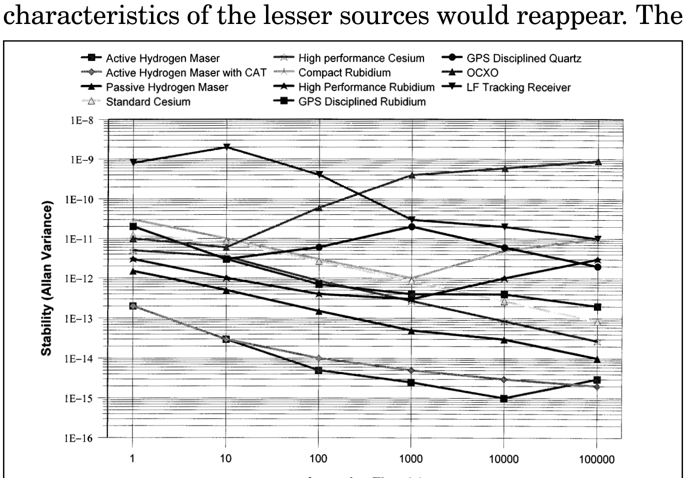

▲ **Figure 12. Comparison of high performance standards.**

The future

A great deal of research is underway to improve present frequency standards. Many devices already exist in laboratories that outperform anything commercially available. They include such exotic sounding devices as trapped ion standards [Tjoelker et al, Fisk], cesium fountains [Weyers et al], oscillators stabilized with cooled sapphire resonators, optical frequency standards [Hamouda et al] and flywheel oscillators. Such high performance devices are almost certain to replace the existing primary standards in laboratories around the world but will have a somewhat restricted niche market. The real challenge lies in the ability to improve crystal oscillators and low-cost atomic standards such as rubidium. This is crucial for 21st Century communication systems, computer networks, navigation and transportation systems including avionics, electric power systems [Martin], space exploration [Emma, Hartl], astronomy and astrometry, geodesy, geology, earthquake monitoring and many others. Performance of these new systems is noted in Figure 10. Figure 12 reviews the relative performance of the various time and frequency standards. ■

References

It was impractical to print the extensive list of references accompanying this article. They have been included with the archived version of this article, which is available via the Internet at www.amwireless.com.

Author information

Richard Percival can be reached at Quartzlock Ltd., Gothic, Plymouth Road, Totnes, Devon TQ9 5LH, UK; tel: +44 1803 862062; fax: +44 1803 86796; e-mail: richard@quartzlock.freeserve.co.uk. The Quartzlock web site is: www.quartzlock.com.