Precise Time and Time Interval
Clocks, Time Frames and Frequency

James R. Clynch
Department of Oceanography
Naval Postgraduate School
February 2002, January 2003

I. Introduction

In the Naval Postgraduate School course on Mapping, Charting, and Geodesy (OC 3902) an overview of precise time and time interval is given. This note is a summary of that material. There is additional material, with an emphasis on GPS and the GPS time frame at the web site www.nps.navy.mil/~jclynch .

Both the civilian and military world make heavy use of precision clocks and accurate time. The US official timekeeper for the civilian world is the National Institutes of Science and Technology (NIST) (www.boulder.nist.gov/timefreq). For the entire US military world the US Naval Observatory (USNO) (www.usno.navy.mil) is the precision time keeper. One of the USNO tasks is to provide the long-term time frame for the Global Positioning Satellites (GPS). Because GPS provides time extremely accurate time inexpensively to any user, in effect USNO is the time provider to a large fraction of the world.

The US has become dependent on the presence of precise time from GPS. If GPS went away, the largest immediate impact on both the civilian and military worlds would be caused by the lack of accurate time, not positional information. Communications networks - cell telephones as well as some military systems - would quickly have parts fall out of synchronization and fail. The power grid is also dependent on GPS to synchronize generators and sub-grids. And the uses of GPS to provide accurate time are growing.

This note will discuss what constitutes a clock, give a few examples, and then discuss the way we characterize the quality of clocks. It is intended to be a midlevel summary, but should be useful to persons without any background in physics or engineering.

II. Clocks - What and How

Basically a clock consists of four items: something that generates events at a regular interval (the oscillator), a counting mechanism, some method to calibrate the rate of the events, and a time setting mechanism. In a fundamental sense, we do not tell time, but only count events. On a low level, we deal in time intervals, not time. The zero of our time system is arbitrary and set by convention – or some committee. This origin, along with a definition of the second, defines a time frame.

The rate that the events occur must be calibrated. This means that there must be standards, and the time from these must get to the user, at least at some point. This is often done at the factory, but very accurate clocks need periodic re-calibration. Both calibration and setting fall under the domain of time dissemination. In days gone by, clocks to be set or calibrated were physically brought to a time laboratory or calibration facility. Now GPS is used for this function in most cases.

Errors in the time a clock provides come from many sources. The usual limiting feature is the failure of a clock to be perfectly regular in it's cycle. These errors in clock's time are normally a function of what time span is of interest. The error characteristic over a short interval is quite different from the form for hundreds of thousands, millions or larger number of events. In addition the period of the clock, the time interval between events, often depends on environmental conditions. Things like temperature, air pressure, the local gravity acceleration and even relative humidity have been shown to affect some clocks. The environmental effects often lead to systematic errors. Then the time error grows large because there is the same error in each event count.

Depending on the accuracy required many things can serve as clocks. The calendar, which counts earth revolutions with respect to the sun, is a clock. For very precise applications the oscillations of atoms in very specific states is used. In between there are many different possibilities.

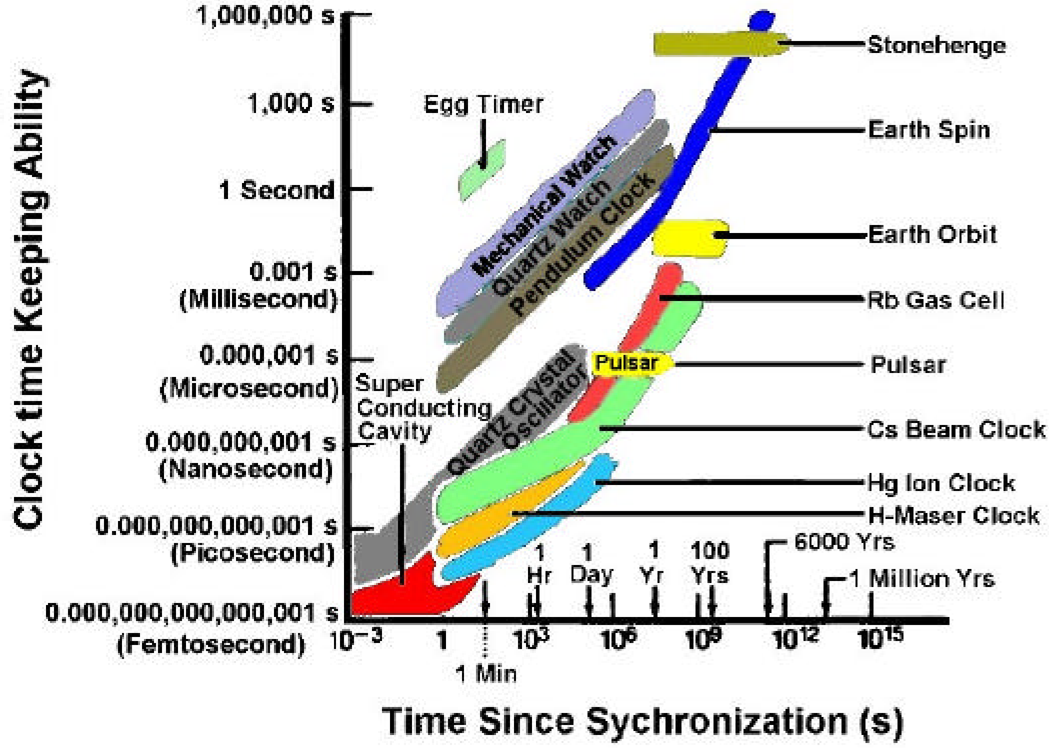

Figure 1

In Figure 1, derived from a plot supplied by David Allan, the expected time error at some time after a clock has been set is shown for a wide variety of "clocks". The widths of the bands are intended to cover "good" examples of each type. In some cases the width may indicate our uncertainty of how things behave at long time intervals. Notice that the error grows with time, and sometimes has bends in the curve.

