Frequency Standards and Clocks: A Tutorial Introduction

H. Hellwig

Time and Frequency Division
Institute for Basic Standards
National Bureau of Standards
Boulder, Colorado 80302

U.S. DEPARTMENT OF COMMERCE, Juanita M. Kreps, Secretary
Sidney Harman, Under Secretary
Jordan J. Baruch, Assistant Secretary for Science and Technology

NATIONAL BUREAU OF STANDARDS, Ernest Ambler, Acting Director

Issued June 1977

CONTENTS

LIST OF FIGURES

LIST OF TABLES

FREQUENCY STANDARDS AND CLOCKS: A TUTORIAL INTRODUCTION

The topic of frequency standards and clocks is treated in a tutorial and non-mathematical way. The concepts of time, frequency, frequency stability, and accuracy are introduced. The general physical principles and design features of frequency standards and clocks are described. The design, performance, and limitations of quartz crystal oscillators and atomic devices (cesium, hydrogen, rubidium) are discussed in detail and critically compared for laboratory devices as well as for devices intended for field usage.

Key words: Cesium beam; clocks (atomic); crystal oscillator; frequency accuracy; frequency stability; frequency standards; hydrogen maser; quartz crystal; rubidium gas cell; timekeeping.

1. INTRODUCTION

Frequency standards and clocks: what do they have in common? A more complete answer will be given later. We note for the moment that most clocks and in particular the very accurate and precise ones are based on frequency standards. The reason for this is the intimate relationship between frequency (symbol ν, "nu") and time (symbol t). If we look at a series of events which are occurring in a somewhat regular fashion, e.g., the rise of the sun every morning, we can state how many of these events occur in a given time period: this number would be the frequency of this series of events. In our example we could say that the frequency of sunrises is $\nu = 7$ events per week or $\nu = 365$ events per year. "Events per week" or "events per year" would be called the unit which we used for our frequency number; this frequency number is different for different units. In our example we assumed that we know somehow what a week or a year is, i.e., we relied on some external definition for our unit of time.

We can now ask, what is the time between the events? The answer for our example is simple, one sunrise succeeds the other after $t = \frac{1}{7}$ week or $t = \frac{1}{365}$ year where we used "week" and "year" as two possible choices for our unit of time. We can rely on the regularity and precision of the occurrence of the events, in other words, on their precise periodicity. Hence, we could define: the unit of time is one week; one week is the time elapsed at the seventh sunrise following an initial sunrise.

We learned two things: (a) For periodic events, the time between the events t is related to the frequency ν of their occurrence in the following simple way

$$\nu = \frac{1}{t} \tag{1}$$

and (b) that periodic events can be used to define time, i.e., the generator of the periodic events - the frequency standard - can be used as a clock. The frequency standard becomes a clock by the addition of a counting mechanism for the events.

In our example above, the frequency standard is the rotating earth. The time between recurring events is one day. This frequency standard served mankind for thousands of years and remained until very recently the source for the definition of time. The counting mechanism which made it a clock was the recording of years and days.

The needs to interpolate time from day to day, to get along for many days without celestial observations, and to more precisely measure time-intervals which are very much shorter than a day brought about the invention of clocks. Although there are other types of clocks like the sand-clock or the decay in radiation intensity of a radioactive substance, we shall confine ourselves to the discussion of clocks based on frequency standards (i.e., systems based on <u>periodic</u> events).

The first clocks based on a frequency standard (a pendulum) were invented about 400 years ago. This type of clock is most widely used today. The pendulum may be a suspended weight (gravitational pendulum) like in "grandfather" clocks or the balance (torsion pendulum) or quartz

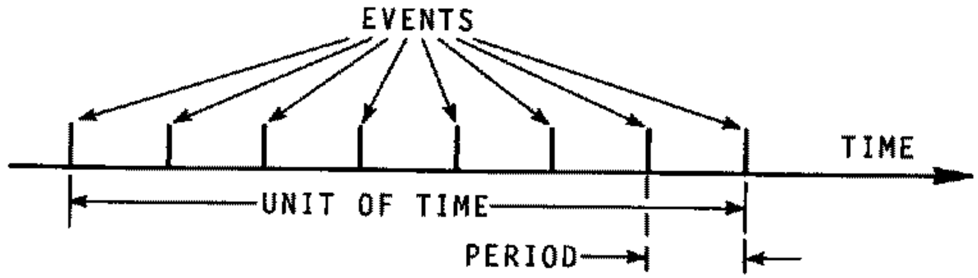

FREQUENCY = NUMBER OF EVENTS PER UNIT OF TIME

$$\text{ACCUMULATED CLOCK TIME} = \frac{\text{TOTAL NUMBER OF EVENTS}}{\text{NUMBER OF EVENTS PER UNIT OF TIME}}$$

POSSIBLE DEFINITION:

UNIT OF TIME = A SPECIFIC NUMBER OF PERIODS OF A WELL-DEFINED EVENT GENERATOR

Fig. 1. Definition of time and frequency.

crystal of modern wristwatches. The objects of our discussion in this report are today's most advanced frequency standards and clocks; however, a close look at our traditional clocks will show all the essential features which we will recognize again in our later discussion of quartz crystal and atomic clocks.

The pendulum in our clock is the frequency determining element. In order to arrive at a frequency standard the pendulum has to be set in motion and to be kept in motion: A source of energy is necessary together with means to transfer this energy to the frequency determining element. In a wristwatch this source of energy is typically the winding spring, and the energy is transferred by mechanical means which are controlled by the pendulum itself in order to cause energy transfer in the proper amount at the proper time in synchronism with the movement of the pendulum. This is called "feedback" in electronic systems. We now have a frequency standard; the tick-frequency of its pendulum could be picked up acoustically, for example, and used as a standard frequency. This is actually being done commercially when adjusting the rate of a clock: the tick-frequency is compared to some (better) standard frequency. In order to arrive at a clock, a read-out mechanism is necessary which counts and accumulates the ticks (more accurately: the time between the ticks) and displays the result. In our example of a wristwatch, this is accomplished by a suitably dimensioned set of gears and the moving hands on the clockface.

The unit of time today is the second (symbol s). Very much in analogy to our sunrise example, the second is defined in reference to a frequency determining element. Since 1967 by international agreement this "natural pendulum" is the cesium atom. One second is defined in the official wording as "the duration of 9192 631 770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the cesium-133 atom". Accordingly, the frequency of the cesium pendulum is 9192 631 770 events per second (the cesium atom is a very rapidly oscillating pendulum). Following our eq (1) the unit of frequency is then defined as hertz (symbol Hz). One hertz equals the

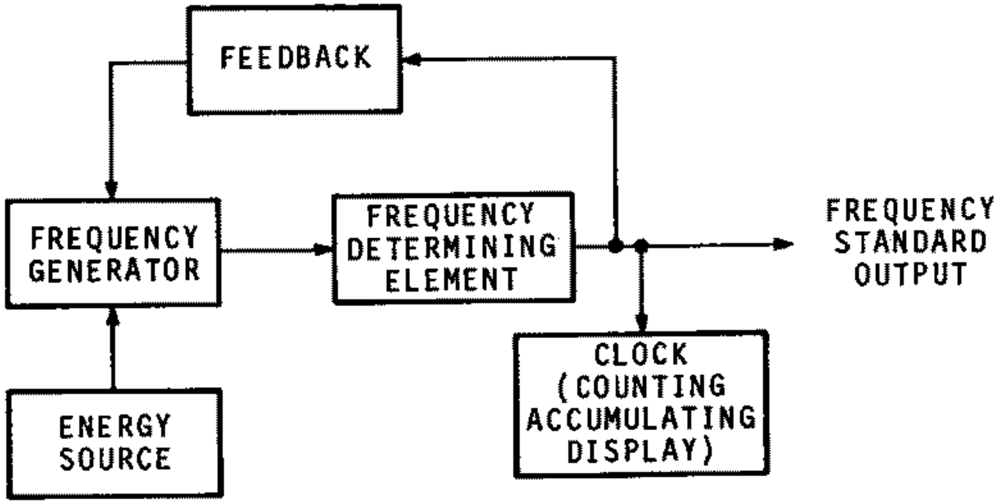

Fig. 2. Frequency standard and clock.

repetitive occurrence of one event per second. The use of "hertz" is preferred to the older term "cycle per second", cps, and was assigned in recognition of the exceptional contributions of the 19th century physicist Heinrich Hertz.

