②

Research and Development Technical Report
SLCET-TR-91-2

Frequency Standards for Communications

Samuel R. Stein
Ball Communication Systems Division

and

John R. Vig
U. S. Army Electronics Technology and Devices Laboratory

January 1991

DISTRIBUTION STATEMENT

Approved for public release.
Distribution is unlimited.

U. S. ARMY LABORATORY COMMAND
Electronics Technology and Devices Laboratory
Fort Monmouth, NJ 07703-5000

NOTICES

Disclaimers

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
January 1991

3. REPORT TYPE AND DATES COVERED
Technical Report

4. TITLE AND SUBTITLE
Frequency Standards for Communications

5. FUNDING NUMBERS
PE# 1L1
PR# 62705
TA# AH94

6. AUTHOR(S)
Samuel R. Stein, Ball Communication Systems Division
John R. Vig, U.S. Army Electronics Technology & Devices Lab

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
US Army Laboratory Command (LABCOM)
Electronics Technology and Devices Laboratory (ETDL)
ATTN: SLCET-EQ
Fort Monmouth, NJ 07703-5000

8. PERFORMING ORGANIZATION REPORT NUMBER
SLCET-TR-91-2

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES This report is a preprint of a chapter that is to be published in The Froehlich/Kent ENCYCLOPEDIA OF TELECOMMUNICATIONS, Fritz E. Froehlich and Alan Kent editors, Marcel Dekker, Inc. Samuel R. Stein, is at Ball Communication Systems Div., P.O. Box 1235, Broomfield, CO 80038-1235.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The fundamentals of quartz and atomic frequency standards are reviewed. The subjects discussed include: crystal resonators and oscillators, atomic oscillators, oscillator types, and the characteristics and limitations of temperature-compensated crystal oscillators (TCXO), oven-controlled crystal oscillators (OCXO), rubidium frequency standards, cesium beam frequency standards and hydrogen masers. The oscillator instabilities discussed include: aging, noise, frequency vs. temperature, warmup, acceleration effects, magnetic field effects, atmospheric pressure effects, radiation effects, and interactions among the various effects. Guidelines are provided for oscillator comparison and selection. A discussion of time transfer techniques, and specifications are also included, as are references and suggestions for further reading.

14. SUBJECT TERMS Oscillator, clock, frequency standard, frequency control, frequency stability, time, timing devices, quartz, quartz crystal, quartz oscillator, atomic clock, atomic frequency standard. rubidium standard. cesium standard, hydrogen maser, stability, aging, noise, phase noise

15. NUMBER OF PAGES
72

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT
Unclassified

18. SECURITY CLASSIFICATION OF THIS PAGE
Unclassified

19. SECURITY CLASSIFICATION OF ABSTRACT
Unclassified

20. LIMITATION OF ABSTRACT
UL

Table of Contents

Figures

Tables

I. The Role of Frequency Standards in Communications

Stable oscillators control the frequencies of communication systems. As communication technology evolved, improvements in oscillator technology allowed improved spectrum utilization. (The evolution of spectrum utilization in commercial two-way communication systems is reviewed elsewhere in this* Encyclopedia [1].) In modern systems, stable oscillators are used not only for frequency control but also for timing.

In digital telecommunication [2, 3] and spread-spectrum [4] systems, synchronization plays critically important roles. In digital telecommunication systems, stable oscillators in the clocks of the transmitters and receivers maintain synchronization, and thereby ensure that information transfer is performed with minimal buffer overflow or underflow events; i.e., with an acceptable level of slips. Slips cause problems such as missing lines in FAX transmission, clicks in voice transmission, loss of encryption key in secure voice transmission, and the need for data retransmission. In AT&T's network [2] for example, timing is distributed down a hierarchy of nodes. A timing source-receiver relationship is established between pairs of nodes containing clocks, which are of four types, in four "stratum levels." The long-term accuracy requirements of the oscillators range from 1×10^{-11} at stratum 1 to 3.2×10^{-5} at stratum 4.

The phase noise of the oscillators can lead to erroneous detection of phase transitions; i.e., to bit errors, when phase shift keyed digital modulation is used. For example, assuming a normal distribution of phase deviations and an RMS phase deviation of ±4.5°, the probability of exceeding a ±22.5° phase deviation is 6×10^{-7}. Spread spectrum techniques are being used increasingly in both military and civilian communication systems. The advantages of spread spectrum use can include: 1) rejection of intentional and unintentional jamming, 2) low probability of interception, 3) selective addressing, 4) multiple access, and 5) more efficient use of the frequency spectrum. As an illustration of the importance of accurate clocks in such systems, consider one type of spread spectrum modulation - frequency hopping. In a frequency hopping system, accurate clocks must insure that the transmitter and receiver hop to the same frequency at the same time. The faster the hopping rate, the higher the jamming resistance, and the more accurate the clocks must be. For example, for a hopping rate of 1,000 hops per second, the clocks must be synchronized to about 100 μs. System parameters such as the autonomy period (radio silence interval) and the time required for signal acquisition (net entry) are also closely dependent on clock accuracy.

II. Frequency Standards

Frequency standards can be divided into two major types: quartz crystal oscillators

---

*This report is a preprint of a chapter that is to be published in Vol. 3 of The Froehlich/Kent ENCYCLOPEDIA OF TELECOMMUNICATIONS, Fritz E. Froehlich and Alan Kent editors, Marcel Dekker, Inc., in late 1991.

and atomic frequency standards. All atomic frequency standards contain a crystal oscillator.

A. The Generalized Crystal Oscillator

1. Description

Figure 1 is a greatly simplified circuit diagram that shows the basic elements of a crystal oscillator [5, 6]. The amplifier section consists of at least one active device, the necessary biasing networks, and may include other elements for band-limiting, impedance matching, and gain control. The feedback network consists of the crystal unit, and may contain other elements, such as a variable capacitor for tuning.

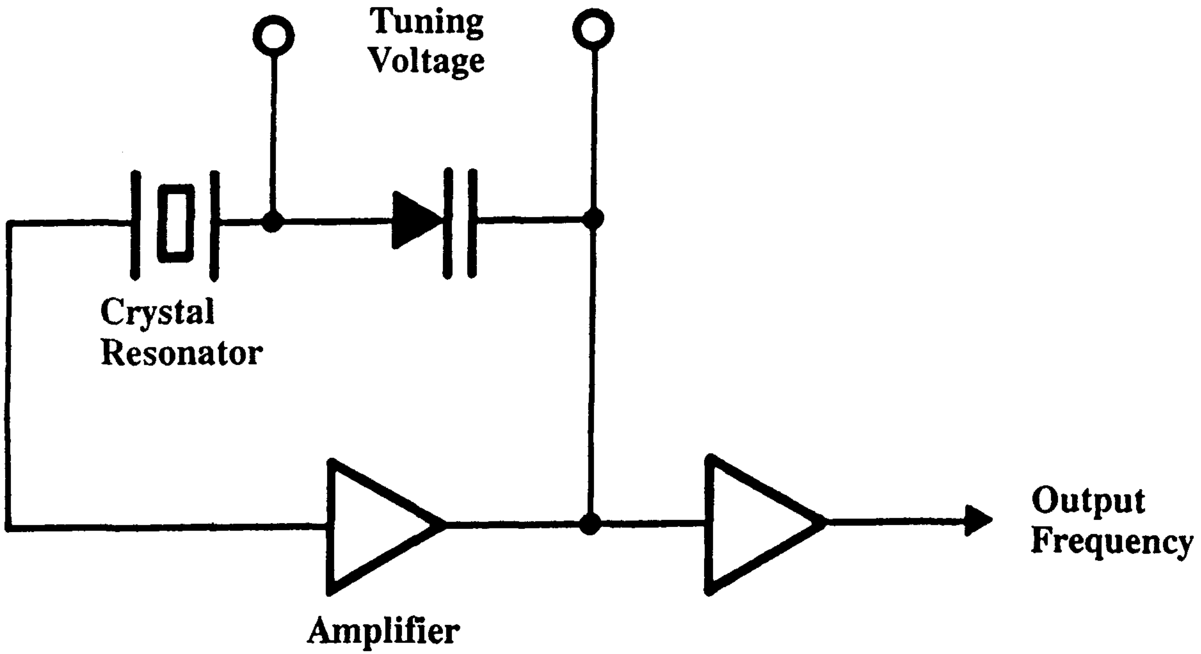

Figure 1. Crystal oscillator - simplified circuit diagram.

The frequency of oscillation is determined by the condition that the closed loop phase shift = $2n\pi$, where n is an integer, usually 0 or 1. When the oscillator is initially energized, the only signal in the circuit is noise. That component of noise, the frequency of which satisfies the phase condition for oscillation, is propagated around the loop with increasing amplitude. The rate of increase depends on the excess loop gain and on the bandwidth of the crystal network. The amplitude continues to increase until the amplifier gain is reduced, either by the nonlinearities of the active elements ("self-limiting") or by an external level-control method.

At steady-state, the closed-loop gain = 1. If a phase perturbation $\Delta\phi$ occurs, the frequency of oscillation must shift by a Δf in order to maintain the $2n\pi$ phase condition. It can be shown [5] that for a series-resonance oscillator

$$\frac{\Delta f}{f} = -\frac{\Delta\phi}{2Q_L},$$

where Q_L is the loaded Q of the crystal in the network. ("Crystal" and "resonator" are often used interchangeably with "crystal unit," although "crystal unit" is the official name. See references [7] and [8], and chapter 3 of reference [6] for further information about crystal units.)

A quartz crystal unit is a quartz wafer to which electrodes have been applied, and which is hermetically sealed in a holder structure. Although the design and fabrication of crystal units comprise a complex subject, the oscillator designer can treat the crystal unit as a circuit component and just deal with the crystal unit's equivalent circuit. Figure 2 shows a simplified equivalent circuit, together with the circuit symbol for a crystal unit. A "load capacitor" C_L is shown in series with the crystal. The mechanical resonance of the crystal is represented by the "motional parameters" L_1, C_1, and R_1. Because the crystal is a dielectric with electrodes, it also displays a "static capacitance" C_0. Figure 3 shows the reactance vs. frequency characteristic of the crystal unit. When the load capacitor is connected in series with the crystal, the frequency of operation of the oscillator is increased by a $\Delta f'$, where $\Delta f'$ is given by

$$\frac{\Delta f'}{f} \approx \frac{C_1}{2(C_0+C_L)}.$$

When an inductor is connected in series with the crystal, the frequency of operation is *decreased*.

2. Stability vs. Tunability

In most crystal oscillator types, a variable load capacitor is used to adjust the frequency of oscillation to the desired value. Such oscillators operate at the "parallel resonance" region of Fig. 3, where the reactance vs. frequency slope, i.e., the "stiffness,"is inversely proportional to C_1. For maximum frequency stability with respect to reactance (or phase) perturbations in the oscillator circuit, the reactance slope (or phase slope) must be maximum. This requires that the C_1 be minimum, and that the Q_L be maximum. The smaller the C_1, however, the more difficult it is to tune the oscillator;

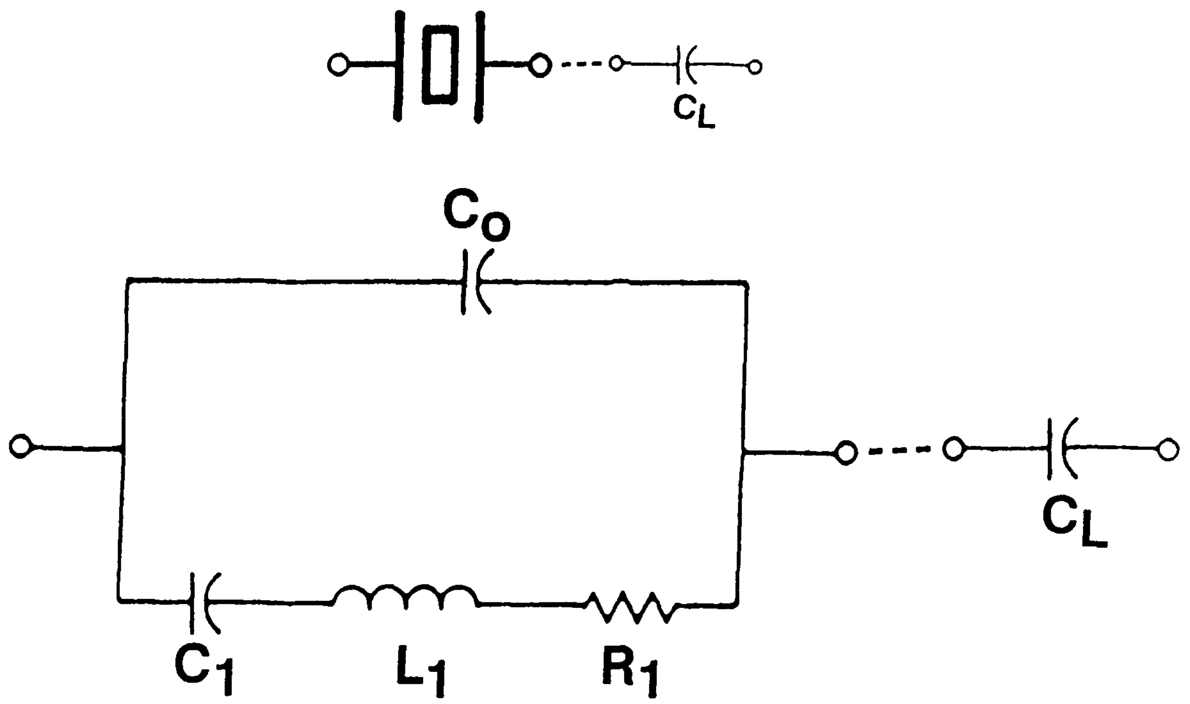

$$\frac{\Delta f}{f_S} \approx \frac{C_1}{2(C_O + C_L)} \longrightarrow \left\{ \begin{array}{l} \text{1. Voltage control (VCXO)} \\ \text{2. Temperature compensation (TCXO)} \end{array} \right.$$

Figure 2. Crystal unit with load capacitor - simplified equivalent circuit.

i.e., the smaller is Δf' for a given change in C_L. The highest stability oscillators use crystal units that have a small C_1. Since C_1 decreases rapidly with overtone number, high stability oscillators generally use third or fifth overtone crystal units. Overtones higher than fifth are rarely used, because R_1 also increases rapidly with overtone number, and some tunability is desirable in order to allow setting the oscillator to the desired frequency.

Wide tuning range voltage-controlled crystal oscillators (VCXOs) use fundamental mode crystal units of large C_1. Voltage control is used: to frequency or phase lock two oscillators; for frequency modulation; for compensation, as in a TCXO (see below); and for calibration; i.e., for adjusting the frequency to compensate for aging. Whereas a high-stability, ovenized 10 MHz VCXO may have a frequency adjustment range of ±5 X 10^{-7} and an aging rate of 2 X 10^{-8} per year, a wide tuning range 10 MHz VCXO may have a tuning range of ±50 ppm and an aging rate of 2 ppm per year.

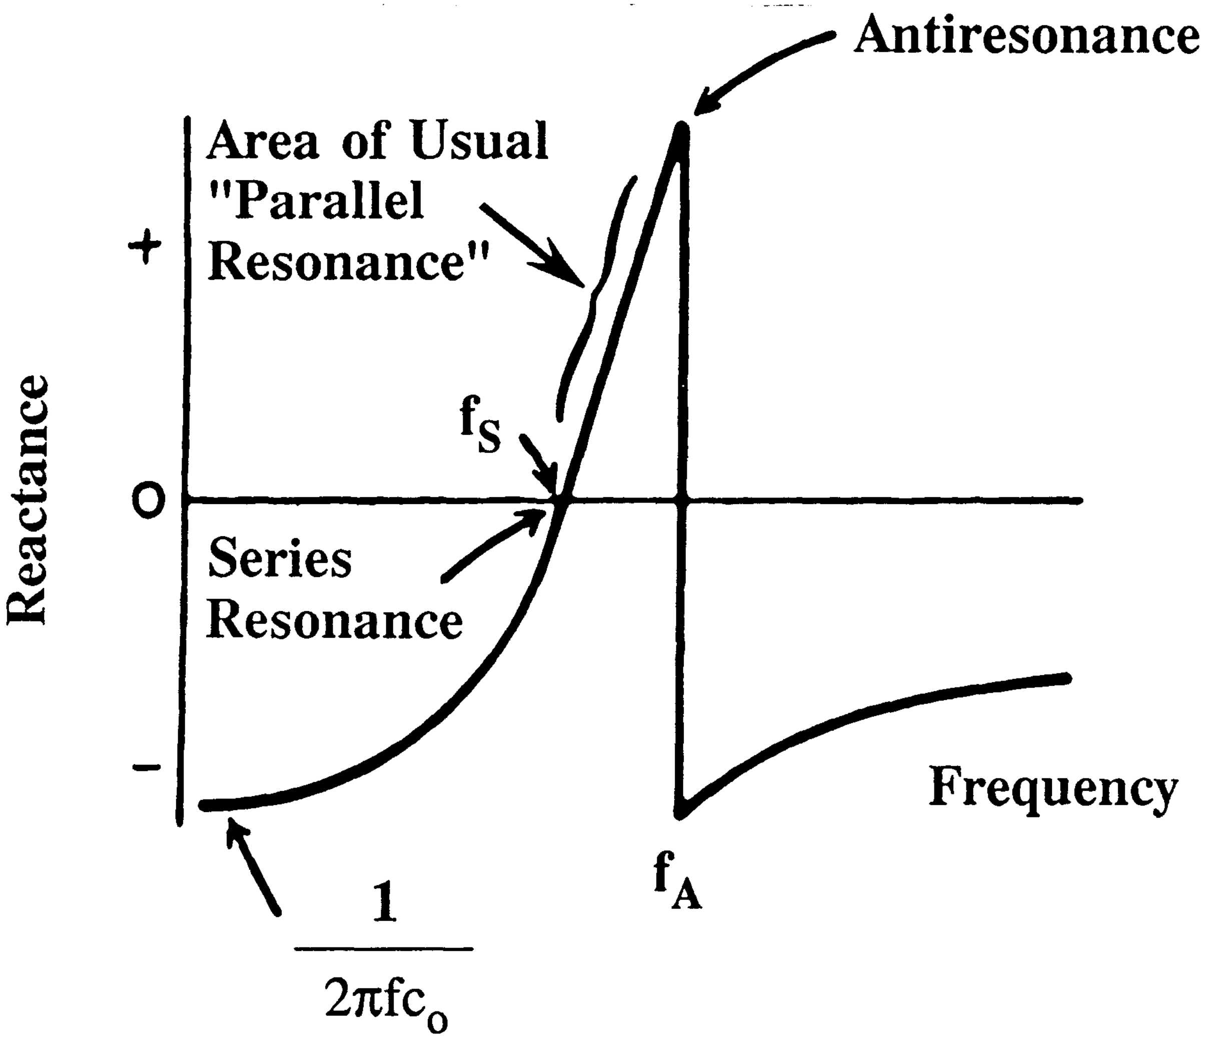

Figure 3. Reactance vs. frequency of a crystal unit.

In general, making an oscillator tunable over a wide frequency range degrades its stability because making an oscillator susceptible to intentional tuning also makes it susceptible to factors that result in unintentional tuning. For example, if an OCXO is designed to have a stability of 1 X 10^{-12} for some averaging time and a tunability of 1 X 10^{-7}, then the crystal's load reactance must be stable to at least 1 X 10^{-5} for that averaging time. Achieving such load reactance stability is difficult because the load reactance is affected by stray capacitances and inductances, by the stability of the varactor's capacitance vs. voltage characteristic, and by the stability of the voltage on the varactor. Moreover, the 1 X 10^{-5} load reactance stability must be maintained not only under benign conditions, but also under changing environmental conditions (temperature, vibration, radiation, etc.). Therefore, the wider the tuning range of an oscillator, the more difficult it is to maintain a high stability.

3. The Quartz Crystal Unit

A quartz crystal unit's high Q and high stiffness makes it the primary frequency, and frequency stability, determining element in a crystal oscillator. The Q values of crystal units ($Q^{-1} = 2\pi f_s R_1 C_1$) are much higher than those attainable with other circuit elements. In general purpose crystal units, Q's are generally in the range of 10^4 to 10^6. A high-stability 5 MHz crystal unit's Q is typically in the range of two to three million. The intrinsic Q, limited by internal losses in the crystal, has been determined experimentally to be inversely proportional to frequency; i.e., the Qf product is a constant. The maximum Qf = 16 million when f is in MHz.

Quartz (which is a single-crystal form of SiO_2) has been the material of choice for stable resonators since shortly after piezoelectric crystals were first used in oscillators - in 1918. Although many other materials have been explored, none has been found to be better than quartz. Quartz is the only material known that possesses the following combination of properties: 1) it is piezoelectric ("pressure electric"; piezein means to press in Greek); 2) zero temperature coefficient resonators can be made from quartz plates when the plates are cut properly with respect to the crystallographic axes of quartz; 3) of the zero temperature coefficient cuts, one, the SC-cut, is "stress compensated"; 4) it has low intrinsic losses, i.e., quartz resonators can have very high Q's; 5) it is easy to process because it is hard but not brittle, and, under normal conditions, it has low solubility in everything except the fluoride etchants; and 6) it is abundant in nature, it is easy to grow in large quantities, at low cost, and with relatively high purity and perfection. Of the man-grown single crystals, quartz, at more than 2000 tons per year (in 1990), is second only to silicon in quantity grown.

Quartz crystals are highly anisotropic; i.e., the properties vary greatly with crystallographic direction. For example, when a quartz sphere is etched in hydrofluoric acid, the etching rate is more than 100 times faster along the fastest etching rate direction, the Z-direction, than along the slowest direction, the slow-X-direction. The "constants" of quartz, such as the thermal expansion coefficient and the temperature coefficients of the elastic constants, also vary with direction. That crystal units can have zero temperature coefficients of frequency is a consequence of the temperature coefficients of the elastic constants ranging from negative to positive values.

The locus of zero temperature coefficient cuts in quartz is shown in Fig. 4. The X, Y, and Z directions are chosen to make the description of properties as simple as possible. The Z-axis in Fig. 4 is an axis of threefold symmetry in quartz; i.e., the physical properties repeat every 120° as the crystal is rotated about the Z-axis. The cuts usually have two-letter names, where the "T" in the name indicates a temperature compensated cut; e.g., the AT-cut was the first such cut discovered. (The SC-cut is also temperature compensated; the name of this cut resulted from it being the only temperature and stress-compensated cut known.) The highest stability crystal oscillators employ SC-cut or AT-cut crystal units.

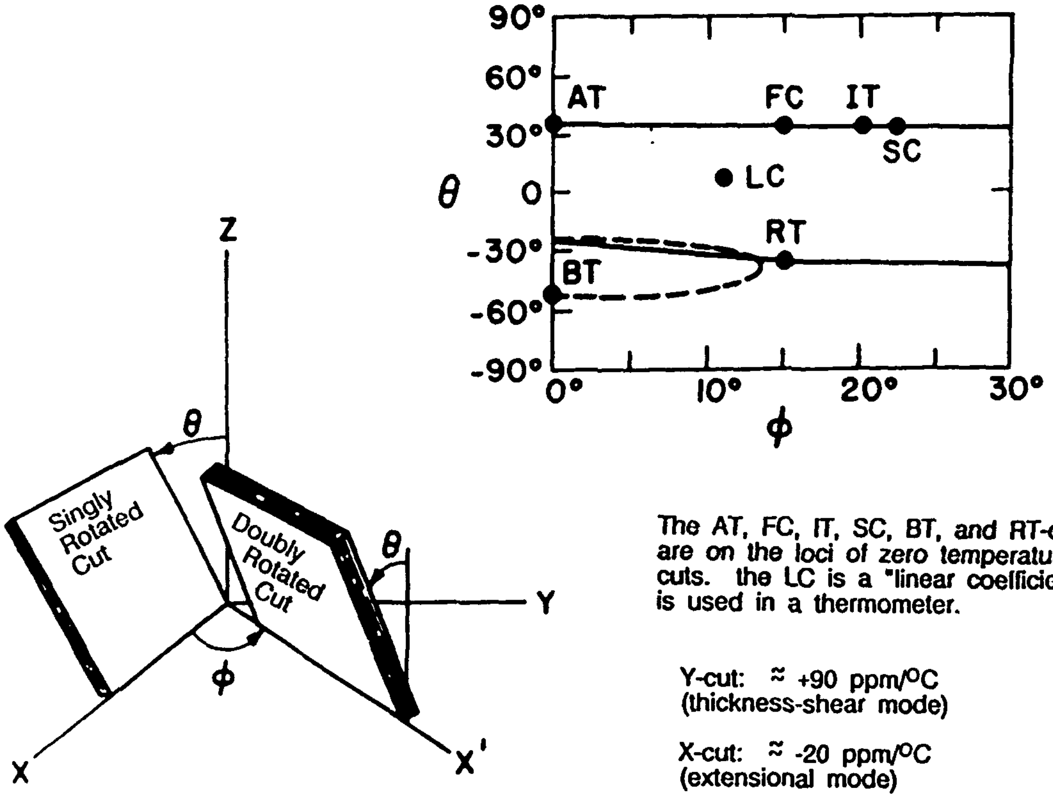

The AT, FC, IT, SC, BT, and RT-cuts are on the loci of zero temperature coefficient cuts. the LC is a "linear coefficient" cut that is used in a thermometer.

Y-cut: ≈ +90 ppm/°C
(thickness-shear mode)

X-cut: ≈ -20 ppm/°C
(extensional mode)

Figure 4. Zero temperature coefficient cuts of quartz.

Because the properties of a quartz crystal unit depend strongly on the angles of cut of the crystal plate, in the manufacture of crystal units, the plates are cut from a quartz bar along precisely controlled directions with respect to the crystallographic axes. After shaping to required dimensions, metal electrodes are applied to the wafer which is mounted in a holder structure. Figure 5 shows the two common holder structures used for resonators of higher than 1 MHz frequency. (The 32 kHz tuning fork resonators used in quartz watches are, typically, packaged in small tubular enclosures.)