The scales on both axes are logarithmic. Logarithmic axes are common in the field of time and clocks. Accuracy's of microseconds or better over time scales ranging from seconds to centuries are involved. In Figure 1, eighteen orders of magnitude (powers of 10) are covered on the time spans of the x-axis. The error, on the vertical axis, covers 9 orders of magnitude. The lower left corner is the province of experimental physicists. The upper right is for astronomers and astrophysics. The rest of us fit in-between. However we use milliseconds and microseconds more than we think. Some automatic teller machines will quit giving cash if they get a few milliseconds off and cannot communicate securely with the bank.

Notice that the spinning earth is somewhat irregular. Over a year the error can be seconds and over a human lifespan can grow to minutes. However to do better (get lower on the graph) we need to use star measurements made throughout out the year (the earth's orbit curve limits this) or some form of atomic clock.

III Examples of Clocks

A. The Earth as a Clock

The first clock man used was probably the spinning earth. The spin is irregular because the moment of inertial changes. This is a function of the mass distribution of the earth. Small changes cause small spin rate variations that add up over time. Occasionally changes come quickly in earthquakes. There are also slower changes due to motions of the fluid within the earth. Recently the distribution of water - how much is locked in glaciers, lakes and reservoirs above sea level, is thought to have made measurable changes.

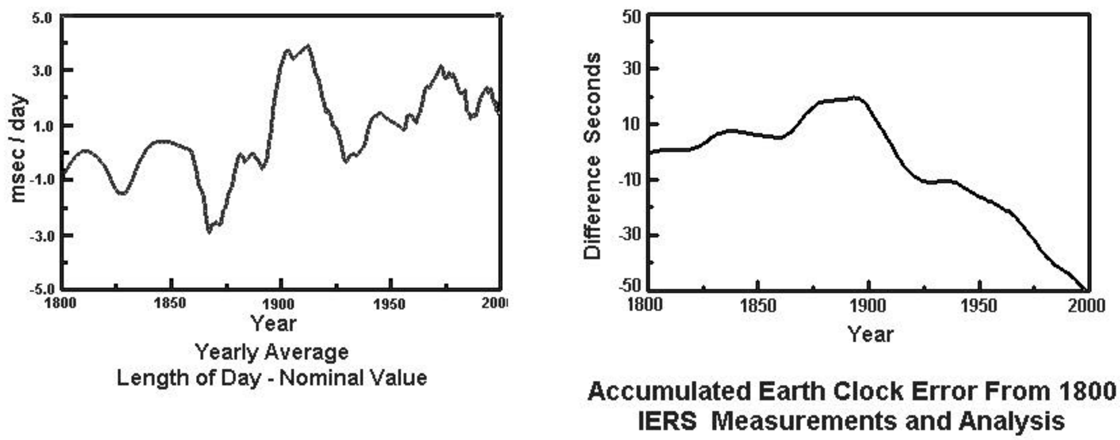

Figure 2

The variation in the length of the day over the past two hundred years is shown in Figure 2. (IERS is the International Earth Rotation Service, the successor to the Bureau International de l'Heure or BIH in Paris.) These data come mainly from astronomical records. The values are a few milliseconds (1/1000 th of a second) per day. The zero of this error curve is fixed by the definition of the second. This has evolved over the years as discussed below.

There are about 100,000 seconds in a day. The relative error in the spin rate is therefore about a part in 100 million. We say this is 10 ppb or 10 parts per billion. However it only takes 3 years to give 1000 days. Therefore the time error accumulations quickly adds up to levels easily observed by humans. The accumulated time errors are also shown in Figure 2.

The rate was both positive and negative over the 19 th century, but has been fairly consistently positive (earth slower than nominal) in the 20 th century. Is this rising sea level? The answer is probably not or only partly. This is an active area of debate.

B. Mechanical and Electronic Clocks

The first fairly accurate short-term clock was the pendulum clock. A diagram of the key timekeeping elements of a pendulum clock is shown in Figure 3. In this case gravity is the restoring force that pulls the weight down each swing. The same elements for a spring and balance wheel clock used in mechanical pocket watches are also shown in Figure 3. In this case the wheel oscillates back and forth with the spring serving as the restoring force.