2. BASIC CONCEPTS

2.1 Accuracy, Reproducibility, Stability

The performance of frequency standards is usually described in terms of accuracy, reproducibility, and stability. We will use these terms in the sense of the following definitions:

Accuracy: the degree of conformity of a measured and/or calculated value to some specified value or definition.

Reproducibility: the degree of agreement across a set of independent devices of the same design after adjustment of appropriate specified parameters in each device. Alternately, it is the ability to reproduce, independently, a previous frequency value.

Stability: the frequency and/or time domain behavior[1] of a process. In the time domain (i.e., the measurement time or duration is the varied quantity) a frequently used measure of stability is the pair-variance (or Allan variance) (to be explained later) or its square root (the pair-stability).

It is obvious from these definitions that frequency accuracy will be largely of interest in scientific measurements and in the evaluation and intercomparison of the most advanced devices, but of little or no interest to the average user of frequency standards. A good reproducibility is an asset in applications where it is of importance to rely on some degree of conformity of the output frequency of several devices such as a factory-guaranteed frequency value. The characterization of the stability of a frequency standard is usually the most important information to the user. The frequency stability (symbol σ, "sigma") of a fre-

---

[1]In this paper, for tutorial purposes, only the time domain stability is used. We note, however, that for many scientific applications the frequency domain stability measure is more useful.

quency standard will depend on a variety of physical and electronic influences both internal and external to the device which cause frequency fluctuations. The frequency stability depends also on the exact measurement procedure which was used to measure the stability.

We shall explain this in the following: Frequency stability can be measured by taking a reasonably large number of successive readings of an electronic counter which counts the frequency of the device to be evaluated. Each counter reading (in hertz) is obtained by measuring or sampling the counted frequency for some specified time, the sampling time (symbol τ, "tau"). This sampling time can usually be chosen by simply adjusting a knob on the counter; for example, a sampling time of 0.1s or 1s or 10s may be chosen. Everyone has had the experience that fluctuations tend to average out if observed long enough; however, this is not always so. Sigma will therefore usually depend on the sampling time of the measurement and tends to get smaller with longer sampling times; again, there are many exceptions to this.

It may be that the fluctuations at some later time are partially caused by, or depend to some degree on, the previous fluctuations. In this case, the actual value of σ will also depend on the particular way in which the many counter readings are averaged and evaluated. Also, it will be of influence whether the counter starts counting again immediately after completion of the preceeding count or if some time elapses ("dead-time") before counting commences again.

Finally, electronic circuits will have a finite response time, e.g., they cannot follow fluctuations faster than some given rate. For example, our eye can not register light fluctuations which occur faster than about every $\frac{1}{10}$ of a second; using eq (1), we say that the eye has a frequency response of 10 Hz, or that its bandwidth is only 10 Hz, i.e., the eye can not follow frequencies higher than 10 Hz. In order to measure frequency stabilities for sampling times larger than some value , we have to provide for an electronic frequency bandwidth which is larger than about $\frac{1}{\tau}$.

We summarize: a recommended way of properly measuring and describing frequency stability is the following: (a) make sure that the frequency bandwidth of the total measuring set-up is larger than $\frac{1}{\tau}_{min}$ where τ_{min} is the smallest desired sampling time; (b) use a counter with a dead-time as small as possible[2]; (c) take a sufficiently large number of readings at a given sampling time which is held constant and compute[3]

$$\sigma = \sqrt{\frac{\text{addition of the squares of the differences between successive readings}}{2 \times \text{total number of differences used}}}$$

In the scientific literature this σ is called the square root of the pair (or Allan) variance; (d) repeat (c) at other sampling times τ and tabulate or plot σ as it depends on τ. See Appendix I (fig. 26).

Commonly, σ will be given as a fractional or normalized value, i.e., the value obtained for the frequency stability is divided by the carrier frequency. For example, if a frequency stability of $\sigma_{\delta\nu} = 10$ Hz were measured at a carrier frequency of $\nu = 5$ MHz (MHz = megahertz = million Hz) then the fractional frequency stability would be

$$\sigma_y = \frac{10}{5 \times 10^{-6}} = 2 \times 10^{-6}.$$

We denote the kind of σ by a subscript, $\delta\nu$ ("delta nu") referring to frequency fluctuations (measured in Hz) or $\frac{\delta\nu}{\nu} = y$ referring to fractional or normalized frequency fluctuations (dimensionless). A stability of one part in a million is thus

$$\sigma_y = 1 \times 10^{-6},$$

and one part in a trillion is written as

$$\sigma_y = 1 \times 10^{-12}.$$

---

[2]The dead-time should be less than the reciprocal bandwidth; if not, computation procedures exist to account for larger dead-times.

[3]The counter readings can be taken in Hz, σ will then have the dimension of Hz.

The common usage of the fractional frequency stability σ_y instead of the frequency stability $\sigma_{\delta\nu}$ (given in Hz) has its good reasons. In many applications of frequency standards their nominal output frequency is used after multiplication or division, i.e., the standard frequency is used to synthesize electronically other frequencies. Also, frequency standards themselves already synthesize several output frequencies, and different frequency standards offer different output frequencies. In such a usage of frequency standards as well as in the comparison of their performances it would be extremely inconvenient to state $\sigma_{\delta\nu}$ because this number (in Hz) would change with any alteration in frequency due to synthesis. This is not the case with σ_y because $\delta\nu$ as well as ν change in the same way in synthesizing procedures. Of course, we assume here that synthesis does not cause additional instabilities, an assumption which is mostly valid with today's electronic capabilities.

2.2 Time Accuracy

In clocks, time accuracy is of importance. In accordance with the definition for accuracy at the beginning of this section, the accuracy of a clock is the conformity of its reading (the date shown) with the reading of a standard time scale (a standard or reference date). A clock can be initially set (synchronized) with respect to a reference clock with some accuracy T_o. Its reading then gradually deviates by an amount $T(t)$ from the reading of the reference clock. $T(t)$ indicates that the deviation or error T is a function of the time t which elapsed after synchronization with the reference. This deviation, i.e., the accuracy of the clock, depends on the stability performance $\varepsilon(t)$ ("epsilon of t") of the frequency standard in the clock, the fractional frequency offset R_o of the clock[4] with respect to the reference clock, and on the change D of the fractional frequency of the clock with time (D = fractional frequency drift). Other parameters may be important; however, those mentioned are most commonly encountered. The time reading of the clock then has a degraded accuracy $T(t)$ which is given by

[4]The counter readings can be taken in Hz; σ will then have the dimension of Hz.