Because quartz is piezoelectric, a voltage applied to the electrodes causes the quartz plate to deform slightly. The amount of deformation due to an alternating voltage depends on how close the frequency of the applied voltage is to a natural mechanical resonance of the crystal. To describe the behavior of a resonator, the differential equations for Newton's laws of motion for a continuum, and for Maxwell's equations, must be solved with the proper electrical and mechanical boundary conditions at the plate surfaces. Because quartz is anisotropic and piezoelectric, with ten independent linear constants and numerous higher order constants, the equations are complex, and have never been solved in closed form for physically realizable three-dimensional resonators. Nearly all theoretical works have used approximations. The nonlinear elastic constants, although small, are the source of some of the important instabilities of crystal oscillators; i.e., the acceleration sensitivity, the thermal-transient effect, and the amplitude-frequency effect, each of which is discussed later.

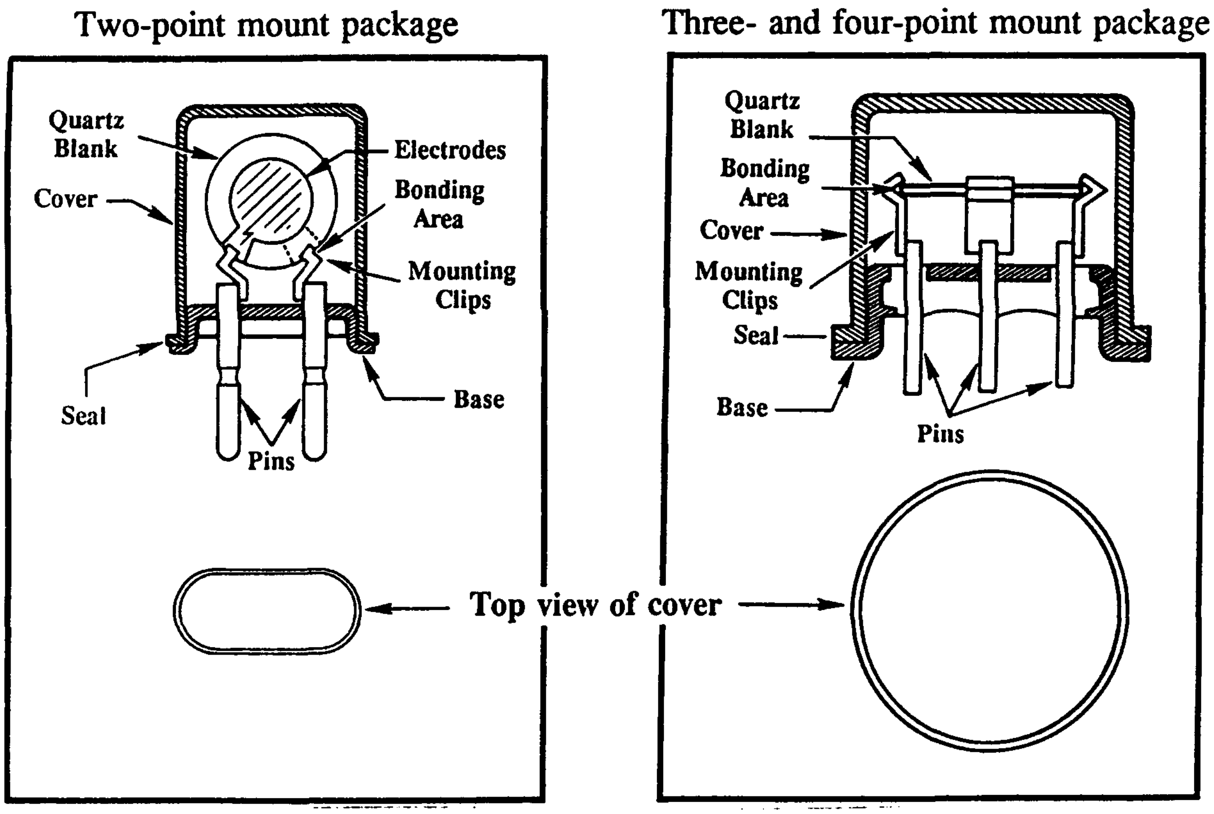

Figure 5. Typical constructions of AT-cut and SC-cut crystal units.

As the drive level, i.e., the current through a crystal, increases, the crystal's amplitude of vibration also increases, and the effects due to the nonlinearities of quartz become more pronounced. Among the many properties that depend on the drive level are: the resonance frequency, the motional resistance R_1, the phase noise, and frequency vs. temperature anomalies, called "activity dips" (which is discussed later). The drive level dependence of the resonance frequency is called the "amplitude-frequency effect." The frequency change with drive level is proportional to the square of the drive current. Because of the drive level dependence of frequency, the highest stability oscillators usually contain some form of automatic level control in order to minimize frequency changes due to oscillator circuitry changes. At high drive levels, the nonlinear effects also result in an increase in the resistance. Crystals can also exhibit anomalously high starting resistance when the crystal surfaces possess imperfections such as scratches and particulate contamination. Under such conditions, the resistance at low drive levels can be high enough for an oscillator to be unable to start when power is applied.

Bulk-acoustic-wave quartz resonators are available in the frequency range of about 1 kHz to 500 MHz. Surface-acoustic-wave (SAW) quartz resonators are available in the range of about 150 MHz to 1.5 GHz. To cover the wide range of frequencies, different cuts, vibrating in a variety of modes, are used. The bulk-wave modes of motion are shown in Fig. 6. The AT-cut and SC-cut crystals vibrate in a thickness shear mode.

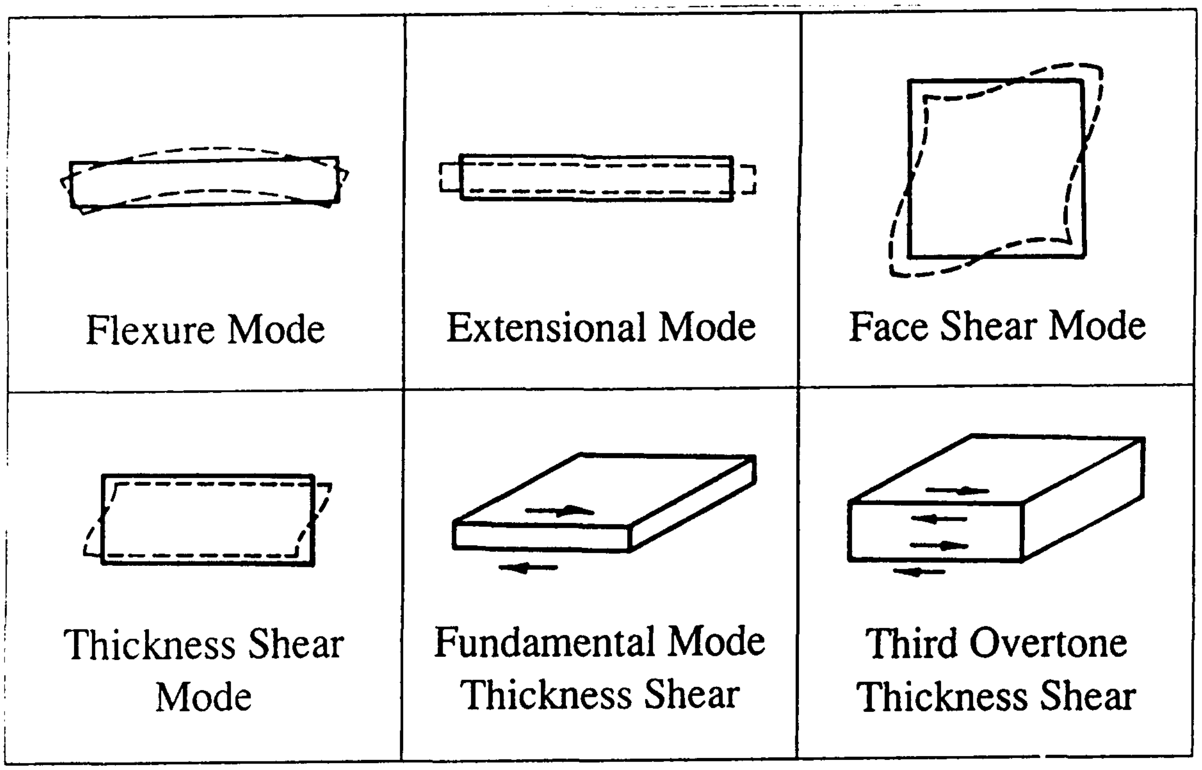

Figure 6. Modes of motion of a quartz resonator.

Although the desired thickness shear mode will exhibit the lowest resistance, the mode spectrum of even properly designed crystal units exhibit unwanted modes above the main mode. The unwanted modes, also called "spurious modes" or "spurs," are especially troublesome in filter crystals, in which "energy trapping rules" are employed to maximize the suppression of unwanted modes. These rules specify certain electrode geometry to plate geometry relationships. In oscillator crystals, the unwanted modes may be suppressed sufficiently by providing a large enough plate diameter to electrode diameter ratio, or by contouring; i.e., generating a spherical curvature on one or both sides of the plate.

Above 1 MHz, the AT-cut is commonly used. For high-precision applications, the SC-cut has important advantages over the AT-cut. The AT-cut and SC-cut crystals can be manufactured for fundamental mode operation up to a frequency of about 200 MHz. (Higher than 1 GHz units have been produced on an experimental basis.) Above 100 MHz, overtone units that operate at a selected harmonic mode of vibration are generally used. Below 1 MHz, tuning forks, X-Y and NT bars (flexure mode), +5° X-cuts (extensional mode), or CT-cut and DT-cut units (face shear mode) can be used. Tuning forks have become the dominant type of low frequency units due to their small size and low cost. Hundreds of millions of quartz tuning forks are produced annually for quartz watches and other applications.

B. The Generalized Atomic Oscillator

1. Description

An atomic frequency standard is a device the frequency of which is determined by some property of a simple atomic system. The term is not restricted to devices where the frequency derives from neutral atoms, but is also applied to devices based on molecules and ions. The terms "atomic clock" and "atomic oscillator" are often used synonymously. However, sometimes the term "clock" is used to mean a frequency standard connected to a counter - a device that provides time as well as frequency. There are several good reviews of atomic frequency standards [9, 10, 11].

Atomic frequency standards must be understood in terms of the concepts of quantum mechanics. The properties of simple atomic systems cannot assume arbitrary values. For example, the energies of the bound states of an atomic system are constrained to discrete values called energy levels. When an atomic system changes energy from an excited state to a state with lower energy, it emits a quantity of electromagnetic energy called a photon, the frequency of which is determined by the energy difference between the two states. If the energy of the upper state is E_2 and the energy of the lower state is E_1, the photon frequency is given by Planck's law

$$\nu = \frac{E_2 - E_1}{h},$$

where h is Planck's constant. The atomic standard produces an output signal the frequency of which is determined by this intrinsic atomic frequency, rather than some property of a bulk material.

The advantages of atomic oscillators all stem from this feature. Intrinsic atomic properties are more easily reproduced than collective properties, endowing atomic frequency standards with the property of accuracy. Atomic systems are easy to isolate from unwanted perturbations, which result in very small sensitivities to temperature, pressure, and other environmental conditions. The low level of interaction also results in extremely sharp resonance features, and reduces errors due to imperfections in the electronics. All atoms of an element are identical, and atomic properties are time invariant, which makes it possible to build very stable devices. Finally, it is surprisingly easy to measure atomic properties and build practical devices suitable for a wide variety of applications.

Atomic frequency standards are categorized in several ways; most often, they are referred to by the type of atom: hydrogen, rubidium, or cesium. Actually, these three devices are based on the same type of atomic interaction, but there are great practical differences in their implementation. Some atomic oscillators are active; in this case, the output signal is derived from the radiation emitted by the atom; others are passive. The atoms are then employed as a discriminator to measure and control the frequency of an electronic oscillator, such as a quartz oscillator. The third classification follows the method of interaction. In atomic beams, the atoms are observed "on the fly"; they pass through the interaction region and are not used again. In contrast, storage devices contain some type of cell that holds the atoms to be observed for a much longer time. In some cases the atoms are recycled.

2. Atomic Spectroscopy

The energy levels of an atom are generally classified according to their physical origin. For example, the principal levels of an atom are associated with the radius of the 'orbit' of an electron about the nucleus. These levels have the largest atomic energy separations. The principal energy levels are subdivided as a result of the quantization of the angular momentum of the atom. The angular momentum due to the motion of a particle, such as an electron, is called "orbital angular momentum". Even when their motion is such that there is no orbital angular momentum, atomic particles may possess an intrinsic angular momentum or spin and a proportional intrinsic magnetic moment. This is another concept unique to quantum mechanics. The principal levels are first divided according to the shape of the electron 'orbits.' Still finer division occurs as a consequence of the particular orientation of the electron's spin and the spin of the nucleus.

The photons emitted when atoms change states between the principal energy levels are usually in the infrared and higher energy regions of the electromagnetic spectrum. The frequencies of these very energetic photons are too high for practical electronic devices. However, very narrow spectroscopic features associated with principal energy levels have been obtained in the laboratory and are useful for relative measurements. Atomic frequency standards are feasible because of the splitting of the ground state of the atom. Next lower, in terms of energy, is the fine structure of the atom, which results from the interaction of the spin of the electron with the magnetic field due to the motion of the electron through the nuclear electric field. This structure is thousands of times smaller than the separation of the principal energy levels. Laboratory atomic frequency standards based on fine structure in calcium and magnesium have been built, but the fundamental frequencies of the atomic transitions are higher than 600 GHz, which is very difficult to synthesize [12].

A finer energy splitting than the spin-orbit coupling is produced by the interaction of the electron and nuclear spins; this is called the hyperfine structure. The ground state of a hydrogen-like atom, for example H, Li, Na, K, Rb, Cs, and singly ionized Be, has a single unpaired electron in a symmetric orbit. In this case, there is no orbital angular momentum and *no fine structure.* The energy splitting due to the intrinsic magnetic moments of the electron and the nucleus can be a million times smaller than the separation of the principal energy levels. The transition frequencies are quite convenient - 1.4 GHz for hydrogen, 6.8 GHz for rubidium, and 9.2 GHz for cesium. All commercial atomic frequency standards are based on the hyperfine spectroscopy of one of these three atoms.

Because the frequency of such a device is determined by the energy of interaction of a pair of magnetic moments, it is generally altered by any background magnetic field. The hyperfine states are also split into multiple energy levels, called Zeeman sublevels, depending upon the component of the total angular momentum in the direction of an applied magnetic field. Figure 7 shows the hyperfine structure of the hydrogen atom in an applied magnetic field. The two states which have no angular momentum component along the direction of the applied field have quadratic dependence of energy on magnetic field for small field values. Thus, the frequency corresponding to a transition between these two levels has very small dependence on the magnetic field and is suitable for use in an atomic frequency standard. Atomic frequency standards utilize a small background magnetic field to allow the selection of the desired magnetic sublevels. The transitions between other levels have too high magnetic field dependence for frequency standard applications - variations in external magnetic fields, such as the earth's field, cannot be shielded sufficiently well to prevent them from perturbing the frequency of the device.

The magnetic field dependent transitions are used to manufacture sensitive magnetometers rather than frequency standards.

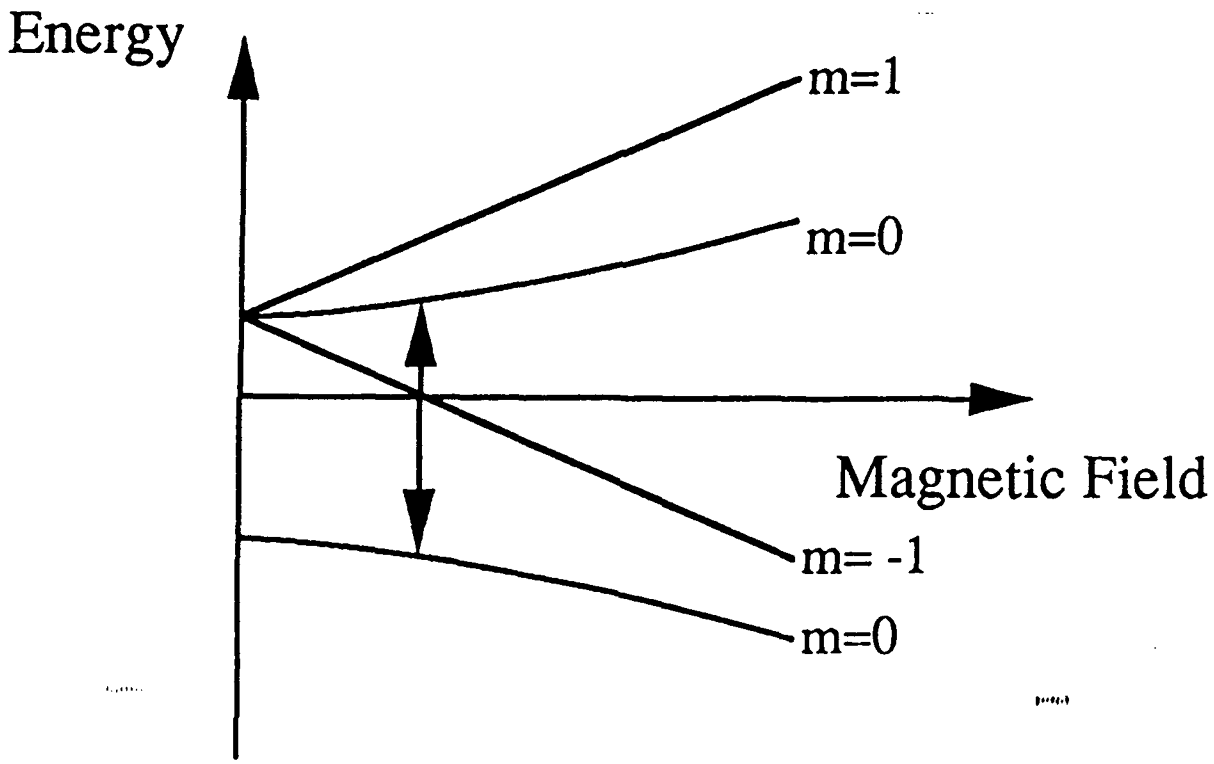

Figure 7. Magnetic field dependence of the hyperfine doublet in the ground state of hydrogen.

The process of measuring the frequency of atomic transitions is called atomic spectroscopy. Atoms are prepared in one atomic state and then stimulated to change to the second state while some method of observing the transition is employed. The atoms must be stimulated to make the transition because the lifetime of the upper hyperfine state is very long; it is an appreciable fraction of the age of the universe in the case of hydrogen. By applying a field at the transition frequency, atoms can be stimulated to make a rapid transition between hyperfine levels. Typically, the atoms change state in a few milliseconds to one second, during which time they are observed.

Since the hyperfine energy separation is small compared to the thermal energy of atoms in a gas at room temperature, one expects to find nearly equal number of atoms in the upper and lower hyperfine states, with slightly more atoms in the lower state. The number of atoms in the lower hyperfine level of room temperature hydrogen gas is 0.01 percent higher than the upper-state population. The number of atoms making transitions from the upper to the lower state and radiating photons is nearly equal to the number making transitions from the lower to the upper state and absorbing photons. Any effects of one process are nearly canceled by the other. In order to observe the process it is necessary to produce a larger discrepancy in the populations of the two levels. After this state selection is performed, it is possible to observe the atomic transition in many ways. In cesium atomic frequency standards, the number of atoms making a transition is measured; whereas in hydrogen masers, the radiation emitted by the atoms is detected.

3. Practical Atomic Spectrometers

Conceptually, the simplest atomic frequency standard is the active oscillator, shown in Fig. 8. Atoms in the upper hyperfine level of the ground state are stored in a microwave resonator. A fraction of the microwaves is coupled out and amplified for the output signal. In order for oscillation to take place, it is necessary for the population of upper state atoms to be increased substantially compared to the lower state atoms, a condition called population inversion. In the presence of microwave fields due to noise, the inversion causes the atoms to emit more power than they absorb. When the gain provided by the atoms exceeds the losses, the microwave fields build up until saturation limits the gain and a steady state condition is reached. Each atom is exposed to a microwave field produced by previously emitting atoms and all the atoms are stimulated to emit approximately in phase with one another. The output signal is highly coherent; the dominant short-term noise is white phase noise. The linewidth of the atomic transition *is determined* by the observation time, which is limited by processes that destroy either the population inversion or the coherence of the atoms. The active hydrogen maser is an example of the atomic oscillator; it is available commercially and is the primary clock used for radio astronomy. In order to achieve high resolution, radio astronomers use long baseline interferometry - two or more radio telescopes that are widely separated make simultaneous observations, each one using a maser to provide the time base. The clocks must stay coherent to a small fraction of an rf cycle during the observation period. For X-band observations, this implies time errors less than 25 ps for observation times from 15 minutes to several hours.

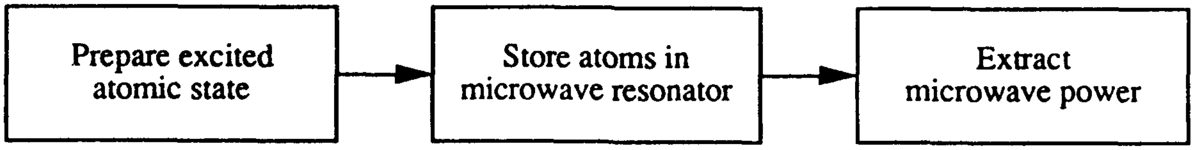

Figure 8. Block diagram of a general active atomic oscillator.

Passive atomic frequency standards are somewhat more complicated than active oscillators; they employ an atomic spectrometer the gain of which is insufficient to sustain oscillation. The spectrometer may even have a net loss. As shown in Fig. 9, a microwave oscillator stimulates the atomic transitions and a control loop provides feedback to tune the oscillator to the frequency that maximizes the transition rate. The peak transition rate is detected by frequency modulating the microwave signal applied to the atoms. The population of one of the hyperfine states is synchronously detected using the modulating signal as the reference. A signal is obtained proportional to the difference between the microwave carrier frequency and the center of the atomic resonance. The passive approach is used when other constraints make it impossible to achieve self oscillation. This happens, for example, when one tries to approach as closely as possible to the ideal of an isolated atom at rest in free space, as in the case of the cesium beam frequency standard, and when it is desirable to minimize the size of the device as in the rubidium frequency standard and the passive hydrogen maser.

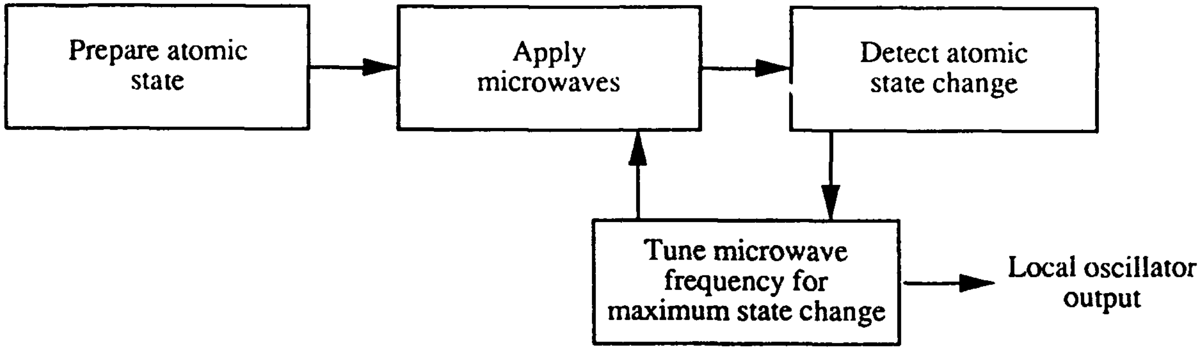

Figure 9. Block diagram of a general passive atomic frequency standard.

Since all atomic frequency standards, both passive and active, derive their output signal from quartz oscillators, the performance of the atomic standards is significantly affected by the capabilities of those oscillators. In particular the very short-term frequency stability, the vibration sensitivity, the radiation sensitivity, and the sensitivity to thermal transients depend principally on the performance of the quartz local oscillator.

a. State Selection

Magnetic state selection and optical pumping are both used to manipulate the hyperfine state populations. Magnetic state selection is based on the concepts which have already been discussed. Examination of Fig. 7 shows that the energy of some of the Zeeman sublevels increase with increasing magnetic field; the energy of the remainder decrease. When the atoms are passed through a region of strongly varying magnetic field, they experience a force proportional to the rate of change of the field with distance. The atoms that increase in energy with increasing magnetic field are deflected towards the region of strong field; the atoms that decrease in energy with increasing magnetic field are deflected in the opposite direction. The method is illustrated schematically in Fig. 10. Dipole, quadrupole, and hexapole optics are all used in atomic frequency standards to achieve the desired state population.

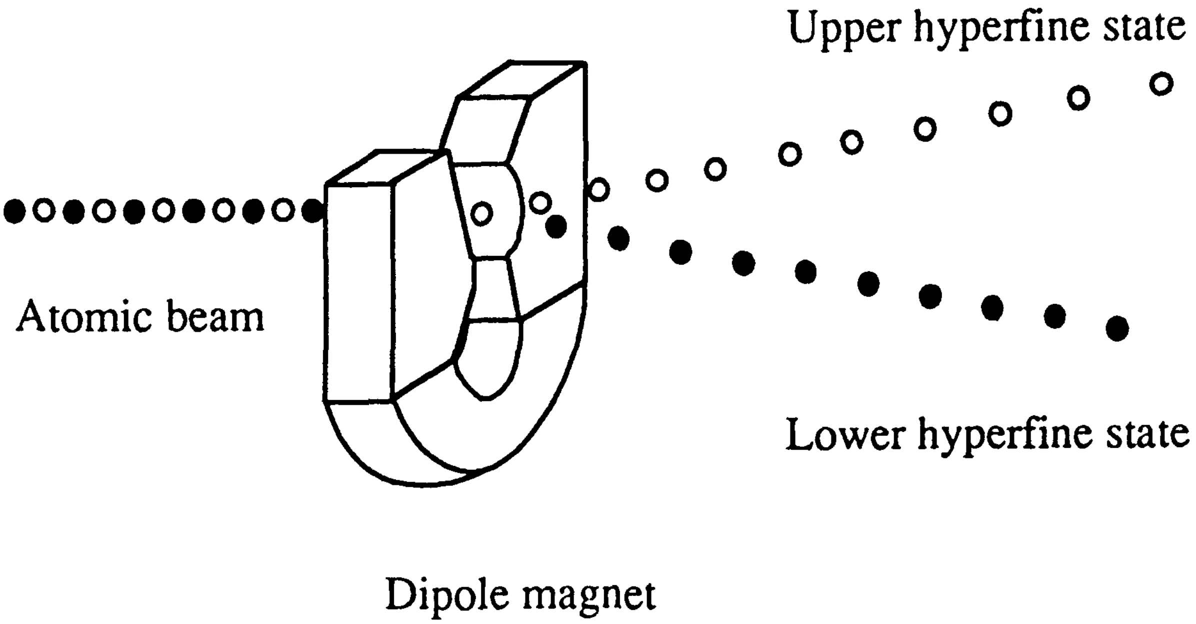

Figure 10. Concept for the magnetic state selection of the upper and lower hyperfine states in an inhomogeneous magnetic field.