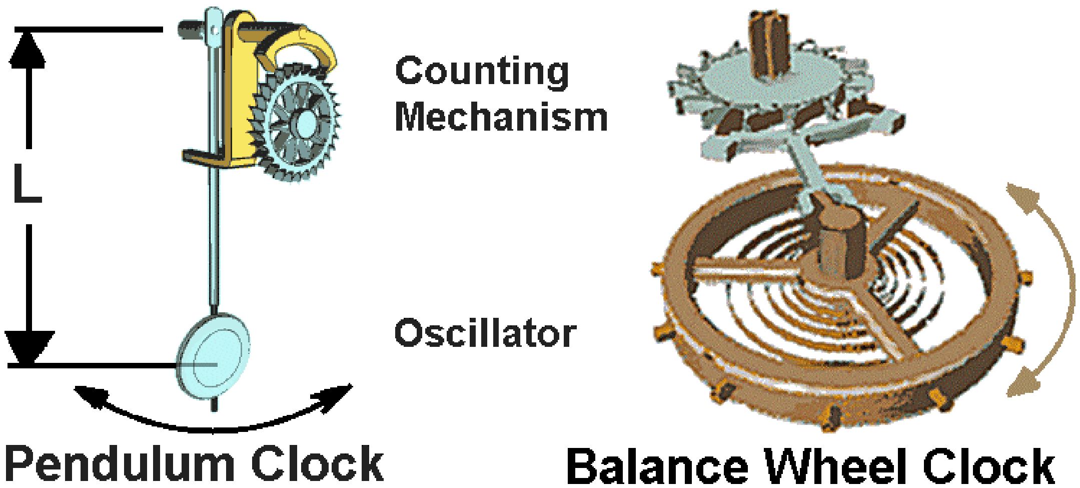

Figure 3

The period, T, of a simple pendulum clock is theoretically given by

$$T = 2\pi\sqrt{\frac{L}{g}}\ ,$$

where L is the length and g is the acceleration of gravity. The value of g is nominally 980 cm/s/s, but varies with both height and latitude. These effects are large enough to be easily seen. In the mid-1600's a clock calibrated in Paris was sent to French Guinea in northern South America on a French astronomical mission. This is very near the equator. It showed errors of over 2 minutes per day, easily seen in the astronomical observation. This lead Newton to conclude the earth is bulged out at the equator. The equator was further from the center of the earth and the value of g less.

The spring/balance wheel watch mechanism, shown in Figure 3, is subject to temperature variations. To overcome this, John Harrison invented the bimetallic strip in the mid 1700's. He invented the chronometer, a very accurate watch for use on ships. This was used to fix longitude at sea. His most famous clock, called H4, lost only a 5 seconds on a 81-day voyage from

England to Jamaica. H4 had many complex mechanisms to compensate for environmental variations.

To go beyond the pendulum clock we must use electronic systems. Quartz is a piezoelectric crystal. This means that it generates a voltage if it is squeezed. Setting it vibrating will generate a oscillating voltage. This is a physical oscillation, for example in the thickness of the crystal. The crystal is ringing. There are also bending, shear, and torsion modes of crystals used for clocks. The thickness mode is most common. These are the basis of crystal oscillators.

A diagram of a quartz crystal and showing how pieces are cut from it is shown in Figure 4. One of the mounting configurations for a crystal oscillator is shown in Figure 5. The thickness mode is essentially a pressure wave going across the crystal. The period is closely related to the physical thickness and the speed of sound in the material.

The frequency of a quartz oscillator varies mainly with temperature. Inexpensive oscillators, such as those found in PC's can show seconds of variation in a day. The variations in simple crystals often has a diurnal curve, due to the ambient temperature variations. There is also an ageing effect where the period gets longer over weeks and years. Sometimes this is discontinuous as some parts of the crystal re-adjust some of the internal stress.

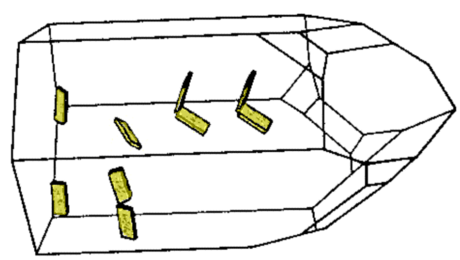

Quartz Crystal
Planes Used for Oscillators

Figure 4

The different "cuts" have different sensitivity to temperature. The better cuts change frequency by 1 part in 10^{11} per degree C. To make the best quartz oscillators, the best cut must be used and the temperature has to be stable to under 0.01 degrees centigrade. To accomplish this a “double oven” is used, a temperature controlled oven inside a second temperature controlled oven.

Crystal oscillators are sensitive to accelerations, even the gravitational field of the earth. A standard test is to change the orientation with respect to up by 90 degrees. This is called a “tip over test”. It usually changes the frequency by over 1 part in 10^8. This means crystal oscillators are “noisy” in a vibrating environment unless well mechanically isolated.

Modern quartz watches have a different geometry. A small “tuning fork” is micro-machined. It is the oscillations of this very small piece of quartz that produces the events counted. This is usually much smaller than the crystal used in a radio. The tuning fork is less accurate than a true quartz oscillator, as seen in Figure 1.

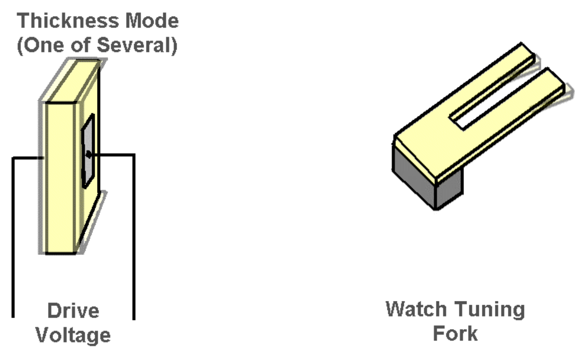

Crystal Oscillator Configurations
Not to Scale

Figure 5

C. Atomic Clocks

The next step from temperature-controlled crystal oscillators is atomic clocks. The energy levels of atoms are discrete. The emission of a bundle of electromagnetic energy comes from the atom going from one energy level to a lower one. Absorption is the inverse process, the atom goes from a lower level to a higher one. The frequency of the light is proportional to the energy level difference. Usually the difference is large enough to produce light or x-rays. However there are some close levels that produce emission in the radio range. These are "fine structure" levels of the outer electrons. Only a few elements can be used, usually specific isotopes of heavy metals. There are several different configurations used depending on the chemistry of the element.