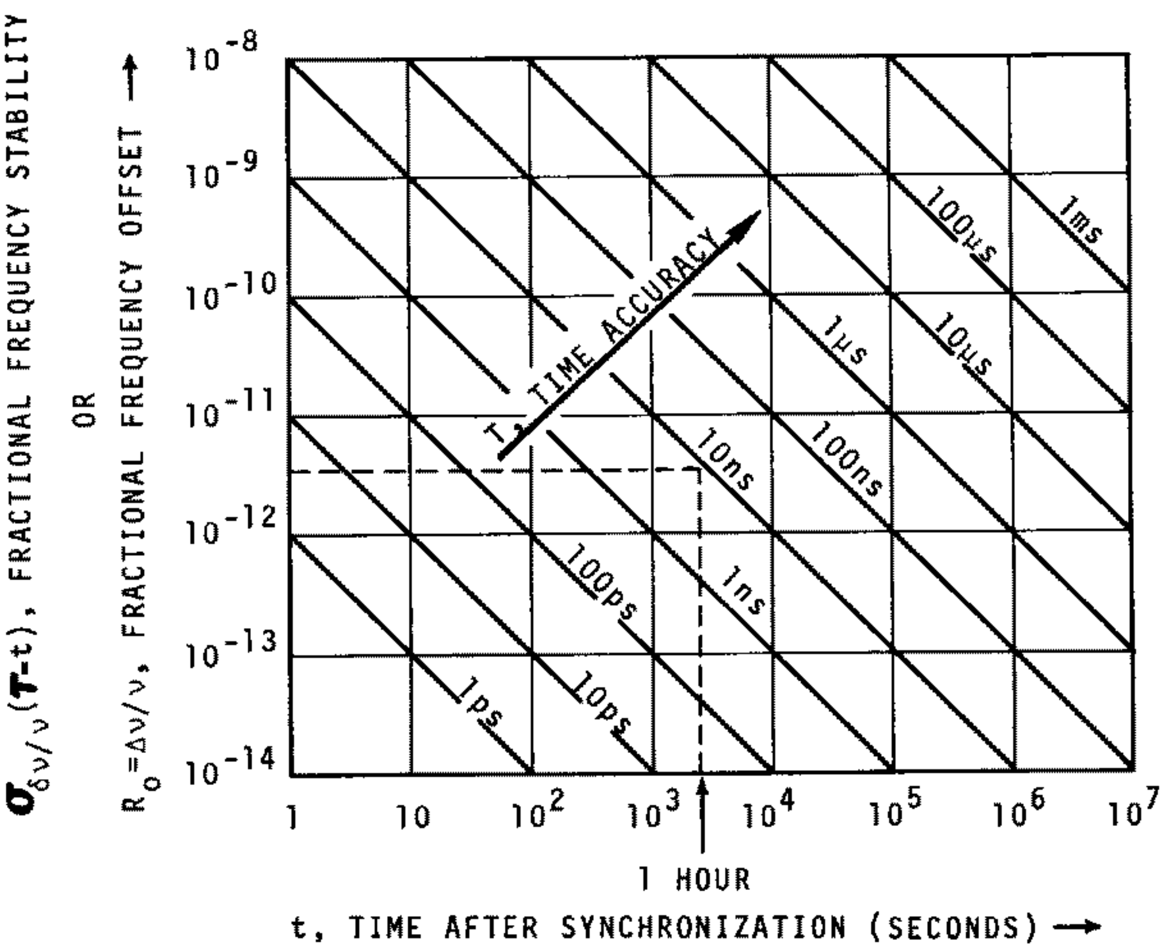

Fig. 3. Relationships between clock accuracy, frequency stability, and frequency offset.

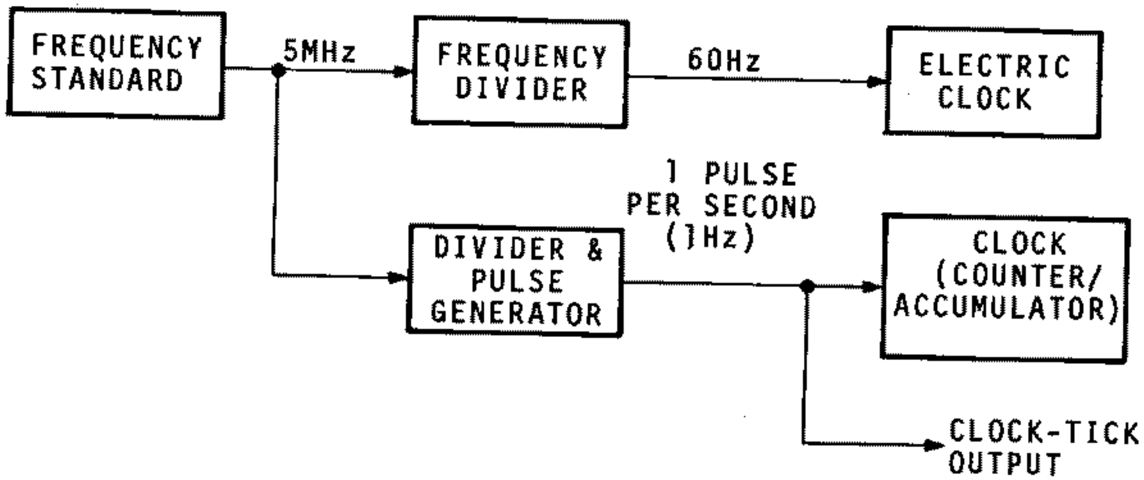

Fig. 4. Example of a clock system.

$$T(t) = T_o + R_o t + \frac{1}{2} Dt^2 + \varepsilon(t) \qquad (2)$$

The presence of a frequency drift D can rapidly deteriorate clock accuracy because the corresponding time error "accelerates" following the square of the elapsed time t. If T_o, R_o and D are not present (i.e., zero), only the random fluctuations can cause a time error $T(t) = \varepsilon(t)$. $\varepsilon(t)$ of course can be calculated from the frequency stability σ_y. This can be a rather complex mathematical problem and is often called "clock modeling" or "time prediction". For most practical uses, however, a very simple, quite good approximation may be used:

$$\varepsilon(t) \simeq t \times \sigma_y(\tau = t) \qquad (2a)$$

where $\sigma_y(\tau = t)$ means the fractional frequency stability evaluated for a sampling time τ equal to the time t which elapsed after the last setting of the clock. For example, a clock with a fractional frequency stability of $\sigma_y = 10^{-11}$ for sampling times of 10^5 seconds (about 1 day) would be accurate to $T = \pm 10^5 \times 10^{-11}$ second = ± 1 microsecond per day (again we assume T_o, R_o and D to be zero).

A clock with a fractional frequency offset of $R_o = \Delta\nu/\nu = 10^{-9}$ would be accurate to T = 100 microseconds after 1 day (assuming T_o, D and $\varepsilon(t)$ to be zero). And, finally, a clock with $D = 10^{-9}$ per day would be accurate to only $T = 5 \times 10^{-6}$ day $\approx$ 0.4 second after 100 days.

2.3 Clocks

In the Introduction we discussed that the addition to a frequency standard of a mechanism, which counts and accumulates and possibly displays the result, creates a clock.

This task can be performed by a frequency divider which, for example, derives a frequency of 60 Hz directly from a 5 MHz crystal oscillator. The 60 Hz voltage can be used to drive an electric clock similar to those driven by the 60 Hz power line frequency which we use at home or at work. Or, an additional electric pulse generator may be used which generates one very sharp electrical pulse per second. The time interval of 1 second between the pulses (corresponding to a frequency of 1 Hz) is directly derived from the output of our frequency standard. The pulses can be used in time comparisons with those of other, similar clocks; or a counter/accumulator/display (clock-face) can be driven by them.

3. FREQUENCY STANDARDS, GENERAL ASPECTS

3.1 Resonators

The performance of a frequency standard is to a considerable degree, but not exclusively, given by its frequency determining element. It determines the frequency by its resonance behavior. Some examples for resonance phenomena are (a) a rod, clamped only at one end, which can vibrate, (b) a block of solid material which can contract and expand and thus vibrate, (c) a capacitor-coil combination (tank circuit) in which the electro-magnetic energy can oscillate back and forth between these two elements, (d) an antenna (a dipole) where the distribution of electric charges can oscillate back and forth, and (e) a coil in which an electric current can create a magnetic field which can oscillate between its two possible polarities (a magnetic dipole). All these devices have in common that they can vibrate or oscillate if they are excited. The method of excitation may be a mechanical pulse for (a) and (b) or an electrical pulse for (c) or a sudden surge of an electric field for (d) or of a magnetic field for (e). The devices exhibit a resonance, i.e., they are resonators with a well-defined frequency which is characteristic of the physical dimensions of the device: the length of the rod, the thickness of the block, the size of the capacitor and coil, the length of the antennas. Once excited, the oscillations will die out gradually with a decay time which is determined by the losses of the resonator. Some of these losses are internal friction as in our cases (a) and (b), and electrical resistance for (c) through (e); in any case, the oscillation energy is ultimately transformed into heat. If there were no losses, the oscillations would never stop; we would have an ideal resonator. The more losses, the faster the oscillations decay and the resonance is less pronounced. It is now obvious that we could use the decay t_d of the oscillations to describe the quality of our resonator. The larger t_d the better the resonator.