An alternative to magnetic state selection is optical pumping. This technique manipulates the populations in the hyperfine levels of the ground state by exciting transitions to higher principal quantum states with infrared, or higher frequency, light. As shown in Fig. 11, the atoms in one hyperfine level are excited optically to a higher state from which they decay spontaneously to both ground state hyperfine levels. The population of the hyperfine state involved in the stimulated transition is rapidly depleted; the population of the second hyperfine level is enhanced. Optical pumping has both advantages and disadvantages compared to magnetic state selection. On the positive side, it can be accomplished in a more compact device and it can enhance the number of atoms in the desired state rather than just rejecting the atoms in the undesired state. On the negative side are increases in complexity and some additional performance degrading mechanisms.

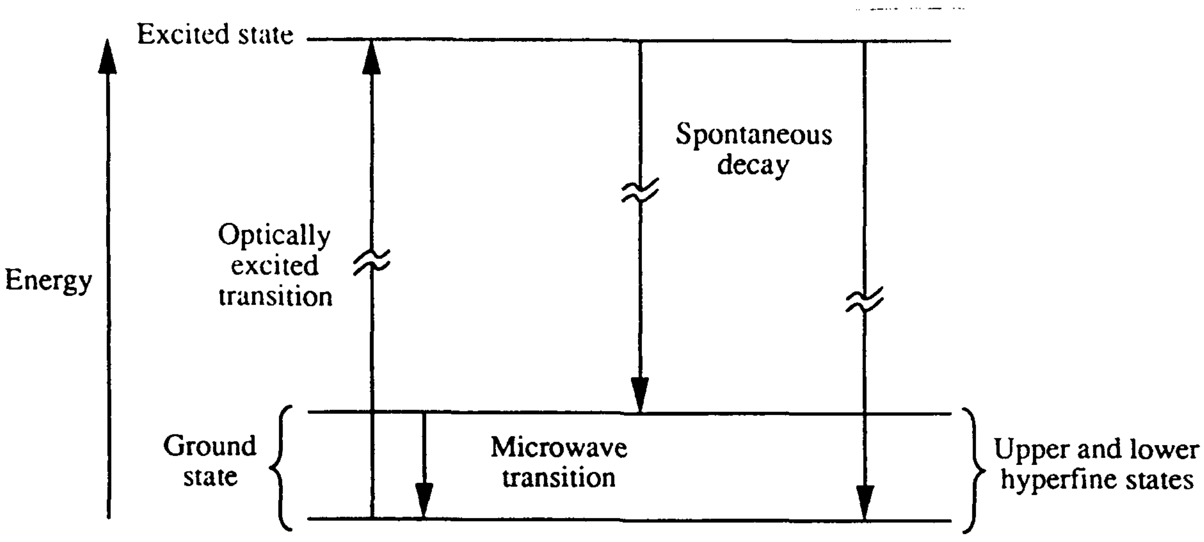

Figure 11. Concept for optical pumping using a three level system.

b. Detection of the Atomic Resonance

Passive atomic frequency standards use a variety of methods to detect the atomic resonance. That is, the standards detect the atomic transition probability as a function of the frequency of the applied radiation. The earliest approach was the direct detection of the state populations. The Stern-Gerlach experiment, which first demonstrated the quantization of angular momentum, used this method. A small background magnetic field established a reference direction. According to classical physics, the angular momentum of the silver atoms along this direction could take on any value between plus and minus the total angular momentum. A beam of silver atoms emitted from an oven passed through a dipole magnet and fell on a glass plate. The silver atoms were expected to deposit along a continuous line on the plate parallel to the magnetic field. However, only two spots were observed, demonstrating for the first time that the angular momentum along the field could have only two values. The amount of silver deposited indicated the populations of the two states.

An extension of this technique is still used. In a cesium beam frequency standard, upper state atoms may initially be selected using magnetic state selection. The application of microwaves near 9.2 GHz causes transitions to the lower state. A second magnet selects the lower state atoms, which subsequently fall on a hot ribbon in the detection region. The cesium atom's ionization potential is significantly less than the work function for platinum or tungsten. Thus the cesium atoms give up an electron and are reevaporated from the wire as positive ions. Laboratory frequency standards collect the ions to indicate the number of atoms arriving at the detector. Commercial cesium beam frequency standards collect the cesium ions using a mass spectrometer to exclude impurities emitted from the ionizer. The cesium ion current is converted to an electron current by a specially treated surface on the first dynode of an electron multiplier. Since the conversion of cesium atoms to ions by the ionizer ribbon is near 100 percent, this method of detection provides nearly ideal performance.

The effect of the microwave frequency on the transition probability can also be detected by its effect on the absorption of the light used to optically pump the atoms, such as in a Rb gas cell atomic frequency standard. With no microwaves applied, all the atoms are removed from one of the hyperfine states and the stored atoms become nearly transparent at the optical frequency. The transmission is maximum as depicted in Fig. 12(a). Tuning the microwaves near the transition frequency causes atoms to be transferred back to the depleted level; the absorption then increases. Maximum absorption occurs at the peak of the transition probability as shown in Fig. 12(b). The atomic resonance may also be detected by its effect on the microwaves themselves. The state selected atoms behave very much like a filter. If upper state atoms predominate, the microwave signal will experience amplification and phase shift. The microwave frequency which maximizes the transition probability may be determined by comparison of the signal passed through the atomic discriminator with a reference. This is the method used in a passive hydrogen maser frequency standard.

4. Systematic Limitations of Atomic Frequency Standards

There are fundamental and practical limitations on the ability of a device to reproduce an atomic frequency [13, 14]. Quantum mechanics and thermal noise limit the quality of the measurements by producing stochastic frequency variations. Imperfections in the electronics, stray electric and magnetic fields, and interactions with other atoms cause the frequency standard to exhibit aging and temperature and pressure sensitivities.

Atomic frequency standards must be designed to prevent performance limitation by the Doppler effect. The first order Doppler effect is the change in the observed frequency when a source is in motion relative to an observer. It is often observed as the change in pitch of a train whistle when the train overtakes and passes a stationary observer. Since many atoms with different velocities must be observed in an atomic frequency standard, the first order Doppler effect would produce broadening of atomic resonance as well as average frequency shifts. The observed Doppler frequency shift is equal to the ratio of the atomic velocity to the speed of light, approximately 1×10^{-6} for room temperature atoms, and would limit the maximum atomic Q to 1×10^{6}.

All atomic frequency standards use some form of first order Doppler shift cancellation in order to achieve Q's in the range of 10^{7} to 10^{9}. One method that is always used in microwave standards is to excite the atomic resonance with a standing wave produced by a microwave cavity. The atoms interact with microwaves of approximately equal intensity traveling in opposite directions and the Doppler shifts for

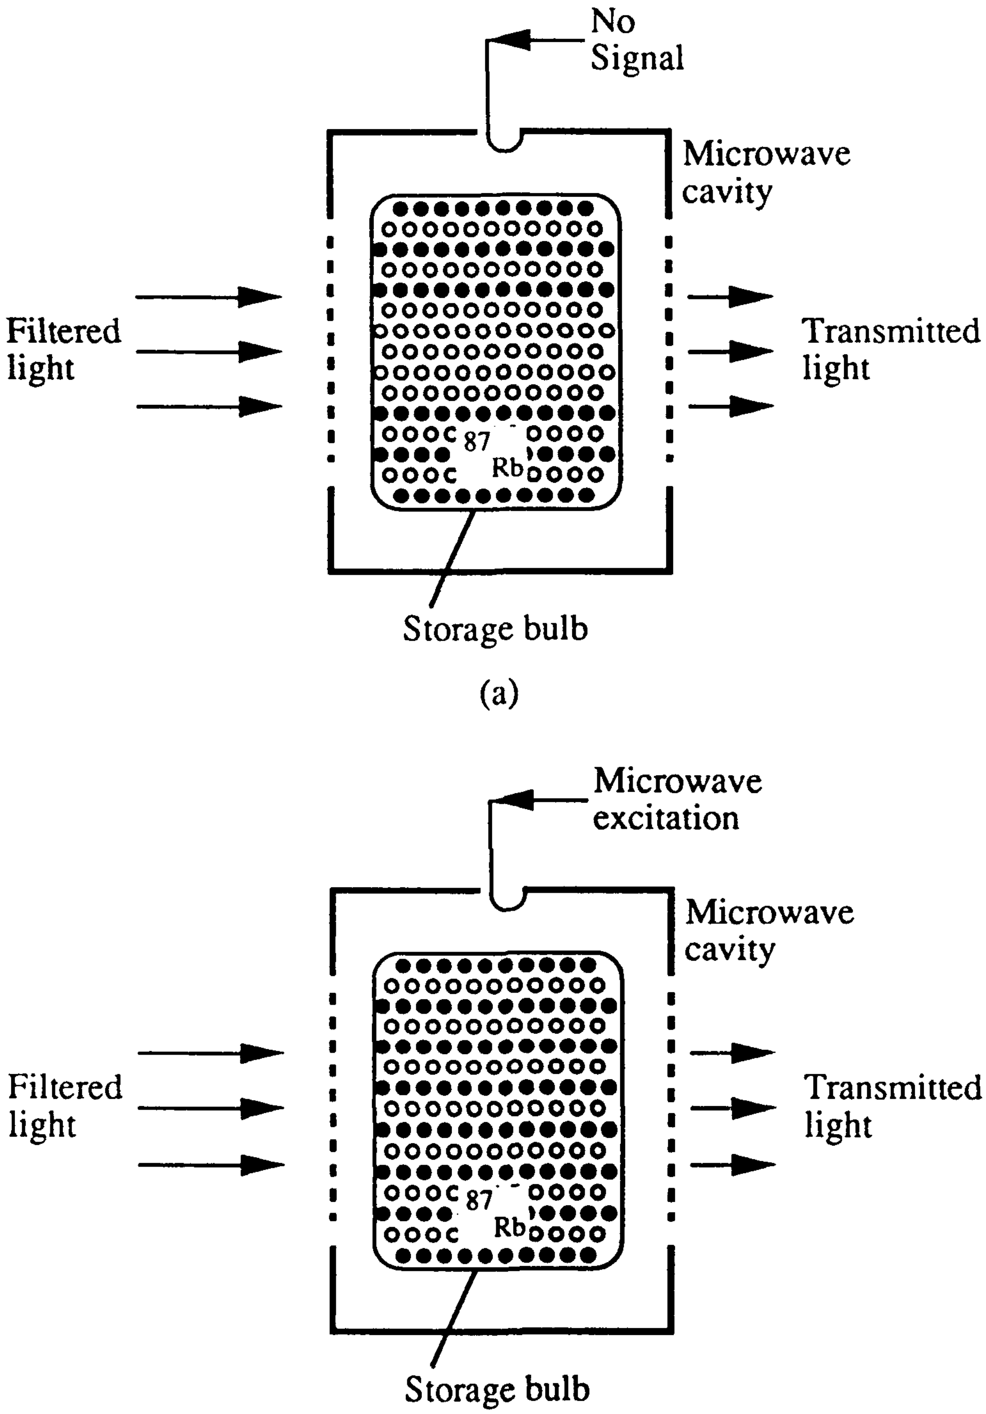

Figure 12. Concept of microwave resonance detection using the optical density of the atomic medium.

these two interactions nearly cancel. Further cancellation is obtained when storage cells are used. The atoms in these devices change direction thousands of times while interacting with the microwave field. However, atom storage within a microwave resonator is insufficient to guarantee narrow, unshifted atomic lines. A more detailed analysis shows that the atoms must be confined to a region smaller than one-half wavelength. This restriction, known as the 'Dicke regime,' prevents frequency modulation due to the Doppler shifts from absorbing all the power from the carrier [15]. Atomic beam frequency standards are always designed so that the residual power flow in the microwave cavity is orthogonal to the atomic beam direction. This construction takes the place of the velocity averaging of storage cell devices. Atomic beam standards must also operate under similar restrictions.

There are many other less fundamental perturbations to the observed atomic frequency. The atomic resonance is observed through the interaction of the atoms with the microwave field of a resonant circuit. If the cavity resonance frequency is not equal to the atomic frequency, the stronger fields at the cavity frequency increase the probability of transitions at a frequency different from the true atomic resonance frequency, and the measured value is not equal to the true atomic frequency. In an active oscillator, the 'pulling' is equal to the error in the cavity resonance reduced by the ratio of the cavity Q to the atomic Q. In a passive atomic frequency standard, the pulling may be reduced by the square of the Q ratio.

In optically pumped frequency standards, the atomic energy levels are affected by the light used for the optical pumping. This problem is a specific instance of a phenomenon called the Stark effect, which is the variation of atomic energy levels with applied electric fields. Other electric fields can usually be reduced to negligible levels, but the fields required for optical pumping are very intense and the resulting 'light shifts' are significant.

Frequency shifts are also caused by interactions of the atoms with other matter. These interactions change the hyperfine interaction energy. 'Pressure shifts' result from collisions with gas molecules. 'Spin exchange shifts' are caused by collisions of the subject atoms with each other. 'Wall shifts' result from collisions with the coating on a storage vessel. Additional changes in the measured maximum of the transition probability curve result from the presence of nearby atomic transitions. The tails of a neighboring transition result in a background slope, which distorts the measured shape of the desired transition.

Imperfections in the electronics cause additional errors in the measurement of the atomic transition frequency. Some of the problems are spurious microwave signals, voltage offsets and drifts, insufficient gain, and insufficient dynamic range. In the end, performance is always limited by the ability of the electronics to find the center of the atomic line, and a major thrust of research and development is towards achieving higher atomic line Q.

C. Oscillator Categories

1. Quartz Oscillator Types

A crystal unit's resonance frequency varies with temperature. Typical frequency vs. temperature (f vs. T) characteristics for crystals used in stable oscillators are shown in Fig. 13. The three categories of crystal oscillators, based on the method of dealing with the crystal unit's f vs. T characteristic, are: crystal oscillator (XO), temperature

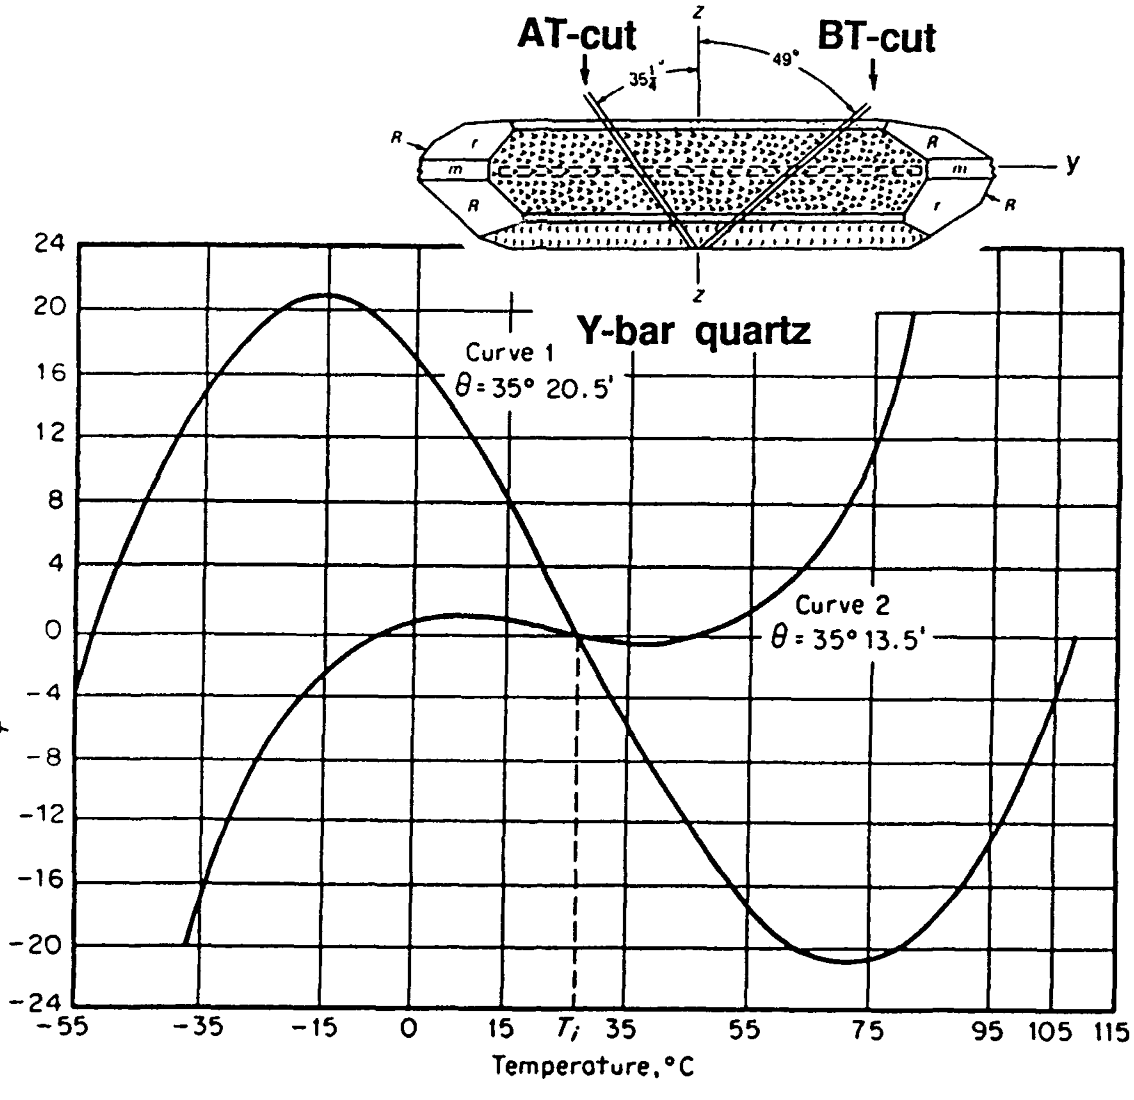

Figure 13. Frequency vs. temperature characterstics of AT-cut crystals, showing AT and BT-cut plates in Y-bar quartz.

compensated crystal oscillator (TCXO), and oven controlled crystal oscillator (OCXO), as illustrated in Fig. 14. A simple XO does not contain means for reducing the crystal's f vs. T variation. A typical XO's f vs. T stability may be ±25 ppm for a temperature range of -55°C to +85°C.

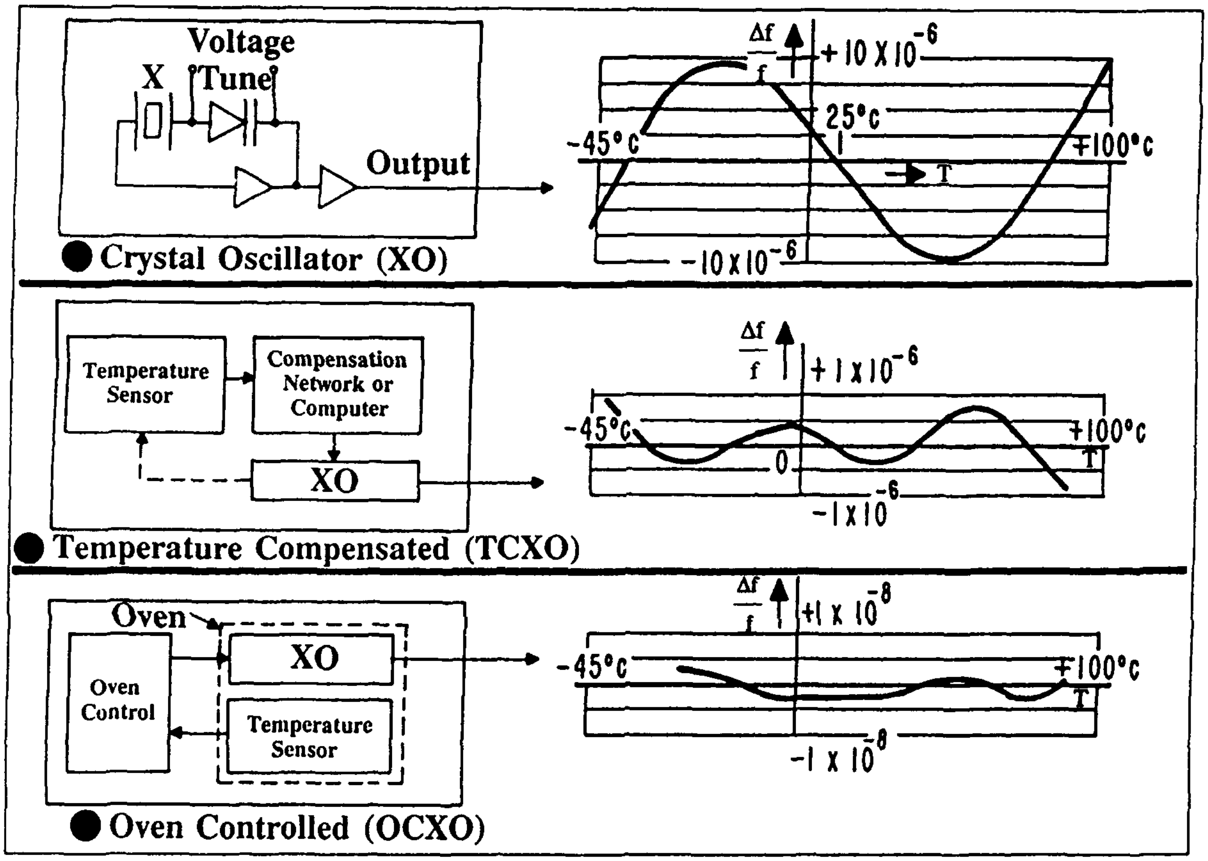

Figure 14. Crystal oscillator categories: XO, TCXO, and OCXO.

In a TCXO [16], the output signal from a temperature sensor (a thermistor) is used to generate a correction voltage that is applied to a voltage variable reactance (a varactor) in the crystal network. The reactance variations produce frequency changes that are equal and opposite to the frequency changes resulting from temperature changes; i.e., the reactance variations compensate for the crystal's f vs. T variations. Analog TCXOs can provide about a 20-fold improvement over the crystal's f vs. T variation. A "good" TCXO may have an f vs. T stability of ±1 ppm for a temperature range of -55°C to +85°C.

In an OCXO [16], the crystal unit and other temperature sensitive components of the oscillator circuit are maintained at a constant temperature in an oven. The crystal is manufactured to have an f vs. T characteristic which has zero slope at the oven temperature. To permit the maintenance of a stable oven temperature throughout the OCXOs temperature range (without an internal cooling means), the oven temperature is selected to be above the maximum operating temperature of the OCXO. OCXOs can provide more than a 1000-fold improvement over the crystal's f vs. T variation. A "good" OCXO may have an f vs. T stability of $\pm 5 \times 10^{-9}$ for a temperature range of -55°C to +85°C. OCXOs require more power, are larger, and cost more than TCXOs.

A special case of a compensated oscillator is the microcomputer compensated crystal oscillator (MCXO) [17]. The MCXO overcomes the two major factors that limit the stabilities achievable with TCXOs: thermometry and the stability of the crystal unit. Instead of a thermometer that is external to the crystal unit, such as a thermistor, the MCXO uses a much higher accuracy "self-temperature sensing" method. Two modes of the crystal are excited simultaneously in a dual mode oscillator. The two modes are combined such that the resulting beat frequency is a monotonic (and nearly linear) function of temperature. The crystal thereby senses its own temperature. To reduce the f vs. T variations, the MCXO uses digital compensation techniques: pulse deletion in one implementation, and direct digital synthesis of a compensating frequency in another. The frequency of the crystal is not "pulled," which allows the use of high stability (small C_1) SC-cut crystal units. A good MCXO may have an f vs. T stability of $\pm 2 \times 10^{-8}$ for a temperature range of -55°C to +85°C.

2. Atomic Frequency Standard Types

a. Atomic Standards in Production

Three types of atomic standards are available to meet frequency stability requirements which exceed the performance capabilities of quartz crystal oscillators: rubidium, cesium, and hydrogen frequency standards. Rubidium oscillators offer better stability in the range from one-hundred seconds to a few hours, a factor of ten improvement in long-term aging, superior reproducibility, and lower sensitivity to the environment. Rubidium frequency standards are often employed in tactical applications or any time the clock must free run for a few hours to a day. Cesium standards offer further improvements. Compared to rubidium oscillators, they perform better for times longer than a few hours, generally suffer from no frequency aging, are more reproducible, and have lower sensitivity to the environment. Cesium standards are used in strategic applications and wherever a clock is required to keep time autonomously for much longer than one day. The third type of standard, the active hydrogen maser, is used in applications requiring extreme coherence for periods of time from minutes to hours. Very long baseline interferometry and spacecraft tracking are the two principal applications.