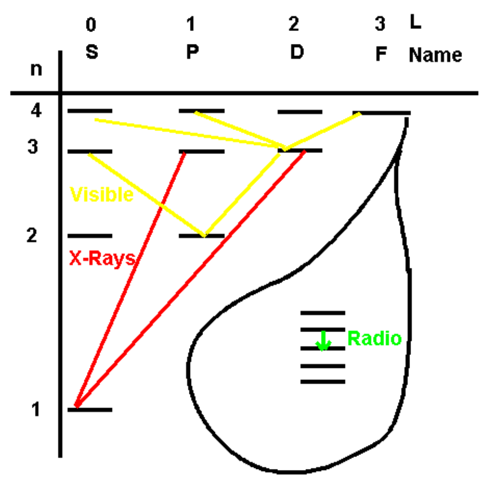

Energy Levels Complex Atom
n = primary quantum number
L = angular momentum number

Atomic oscillators are much more accurate because the electrons that generate the levels are isolated from any physical mounting. There are environmental effects though. Typically magnetic fields and temperature cause frequency shifts or line broadening. The sensitivity is down by several orders of magnitude from the effects on mechanical clocks and crystal oscillators. A magnetic shield is commonly used.

The inherent accuracy of an atomic clock is related to the width of the energy level which generates the width of the observed line. This is related to two things. First the clock is sensing a very large number of atoms. These may be in motion with different velocities. The different Doppler shifts cause line broadening. Other environmental effects enter in the same way. Second the time the atom spends in a state, T, limits the width to 1/T or more. This is important for the beam clocks.

The most common "inexpensive" atomic clock is the Rubidium (atomic symbol Rb). This is a light abortion cell shown in Figure 6. A Rb light shines pure light through a cell of dilute Rb gas. There is an oscillating microwave field applied to this gas at 6.8 GHz. When the frequency is correct, the light is greatly absorbed. A feedback loop tunes the frequency to minimum transmission. The frequency of the microwaves is counted and used as the events of the atomic clock. In the diagrams that follow, only the key items are highlighted.

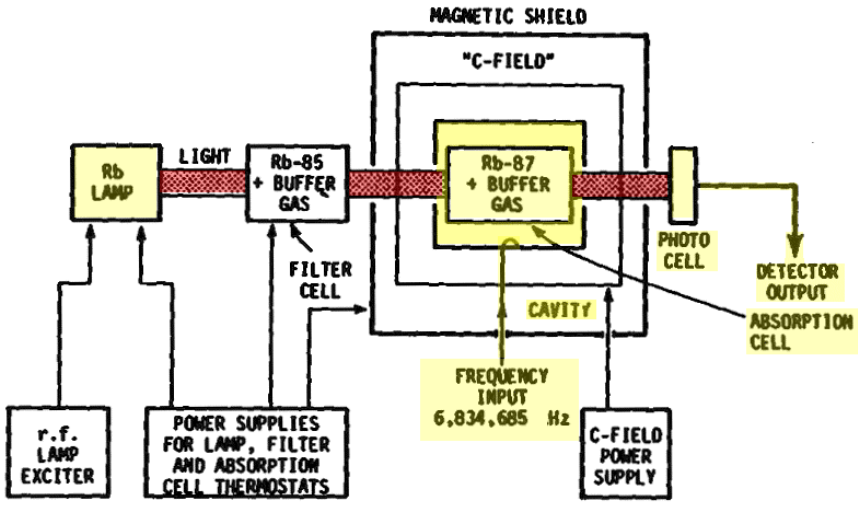

Rb Gas Cell Atomic Oscillator

Figure 6

The Cesium (Cs) atomic clock, shown in Figure 7, actually sends a beam of Cs atoms though two sections of a microwave cavity. (This is called a Ramsey Cavity. Norman Ramsey got the Nobel Prize for his work on atomic clocks.) The beam then goes by a strong magnet. If the microwave frequency is correct, about 9.2 GHz, the atoms change state, are bend more by the exit magnet and are deflected into a detector. The atoms enter the Ramsey cavity twice in the two forks. The line width is proportional to the time the atoms spend in between these two passages. For commercial Cs clocks, the forks are about 10 cm apart. For the national time laboratories, such as NIST, the length can be over 2 m.

The cesium atomic is better than the rubidium. Precisely by how much and how to specify the accuracy of a clock are the subject of the next section.

The Cesium atom of atomic weight 133 is now used to define the second. The second is 9,129,631 periods of the radiation corresponding to the transition between two specific hyperfine levels of the ground state. Up until 1960, the seconds was defined as 1/86400 th of a mean solar day. It was realized that the earth's day varied much more that the clocks being used. So from

1960 to 1967 it was the mean solar day of year 1905. In 1967 the atomic second was substituted. In 1972 a very specific atomic transition of Cs-133 under a specified environment was adopted.