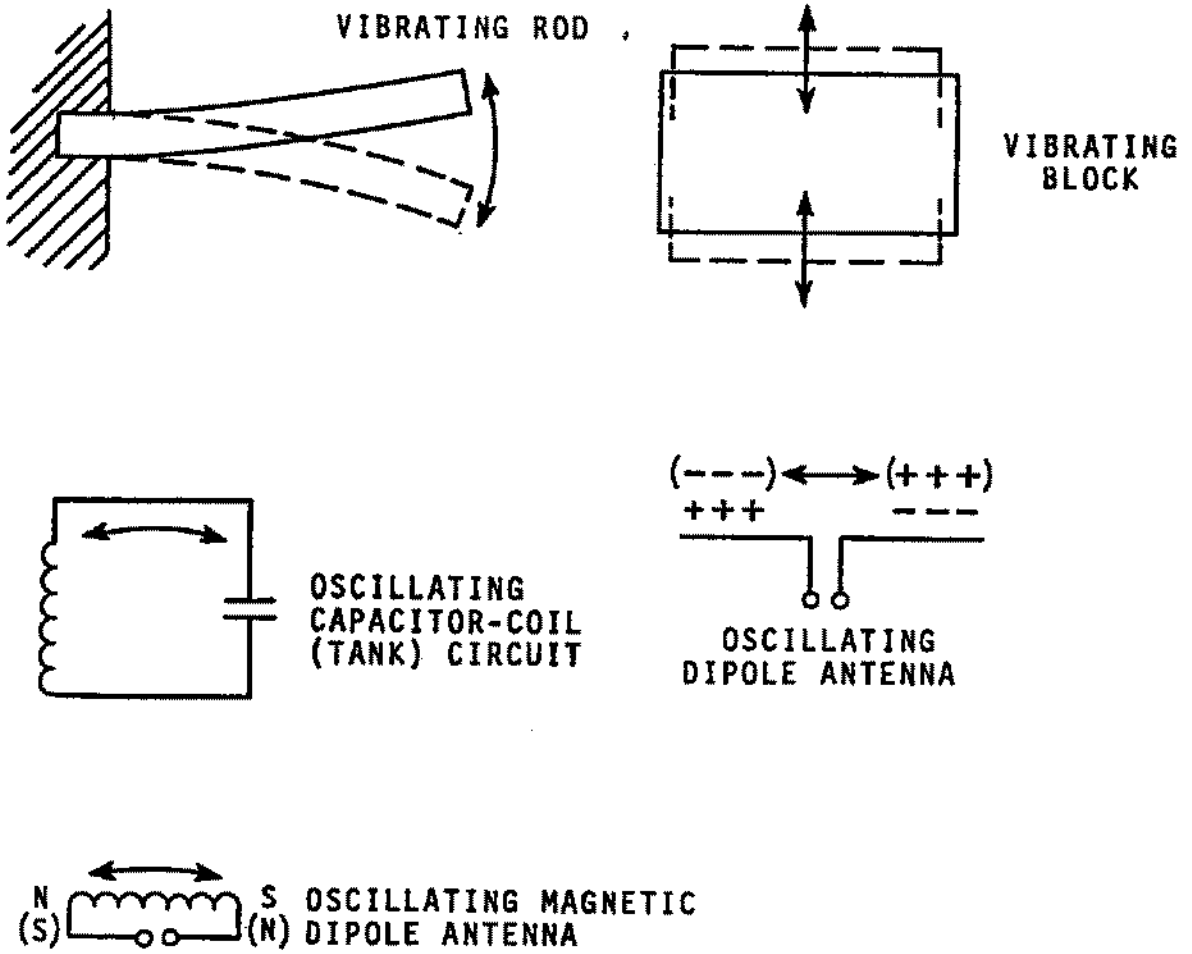

Fig. 5. Examples of resonators.

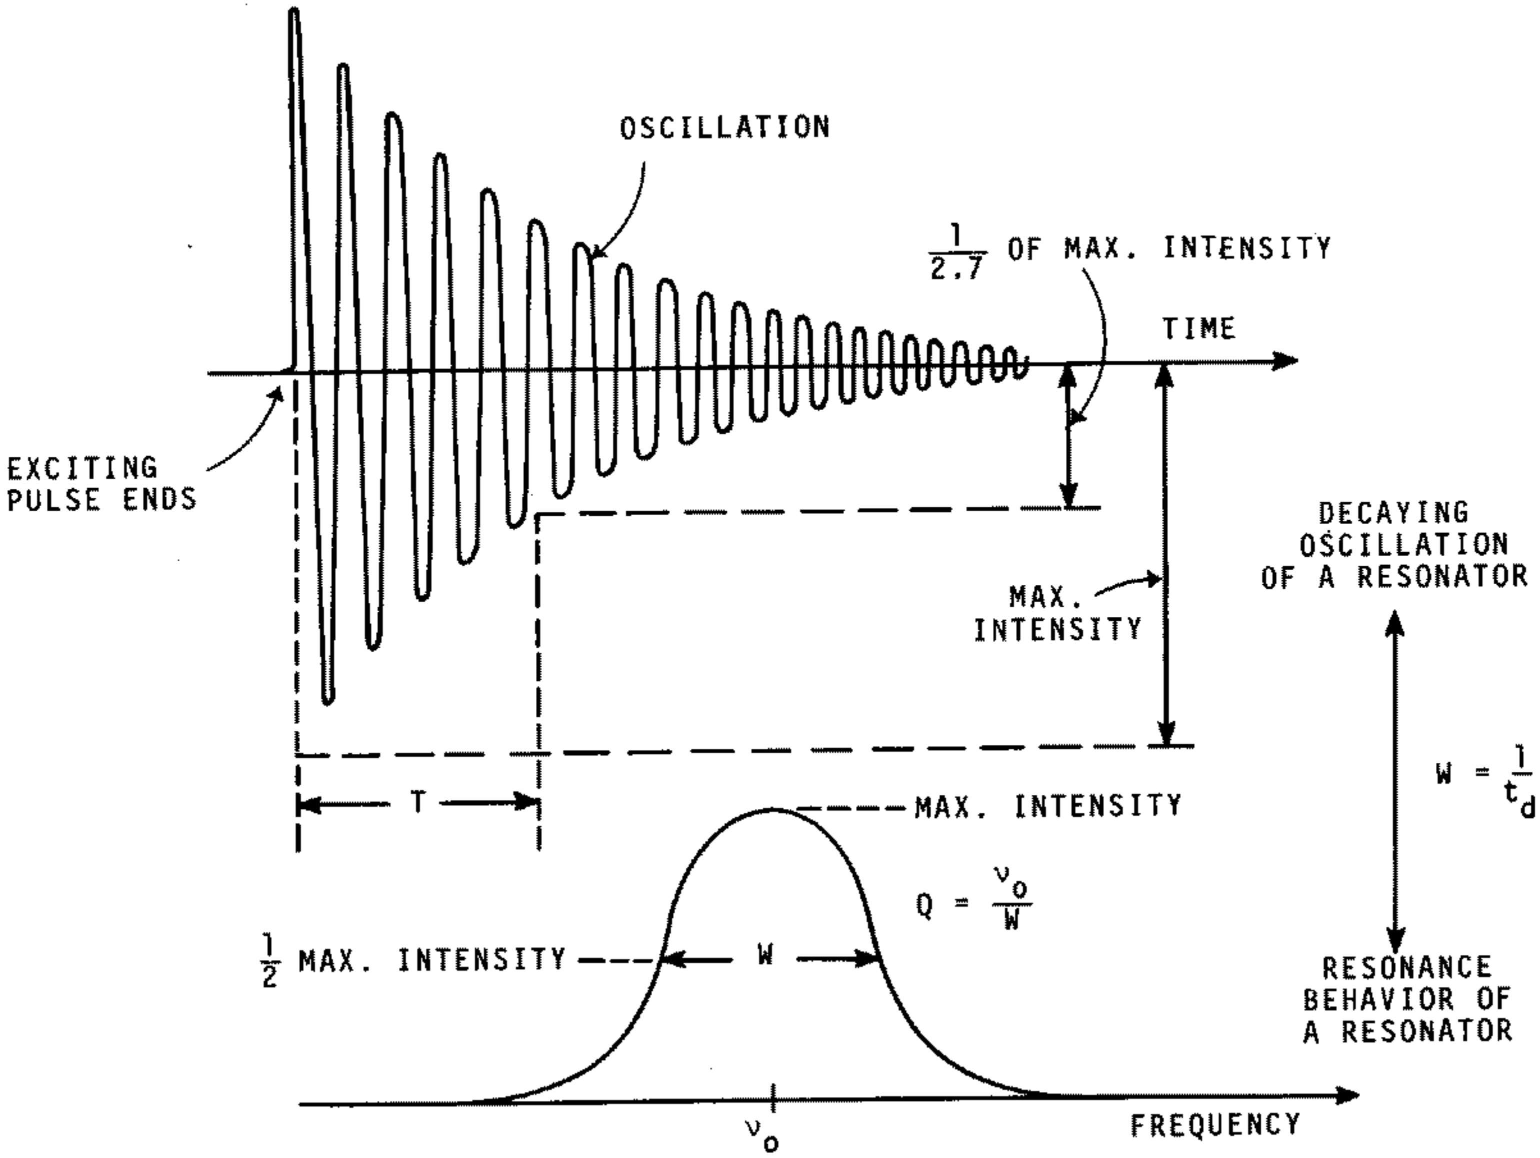

Fig. 6. Decay time, linewidth, and Q-value of a resonator.