Figure 15 is a schematic of a cesium beam tube. Spectroscopy is performed on cesium atoms in free flight through the microwave cavity in order to minimize environmental influences. A magnetic state selector prior to the cavity rejects atoms in the unwanted hyperfine level. A magnetic state selector following the cavity passes atoms having the state that was rejected in the first region. Thus, atoms must make a hyperfine transition in order to reach the hot wire ionizer. A feature of this beam tube, not previously discussed, is the U-shaped microwave cavity also called a Ramsey cavity. In this design, atoms are exposed to the microwave field, pass through a microwave field free region, and are then exposed to a second microwave field, in phase with the first. As a result, the atomic line is an interference pattern that maximizes the resolution of the spectrometer. Even more important, the design of the Ramsey cavity reduces the sensitivity to magnetic field inhomogeneities. Without it, the magnetic shielding requirement would be impractical.

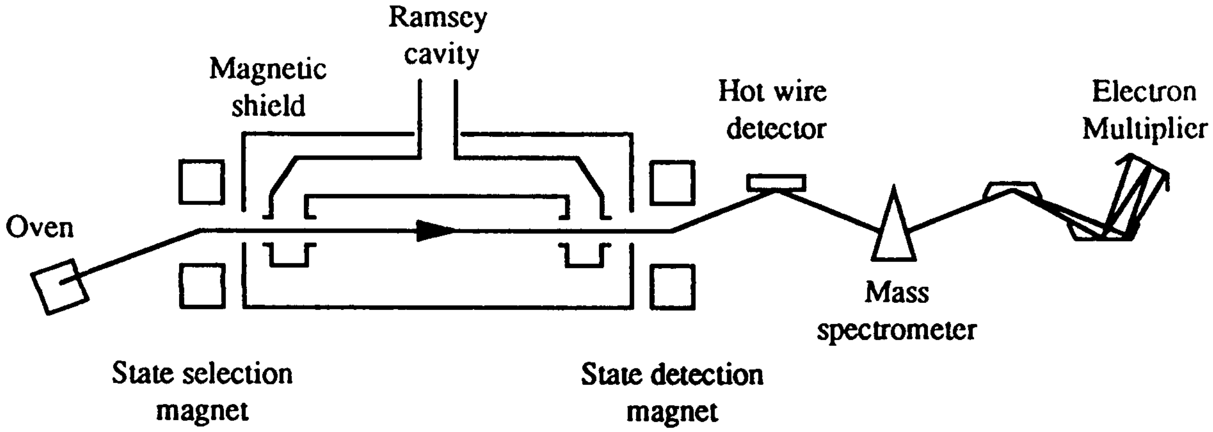

Figure 15. Schematic drawing of a cesium beam tube.

Figure 16 shows a typical electronic schematic for the cesium frequency standard. A 5-MHz VCXO provides the output signal. After audio frequency modulation, it is used to synthesize the microwave signal at 9192.631 MHz, which is then used to excite the hyperfine transition. The detected atomic beam current from the beam tube is applied to a phase sensitive detector using the audio modulation as reference. After integration and additional loop filtering, the phase detector output is used to control the frequency of the 5-MHz VCXO. Figure 17 illustrates how phase sensitive detection generates a control signal that passes through zero at the resonance frequency.

Cesium technology has been optimized for long-term timekeeping and frequency reproducibility at the expense of size, weight, power, and warm-up time. A healthy device shows no frequency aging, and the reproducibility is on the order of 2 X 10^{-12} for the life of the cesium beam tube. The dominant long-term noise is random walk frequency, which limits the timekeeping to approximately 0.1 microsecond per month.

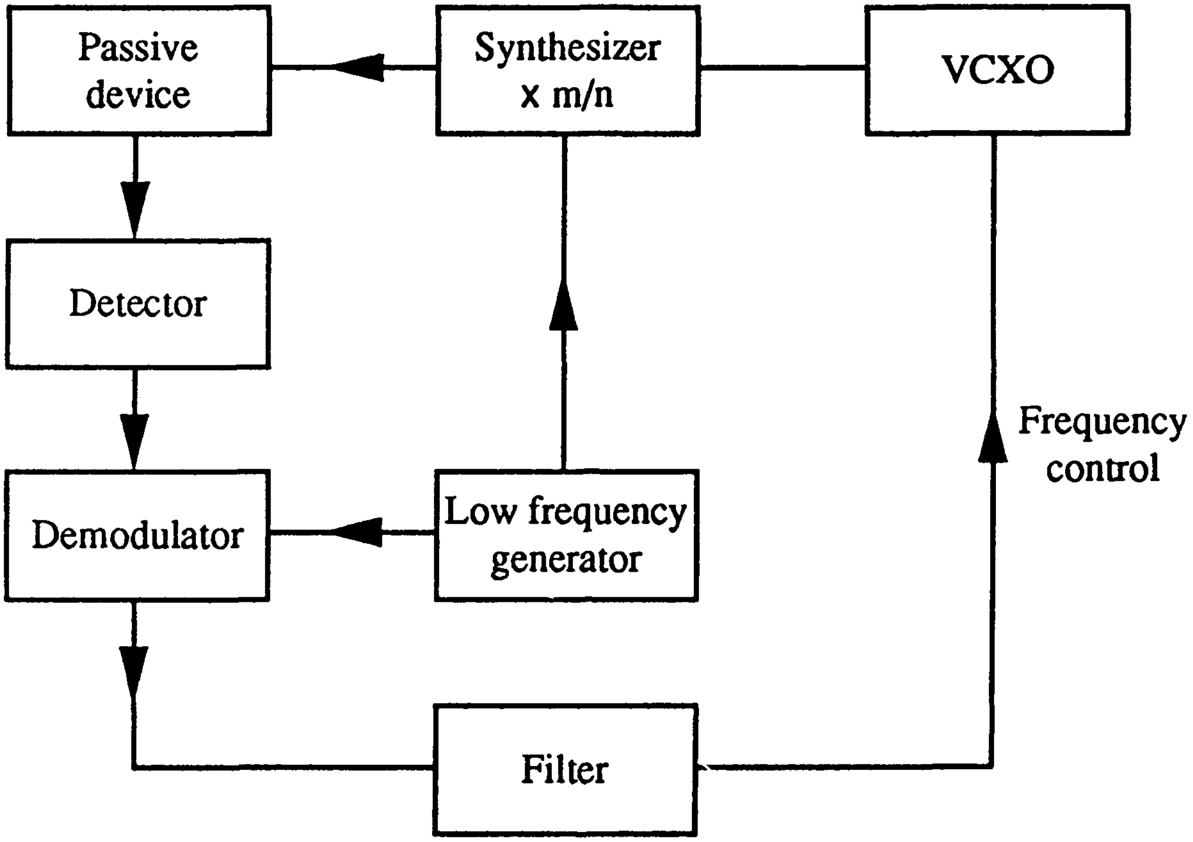

Figure 16. Electronic schematic showing the frequency lock of a voltage controlled quartz oscillator (VCXO) to an atomic (or other) resonance.

The atomic spectrometer of a typical rubidium standard is shown in Fig. 18. Commercial rubidium standards, based on the hyperfine transition in the ground state of Rb 87, are optically pumped using the light emitted from a Rb 87 discharge lamp. The emission from this lamp contains light with wavelengths corresponding to the transitions from the two hyperfine ground states to the excited state. The light from the Rb 87 passes through a region containing Rb 85. What makes rubidium standards practical is the fact that Rb 85 strongly absorbs the wavelength which would excite the upper hyperfine state and transmits most of the lower state transition light. This filtered light is selectively absorbed by the lower hyperfine state atoms in the resonance cell. These atoms are excited to a third, optical, state, which decays to both the ground state hyperfine levels. Atoms which return to the lower hyperfine level are excited again by the

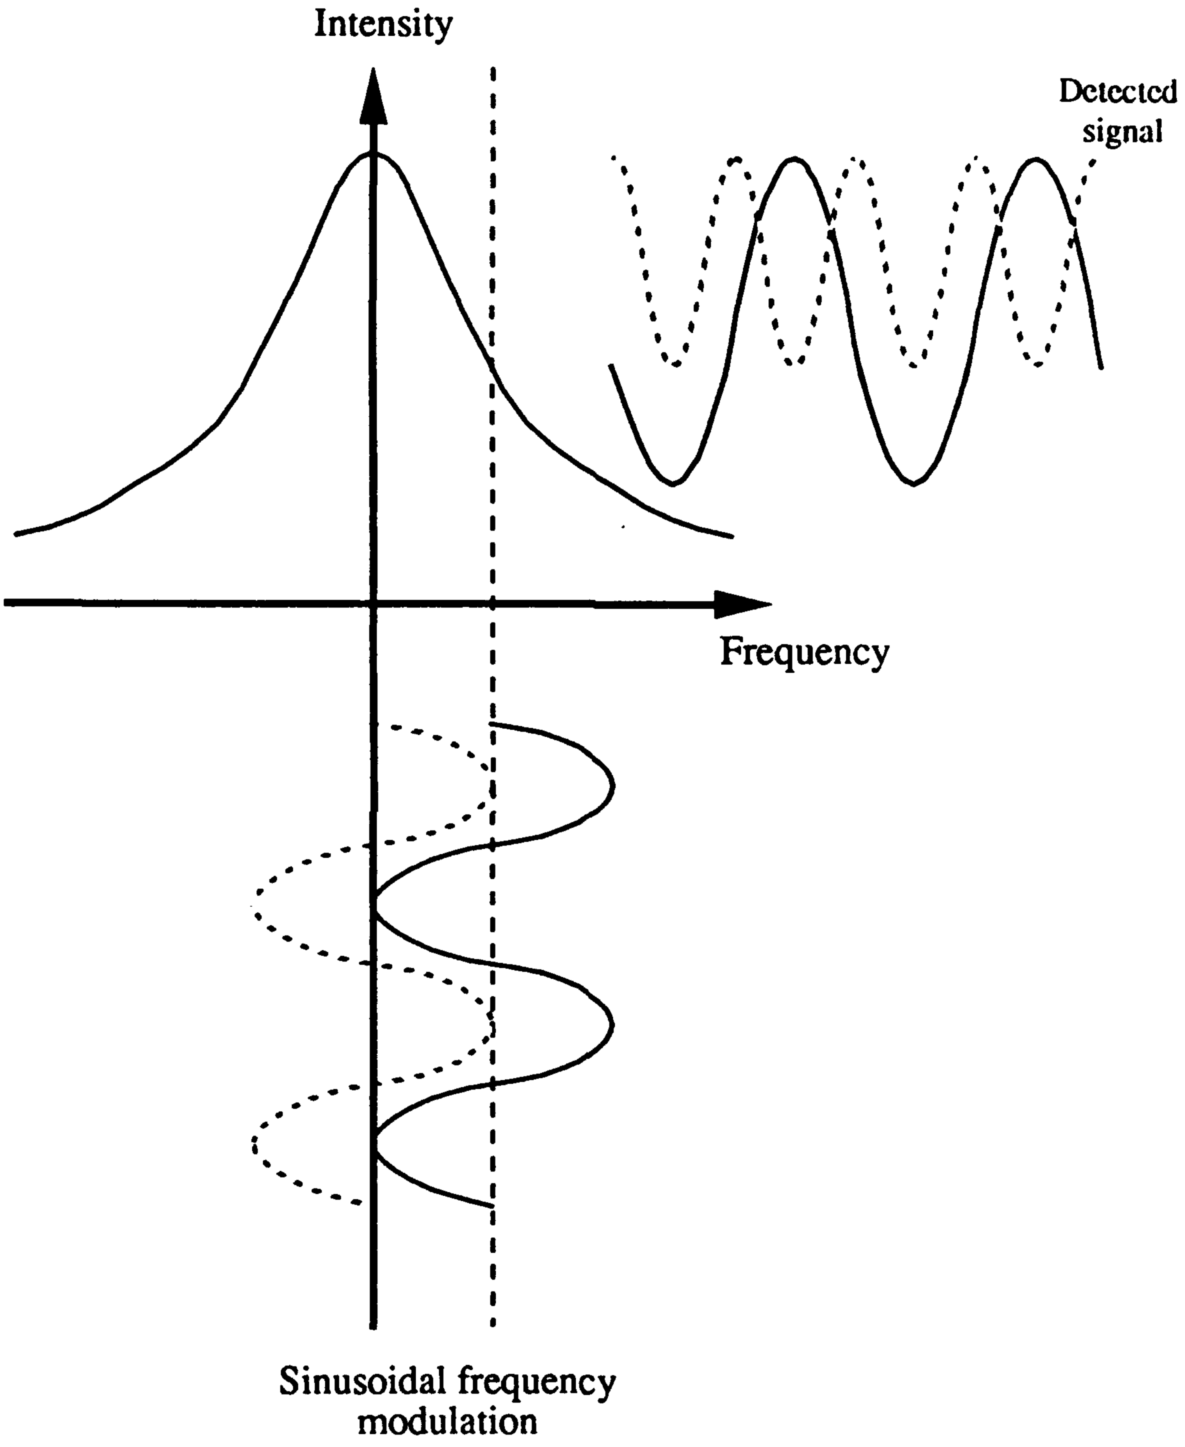

Figure 17. Concept for the detection of a passive resonance using sinusoidal frequency modulation (FM).

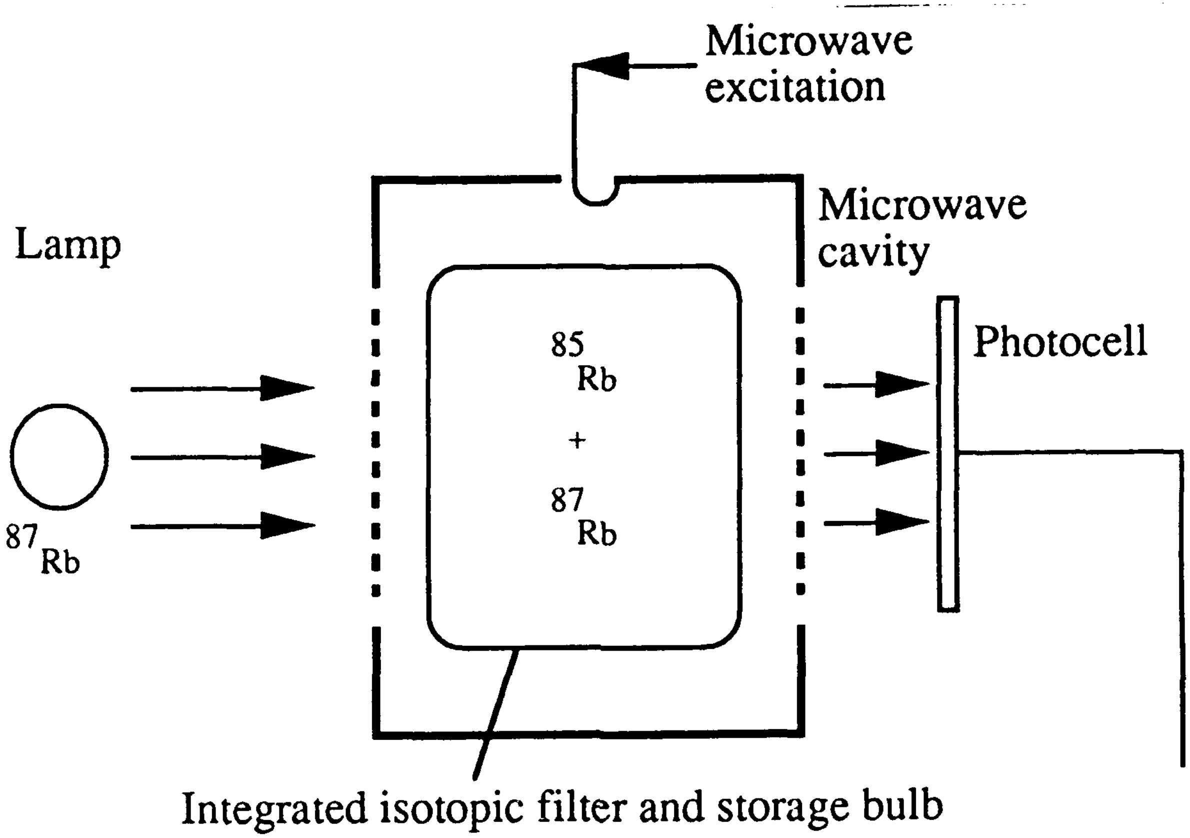

Figure 18. Schematic drawing of the atomic spectrometer for a passive rubidium frequency standard.

filtered light until they are converted into upper hyperfine level atoms. Thus, the filtered light optically pumps Rb 87 atoms contained in the microwave resonance cell into the upper hyperfine state. In order to make a compact device, the resonance cell completely fills a small microwave cavity. The filter can be implemented separately, but most often is integrated with the cell by adding Rb 85. When microwaves are applied to the resonator cell, atoms are connected back from the upper to the lower hyperfine states. This increases the absorption of the filtered light which passes through the resonance cell. The microwave transition frequency is observed by detecting the filtered light which passes through the resonance cell. The Dicke criterion is satisfied within the resonance cell by using a buffer gas to localize the Rb 87 atoms during interaction with the microwave field. The buffer gas prevents the rubidium atoms from experiencing wall collisions that would introduce major perturbations. However, collisions of the Rb atoms with the buffer gas cause pressure shifts of the frequency. The resulting pressure and temperature coefficients can be set to zero at the nominal operating temperature by using a mixture of gases, such as nitrogen and argon, that have opposite sign pressure shifts.

The rubidium standard control electronics is very similar to that of the cesium standard. The transition frequency is 6.834 GHz and the linewidth is several hundred Hz. The detection of the resonance is performed by a photodetector that measures the transmission of pumping light through the cell. There is approximately one to ten percent decrease in transmission when the microwave frequency is at the peak of the transition probability.

Rubidium technology has been optimized for small size and fast warm-up. The smallest is 5 cm by 7.5 cm by 10 cm, has a warm-up time of less than 4 minutes from 25°C for 1×10^{-9} reproducibility, and the long-term frequency aging is 2×10^{-10} per year. Such a device does not have state-of-the-art frequency stability or phase noise. These performance parameters can be improved at the expense of increased size. In fact, a full-size rubidium standard, the same size as a cesium standard, performs almost as well as a cesium standard.

Figure 19 is a schematic representation of an active hydrogen maser. Unlike cesium and rubidium, hydrogen occurs only in molecular form at room temperature. The necessary atomic hydrogen is produced using a radio frequency discharge. Typically, the upper state atoms are focused into a storage bulb using a hexapole state selector magnet. Eventually, spent atoms are pumped away. The bulb surface is Teflon coated to ensure nearly elastic collisions of the hydrogen atoms on the wall. This surface is so

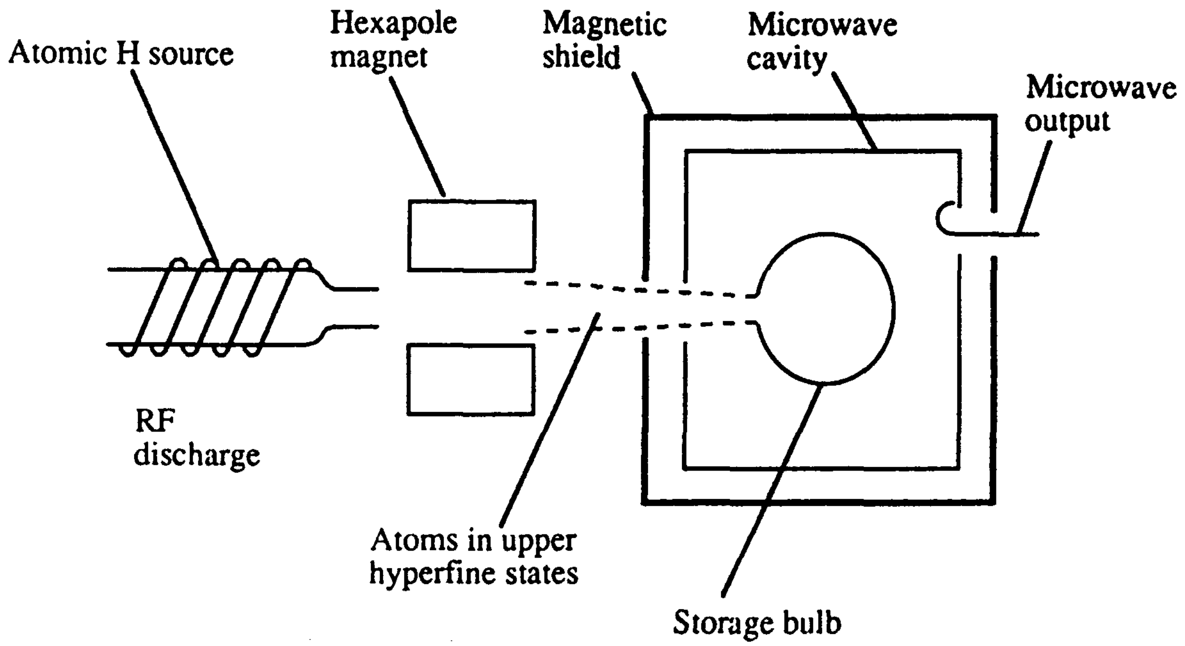

Figure 19. Schematic drawing of an active hydrogen maser.

good that more than ten-thousand bounces occur before the coherence of the hydrogen atoms is destroyed. Consequently, the observation time can exceed one second. As a result, the atomic line Q is greater than 1 X 10^9, the highest of the commercial atomic frequency standards. The density of hydrogen in the storage bulb is limited by hydrogen-hydrogen collisions that alter the hyperfine energy and thus produce frequency shifts. A high Q microwave cavity is required to achieve self oscillation. The large size of most hydrogen masers (typical active hydrogen masers are ten times the volume of cesium frequency standards) results from the 1.4 GHz TE011 mode microwave cavity and the surrounding layers of magnetic shielding and temperature stabilizing ovens.

The control electronics for the active hydrogen maser is shown in Fig. 20. The 1.42 GHz output from the atoms is amplified and mixed with a signal obtained by multiplying the output frequency to the microwave region. A final downconversion of this difference frequency using a reference synthesized from the output frequency completes the conversion of the atomic frequency to baseband. Feedback to the VCXO phaselocks it to the microwave signal originating from the atoms. In contrast to passive atomic frequency standards, which convert thermal noise at their preamplifier inputs to white

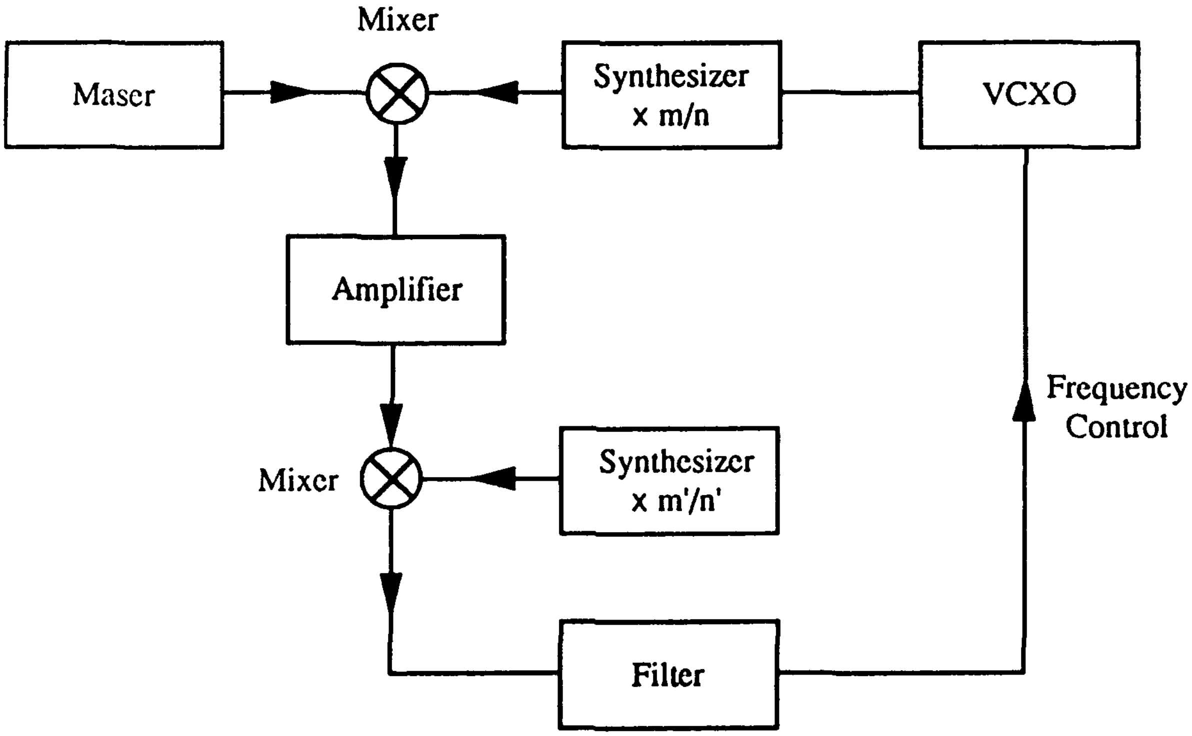

Figure 20. Electronic schematic showing the phase lock of a voltage controlled quartz oscillator (VCXO) to the atomic radiation emitted by an active maser.

frequency noise, the active maser converts this thermal noise to white phase noise. As a result, the short-term stability of active masers improves inversely with the measurement interval, as opposed to the stability of passive standards, which improves as the square-root of the measurement interval.

Active masers have the best short-term frequency stability of all the atomic frequency standards. Typical performance is $5 \times 10^{-13}/\tau$, reaching a floor of better than 1×10^{-15}. However, the long-term performance is not as good as that of cesium frequency standards. Most masers have had substantial frequency aging which has been attributed to cavity pulling. The frequency stability at one day is typically 1×10^{-14}.

b. Experimental Atomic Frequency Standards

At the time of the writing of this section, in 1990, one of the most active areas of research in atomic frequency standards was atomic beam tube technology. The efforts are directed towards two different and contradictory objectives, but both employ optical pumping technology. One goal is to develop a frequency standard with the low environmental sensitivity typical of the cesium standard, but having the size of a miniature rubidium standard. Optical pumping is used to replace the magnetic state selection and detection. Elimination of the high magnetic fields makes it possible to design a more compact device. The linear beam of an optically pumped device makes more efficient use of the atoms emitted from the oven than the usual magnetic optics, and helps achieve good short-term stability. The second goal is to improve the accuracy, reproducibility, and long-term stability of the atomic beam frequency standard [18]. The performance limitations of the existing technology are due to residual first order Doppler shift, second order Doppler shift, and pulling by neighboring lines. The linear beam of the optically pumped device reduces the complicated coupling between the velocity and the position of the atoms in the beam, making it possible to better measure the Doppler effect. Optical pumping can also be used to transfer almost all the atoms to the desired hyperfine state, thus reducing the size of and pulling by neighboring magnetic field sensitive transitions. This technique can be combined with advances made in alignment of the beam tube to reduce the pulling problem to negligible levels.