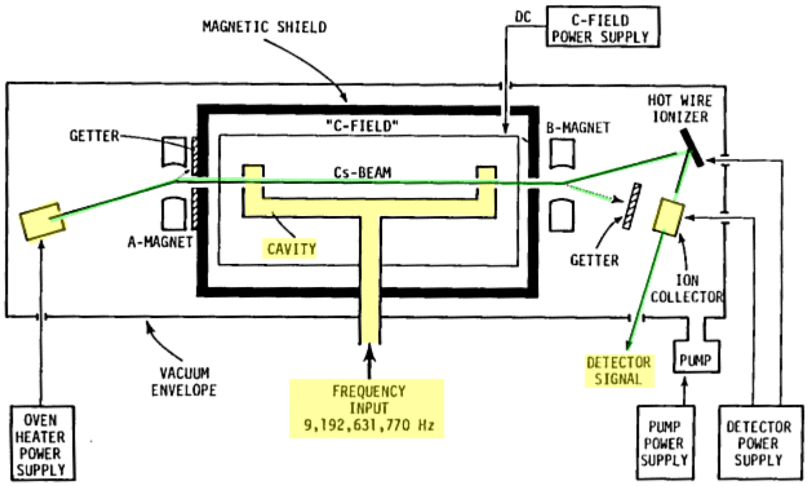

Cesium Atomic Beam Oscillator

Figure 7

There are several newer types of atomic oscillators that have even better behavior. These include the Hydrogen Maser (both active and passive types), the Mercury ion clock, and the Cesium fountain clock. Some of the latest research can be found at the Time Division of USNO website, (tycho.usno.navy.mil) and the NIST site (www.boulder.nist.gov/timefreq).

D. Paper Clocks - Clock Ensembles

Finally we should talk about the clocks that are really used by time standard organizations such as NIST and USNO. Many good clocks are averaged to find an average time. There is usually no physical clock that represents the average value. Therefore this ensemble average is sometimes called a paper clock. In some cases, such as the USNO, one Cs clock is driven to the ensemble average, but there is a small error in this locking process.

Each clock in the ensemble is weighted in the average. This average, which does not represent any physical clock, is called the paper clock. The USNO uses about 50 Cesium clocks and several other more advance clocks in the standard they provide to the US military. (And because this drives the GPS time base, it provides time to much of the world). There are several sets of Cesium clocks in separate clock rooms that are temperature controlled and on different power systems. Each clock is in its own temperature chamber inside these rooms. NIST has a similar set up.

III. Characterizing Clocks

A. Problem Defining Clock Error Standard Statistics

How do we know that one type of clock, or one specific clock, is better than another? What do we measure to determine this. The answer turns out to be a little complicated. Ordinarily one would measure many time errors and compute the standard deviation. Any average value can be removed by proper calibration so the standard deviation would indicate the quality of the clock. However it is found that the value of the standard deviation of the clock time error is not well defined. It depends on the length of time that is used to make the measurement and the rate at which the measurements are made.

This problem arises because of the characteristics of the time error. One type of error common in good clocks is a random walk in frequency. If you measure a clock's frequency and wait a small interval of time, it will change slightly. This is a random number in size and direction. Because this is a walk in frequency, you have to integrate (add up) the changes to get the error in time. It turns out the standard deviation of the time error from this process grows with the length of the number of steps or samples. It is ill defined.

David A. Allan came up with a solution to this. He defined what is now called the Allan Variance. It is often denoted σ_τ^2, or $\sigma_y^2(\tau)$. The square root of the Allan variance is usually reported, which is a standard deviation like value. The Allan variance uses a ratio of time error to time interval as the fundamental variable. Therefore the Allan variance is a unit-less quantity.

A diagram of the normal form of the Allan variance plot is shown in Figure 8. The plot is on a log-log scale. The line would be flat - independent of sample interval - if the standard deviation was well defined. However the plots are not flat, but usually decrease (get better) with increasing time sample length, flatten out and then rise. Several orders of magnitude are shown on each axis. Higher values are worse.

The nominal Allan variance curve has three areas. In the first area, the Allan variance decreases as a power of the sample time. (On a log-log plot, a power law, $y = x^n$, is a straight line of slope n.) The power depends on the type of oscillator.

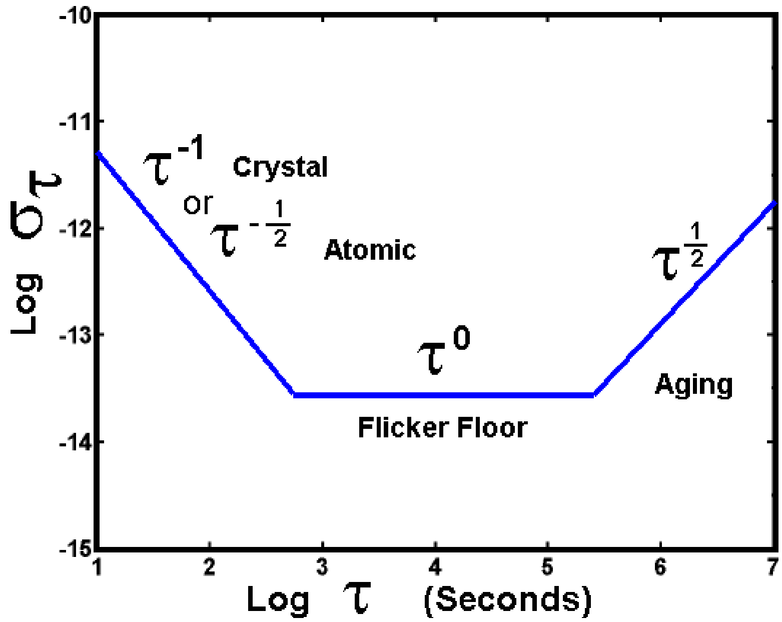

Allan Variance Plot

Figure 8

In the center there is a region where the curve is relatively flat. This lowest level is sometimes called the "flicker floor". The noise type there is called flicker noise. This type of noise has no easy description outside of mathematics.