An alternate way of measuring the resonance behavior is to use an external oscillator, to couple it to the resonator, and to sweep slowly the frequency of the external oscillator across the resonance. We will find again a resonance frequency at which the resonator will oscillate most pronounced, i.e., with the greatest intensity. On both sides of the resonance frequency, the response of the resonator will lessen until it ceases to respond. We can define a frequency interval around resonance in which the resonator response is relatively strong; we call it the resonance linewidth (symbol W). There is a simple relationship[5] between W and the decay time t_d

$$W = \frac{1}{t_d} . \qquad (3)$$

Again, as in the case of σ, it is of advantage to state the linewidth in a fractional way. The fractional linewidth would be $\frac{W}{\nu_o}$ where the symbol ν_o is used for the resonance frequency. More widely used is the quality factor of the resonance (symbol Q) which is defined as the reciprocal fractional linewidth

$$Q \equiv \frac{\nu_o}{W} . \qquad (4)$$

In frequency standards we obviously like to have large values for the Q of the frequency determining element. As an example, if $Q = 10^6$, then a fractional accuracy of 10^{-10} would imply that we can determine the center of our resonance curve to a small fraction (10^{-4} or one hundredth of a percent) of its width. In the same example, a stability of 10^{-12} for some sampling time would correspond to an ability to keep the frequency to within 10^{-6} (one millionth) of the resonance linewidth around a given value. It is clear, therefore, that the frequency stability and accuracy of a frequency standard may be expected to become the better the higher the Q-value of the frequency determining element.

Many kinds of frequency determining elements have been and are being used in frequency standards. They can be grouped into three classes:

---

[5]Equation (3) is only an approximation. Depending on the particular shape of the resonance curve we would have to insert a factor into this equation; however, this factor is never far from unity. t_d is defined as the decay to $\frac{1}{e}$ of the original amplitude. e is a number, $e \simeq 2.7$.

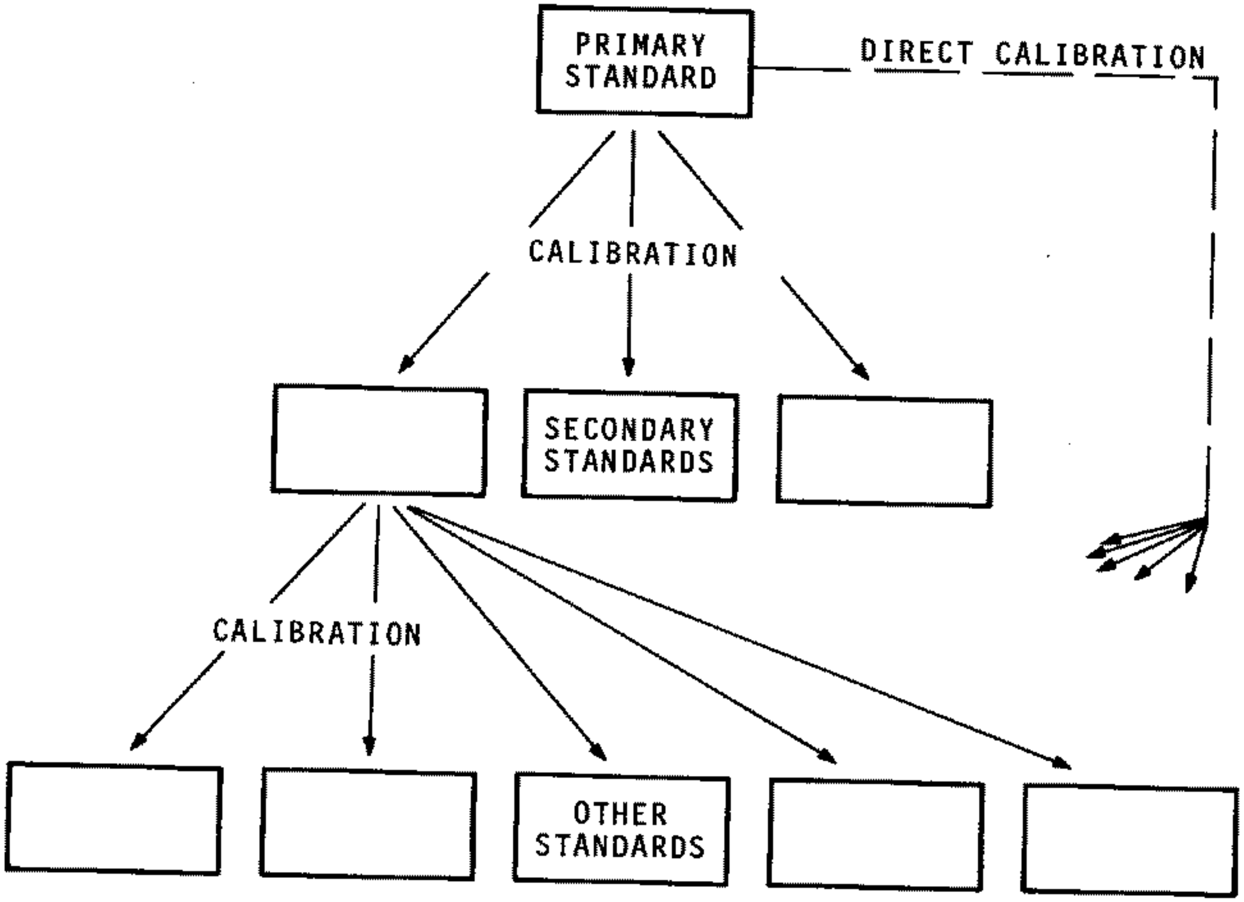

Fig. 7. Hierarchy of frequency standards.

Mechanical resonators;

electronic resonators;

atomic resonators.

As far as mechanical resonators are concerned we will only discuss one group in detail, the quartz crystals. Other mechanical resonators like the pendulum and the tuning fork are of no importance in today's high performance frequency standards although they have been historically very important and are still used in low performance devices (e.g., in watches). For similar reasons we will also omit the discussion of electronic resonators like the tank circuits (our device (c) of above) and microwave cavities.[6] Atomic resonators form the heart of our most accurate frequency standards and clocks and will, therefore, be extensively discussed.

3.2 Primary and Secondary Standards

At this point, we should briefly discuss the frequently used terms "primary frequency standard" and "secondary frequency standard". These terms should refer to the systems-use of the devices; any frequency standard, regardless of its accuracy or stability, can be a primary frequency standard, if it is used as the sole calibration reference for other frequency sources. A secondary frequency standard is a device which is occasionally calibrated against a primary frequency standard but operationally serves as the working reference for other frequency sources. The use of the terms "primary" and "secondary" to describe the performance and/or the design of a frequency standard itself is not helpful, often misleading and therefore discouraged. These aspects can be described adequately and accurately by stating accuracy, reproducibility, and design features.

One class of frequency standards can be separated from the rest. We call those the evaluable or "primary" frequency standards, which, by virtue of their basic physical operation as well as their design, allow the experimental evaluation of all known influences which might alter the output frequency. The accuracy of the evaluable frequency standard can thus be stated for the single device without reference to any other frequency

---

[6]We will touch on one very interesting example of a standard using a microwave cavity, the superconducting cavity oscillator (See Section 8)

standard and the term "primary" is then permissible because they are by virtue of their design the top of a calibration hierarchy. Some (but not all) cesium and hydrogen standards fall in this category.