The most promising developments for improving long-term stability of atomic frequency standards are in the field of ion storage. The goal of this work is to achieve the low interactions of the beam type frequency standard with the long observation times possible in a storage device. Singly ionized atoms, such as beryllium and mercury, have been stored using ion trap technology [19, 20]. The ions experience no collisions due to the storage mechanism, but are confined by dc or rf electromagnetic fields. Mercury ion frequency standards have been built using the rf ion trapping technique. The technology appears practical for commercial application in the near term. Other ion frequency standard research uses Penning traps which employ dc trapping fields. This approach is compatible with laser cooling of the ions - a technique that extracts energy and momentum from the ions by scattering a laser beam. Cooling the ions below one Kelvin reduces the magnitude of the Doppler effects and improves the accuracy and reproducibility of the frequency standard. The large magnets and complicated lasers used for laser cooling will relegate this approach to the laboratory for some time to come.

III. Oscillator Instabilities

A. Accuracy, Stability and Precision

Figure 21 illustrates the meanings of accuracy, stability and precision for a marksman and a frequency source. (For the marksman, each bullet hole's distance to the center of the target is the "measurement.") ***Accuracy*** is the extent to which a given measurement, or the average of a set of measurements on one sample, agrees with the definition of the quantity being measured. It is the degree of "correctness" of a quantity. Atomic frequency standards have varying degrees of accuracy. The SI units for time and frequency, the second and the Hz, are obtained in laboratories using very accurate frequency standards called 'primary' standards. A 'primary' standard operates at a frequency calculable in terms of the SI definition of the second - "the duration of

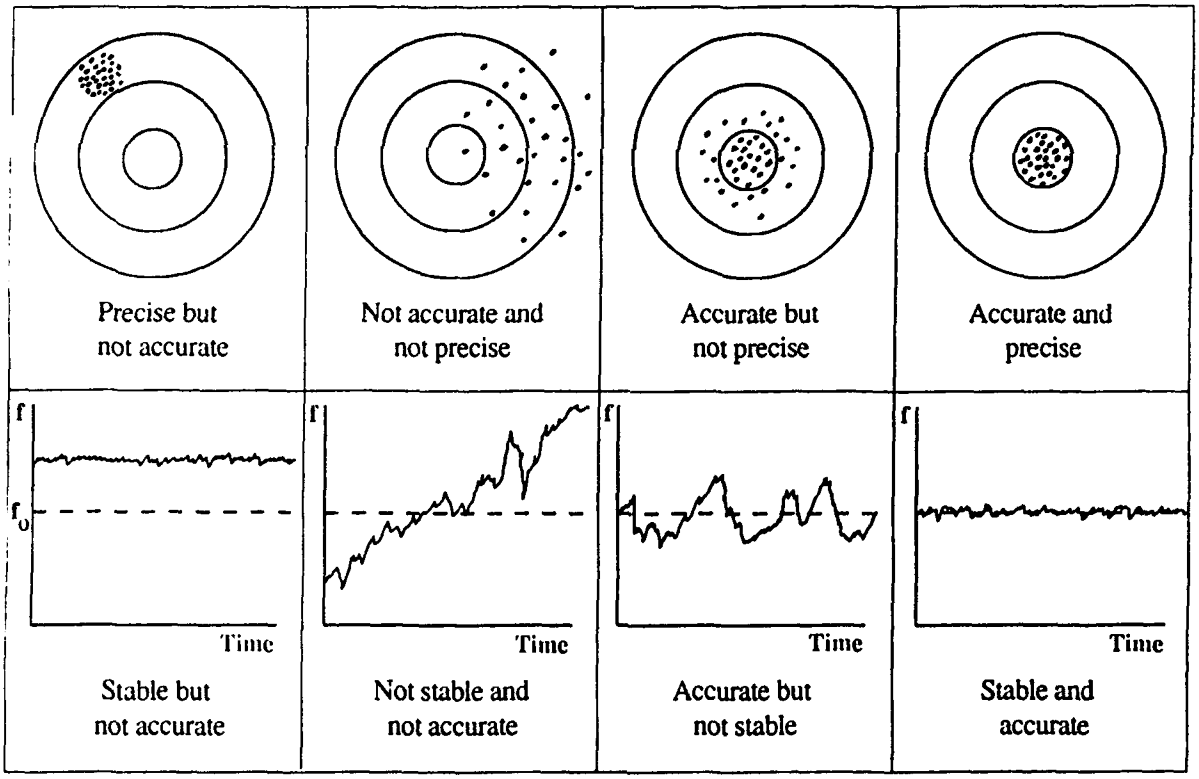

Figure 21. Accuracy, stability and precision.

9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the cesium atom 133." ***Reproducibility*** is the ability of a single frequency standard to produce the same frequency, without adjustment, each time it is put into operation. From the user's point of view, once a frequency standard is calibrated, reproducibility confers the same advantages as accuracy. ***Stability*** describes the amount something changes as a function of a parameter such as time, temperature, shock, etc. ***Precision*** is the extent to which a given set of measurements, on one sample, agrees with the mean of the set. (A related meaning of the term is used as a descriptor of the quality of an instrument, as in a "precision instrument." In that context, the meaning is usually accurate and precise, although a precision instrument can also be inaccurate and precise, in which case the instrument needs to be calibrated.)

B. Aging

"Aging" and "drift" have occasionally been used interchangeably in the literature. However, in 1990, recognizing the "need for common terminology for the unambiguous specification and description of frequency and time standard systems," the CCIR adopted a glossary of terms and definitions [21]. According to this glossary, aging is "the systematic change in frequency with time due to internal changes in the oscillator," and drift is "the systematic change in frequency with time of an oscillator." Drift is due to aging plus changes in the environment and other factors external to the oscillator. Aging is what one specifies in a specification document and what one measures during oscillator evaluation. Drift is what one observes in an application. For example, the drift of an oscillator in a spacecraft might be due to (the algebraic sum of) aging and frequency changes due to radiation, temperature changes in the spacecraft, and power supply changes.

1. Quartz Oscillator Aging

Aging can be positive or negative. Occasionally, a reversal in aging direction is observed. Typical (computer simulated) aging behaviors are illustrated in Fig. 22. The curve showing the reversal is the sum of the other two curves. A reversal indicates the presence of at least two aging mechanisms. The aging rate of an oscillator is highest when it is first turned on. At a constant temperature, aging usually has an approximately logarithmic dependence on time. When the temperature of a crystal unit is changed, e.g., when an OCXO is turned off, and turned on at a later time, a new aging cycle starts. (See "retrace" below).

The primary causes of crystal oscillator aging are mass transfer to or from the resonator's surfaces due to adsorption or desorption of contamination, stress relief in the mounting structure of the crystal, and, possibly, changes in the quartz material. Because the frequency of a thickness shear crystal unit, such as an AT-cut or an SC-cut, is inversely proportional to the thickness of the crystal plate, and because a typical 5 MHz

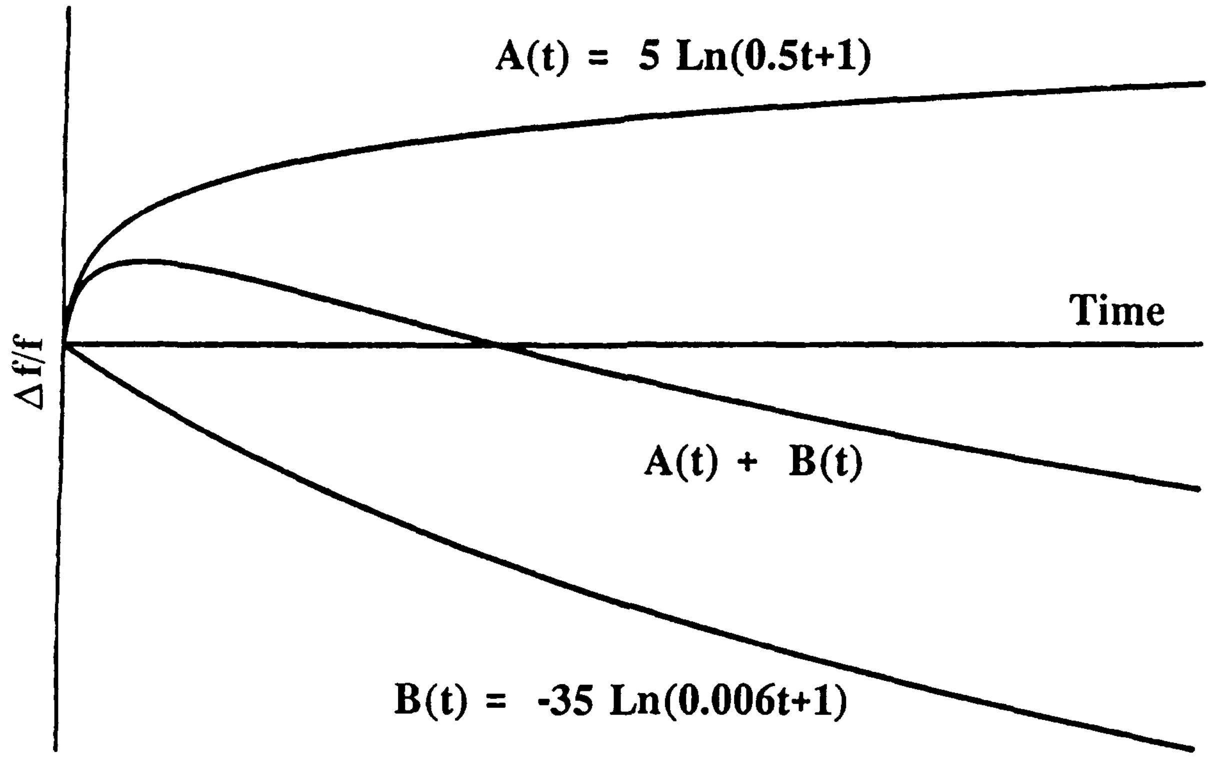

Figure 22. Typical aging behaviors.

plate is on the order of 1 million atomic layers thick, the adsorption or desorption of contamination equivalent to the mass of one atomic layer of quartz changes the frequency by about 1 ppm. Therefore, in order to achieve low aging, crystal units must be fabricated and hermetically sealed in an ultraclean, ultrahigh vacuum environment. The aging rates of typical commercially available crystal oscillators (in 1990) range from 5 to 10 ppm per year for an inexpensive XO, to 0.5 to 2 ppm per year for a TCXO, to 0.05 to 0.1 ppm per year for an OCXO. The highest precision OCXOs can age less than 0.01 ppm per year.

2. Atomic Standard Aging

One mechanism for frequency aging potentially affects all passive atomic frequency standards, that is, failure to control the aging of the internal VCXO. This source of aging is a consideration for passive frequency standards that use analog control electronics, in which case, the integrators in the control loop have finite dc gain. The open-loop VCXO aging divided by the dc gain must be less than the noise-induced frequency changes, otherwise, frequency aging will be detectable. This requirement must not only be met in the new device, but also after the gain has been degraded by operation and exposure to the environment. For example, the gain of the electron multiplier used in a cesium standard degrades as a result of the unavoidable exposure to cesium, and the gain of a rubidium standard degrades due to the exposure of the photodetector to ionizing radiation. In the past, cesium standards exhibited this form of frequency aging. Today, 'healthy' cesium standards exhibit no measurable frequency aging. Hydrogen masers exhibit frequency aging in the range from 1 to 10 X 10^{-12} per year. This aging is often attributed to pulling by the microwave cavity. In the past, it has been controlled through the use of very stable microwave cavities manufactured from ceramic materials. More recently, electronic control has been used to minimize the cavity frequency error in miniature active and passive hydrogen masers [22,23]. Rubidium frequency standards exhibit the most aging, on the order of 1 X 10^{-10} per year. The largest effects result from changes in the light shift. The electrical discharge and elevated temperatures in the lamp cause diffusion of Rb into the glass and change its electrical properties. Over time, large changes in the lamp spectrum and intensity can take place with associated changes in the output frequency. Additional aging can be produced by the diffusion of helium into the resonance cell. The increasing helium partial pressure changes the pressure shift. As a result, all manufacturers use glass chosen for relatively low helium permeation rate.

C. Noise in Frequency Standards

1. *The Effects of Noise*

Sometimes the suitability of oscillators for an application is limited by deterministic phenomena. In other instances, stochastic (random) processes establish the performance limitations. Except for vibration, the short-term instabilities almost always result from noise. Long-term performance of quartz and rubidium standards is limited by the temperature sensitivity and the aging, but the long-term performance of cesium and hydrogen standards is limited by random processes.

Noise can have numerous adverse effects on system performance. Among these are: 1) it limits the ability to determine the current state and the predictability of precision oscillators - for example, the noise of an oscillator produces time prediction errors of $\sim \tau\sigma_y(\tau)$ for prediction intervals of τ, 2) it limits synchronization and syntonization accuracies, 3) it can limit a receiver's useful dynamic range, channel spacing, and selectivity, 4) it can cause bit errors in digital communication systems, 5) it can cause loss of lock, and limit acquisition and reacquisition capability in phase-locked-loop systems, and 6) it can limit radar performance, especially when it is a Doppler radar.

2. The Characterization of Noise

It is important to have appropriate statistical measures to characterize the *random*-component of oscillator instability. The voltage from a precision oscillator can be written

$$V(t) = [V_0 + \varepsilon(t)] \sin[\omega_0 t + \phi(t)] ,$$

where V_0 is the nominal amplitude, $\varepsilon(t)$ represents amplitude fluctuations, ω_0 is the nominal angular frequency, and $\phi(t)$ represents phase fluctuations. The amplitude noise is usually small compared to the phase noise, and will not be discussed further, but the validity of this assumption should always be verified. Two methods will be described that are appropriate for characterizing the stochastic variations in the phase or frequency of an oscillator [24]. If they are present, other methods must be used to minimize or remove deterministic effects. Normally, short-term frequency stability measurements are made with the oscillator in a benign environment to reduce extraneous frequency perturbations. Building vibrations, and temperature, pressure, and humidity fluctuations can affect the stability measurements. The effect of aging must usually be removed during stability analysis. Simple methods exist for estimating frequency aging in the presence of random walk frequency noise, which is the usual case.

One method of describing the phase noise of an oscillator is the spectrum of the noise. The spectral density $S_x(f')$ of a quantity x is its mean square value per Hz Fourier *frequency f'. The Fourier frequency is a fictitious frequency used in* Fourier analysis of the signal. Zero Fourier frequency corresponds to the carrier, and a negative Fourier frequency refers to the region below the carrier. The integral of the spectral density over all Fourier frequencies from minus infinity to infinity is the mean square value of the quantity. The spectral density of phase noise $S_\phi(f')$ is very important because it is directly related to the performance of oscillators in RF signal processing applications. The single-sideband noise power per Hz to total signal power ratio is often specified for oscillators instead of the phase spectral density. This ratio has been designated $\mathcal{L}(f)$. Recently, the definition of $\mathcal{L}(f)$ has been changed to one-half $S_\phi(f')$ [25]. When defined this way, $\mathcal{L}(f)$ is equal to the SSB noise to signal ratio only as long as the integrated phase noise from f' to infinity is small compared to one rad^2. The phase spectral density depends on carrier frequency. When the signal from an oscillator is multiplied by n in a noiseless multiplier, the FM sidebands increase in power by n^2, as does the spectral density of phase. Consequently, it is important to state the oscillator frequency together with any measurement of the spectral density of phase. Sometimes oscillator noise is described in terms of frequency rather than phase. The instantaneous angular frequency $\omega(t)$ of an oscillator is the derivative of the total phase.

$$\omega(t) = \omega_0 + \frac{d\phi(t)}{dt}$$

The phase noise in precision oscillators is usually described in terms of a dimensionless instantaneous frequency fluctuation, y, which is defined in terms of the angular frequency.

$$y \equiv \frac{\omega(t) - \omega_0}{\omega_0} = \frac{1}{\omega_0} \frac{d\phi(t)}{dt}$$

Since the frequency is the derivative of the phase, the spectral density of y is simply related to the phase spectral density.

$$S_y(f') = (\frac{2\pi f'}{\omega_0})^2 S_\phi(f')$$

A relatively simple model is adequate to describe the noise in the precision oscillators discussed in this article. The model spectral density consists of a finite sum of terms proportional to the Fourier frequency raised to a positive or negative integer power.

$$S_y(f') = \sum_\alpha h_\alpha (f')^\alpha$$

Six processes are sufficient:

white phase modulation (PM)	$\alpha = +2$
flicker PM	$\alpha = +1$
white frequency modulation (FM)	$\alpha = +0$
flicker FM	$\alpha = -1$
random-walk FM	$\alpha = -2$
random-walk frequency aging	$\alpha = -4$

The even terms can all be produced by the integration or differentiation of white noise. Thus, models containing just these terms have simple Kalman filter or ARIMA model equivalents. Under these circumstances, it is easy to design optimum filters for using the clocks in systems and control loops. If either of the flicker noise processes appears to be present over a substantial Fourier frequency interval, many lag-lead filters may be necessary to even approximately produce the noise from a white noise source, making optimal system design very difficult.

The origin of white PM and white FM in a clock can usually be understood in terms of the physical and electronic design. Additive thermal noise in the buffer amplifiers is the usual source of white PM. Thermal noise in the control loop of a passive atomic frequency standard is a common source of white FM. The physical sources of the other noise processes are usually not known. The random-walk FM noise may only be apparent, that is, it may be the result of frequency fluctuations caused by changes in temperature, temperature gradient, pressure, humidity, magnetic field, or mechanical stress. If the spectrum were truly random walk in nature, the frequency would be unbounded, whereas the frequency of a precision oscillator usually doesn't change by more than a few linewidths. Nevertheless, this term should be included in the oscillator model if the mean frequency of the oscillator appears to change. Similarly, random-walk frequency aging can be included in the model of an oscillator that has a variable frequency aging.

Although the spectrum is a powerful method of characterizing an oscillator, it is not very directly related to its timekeeping ability. The Allan variance, $\sigma_y^2(\tau)$, is an alternate measure of frequency stability that is quite useful for this purpose. It is defined by

$$\sigma_y(\tau) = \frac{1}{2} E\{[\frac{\phi(t+2\tau)-2\phi(t+\tau)+\phi(t)}{\omega_0 t}]^2\}$$

where E{} refers to the expectation value over the ensemble of possible observations. The traditional variance describes the deviation of a set of observations from the mean, but is not defined for noise processes more divergent than white frequency noise. On the other hand, the Allan variance describes the variation of the frequency from one measurement interval to the next with no dead time between intervals. As a result, it converges for both flicker frequency and random walk frequency noise, but is not defined for random walk frequency aging noise. The rms time error of a clock after a free running interval τ is approximately $\tau\sigma_y(\tau)$. Thus, estimating the Allan variance provides a nonparametric method of characterizing the timekeeping ability of a clock.

The terms jitter and wander are used in characterizing timing instabilities in digital communications. Jitter refers to the high-frequency timing variations of a digital signal; wander refers to the low-frequency variations. The dividing line between the two is often taken to be 10 Hz. Wander and jitter can be characterized by the appropriate measurement of the rms time error of the clock. A 10 Hz low pass filter should be used to remove the effects of jitter, if necessary. For very high Fourier frequencies or short integration times, it may be necessary to calculate the jitter from the spectrum rather than measure it directly. For example, the mean-square phase jitter during a one-tenth second interval is the integral of the spectral density of phase over the Fourier frequency range from 10 Hz to infinity.

3. Noise in Crystal Oscillators

Although the causes of noise in crystal oscillators are not fully understood, several causes of short term instabilities have been identified. Temperature fluctuations can cause short term instabilities via thermal transient effects (see section III.D.1 below), and via activity dips at the oven set-point in OCXOs. Other causes include Johnson noise in the crystal unit, random vibration (see section III.E.1 below), noise in the oscillator circuitry (both the active and passive components can be significant noise sources), and fluctuations at various interfaces on the resonator, e.g., in the number of molecules adsorbed on the resonator's surface.

In a properly designed oscillator, the resonator is the primary noise source close to the carrier and the oscillator circuitry is the primary source far from the carrier. The noise close to the carrier, i.e., within the bandwidth of the resonator, has a strong inverse relationship with resonator Q; $S_\phi(f') \propto 1/Q^4$. In the time domain, $\sigma_y(\tau) \approx 1/Q(2 \times 10^{-7})$ at the noise floor. In the frequency domain, the noise floor is limited by Johnson noise, the noise power of which is kT = -174 dBm/Hz at 290°K. A higher signal, i.e., a higher resonator drive current, will improve the noise floor but not the close-in noise. In fact, for reasons that are not fully understood, above a certain point, higher drive levels generally degrade the close-in noise. For example, the maximum "safe" drive level is about 100 μa for a 5 MHz 5th overtone AT-cut resonator with $Q \approx 2.5$ million. The safe drive current can be substantially higher for high frequency SC-cut resonators. For example, $\mathcal{L}(f) \approx -180$ dBc/Hz has been achieved with 100 MHz 5th overtone SC-cut resonators at drive currents $\approx$ 10 ma. However, such a noise capability is useful only in a vibration-free environment, for if there is vibration at the offset frequencies of interest, the vibration-induced noise will dominate the quiescent noise of the oscillator - see section III.E.1 below.

When low noise is required in the microwave (or higher) frequency range, surface acoustic wave (SAW) oscillators and dielectric resonator oscillators (DRO) are sometimes used. When compared with multiplied-up (bulk-acoustic wave) quartz oscillators, these oscillators can provide lower noise far from the carrier, at the expense of poorer noise close to the carrier, poorer aging, and poorer temperature stability. SAW oscillators and DROs can provide lower noise far from the carrier because these devices can be operated at higher drive levels, thereby providing higher signal to noise ratios, and because the devices operate at higher frequencies, thereby minimizing the "20 log N" losses due to frequency multiplication by N. $\mathcal{L}(f) = -180$ dBc/Hz noise floors have been achieved with state-of-the-art SAW oscillators. Of course, as is the case for high frequency bulk-wave oscillators, such noise floors are realizable only in environments which are free of vibrations at the offset frequencies of interest.

4. Noise in Atomic Frequency Standards

Most devices function in a way that can be described by classical physics. In this regime, the quality of measurements is limited only by thermal noise and the capability of the measurement equipment. Atomic frequency standards, on the other hand, depend on the quantum nature of the atom. This means that when a microwave field is applied to an atom, there is a probability that the atom will make a transition. Since all atomic frequency standards function by observing the effect of microwaves on the atoms, there is a variability in the outcome of each 'observation' that limits the precision of the frequency standard. This noise enters into active and passive frequency standards somewhat differently, but results in white frequency noise in both cases. In an active oscillator, the noise outside the bandwidth of the atomic resonance integrates to produce phase shocks. Since the transition rate of the atoms is independent of the phase of the stimulating field, there is no correction of these phase errors, and the total phase displays a random walk, which is the same as white frequency noise. In a passive atomic frequency standard, the correction signal is proportional to the number of atoms and the frequency error of the microwave source. The fluctuations in detected signal due to transition probability have a white spectrum and are proportional to the square root of the number of atoms. Thus, even when there is no error in the microwave frequency, the control impresses frequency corrections on the VCXO. This results in white frequency noise inversely proportional to the square root of the number of atoms.

In addition, all atomic frequency standards have white phase noise due to additive thermal noise in the buffer amplifiers used to provide the output signal. They also display more divergent noises. All atomic frequency standards suffer from random walk frequency noise. However, the source of this noise cannot be explained in the same physical manner as the white phase and frequency modulations. The random walk frequency noise most likely results from the sum of many different frequency perturbations due to variations in temperature, pressure, humidity, and magnetic field. It is likely that a more complete understanding of each atomic frequency standard would show that the frequency variations described as random walk frequency noise are actually deterministic. However, this is not possible today, and the stochastic description of the long term frequency variations makes it possible to design appropriate filters for using the frequency standards optimally. Some atomic frequency standards suffer from (variable) frequency aging due to even closer coupling of the atoms to the environment.

D. Frequency vs. Temperature Stability

1. Frequency vs. Temperature Stability of Quartz Oscillators

a. Static f vs. T Stability

As an illustration of the effects that temperature can have on frequency stability, Fig. 23 shows the accuracy vs. temperature of a typical quartz wristwatch. Near the wrist temperature, the watch can be very accurate because the frequency of the crystal (i.e., the clock rate) changes very little with temperature. However, when the watch is cooled to -55°C or heated to +100°C, it loses about 20 seconds per day, because the typical temperature coefficient of frequency of the tuning fork crystals used in quartz watches is -0.035 ppm/°C^2.