Finally at long times the Allan variance increases as the square root of time. The values on the axis in Figure 8 will vary greatly from clock type to clock type.

B. Conversion of Allan Variances to Expected Clock Error

The conversion of "sigma tau" or the Allan variance into the expected time error is straightforward. The Allan Variance curve of the clock is measured in a laboratory. Then if you set a clock at some time, the standard deviation of the expected time error is just

$$\sigma_\tau(\tau) = \tau \sigma_\tau$$

where τ is the time since clock setting and σ_τ is the square root of the Allan variance. You just have to read the values off the curve and multiply the two coordinates to get the expected time error. Figure 9 shows an example of this process for a Cesium clock. The resultant "error" is the expected variation standard deviation.

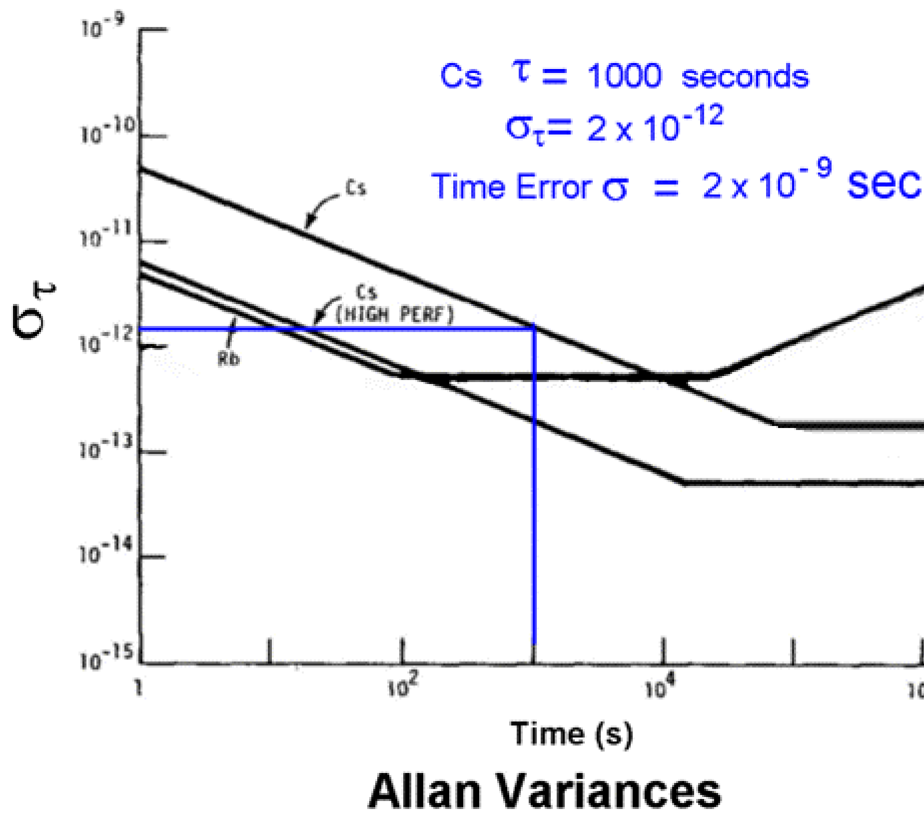

Figure 9

C. Examples of Allan Variances

A set of typical Allan variances (Allan sigmas really), are shown in Figure 10. A wide variety of clock types are shown. The highest curves (worst clocks) are the crystals. The floor is at short times, between 1 and 100 seconds. At very short times, crystals are better than most atomic clocks. Therefore atomic clocks usually have good crystals that are used over time frames of a few seconds. The output loop is driven at long times by the atomic oscillator.

The Cesium clocks (green lines) are the most common high accuracy used in national time standards laboratories. The larger time laboratories and some academic institutions have H-masers (gold line) and other newer atomic clocks. Cesium clocks are also on most larger US Navy ships. Their noise floors covers time spans from a few days to a few weeks.

Rubidium clocks (red lines) are becoming inexpensive and common. However they are being replaced with a form of GPS clock. This is of the form GPS + something else (dashed lines). At short times these are on the "something else" stability diagram. At long times they have the characteristics of GPS time - that is the USNO's clock ensemble. This is seen in the overlapping lines on the right edge of Figure 10.

Notice that the "better" Cs clock is shown as having the same Allan variance as the Rb for times shorter than about 10 sec. If only one number is reported on the quality of a clock, it is usually the lowest value of the Allan variance, the flicker floor. In that sense the Cs clock is better than the Rb.