4. QUARTZ CRYSTAL DEVICES

4.1 Quartz Crystal Resonator

The quartz crystal is a mechanical resonator much like our examples (a) and (b) of section 3.1 The resonator's oscillations have to be excited and sensed externally. In the case of a quartz crystal this is done by taking advantage of its piezoelectric properties. The piezoelectric effect is a special property of a certain class of crystals. Compression or dilatation of the crystal generates a voltage across the crystal, and conversely, the application of an external voltage across the crystal causes the crystal to expand or contract depending on the polarity of the voltage. A crystal is not a homogeneous medium but has certain preferred directions; thus, the piezoelectric effect has a directional dependence with respect to special, preferred directions of the crystal. In order to take advantage of the piezoelectric effect one has to cut the crystal resonator from the crystal block in a well defined way with respect to these crystallographic directions. The raw material today is both natural quartz and synthetic quartz. The crystal is cut out of the raw crystal in the desired orientation with the aid of optical techniques which allow the determination of the crystallographic axes. The high precision final orientation of the cut and the tuning to the desired frequency is then done by grinding and etching under control of x-ray methods.

The quartz crystal can be cut and electrically excited in a variety of ways. The most common types of vibrations (modes) are the longitudinal and thickness modes, the flexure (bending) mode, the torsional mode, and the shear mode. In order to use the piezoelectric effect, electrodes have to be put on some of the crystal surfaces such that the desired mode is excited. The electrodes are typically created as extremely thin metallic coatings by vacuum evaporation of metals. Electric leads are attached to the electrodes, (e.g., by soldering), which usually also serve as the mounting support thus suspending the quartz crystal. In order to least perturb the mechanical vibrations of the crystal, the electrode-support leads are attached at points where no vibrational motion occurs (modes). The crystal is usually encased (in metal or glass) and the enclosure is sometimes filled with a protective gas or is evacuated.

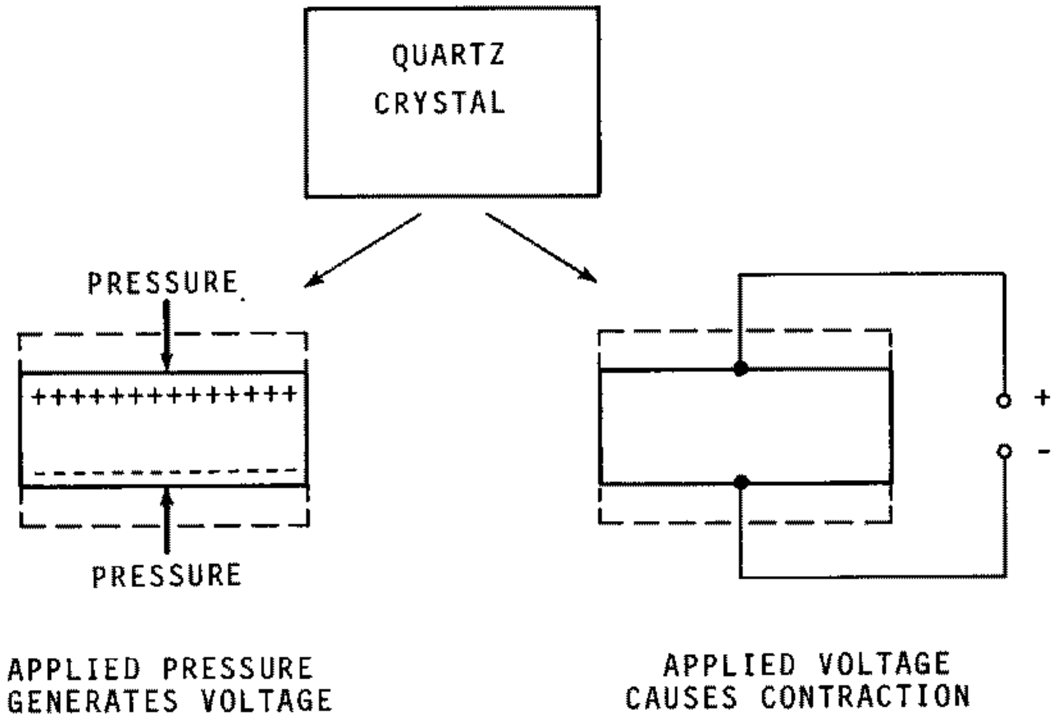

Fig. 8. The piezoelectric effect.

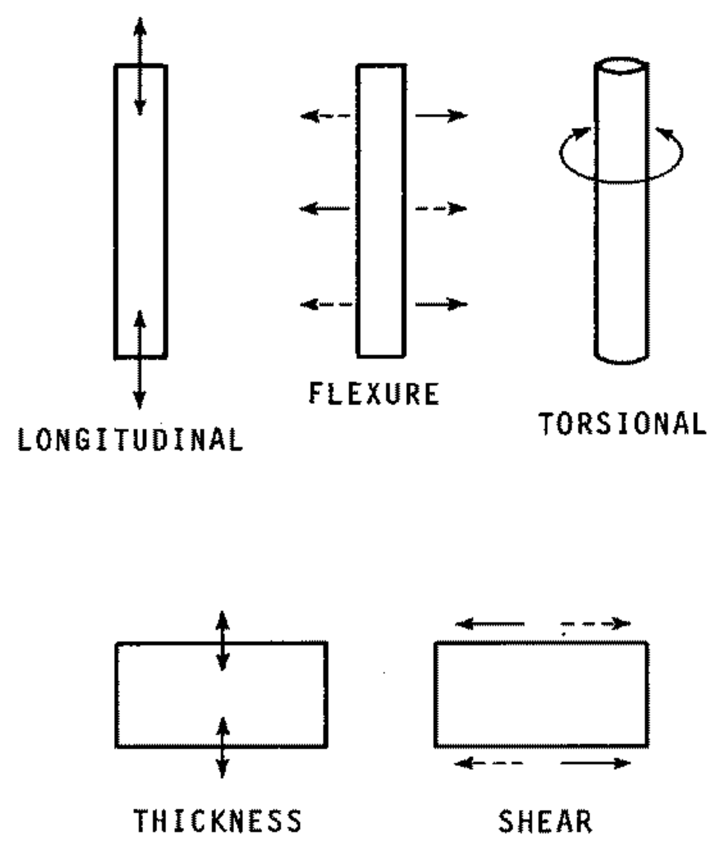

Fig. 9. Principal vibrational modes of quartz crystals.

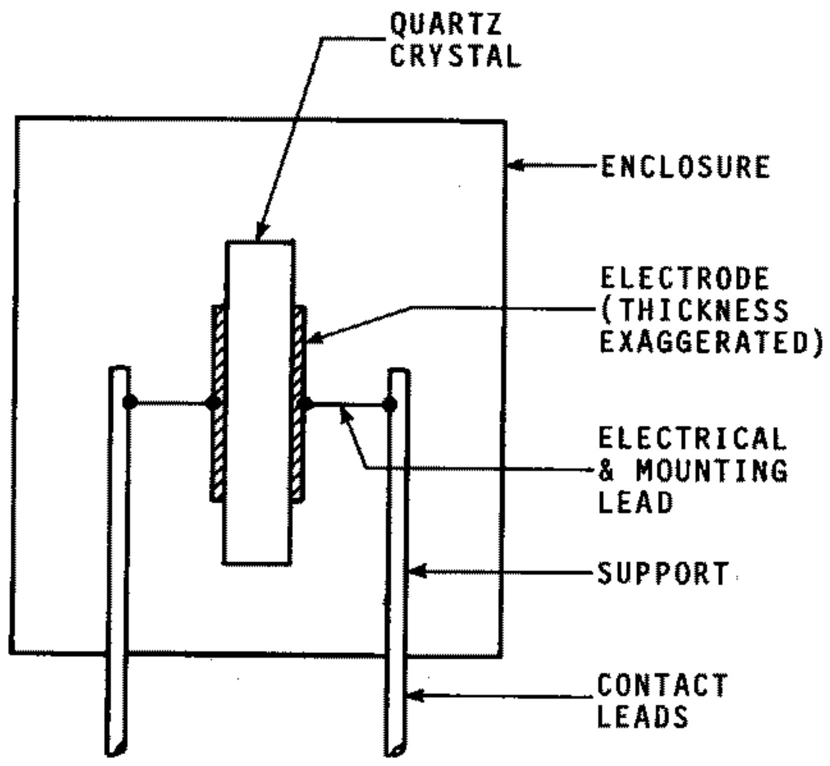

Fig. 10. Typical quartz crystal mount.