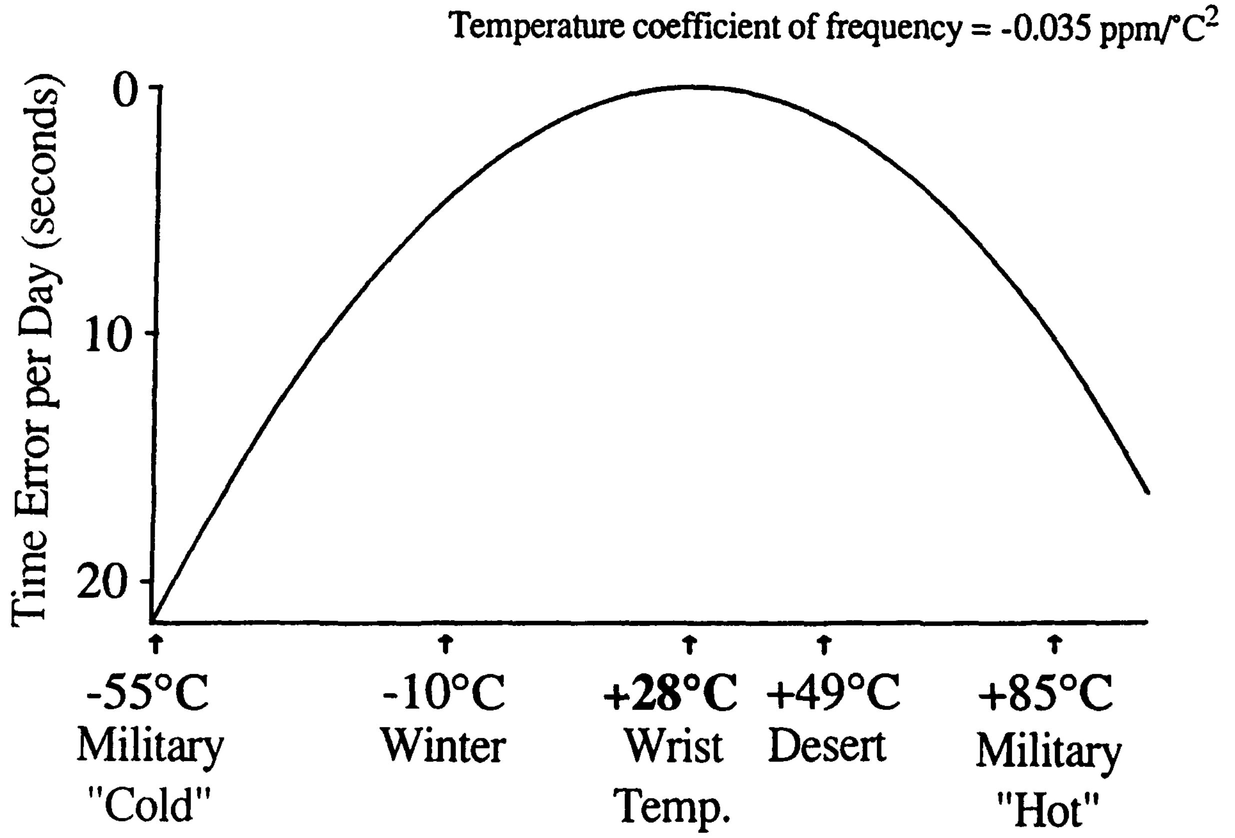

Figure 23. Wristwatch accuracy vs. temperature.

The static f vs. T characteristics of crystal units are determined primarily by the angles of cut of the crystal plates with respect to the crystallographic axes of quartz. "Static" means that the rate of change of temperature is slow enough for the effects of temperature gradients (explained later) to be negligible. As Fig. 13 illustrates for the AT-cut, a small change in the angle of cut (seven minutes in the illustration) can significantly change the f vs. T characteristics. The points of zero temperature coefficient, the "turnover points," can be varied over a wide range by varying the angles of cut. The f vs. T characteristics of SC-cut crystals are similar to the curves shown in Fig. 13, with T_i shifted to about 95°C.

Other factors that can affect the f vs. T characteristics of crystal units include: the overtone; the geometry of the crystal plate; the size, shape, thickness, density and stresses of the electrodes; the drive level; impurities and strains in the quartz material; stresses in the mounting structure; interfering modes; ionizing radiation; the rate of change of temperature (i.e., thermal gradients); and thermal history. The last two factors are important for understanding the behaviors of OCXOs and TCXOs, and are, therefore, discussed separately.

Interfering modes can cause "activity dips," as is illustrated in Fig. 24. Activity dips can be strongly influenced by the crystal's drive level and load reactance. Near the activity dip temperature, anomalies appear in both the f vs. T and R vs. T characteristics. The activity dip temperature is a function of C_L because the interfering mode usually has a large temperature coefficient and a C_1 that is different from that of the desired mode. Activity dips are troublesome in TCXOs, and also in OCXOs when the dip occurs at the oven temperature. When the resistance increases at the activity dip, and the oscillator's gain margin is insufficient, the oscillation stops. The incidence of activity dips in SC-cut crystals is far lower than in AT-cut crystals.

An important factor that affects the f vs. T characteristics of crystal oscillators is the load capacitor. When a capacitor is connected in series with the crystal, the f vs. T characteristic of the combination is rotated slightly from that of the crystal alone. The temperature coefficient of the load capacitor can greatly magnify the rotation.

The f vs. T of crystals can be described by a polynomial function. A cubic function is usually sufficient to describe the f vs. T of AT-cut and SC-cut crystals to an accuracy of ±1 ppm. In the MCXO, in order to fit the f vs. T data to $\pm 1 \times 10^{-8}$, a polynomial of at least sixth order is usually necessary.

b. Dynamic f vs. T Effects

Changing the temperature surrounding a crystal unit produces thermal gradients when, for example, heat flows to or from the active area of the resonator plate through the mounting clips. The static f vs. T characteristic is modified by the "thermal-transient effect." When an OCXO is turned on, there can be a significant thermal-transient effect.

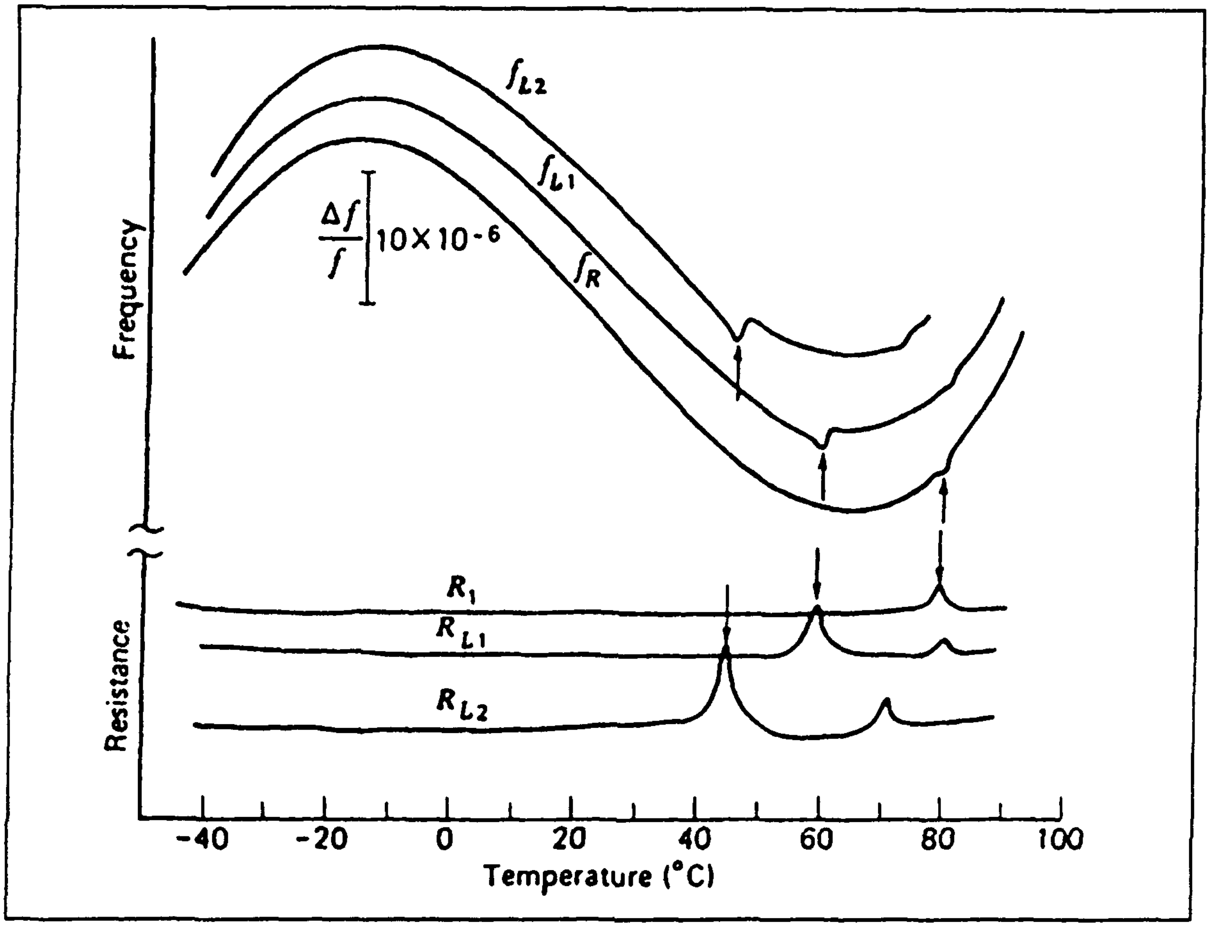

Figure 24. Activity dips in the f vs. T and R vs. T characteristics, with and without C_L.

Figure 25 shows what happens to the frequency output of two OCXOs, each containing an oven that reaches the equilibrium temperature in six minutes. One oven contains an AT-cut, the other, an SC-cut crystal. Thermal gradients in the AT-cut produce a large frequency undershoot that anneals out several minutes after the oven reaches equilibrium. The SC-cut crystal, being insensitive to such thermal-transients, reaches the equilibrium frequency as soon as the oven stabilizes.

In addition to extending the warmup time of OCXOs, when crystals other than SC-cuts are used, the thermal-transient effect makes it much more difficult to adjust the temperature of OCXO ovens to the desired turnover points, and the OCXO frequencies are much more sensitive to oven temperature fluctuations.

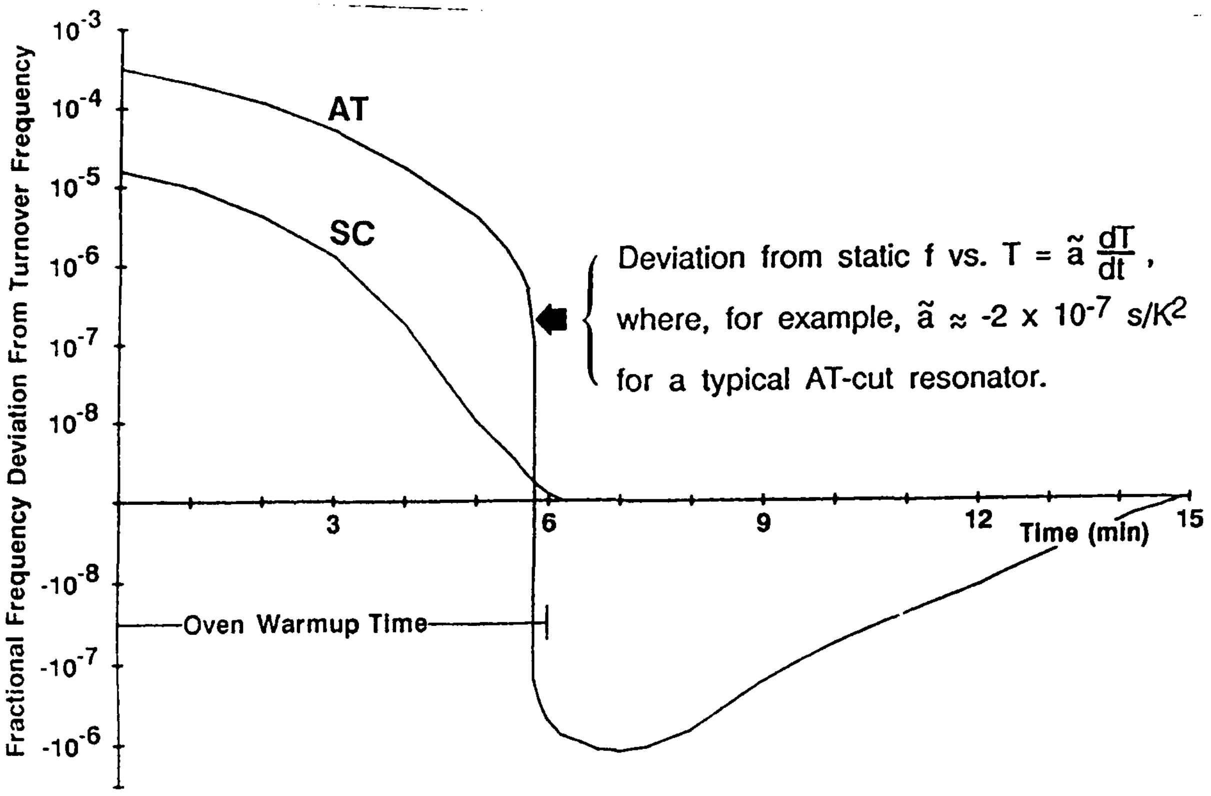

Deviation from static f vs. T = $\tilde{a}\frac{dT}{dt}$, where, for example, $\tilde{a} \approx -2 \times 10^{-7}$ s/K^2 for a typical AT-cut resonator.

Figure 25. Warmup characteristics of AT-cut and SC-cut OCXOs.

The testing and compensation accuracies of TCXOs are also adversely affected by the thermal-transient effect. As the temperature is changed, the thermal-transient effect distorts the static f vs. T characteristic, which leads to apparent hysteresis. The faster the temperature is changed, the larger is the contribution of the thermal-transient effect to the f vs. T performance.

c. Thermal Hysteresis and Retrace

The f vs. T characteristics of crystal oscillators do not repeat exactly upon temperature cycling [26]. The lack of repeatability in TCXOs, "thermal hysteresis," is illustrated in Fig. 26. The lack of repeatability in OCXOs, "retrace," is illustrated in Fig. 27. Hysteresis is defined [27] as the difference between the up-cycle and the down-cycle f vs. T characteristics, and is quantified by the value of the difference at the temperature where the difference is maximum. Hysteresis is determined during a complete quasistatic temperature cycle between specified temperature limits. Retrace is defined as the nonrepeatability of the f vs. T characteristic at a fixed temperature (which is usually the oven temperature of an OCXO) upon on-off cycling an oscillator under specified conditions.

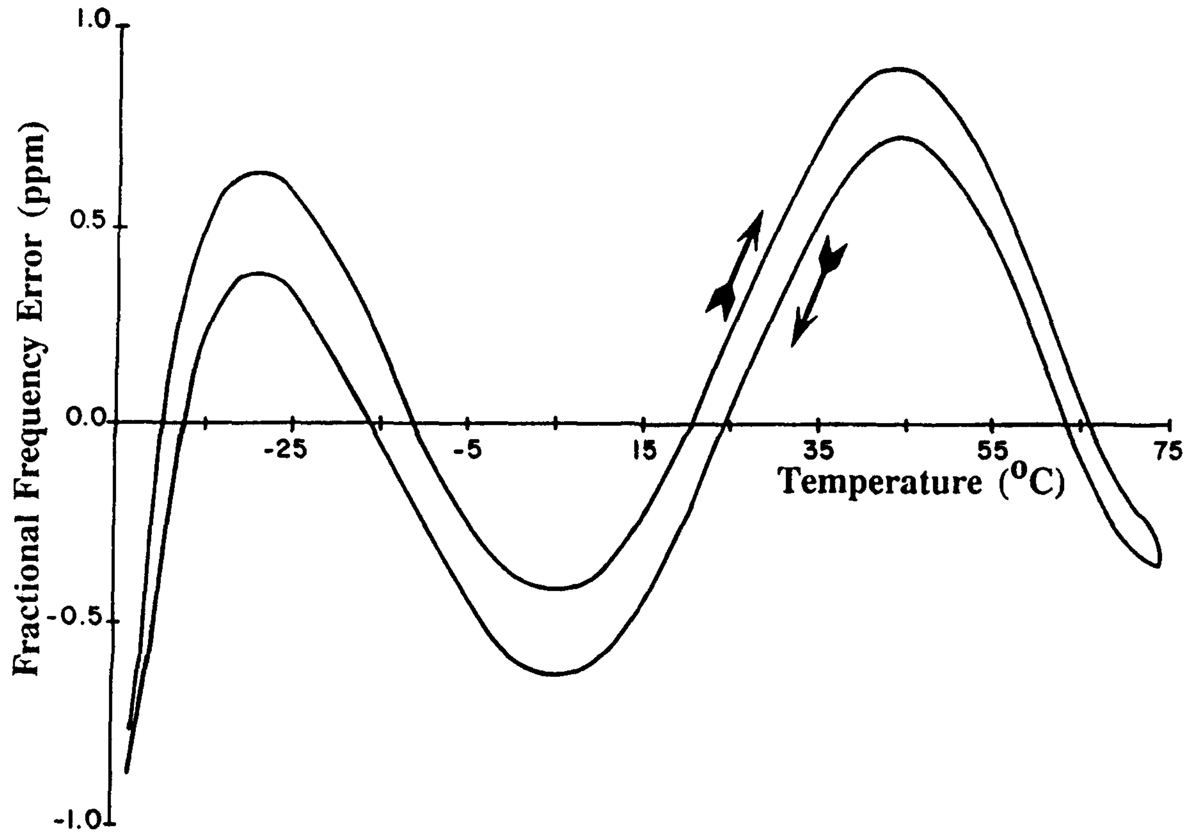

Figure 26. TCXO thermal hysteresis.

Hysteresis is the major factor limiting the stability achievable with temperature compensated crystal oscillators. It is especially so in the MCXO because, in principle, the digital compensation method used in the MCXO would be capable of compensating for the f vs. T variations to arbitrary accuracy if the f vs. T characteristics could be described by a single valued function. Retrace limits the accuracies achievable with OCXOs in applications where the OCXO is on-off cycled. Typical values of hysteresis in TCXOs range from 1 ppm to 0.1 ppm when the temperature cycling ranges are 0°C to 60°C, to -55°C to +85°C. Typical OCXO retrace specifications, after a 24 hour off period at about 25°C, range from 2 X 10^{-8} to 1 X 10^{-9}. Low-temperature storage during the off period, and extending the off period, usually make the retrace worse.

The causes of hysteresis and retrace are not well understood; the experimental evidence to date is inconclusive. The mechanisms that can cause these effects include:

strain changes, changes in the quartz, oscillator circuitry changes, contamination redistribution in the crystal enclosure, and apparent hysteresis or retrace due to thermal gradients.

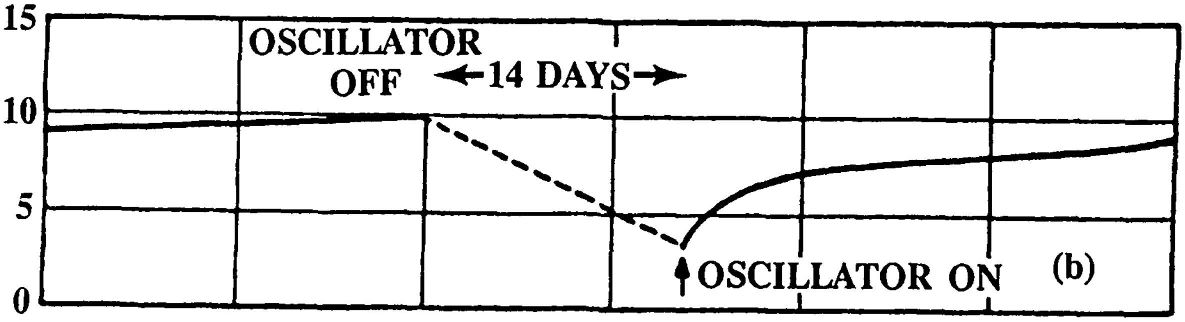

Figure 27. OCXO retrace.

2. Frequency vs. Temperature Stability of Atomic Clocks

The frequency sensitivity of atomic oscillators to temperature changes results from actual atomic frequency changes and electronic effects. The temperature sensitivity of a typical miniature rubidium standard is 3 X 10^{-10} over the temperature range -55° to +65°C. The pressure shift is probably the largest contributor to this temperature sensitivity. The small size of the frequency standard limits the quality of the rubidium cell thermal control, and the temperature variations produce pressure changes of the buffer gas. In addition to the pressure shift, all the effects discussed below in the context of the other atomic frequency standards contribute to the temperature sensitivity of rubidium frequency standards.

The cesium standard has a much lower temperature coefficient. A typical value is 3 X 10^{-12} over the range 0° to 50°C. The cesium atoms are better insulated from the environment than in the case of rubidium. Cavity pulling, residual Doppler shifts and microwave power changes all contribute to the residual variations of the atomic frequency. In addition, imperfections in the electronics result in an offset of the operating frequency from the atomic frequency that changes with temperature as the active component offset voltages and currents change.

The temperature sensitivity of hydrogen masers is comparable to that of cesium standards. In active masers, the most significant contributor is cavity pulling. Stable ceramic cavities, and four or five levels of thermal control are used to limit the thermal sensitivity. Active control of the frequency of the cavity is sometimes used, which reduces the requirement on thermal control. Electronic effects are not a problem for active masers, but do affect the frequency of passive masers.

Retrace in cesium standards varies between 5 X 10^{-13} and 3 X 10^{-12}. In miniature rubidium standards, the typical retrace is 5 X 10^{-11}.

E. Warmup

When power is applied to a frequency standard, it takes a finite amount of time before the equilibrium frequency stability is reached. Figure 25 illustrates the warmup of two OCXOs. The warmup time of an oscillator is a function of the thermal properties of the oscillator, the input power, and the oscillator's temperature prior to turn-on. Typical warmup time specifications of OCXOs and rubidium frequency standards, e.g., from a 0°C start, range from three to ten minutes. The warmup times of cesium standards range from 30 to 60 minutes. Hydrogen masers warm up in 4 hours to one day. Even TCXOs, MCXOs and simple XOs take a few seconds to warm up, although these are not ovenized. The reasons are that it takes a finite amount of time for the signal to build up in any high Q circuit, and the few tens of milliwatts of power which are dissipated in these oscillators do change the thermal conditions within the oscillators.

F. Acceleration Effects

1. Acceleration Effects in Crystal Oscillators

Acceleration changes a crystal oscillator's frequency [28]. The acceleration can be a steady-state acceleration, vibration, shock, attitude change (2-g tipover), or acoustic noise. The amount of frequency change depends on the magnitude and direction of the acceleration $\vec{A}$, and on the acceleration sensitivity of the oscillator $\vec{\Gamma}$. The acceleration sensitivity $\vec{\Gamma}$ is a vector quantity. The frequency change can be expressed as

$$\frac{\Delta f}{f} = \vec{\Gamma} \bullet \vec{A} \, .$$

Typical values of $|\vec{\Gamma}|$ are in the range of 10^{-9} per g to 10^{-10} per g. For example, when $\vec{\Gamma} = 2 \times 10^{-9}$ per g and is normal to the earth's surface, and the oscillator is turned upside down (a change of 2 g's), the frequency changes by 4 X 10^{-9}. When this oscillator is vibrated in the up and down direction, the time dependent acceleration modulates the oscillator's output frequency at the vibration frequency, with an amplitude of 2 X 10^{-9} per g. In the frequency domain, the modulation results in vibration induced sidebands that appear at plus and minus integer multiples of the vibration frequency from the carrier frequency. Figure 28 shows the output of a spectrum analyzer for a 10 MHz, 1.4 X 10^{-9} per g oscillator that was vibrated at 100 Hz and 10 g. When the frequency is multiplied, as it is in many applications, the sideband levels increase by 20 dB for each 10X multiplication. The increased sideband power is extracted from the carrier. Under certain conditions of multiplication, the carrier disappears, i.e., all the energy is then in the sidebands.

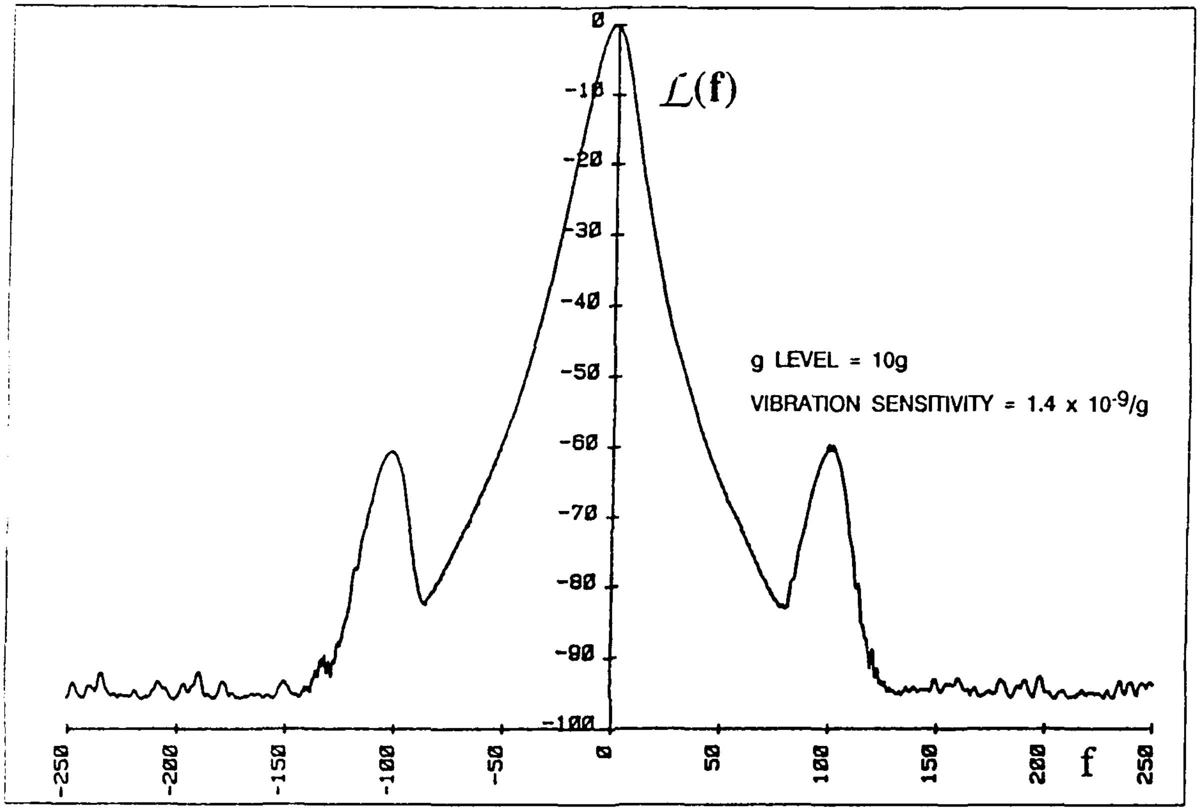

Figure 28. Vibration-induced sidebands.