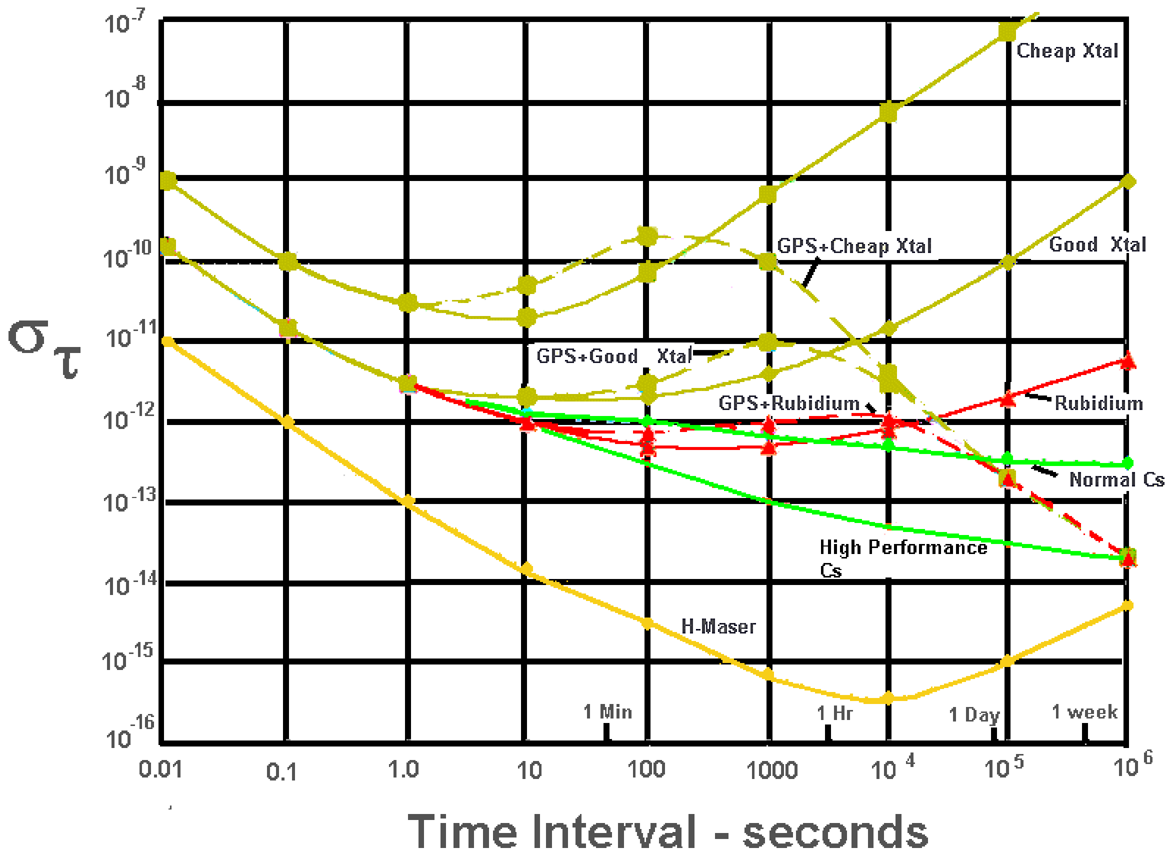

Figure 10

IV. Time Dissemination

These high quality clocks are expensive and generally used only by national time organizations. These time standard organizations, of which there are a few dozen throughout the world, know time very well. But they need to get it to the users. They need to disseminate time.

The simplest method is the telephone. You can dial a number in Washington DC and hear the USNO time "tic" with the time announced on the minute. A similar service is provided by NIST. In addition NIST provides this signal over the radio stations WWV or WWVH at 5, 10,15 and 20 MHz. There is also a NIST station, WWVB at 60 kHz, which contains only digital data used to set a clock. Inexpensive “atomic” clocks in the $50 range that are being sold by several vendors that use the WWVB signal. All these radio signals are a fairly inaccurate means of time distribution by modern standards. Most can be used to set a clock to a few msec at best.

Before GPS came along, many different methods were used for time dissemination. Most still exist. The Figure 11 shows many of these methods and their accuracy limits. For many of the radio systems, the ionospheric conditions are a limiting effect. Notice that the most accurate methods involve GPS.