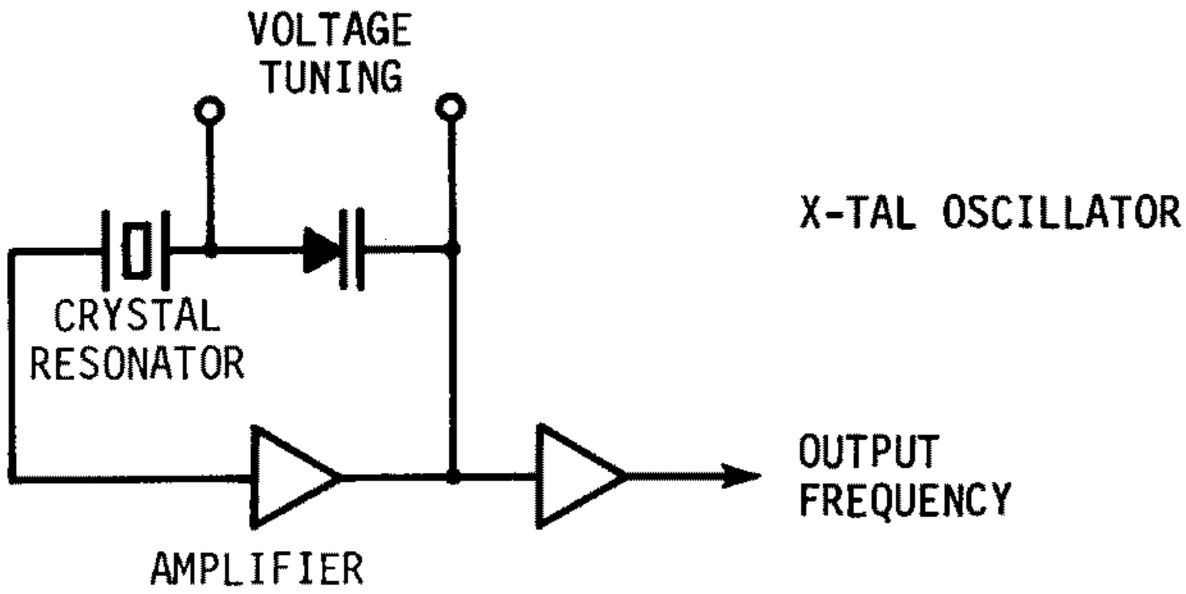

Fig. 11. Block diagram of a quartz crystal oscillator.

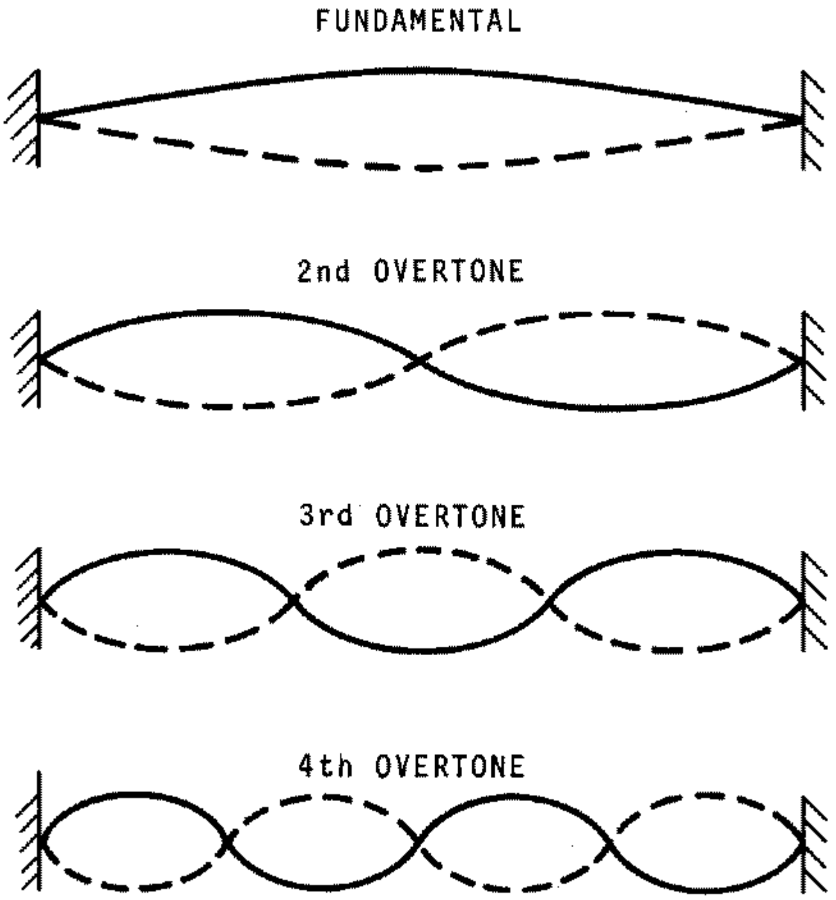

Fig. 12. Fundamental and overtone resonance frequencies.

A frequency standard can now be built by adding an electronic amplifier (energy transfer), and feedback. We call such a device a quartz crystal oscillator. Its output frequency is determined by the quartz crystal resonator whose frequency in turn is determined by the physical dimensions of the crystal together with the properties of crystalline quartz. The resonance frequency thus depends on the orientation of the cut, the particular mode, and the dimensions of the crystal. As an example, we find that the resonance frequency for a longitudinal mode of vibration is approximately given by

$$\nu_o = 2.7 \times 10^3 \times \frac{1}{\ell} , \quad (5)$$

where ℓ is the length of the crystal. Equation (5) is written such that the use of meters to express ℓ will give the resonance frequency ν_o in hertz. Equation (5) allows us to estimate the size of crystals. For example, a 100 kHz crystal has a length which is of the order of a few centimeters and a 10 MHz crystal has a length (thickness) of just a few tenths of a millimeter. We see from this that the production of quartz crystals with resonance frequencies much above 10 MHz is hardly possible. However, one can excite resonators not only in their so-called fundamental mode (which we discussed so far) but also at multiples (overtones) of this fundamental resonance frequency. The best example for this is the violin string which also can be made to oscillate at frequencies which are multiples of the fundamental frequency; the violinist depends on this. Quartz crystals which are designated for the excitation of multiples of their fundamental resonance are called overtone crystals.

4.2 Effects on the Crystal Resonator Frequency

Two deleterious effects, among others, are important in the design of crystals and crystal oscillators and limit their usefulness. The first is the temperature dependence of the quartz crystal resonance frequency, the second is a slow change of the resonance frequency as time goes on (frequency drift or aging).

The temperature dependence is caused by a slight change in the elastic properties of the crystal with temperature. This can easily be imagined from the general behavior of matter: the packing density of atoms increases with lowering the temperature. However, certain cuts, i.e., certain crystallographic orientations of the crystal, minimize this effect over a rather wide range of temperatures, most notably the so-called "AT" and "GT" cuts. Temperature coefficients of less than one part in 100 million per degree (celsius or kelvin) are possible. In other words, the fractional frequency change is less than 10^{-8} with one degree of temperature change. Nevertheless, this effect demands certain precautions in the design of a crystal oscillator if very high frequency stabilities over longer times (hours or days) are desired and/or if large environmental temperature fluctuations are to be tolerated. Hence, crystals are enclosed in electronically regulated ovens which maintain a constant temperature; in certain crystal oscillators this is done to better than $\frac{1}{1000}$ of a degree.

A different solution to the temperature problem is the so-called temperature compensated crystal oscillator or TCXO. An additional frequency determining element in the oscillator, which can essentially be just a small capacitor (see fig. 11 and our example (c) of sec. 3.1), gives the opportunity to tune the oscillator over a limited range by varying this additional element. If a temperature sensor is added which causes a change in this capacitor one can adjust the response in such a way that the change in resonance frequency of the crystal resonator is just compensated by a suitable adjustment of the added capacitor. Bimetal springs have been used which mechanically change the setting of the capacitor much in the same way as the automatic choke in many automobiles works. Today, capacitors whose value changes with an applied voltage (varactors) are used, as shown in figure 11; the applied voltage is derived from a temperature sensing circuit. The TCXO thus does not necessarily require further temperature control by an oven. However, we see the drawback of this approach. In adding a further frequency determining element, the crystal resonator has to relinquish a corresponding part of its control on the output frequency of the whole oscillator. We, therefore, realize that the stability performance of a TCXO will degrade the more, the wider the temperature range of compensation is made.[7] The long-term stability (days) of TCXO's is therefore below that of crystals with a good oven control. We find TCXO's in small, usually portable units of relatively low performance; i.e., for applications where frequency stabilities from day to day and frequency changes over some tens of degrees of temperature of not better than 10^{-9} are needed.