The effect of random vibration is to raise the phase noise level of the oscillator. The degradation of phase noise can be substantial when the oscillator is on a vibrating platform, e.g., on an aircraft. Figure 29 shows a typical aircraft random vibration specification (power spectral density vs. vibration frequency) and the resulting vibration-induced phase noise degradation. Acoustic noise is another source of acceleration that can affect the frequency of oscillators.

During shock, a crystal oscillator's frequency changes suddenly due to the sudden acceleration. The frequency change follows the above expression for acceleration-induced frequency change, except, if during the shock some elastic limits in the crystal's support structure or electrodes are exceeded (as is almost always the case during typical shock tests) the shock will produce a permanent frequency change. If the shock level is sufficiently high, the crystal will break; however, in applications where high shock levels

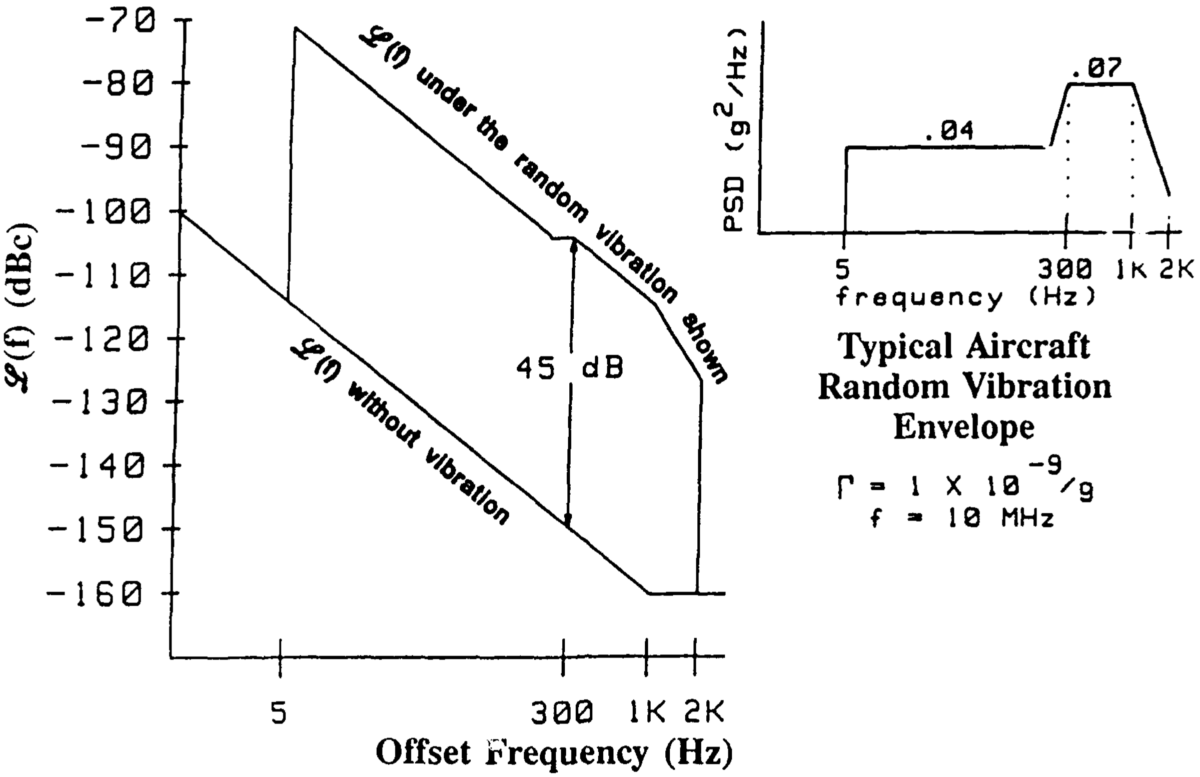

Figure 29. Random vibration-induced phase noise degradation.

are a possibility, crystal units with chemically polished crystal plates can be used. Such crystals can survive shocks in excess of 20,000 g (and have been fired successfully from howitzers [29]).

2. Acceleration Effects in Atomic Frequency Standards

Static acceleration is not much of a problem in atomic frequency standards. For example, miniature rubidium frequency standards exhibit acceleration sensitivity of approximately 2×10^{-12}/g. However, transient acceleration due to vibration can have significant effects. Control loop time constants of a fraction of a second are typical, hence the sensitivity of the frequency to vibration throughout the audio range is primarily determined by the quartz crystal unit. Not only is the quality of the spectrum degraded, but the time errors due to random vibration accumulate. The time (or phase) errors do not average out because the white frequency noise is integrated to produce random walk of the phase. The time errors increase proportionally to the square-root of the elapsed time. An additional vibration induced problem occurs because the microwave signal used to interrogate the atoms is produced by frequency multiplication, a process that transfers power from the carrier to the modulation sidebands. Under intense vibration, the carrier may essentially disappear, and loss of lock occurs.

G. Magnetic Field Effects

1. Magnetic Field Effects in Quartz Oscillators

Quartz is diamagnetic; however, magnetic fields can affect magnetic materials in the crystal unit's mounting structure, electrodes, and enclosure. Time varying electric fields will induce eddy currents in the metallic parts. Magnetic fields can also affect components such as inductors in the oscillator circuitry. When a crystal oscillator is designed to minimize the effects of magnetic fields, the sensitivity can be much less than 10^{-10} per oersted. Magnetic field sensitivities on the order of 10^{-12} per oersted have been measured [30] in crystal units designed specifically for low magnetic-field sensitivity.

2. Magnetic Field Effects in Atomic Frequency Standards

Atomic frequency standards are particularly sensitive to magnetic fields because the hyperfine frequency is proportional to a magnetic interaction energy. All the commercial atomic frequency standards use atomic transitions with small quadratic magnetic field dependence at the operating magnetic field of the unit - typically a few thousandths of an oersted. Several layers of magnetic shielding are used to provide a stable magnetic environment. Typical sensitivity for a miniature rubidium standard is 2×10^{-11} per oersted change in the external field. *Typical cesium standards have* magnetic field sensitivity of 2×10^{-12} per oersted. Both rubidium and cesium are available in high performance versions that have ten times smaller magnetic field sensitivity obtained by using additional shielding. The magnetic field sensitivity of hydrogen masers is comparable to that of cesium standards.

H. Radiation Effects

1. Radiation Effects in Quartz Oscillators

Ionizing radiation changes a crystal oscillator's frequency primarily because of changes the radiation produces in the crystal unit [31]. Under certain conditions, the radiation will also produce an increase in the crystal unit's equivalent series resistance. The resistance increase can be large enough to stop the oscillation when the oscillator is not radiation hardened.

Figure 30 shows a crystal oscillator's idealized frequency response to a pulse of ionizing radiation. The response consists of two parts. Initially, there is a transient frequency change that is due primarily to the thermal-transient effect caused by the sudden deposition of energy into the crystal unit. This effect is a manifestation of the

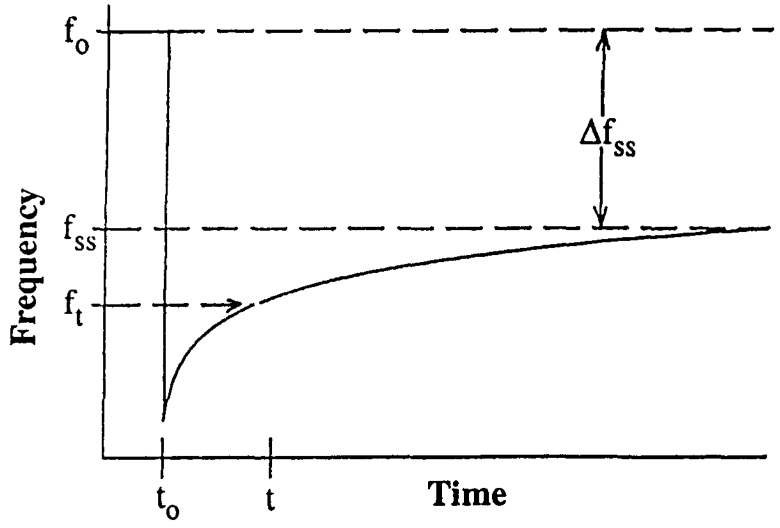

f_o = original, preirradiation frequency

Δf_{ss} = steady-state frequency change (0.2 to 24 hours after exposure)

f_t = instantaneous frequency at any time (t)

$$\Delta f'_{ss}/\mathrm{rad}^* \approx \begin{cases} 10^{-11} \text{ for natural quartz (R increase can stop the oscillation)} \\ 10^{-12} \text{ for cultured quartz} \\ 10^{-13} \text{ for swept cultured quartz} \end{cases}$$

* for 1 Mrad dose

Figure 30. Crystal oscillator's response to a pulse of ionizing radiation.

dynamic f vs. T effect discussed earlier. The transient effect is absent in SC-cut resonators made of high purity quartz.

After steady-state is reached, there is a permanent frequency offset which is a function of the radiation dose and the nature of the crystal unit. The frequency change vs. dose is nonlinear, the change per rad being much larger at low doses than at large doses. At doses above 1 krad (SiO_2), the rate of frequency change with dose is quartz impurity defect dependent. For example, at a 1 Mrad dose, the frequency change can be as large as 10 ppm when the crystal unit is made from natural quartz; it is typically one to a few ppm when the crystal is made from cultured quartz, and it can be as small as 0.02 ppm when the crystal is made from swept cultured quartz.

The impurity defect of major concern in quartz is the substitutional Al^{3+} defect with its associated interstitial charge compensator, which can be an H^+, Li^+, or Na^+ ion, or a hole. This defect substitutes for a Si^{4+} in the quartz lattice. Radiation can result in a change in the position of weakly bound compensators, which changes the elastic constants of quartz and thereby leads to a frequency change. The movement of ions also results in a decrease in the crystal's Q, i.e., in an increase in the crystal's equivalent series resistance. If the oscillator's gain margin is insufficient, the increased resistance stops the oscillation.

Sweeping is a high-temperature, electric-field driven, solid-state purification process in which the weakly bound alkali compensators are diffused out of the lattice and replaced by more tightly bound H^+ ions and holes. In the typical sweeping process, conductive electrodes are applied to the Z-surfaces of a quartz bar, the bar is heated to about 500°C, and a voltage is applied so as to produce an electric field of about 1 kV/cm along the Z-direction. After the current through the bar (due to the diffusion of impurities) decays to some constant value, the bar is cooled slowly, the voltage is removed, and then the electrodes are removed. Crystal units made from swept quartz exhibit neither the radiation-induced Q degradation nor the large radiation induced frequency shifts. Swept quartz (or low aluminum content quartz) should be used in oscillators which are expected to be exposed to ionizing radiation.

At low doses [32], e.g., at a few rads, the frequency change per rad can be as high as 10^{-9} per rad. The low dose effect is not well understood. It is not impurity dependent, and it saturates at about 300 rads. At very high doses, i.e., at >> 1 Mrad, the impurity dependent frequency shifts also saturate because, since the number of defects in the crystal are finite, the effects of the radiation interacting with the defects are also finite.

When a fast neutron hurtles into a crystal lattice and collides with an atom, it is scattered like a billiard ball. A single such neutron can produce numerous vacancies, interstitials, and broken interatomic bonds. The effect of this "displacement damage" on oscillator frequency is dependent primarily upon the neutron fluence. The frequency of oscillation increases nearly linearly with neutron fluence at rates of: $8 \times 10^{-21}/n/cm^2$ at a fluence range of 10^{10} to 10^{12} n/cm^2, $5 \times 10^{-21}/n/cm^2$ at 10^{12} to 10^{13} n/cm^2, and $0.7 \times 10^{-21}/n/cm^2$ at 10^{17} to 10^{18} n/cm^2.

2. Radiation Effects in Atomic Frequency Standards

To the extent that they occur fast compared to the loop time constant, the transient effects of radiation on the quartz oscillator are unaffected by the frequency control loop in the atomic oscillator. Longer term effects are reduced by the loop gain up to a limit imposed by the effect of the radiation on the control electronics. For example, a military rubidium standard irradiated with a dose of 600 rad (Si) at a rate of 1×10^{10} rad (Si) per second suffered an offset of 6×10^{-8} after one second, but recovered to better than 5×10^{-11} in 20 seconds. Both rubidium and cesium standards have been radiation hardened to levels appropriate for use on the GPS satellites. The physics packages are intrinsically insensitive to radiation dose, but the electronics may require latchup protection, additional gain margin and other similar modifications for use in radiation environments.

I. Other Effects on Stability

Ambient pressure change (as during an altitude change) can change a crystal oscillator's frequency if the pressure change produces a deformation of the crystal unit's or the oscillator's enclosure (thus changing stray capacitances and stresses). The pressure change can also affect the frequency indirectly through a change in heat transfer conditions inside the oscillator. Humidity changes can also affect the heat transfer conditions. In addition, moisture in the atmosphere will condense on surfaces when the temperature falls below the dew point, and can permeate materials such as epoxies and polyimides, and thereby affect the properties (e.g., conductivities and dielectric constants) of the oscillator circuitry. The frequency of a properly designed crystal oscillator changes less than 5×10^{-9} when the environment changes from one atmosphere of air to a vacuum.

All atomic frequency standards are indirectly sensitive to pressure to the extent that the change in the thermal conductivity of the air modifies the thermal gradients within the unit. Hydrogen masers can be directly affected by pressure changes through the cavity pulling if the microwave cavity is not isolated from these changes or compensated. Rubidium frequency standards are directly affected by pressure changes as a result of the distortion of the rubidium cell and the consequent change in density of the buffer gas. The sensitivity of rubidium frequency standards to variations in ambient pressure is approximately 1×10^{-13} per torr. The typical cesium standard specification is "altitude: $< 2 \times 10^{-12}$ change up to 12.2 km (40,000 ft.)."

Electric fields can change the frequency of a crystal unit. An ideal AT-cut is not affected by a dc voltage on the crystal electrodes, but "doubly rotated cuts," such as the SC-cut, are affected. For example, the frequency of a fundamental mode SC-cut crystal changes 7×10^{-9} per volt. Direct-current voltages on the electrodes can also cause *sweeping, which can affect the frequencies of all cuts.*

Power supply and load impedance changes affect the oscillator circuitry and, indirectly, the crystal's drive level and load reactance. A change in load impedance changes the amplitude or phase of the signal reflected into the oscillator loop, which changes the phase (and frequency) of the oscillation. The effects can be minimized through voltage regulation and the use of buffer amplifiers. The frequency of a "good" crystal oscillator changes $< 5 \times 10^{-10}$ for a 10% change in load impedance. The typical sensitivity of a high-quality crystal oscillator to power supply voltage changes is 5×10^{-11}/V; that of a rubidium frequency standard is 5×10^{-12}/V.

Gas permeation under conditions where there is an abnormally high concentration of hydrogen or helium in the atmosphere can lead to anomalous aging rates. For example, hydrogen can permeate into "hermetically" sealed crystal units in metal enclosures, and helium can permeate through the walls of glass-enclosed crystal units and through the walls of the glass bulbs of rubidium standards.

J. Interactions Among the Influences on Stability

The various influences on frequency stability can interact in ways that lead to erroneous test results, if the interfering influence is not recognized during testing. For example, building vibrations can interfere with the measurement of short-term stability. Vibration levels of 10^{-3} to 10^{-2} g are commonly present in buildings. Therefore, if an oscillator's acceleration sensitivity is 1×10^{-9} per g, then the building vibrations alone can contribute short term instabilities at the 10^{-12} to 10^{-11} level.

The 2-g tipover test is often used to measure the acceleration sensitivity of crystal oscillators. Thermal effects can interfere with this test, because when an oscillator is turned upside down, the thermal gradients inside the oven can vary due to changes in convection currents. Other examples of interfering influences include: temperature and drive level changes interfering with aging tests; induced voltages due to magnetic fields interfering with vibration-sensitivity tests; and the thermal-transient effect, humidity changes, and load reactance temperature coefficients interfering with the measurement of crystal units' static f vs. T characteristics.

An important effect in TCXOs is the interaction between the frequency adjustment during calibration and the f vs. T stability [33]. This phenomenon is called the "trim effect." In TCXOs, a temperature dependent signal from a thermistor is used to generate a correction voltage that is applied to a varactor in the crystal network. The resulting reactance variations compensate for the crystal's f vs. T variations. During calibration, the crystal's load reactance is varied to compensate for the TCXO's aging. Since the frequency vs. reactance relationship is nonlinear, the capacitance change during calibration moves the operating point on the frequency vs. reactance curve to a point where the slope of the curve is different, which changes the compensation; i.e., compensating for aging degrades the f vs. T stability. Figure 31 shows how, for the same compensating C_L vs. T, the compensating f vs. T changes when the operating point is moved to a different C_L. Figure 32 shows test results for a "0.5 ppm" TCXO that had a ±6 ppm frequency adjustment range (to allow for aging compensation for the life of the device). When delivered, this TCXO met its 0.5 ppm f vs. T specification; however, when the frequency was adjusted ±6 ppm during testing, the f vs. T performance degraded significantly.

IV. Oscillator Comparison and Selection

The following discussion applies to wide-temperature-range frequency standards, i.e., to those which are designed to operate over a temperature range that spans at least 90°C. Laboratory devices that operate over a much narrower temperature range can have better stabilities than those in the comparison below.

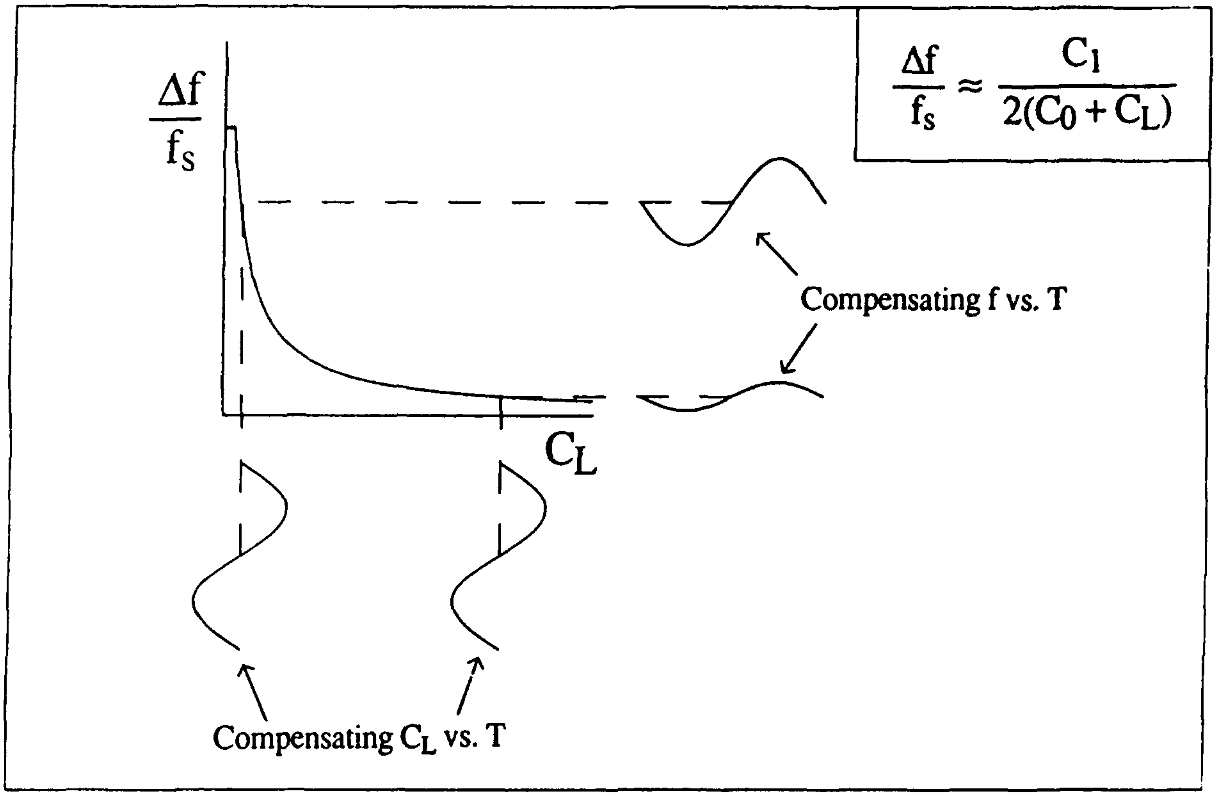

Figure 31. Change in compensating f vs. T due to C_L change.

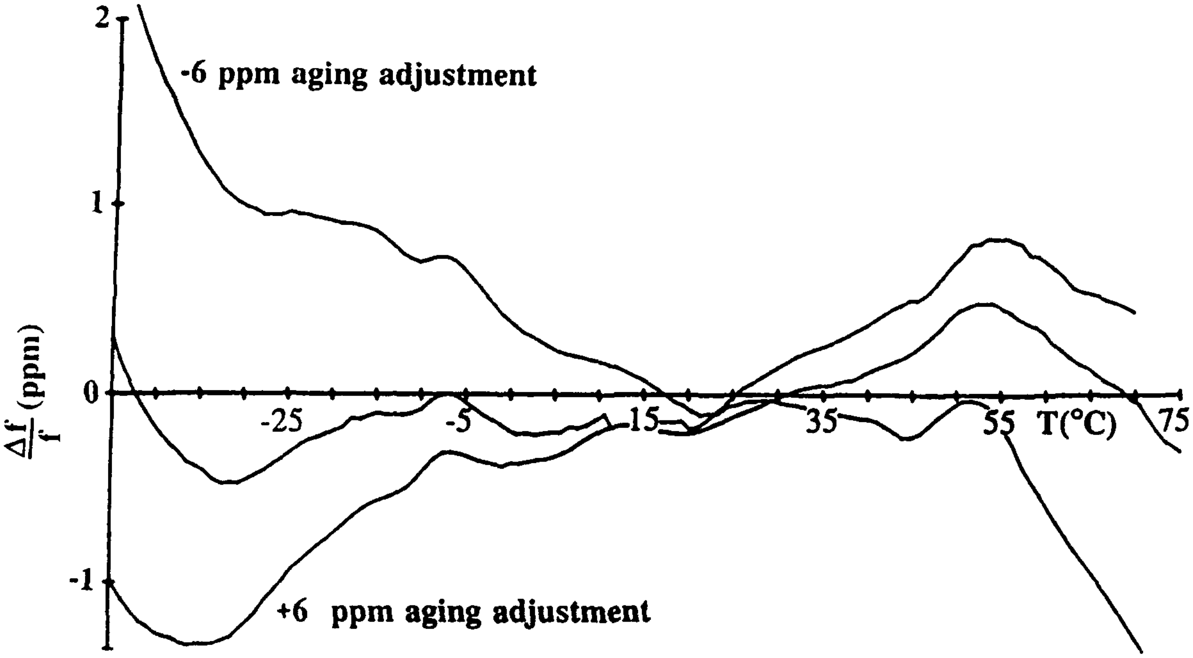

Figure 32. TCXO trim effect.

Commercially available frequency sources cover an accuracy range of several orders of magnitude - from the simple XO to the cesium beam frequency standard. As the accuracy increases, so does the power requirement, size, and cost. Figure 33, for example, shows the relationship between accuracy and power requirement. Accuracy vs. cost would be a similar relationship, ranging from about a dollar for a simple XO to about $40,000 for a cesium standard (1990 prices). Table 1 shows a comparison of salient characteristics of frequency standards. Figure 34 shows the comparison of stability ranges as a function of averaging time, Fig. 35 shows a comparison of phase noise characteristics, and Table 2 shows a comparison of weaknesses and wearout mechanisms.

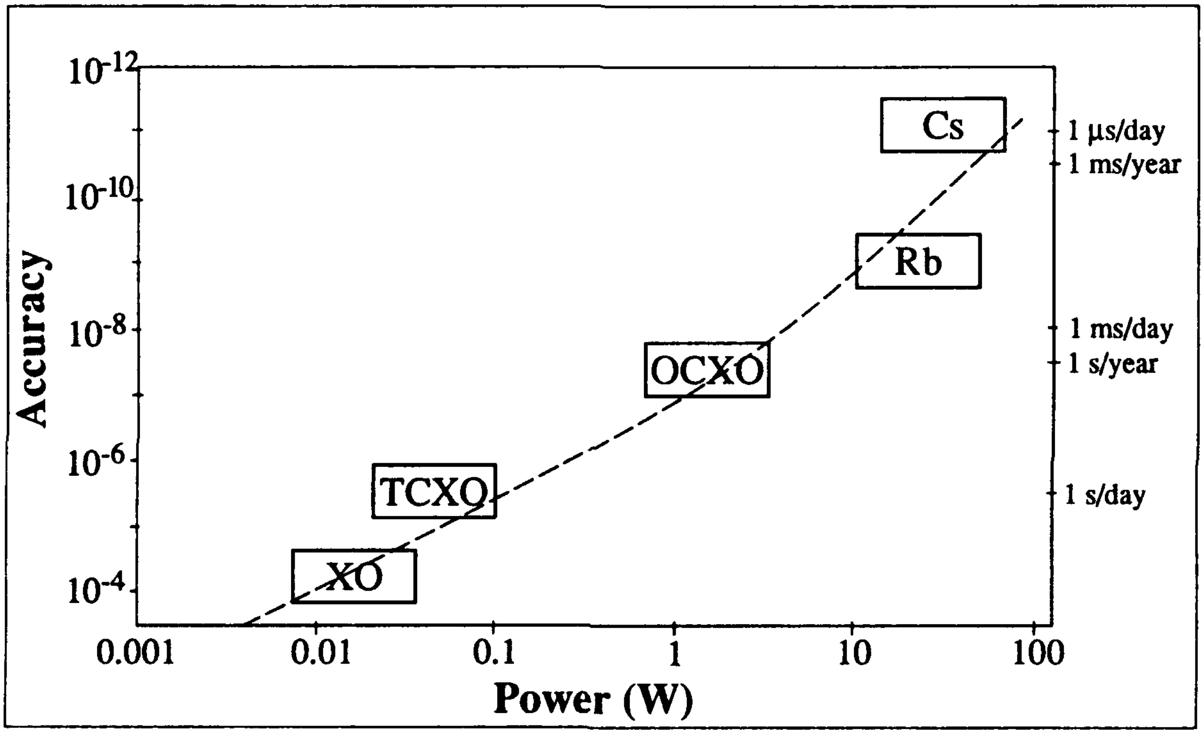

Figure 33. Accuracy vs. power requirement.

The RbXO in Table 1 is the rubidium-crystal oscillator, a device intended for applications where power availability is limited, but where atomic frequency standard accuracy is needed. It consists of a rubidium frequency standard, a low-power and high-stability crystal oscillator, and control circuitry which adjusts the crystal oscillator's frequency to that of the rubidium standard. The rubidium standard is turned on periodically, e.g., once a week, for the few minutes it takes for it to warm up and correct the frequency of the crystal oscillator. With the RbXO, one can approach the long-term stability of the rubidium standard with the low (average) power requirement of the crystal oscillator.

	Quartz Oscillators			Atomic Oscillators		
	TCXO	**MCXO**	**OCXO**	**Rubidium**	**RbXO**	**Cesium**
Accuracy* (per year)	2×10^{-6}	5×10^{-8}	1×10^{-8}	5×10^{-10}	7×10^{-10}	2×10^{-11}
Aging/Year	5×10^{-7}	2×10^{-8}	6×10^{-9}	2×10^{-10}	2×10^{-10}	0
Temp. Stab. (range, °C)	5×10^{-7} (-55 to +85)	2×10^{-8} (-55 to +85)	1×10^{-9} (-55 to +85)	3×10^{-10} (-55 to +68)	5×10^{-10} (-55 to +85)	2×10^{-11} (-28 to +65)
Stability, $\sigma_y(\tau)$ ($\tau = 1$ s)	1×10^{-9}	1×10^{-10}	1×10^{-12}	3×10^{-11}	5×10^{-12}	5×10^{-11}
Size (cm^3)	10	50	20-200	800	1200	6000
Warmup Time (min)	0.1 (to 1×10^{-6})	0.1 (to 2×10^{-8})	4 (to 1×10^{-8})	3 (to 5×10^{-10})	3 (to 5×10^{-10})	20 (to 2×10^{-11})
Power (W) (at lowest temp.)	0.05	0.04	0.25 - 4	20	0.35	30
Price (~$)	100	1,000	2,000	8,000	10,000	40,000

*** Including environmental effects (note that the temperature ranges for Rb and Cs are narrower than for quartz).**

Table 1. Comparison of frequency standards' salient characteristics.

The major questions to be answered in choosing an oscillator include: 1) what frequency accuracy or reproducibility is needed for the system to operate properly; 2) how long must this accuracy be maintained; i.e., will the oscillator be calibrated (or replaced) periodically, or must the oscillator maintain the required accuracy for the life of the system; 3) is ample power available, or must the oscillator operate from batteries; 4) what warmup time, if any, is permissible; 5) what are the environmental extremes in which the oscillator must operate; 6) what is the short term stability (phase noise) requirement; and 7) what is the size constraint.

An important question related to question 2) is: What cost is to be minimized, the initial acquisition cost or the life cycle cost? Often, the cost of recalibration is far higher than the added cost of an oscillator that can provide calibration-free life. A better oscillator may also allow simplification of the system's design.

The frequency of the oscillator is another important consideration, because the choice can have a significant impact on both the cost and the performance. Everything else being equal, an oscillator of standard frequency, such as 5 MHz or 10 MHz, will cost

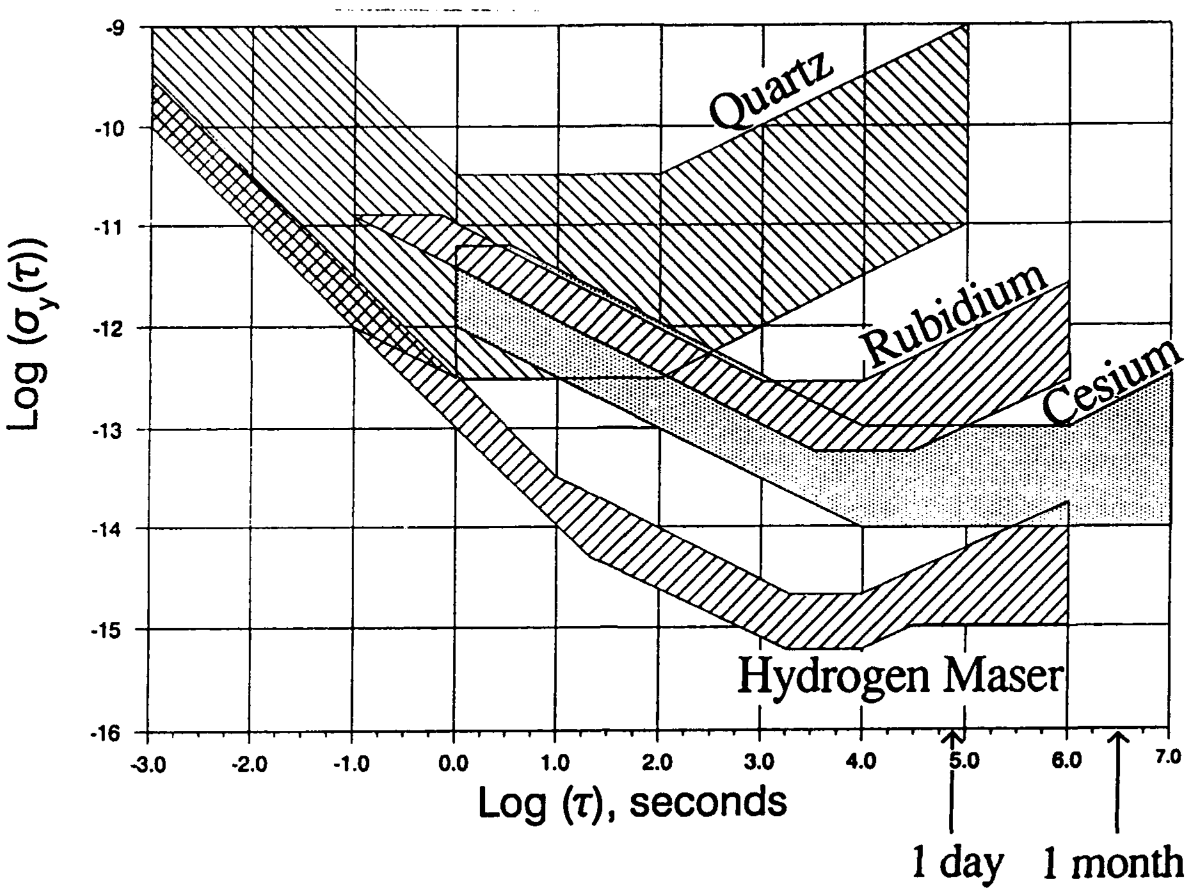

Figure 34. Comparison of frequency standards' stabilities as a function of averaging time.

less than one of an unusual frequency, e.g., 8.34289 MHz. Moreover, for thickness-shear crystals, such as the AT-cut and SC-cut, the lower the frequency, the lower the aging and the noise floor. Since at frequencies much below 5 MHz, thickness-shear crystals become too large for economical manufacturing, and since all the highest stability oscillators use thickness shear crystals, the highest stability commercially available oscillator's frequency is 5 MHz. Such oscillators will also have the lowest phase noise capability close to the carrier. There are also some excellent 10 MHz oscillators on the market; however, oscillators of much higher frequency than 10 MHz have significantly higher aging rates and phase noise levels close to the carrier than do 5 MHz oscillators. For lowest phase noise far from the carrier, where the signal to noise ratio determines the noise level, higher frequency crystals, e.g., 100 MHz, can provide lower noise because such crystals can tolerate higher drive levels, thereby allowing higher signal levels.

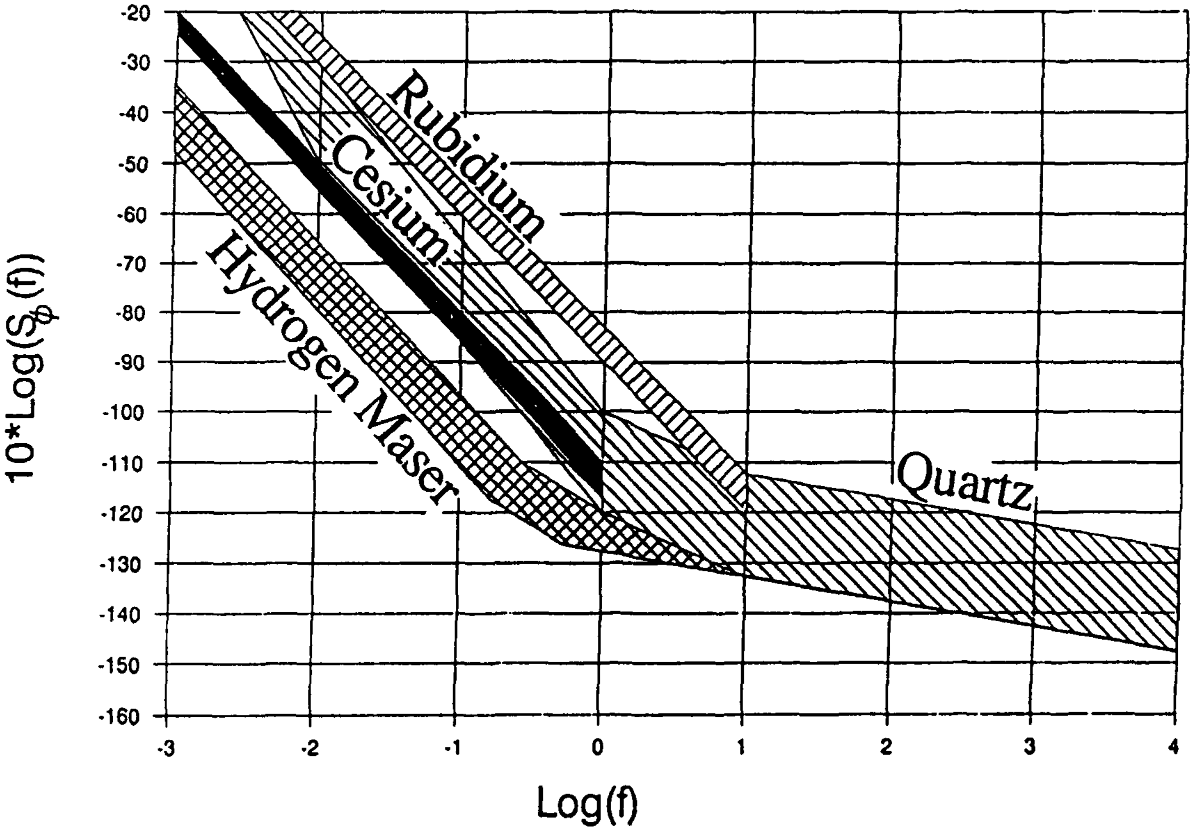

Figure 35. Comparison of frequency standards' phase instabilities.

V. Time Transfer

Time transfer techniques provide an additional method of maintaining synchronization among remote locations which complements the use of independent clocks. In fact, most systems derive time using both external time references and internal clocks. The former provide long-term accuracy and interoperability; the latter provide autonomous capability in the absence of the external references. A variety of time transfer techniques are in frequent use today: telephone, LORAN-C, GPS, WWV/VB, GOES, and two-way satellite [34]. They vary in capability from a few milliseconds to a few nanoseconds.

Radio broadcast services, such as WWV and WWVB in the United States, disseminate time with modest accuracy and stability. The high-frequency broadcasts between 2.5 and 20 MHz are usually received after reflection from the ionosphere. As a result, variability in the path delay limits the accuracy to a few milliseconds for most users. The broadcasts contain time 'ticks,' time code, and voice announcements referenced to the National Institute of Standards and Technology (NIST) time scale. Commercial WWV receivers are available.

	Weaknesses	Wearout Mechanisms
Quartz	Aging Rad hardness	None
Rubidium	Life Power Weight	Rubidium depletion Buffer gas depletion Glass contaminants
Cesium	Life Power Weight Cost Temp. range	Cesium supply depletion Spent cesium gettering Ion pump capacity Electron multiplier

Table 2. Comparison of frequency standards' weaknesses and wearout mechanisms.

Improved performance is provided by the GOES system, the timing signals of which are transmitted from the Geostationary Operational Environmental Satellites and may be received throughout most of the continental United States. Commercial time code receivers are available which provide timing accuracies of approximately 100 microseconds limited by uncertainties in satellite position and receiver delays. The NIST time scale is also the reference for the GOES timing signals.

The LORAN-C navigation system may be used to obtain time transfer accuracies of approximately one microsecond referenced to the United States Naval Observatory (USNO) time scale. These low-frequency broadcasts are propagated via ground wave, which is much more stable than the sky wave propagation of the HF broadcasts. Commercial timing receivers are available, which simplify time recovery from LORAN.

The Global Positioning System (GPS) also disseminates time referenced to the USNO time scale. Because of the precise knowledge of the satellite positions, the GPS Standard Positioning Service's C/A code is capable of disseminating time with approximately one hundred nanoseconds accuracy. Several commercial GPS timing receivers are available, which provide completely automatic operation. GPS may also be used in a differential mode, often called common view, to provide improved synchronization capability [35]. For sites located within several thousand kilometers of one another, timing errors due to errors in the ephemeris and the propagation delay are approximately equal. Thus, when the absolute GPS times of arrival of simultaneously observed satellite signals are subtracted from one another, the differential accuracy improves to several tens of nanoseconds.

The highest accuracy time synchronization is obtained via two-way satellite techniques [36]. Both the propagation errors and the delays through the receiver are calibrated by transmitting time in both directions between two sites. Each site measures the difference between the time of arrival of the pulse from the other site and the time of the local clock. The difference in the measurements made at the two ends provides the relative time of the two local clocks. The effects of the transmitter and receiver delays, the uplink and downlink propagation delays, and the delays through the satellite are substantially canceled. As a result, time synchronization accuracy of a few nanoseconds has been obtained using commercial communication satellites and very small aperture terminals (VSAT). A custom spread-spectrum time transfer modem is necessary.

Clocks and timing receivers can be combined in a timing system to provide a broader range of timing capabilities than either one can provide alone [37]. Such a system uses the received timing signal to calibrate the local clock, and learn its time, frequency, and frequency aging. When the timing signal is unavailable, the local clock acts as a 'flywheel.' Its free running operation starts using the time and frequency corrected by the calibration of the external source and the frequency, and may subsequently be periodically updated for the predicted frequency aging. This procedure produces the minimum possible free running timing errors. Commercial 'disciplined oscillators' now provide all these functions in an integrated package.

Relativistic effects become significant when nanosecond-level time transfer accuracies are desired, and when clocks are widely separated or have high velocities [34]. For example, at latitude 40°, a clock will gain 9.4 ns/day when it is moved from sea level to a 1 km elevation, and the clocks in GPS satellites (12 hr period circular orbits) gain 44 μs/day when compared to their rates on earth prior to launch.

VI. Specifications, Standards, Terms and Definitions

Numerous specifications and standards exist which relate to frequency standards. The major organizations responsible for these documents are: the Institute of Electrical and Electronics Engineers (IEEE), the International Electrotechnical Commission (IEC), the International Radio Consultative Committee (CCIR), and the U. S. Department of Defense, which maintains the Military Specification (MIL-SPEC) System. A listing of "Specifications and Standards Relating to Frequency Control" can be found in the back pages of the Proceedings of the Annual Symposium on Frequency Control. In the 1990 Proceedings [38], for example, 79 such documents are listed. Many of the documents include terms and definitions, some of which are inconsistent. Unfortunately, no single authoritative document exists for terms and definitions relating to frequency standards.

The terms and definitions in the CCIR glossary [21], in IEEE Std. 1139-1988 [25], and in MIL-O-55310's section 6 are the most recent (to 1990), they address different aspects of the field, and together form a fairly good set of terms and definitions for users of frequency standards.

The most comprehensive document dealing with the specification of frequency standards is "MIL-O-55310, Military Specification, Oscillators, Crystal, General Specification for" [27]. The evolution of this document over a period of many years has included periodic coordinations between the government agencies that purchase crystal oscillators and the suppliers of those oscillators. The document addresses the specifications of all the oscillator parameters discussed above, plus many others. This specification was written for crystal oscillators. Because the output frequencies of atomic frequency standards originate from crystal oscillators, and because no comparable document exists that addresses atomic standards specifically, MIL-O-55310 can also serve as a useful guide to specifying atomic standards.

MIL-STD-188-115, on "Interoperability and Performance Standards for Communications Timing and Synchronization Subsystems," specifies that the standard frequencies for nodal clocks shall be 1 MHz, 5 MHz, or 5 X 2^N MHz, where N is an integer. This standard also specifies a 1 pulse per second timing signal of amplitude 10 V, pulse width of 20 µs, rise time < 20 ns, fall time < 1 µs; and a 24 bit BCD time code that provides UTC time of day in hours, minutes and seconds, with provisions for an additional 12 bits for day of the year, and an additional four bits for describing the figure of merit (FOM) of the time signal. The FOMs range from BCD character 1 for better than 1 ns accuracy to BCD character 9 for "greater than 10 ms of fault."

VII. For Further Reading

Reference 8 contains a thorough bibliography on the subject of frequency standards, up to 1983. The principal forum for reporting progress in the field has been the Proceedings of the Annual Symposium on Frequency Control [34]. Other publications that deal with frequency standards include: the IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, the IEEE Transactions on Instrumentation and Measurement, the Proceedings of the Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting [39], and the Proceedings of the European Frequency and Time Forum [40]. Review articles can be found in special issues and publications [41-44].

VIII. References

[1] Kinsman, R., Gailus P., and Dworsky, L., "Communications System Frequency Control," elsewhere in this Encyclopedia.

[2] Abate, J. E., et al., "AT&T's New Approach to the Synchronization of Telecommunication Networks," IEEE Communications Magazine, pp. 35-45, April 1989.

[3] Pan, J., "Present and Future of the Synchronization in the US Telephone Network," IEEE Trans. Ultrasonics, Ferroelectrics, and Frequency Control, Vol. UFFC-34, pp. 629-638, November 1987.

[4] Dixon, R. C., Spread Spectrum Systems, John Wiley and Sons, NY 1976.

[5] Smith, W. L., "Precision Oscillators," in Precision Frequency Control, Vol. 2, (E. A. Gerber and A. Ballato, eds.) Academic Press, Inc., 1985, pp. 45-98.

[6] Parzen, B., Design of Crystal and Other Harmonic Oscillators, John Wiley and Sons, Inc., 1983.

[7] Bottom, V. E., Introduction to Quartz Crystal Unit Design, Van Nostrand Reinhold Co., 1982.

[8] Precision Frequency Control, (E. A. Gerber and A. Ballato, eds.) Academic Press, Inc., 1985.

[9] Hellwig, H. H., "Microwave Frequency and Time Standards," in Precision Frequency Control, Vol. 2, (E. A. Gerber and A. Ballato, eds.) Academic Press, Inc., 1985, pp. 113-176.

[10] Audoin, C., and Vanier, J., Journal of Physics E: Scientific Instruments, 9:697-720 (1976).

[11] Ramsey, N., Journal of Research, NBS, 88:301-320 (1983).

[12] Strumia, F., Minguzzi, P., Francesconi, M., and Benedetti, R., 28th Ann. Symp. Frequency Control, 350-354 (1974).

[13] Holloway, J. H., and Lacey, R. F., Proc. Int. Conf. Chronometry, 317-331 (1964).

[14] Wineland, D. J., Allan, D., Glaze, D. J., Hellwig, H. H., and Jarvis, S., Jr., IEEE Trans. Instrum. Meas. IM-25:453-458 (1976).

[15] Dicke, R. M., Phys. Rev. 89:472-473 (1953).

[16] Frerking, M. E., "Temperature Control and Compensation," in Precision Frequency Control, Vol. 2, (E. A. Gerber and A. Ballato, eds.) Academic Press, Inc., 1985, pp. 99-111.

[17] Schodowski, S. S., Filler, R. L., Messina, J. A., Rosati, V. J., and Vig, J. R., "Microcomputer Compensated Crystal Oscillator for Low Power Clocks," Proc. 21st Ann. Precise Time & Time Interval (PTTI) Applications & Planning Meeting, pp. 445-464, 1989. Proc. available from the U. S. Naval Observatory, Time Services Dept., 34th & Massachusetts Ave., NW, Washington, DC 20392. Details of the MCXO are also described in a series of five papers in the Proc. 43rd Ann. Symp. on Frequency Control, 1989, IEEE Catalog No. 89CH2690-6.

[18] Giordano, V., Hamel, A., Petit, P., Theobald, G., Dimarcq, N., Cerez, P., and Audoin, C., 43rd Ann. Symp. Freq. Control, 130-134 (1989).

[19] Wineland, D. J., Bergquist, J. C., Bollinger, J. J., Itano, W., Heinzen, D., Gilbert, S., Manney, C., and Weimer, C., 43rd Ann. Symp. Freq. Control, 143-150 (1989).

[20] Cutler, L. S., Giffard, R. P., Wheeler, P. J., and Winkler, G. M. R., 41st Ann. Symp. Freq. Control, 12-19 (1987).

[21] Recommendation No. 686, "Glossary," CCIR 17th Plenary Assembly, Volume VII, "Standard Frequencies and Time Signals (Study Group 7)," (1990) Consultative Committee on International Radio (CCIR); copies available from: International Telecommunications Union, General Secretariat - Sales Section, Place des Nations, CH-1211 Geneva, Switzerland.

[22] Walls, F. L. and Persson, K. B., 38th Ann. Symp. Freq. Control, 416-419 (1984).

[23] Peters, H. E., 38th Ann. Symp. Freq. Control, 420-427 (1984).

[24] Barnes, J. A., Chi, A. R., Cutler, L. S., Healey, D. J., Leeson, D. B., McGunigal, T. E., Mullen, J. A., Jr., Smith, W. L., Sydnor, R. L., Vessot, R. F. C., and Winkler, G. M. R., IEEE Trans. Instrum. Meas. IM-20:105-120 (1971).

[25] "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology," IEEE Std 1139-1988, The Institute of Electrical and Electronics Engineers, Inc., 345 East 47th Street, New York, NY 10017, USA.

[26] Kusters, J. A., and Vig, J. R., "Thermal Hysteresis in Quartz Resonators - A Review," Proc. 44th Ann. Symp. on Frequency Control, 165-175, 1990, IEEE Cat. No. 90CH2818-3.

[27] "MIL-O-55310, Military Specification, Oscillators, Crystal, General Specification for,"

the latest revision is available from Military Specifications and Standards, 700 Robbins Ave., Bldg. 4D, Philadelphia, PA 19111-5094.

[28] Filler, R. L., "The Acceleration Sensitivity of Quartz Crystal Oscillators: A Review," IEEE Trans. on Ultrasonics, Ferroelectrics and Frequency Control, Vol. 35, pp. 297-305, 1988.

[29] Filler, R. L., Keres, L. J., Snowden, T. M., and Vig, J. R., "Ceramic Flatpack Enclosed AT and SC-cut Resonators," Proc. 1980 IEEE Ultrasonic Symp., pp. 819-824, 1980, IEEE Cat. No. 80CH1602-2.

[30] Brendel, R., El Hasani, C., Brunet, M. and Robert, E., "Influence of Magnetic Field on Quartz Crystal Oscillators," Proc. 43rd Ann. Symp. on Frequency Control, pp. 268-274, 1989, IEEE Cat. No. 89CH2690-6.

[31] King, J. C. and Koehler, D. R., "Radiation Effects on Resonators," in Precision Frequency Control, Vol. 2, (E. A. Gerber and A. Ballato, eds.) Academic Press, Inc., 1985, pp. 147-159.

[32] Flanagan, T. M., Leadon, R. E., and Shannon, D. L., "Evaluation of Mechanisms for Low-Dose Frequency Shifts in Crystal Oscillators," Proc. 40th Ann. Symp. on Frequency Control, pp. 127-133, 1986, IEEE Cat. No. 86CH2330-9.

[33] Filler, R. L., Rosati, V. J., Schodowski, S. S. and Vig, J. R., "Specification and Measurement of the Frequency Versus Temperature Characteristics of Crystal Oscillators," Proc. 43rd Ann. Symp. on Frequency Control, pp. 253-256, 1989, IEEE Cat. No. 89CH2690-6.

[34] Allan, D. W., "Frequency and Time Coordination, Comparison, and Dissemination," in Precision Frequency Control, Vol. 2, (E. A. Gerber and A. Ballato, eds.) Academic Press, Inc., 1985, pp. 233-273.

[35] Allan, D. W., and Weiss, M. A., 34th Ann. Symp. Freq. Control, 334-346 (1980).

[36] Howe, D. A., 41st Ann. Symp. Freq. Control, 149-160 (1987).

[37] MacIntyre, A. and Stein, S. R., 40th Ann. Symp. Freq. Control, 465-469 (1986).

[38] The Proceedings of the Annual Symposium on Frequency Control have been published since the 10th Symposium in 1956. The earlier volumes are available from the National Technical Information Service; the later volumes, from the IEEE (see ref. 12 above). Ordering information for all the Proceedings can be found in the back of the latest volumes, e.g., the Proceedings of the 44th Annual Symposium on Frequency Control, 1990, is available from the IEEE, Cat. No. 90CH2818-3.

[39] The Proceedings of the Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting are available from the U. S. Naval Observatory, Time Services Department, 34th and Massachusetts Ave., N.W., Washington, DC 20392-5100. The latest volumes are also available from the National Technical Information Service.

[40] The Proceedings of the European Frequency and Time Forum are available from the Swiss Foundation for Research in Metrology (FSRM), Rue de l'Orangerie 8, CH-2000 Neuchatel, Switzerland.

[41] IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, Vol. UFFC-34, November 1987.

[42] IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, Vol. UFFC-35, May 1988.

[43] Frequency Stability: Fundamentals and Measurement, V. F. Kroupa, ed., IEEE Press, New York, 1983, IEEE Order No. PC01644.

[44] Characterization of Clocks and Oscillators, edited by D. B. Sullivan et. al, NIST Technical Note 1337, 1990, National Institute of Standards and Technology, Boulder, CO 80303-3328.