The drift, or aging, is a common behavior of all crystal oscillators. It is a nearly linear (uniform) change in resonance frequency with time, which frequently is negative (i.e., the resonance frequency decreases). A frequency decrease could be interpreted as an increase in the crystal size according to eq (5). Many physical mechanisms have been considered as the cause: contamination of the surfaces (deposition of foreign material); changes in the electrodes or the metallic plating or the mounting; reformation of loose (from grinding and etching) surface material; changes in the internal crystal structure; etc.; all of this possibly caused or enhanced by the vibrating motion of the oscillating crystal. Careful fabrication and electrode design combined with clean vacuum enclosures have led over the years to a reduction of the aging to about 10^{-11} per day and better for the best crystals. This aging corresponds to 10^{-11} fractional thickness change per day, as we can see from eq (5). For a 5 MHz crystal with a thickness of a little less than a millimeter this aging corresponds to an absolute thickness change of only 10^{-11} of a millimeter or less than $\frac{1}{10}$ of one percent of the diameter of an atom. It seems surprising that mechanical resonators can be built which change their dimension by so little.

Three more effects on crystal resonators are to be considered. One is its relative sensitivity to gravitational forces and acceleration:

[7]This is also true for any tuning (frequency adjustment) capability which is added to any frequency standard. The larger the tuning range, the less the stability. The best frequency standards and clocks can not directly be adjusted; however, adjustment can be done external to the device via frequency synthesizers or time steppers (phase shifters).

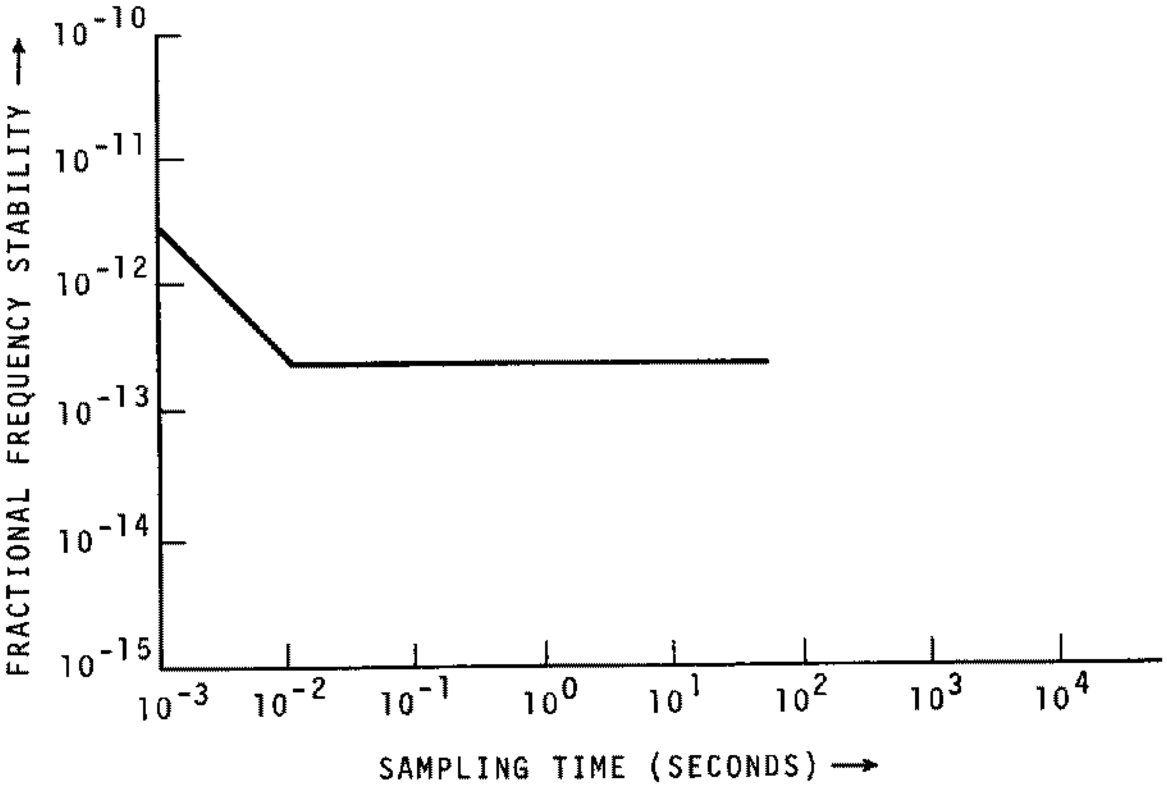

Fig. 13. Frequency stability of the better quartz crystal oscillators.

frequency changes will occur because of the stresses in the crystal caused by these forces. This influence depends on the direction of the force relative to the crystallographic axes and thus can be minimized for certain orientations. The magnitude of the effect is typically of the order of 10^{-9} for accelerations corresponding to the earth's gravitation.

The second effect is related to intermittent operation. If a crystal oscillator is turned off and, after some time, put back into operation it will not oscillate immediately at the original frequency but will exhibit first a "warm-up" due to temperature stabilization of the crystal resonator and its oven and then for some time (as long as many days) a large but diminishing drift until it reaches its previous aging performance. The frequency at which it will then operate might also be substantially different (as compared to its stability and aging performance) from its frequency before the interruption. The third effect is the sensitivity of the crystal to ionizing radiation (x-rays, nuclear radiation): transitory and lasting frequency changes are produced via radiation-induced changes in the crystal lattice.

4.3 Quartz Crystal Oscillator Performance

Crystal resonators have Q-values which are typically in the range from $Q = 10^5$ to almost $Q = 10^7$ (around room temperature; at very low temperatures much higher Q values can be obtained). These are very high Q-values as compared to most other resonators except, most notably, atomic resonators.

These high Q-values are an essential prerequisite for the excellent stability performance of crystal oscillators. The best presently available devices show stabilities of about one part in 10^{13} for sampling times from one second to hundreds of seconds. There is some experimental evidence that some crystal resonators may perform better and that limitations are primarily caused by noise in the electronic components in the oscillator circuits. This noise (flicker noise) may possibly be reduced by a special selection of low noise components (transistors, capacitors, etc.) and by some special circuit design. Thus there is a reasonable chance that the stability may reach values of better than 10^{-13} for sampling times of seconds to hours. For times shorter than one second, the stability is often determined by additive noise in the output amplifiers and can then be reduced by other better electronic circuits or by a (crystal) filter in the output. The long-term stability beyond several hours sampling time is determined by the aging and by external influences such as line voltage variations, temperature fluctuations, etc. The immediate environment of the resonator, i.e., the mounting, enclosure, etc., probably are of substantial influence in this regard and improvements appear possible. A specification of accuracy is not very well possible with crystal oscillators. Without any frequency calibration they possibly can be fabricated via thickness-determination (not very practical) to about 10^{-6} frequency accuracy. If they are calibrated against a high accuracy frequency standard they maintain this calibration (accuracy) according to their long-term stability performance, e.g., a crystal with the low aging of 10^{-11} per day will be accurate to a few parts in 10^{8} for the duration of a year.

We realize, therefore, that crystal oscillators require calibration, which may be rather frequent depending on the requirements. The frequency adjustments are being made with a small added capacitor in much the same way as we discussed before in connection with the TCXO. The most stable ones with the lowest aging rate cost more than $1000, have a volume of a few hundred cubic centimeters, and require a power of several watts. They have a crystal oven for temperature control, well-designed electronics, and usually several output frequencies which are derived from the oscillator frequency with the aid of frequency dividers and multipliers. These high performance devices presently use crystal resonators in the 1 to 10 MHz range with Q-values of more than one million.

Crystal oscillators which are cheaper and/or smaller are available in an immense variety of designs at certain sacrifices in frequency stability and/or environmental insensitivity. Costs can go down to below $100, sizes to a few cubic centimeters, and power requirements to less than 0.1 watt. The reliability of crystal oscillators is usually not limited by the crystal, the mean time between failure (MTBF) being that of any electronic circuit of equivalent sophistication.