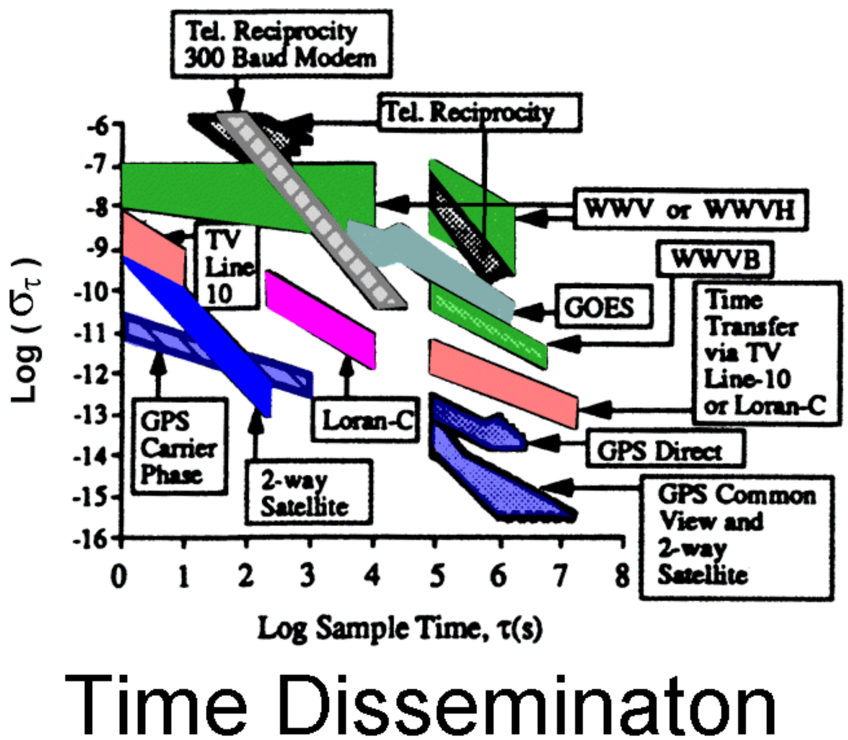

Figure 11

There are two ways GPS can be used to set time. For the vast majority of users, the time is just a side product of finding a location. This time is guaranteed to be within 100 nsec (0.1 microseconds) of the official time called UTC. In fact it is usually under 10 nsec. This is a mode that does not have a “sample” time and hence no Allan variance.

A common type of accurate clock is called a disciplined crystal or disciplined Rubidium. Here a GPS receiver is combined with the other clock. The position of the GPS antenna is generally known and set into the system. Then the only unknown the GPS receiver solves for is time. This value of time is compared with the crystal or Rubidium clock and the error in the local clock computed and stored. Several parameters in a power series may be determined for a crystal. When GPS is unavailable, the time from the local clock, after corrections are applied, is used.

Time services compare their results using GPS with a much more precise method. Two time services observe the same GPS satellites at the same time and the raw data (or their residuals from a known model) are exchanged. All the errors that originate in the spacecraft and the operational control center are common to the two measurements. Differencing the data gives the time difference between the two users. . This is called GPS common view. This provides a comparison of the different national time frames at a level of a few 1/1000 of a nanosecond (picoseconds).

V Annotated Bibliography

There are many books, reports, and web pages devoted to the measurement of time and time interval over the ages. Here only a few items will be listed.

1. NIST/NBS Overviews

There are three items that provide a very good overview of clocks in general and precise time in particular. Two are NIST reports (or NBS reports under their old name). The first is a 1977 compendium of reports.

Hellwig, H, NBS Technical Note 616, "Frequency Standards and Clocks: A Tutorial Introduction", 1977

Hellwig has an excellent tutorial with descriptions of the characteristics of crystal and atomic clocks. This includes basic physics and engineering discussions suitable for a scientist with no background in precise time. It also discusses practical problems such as crystal mountings and system Q etc. Diagrams of different atomic clock configurations are presented that form the basis of some of the diagrams in this report. Copies of review papers by Hellwig and Allan are attached. At the end there is an unpublished note by Allan contains a detailed procedure for computed Allan variances and tables of spectral characteristics.

The second is an updated version of this concept. This contains an introduction paper with an overview of the subject followed by a collection of papers spanning 20 years. The papers are organized by topic, not publish date.

Sullivan, D.B. et al (Editors), NIST Technical Note 1337, "Characterization of Clocks and Oscillators", 1990.

2. Technical Application Notes and Reports

Agilent (old HP Test and Measurement Divisions), has an excellent set of application notes. In addition to notes on how to use their counters, oscillators, and atomic clocks, there is an excellent overview by Allan, Ashby and Hodge published as application note 1289. HP/Agilent application notes in PDF can be obtained on line from their web page going to products and then support. This material includes the tables of clock noise types as well as detailed examples on how to compute both the Allan variance and the modified Allan variance.

HP Application Note 1289, "The Science of Timekeeping"
http://www.tm.agilent.com/classes/MasterServlet?view=applicationnote&apn-ItemID=1000000363&language=eng&locale=US

The Allan note, is also available at the web site that David Allan set up after retiring from NIST. http://www.allanstime.com/ , with the direct link being http://www.allanstime.com/Publications/DWA/Science_Timekeeping/index.html .

Allan, Ashby, and Hodge also published similar material as a supplement to the December 2000 issue of *GPS World*.

For crystal oscillators, John Vig who works for a US Army research laboratory, has written an excellent tutorial that appeared in an IEEE publications. It is also available on line.

“Introduction to Quartz Frequency Standards”, John R. Vig, SLCET-TR-92-1 (Rev. 1) October 1992, http://www.ieee-uffc.org/freqcontrol/quartz/vig/vigtoc.htm .

The Piezo Corporation had an excellent tutorial on crystal oscillators in their older catalogs. This had numeral values for sensitivities to different environmental factors for different crystal cuts. A tutorial is now on line at http://www.piezo.com/edu.html .

3. Internet References

Many national standards laboratories have time divisions. Most have web pages with some introduction material. The two US organizations have a joint page

http://www.time.gov/

which has links to different sites with educational material for the lay person. More wide-ranging material is available at the NIST Time Division Frequently Ask Questions page

http://www.boulder.nist.gov/timefreq/phase/Properties/main.htm .

The main page for the NIST Time Division is www.boulder.nist.gov/timefreq while the USNO Time Division’s page is at (tycho.usno.navy.mil) .

The UK national standards laboratory, often called just NPL, has some very good tutorial material at:

Home Page http://www.npl.co.uk/
Time and Frequency Division Home Page http://www.npl.co.uk/npl/ctm/
Time Tutorial http://www.npl.co.uk/npl/publications/atomic/

4. Papers

Allan, D.W., "Statistics of Atomic Frequency Standards", IEEE, 54, 221,1966
(The is the paper where the Allan variance was introduced.)

Lesage,P, and C. Audoin, "Characterization and Measurement of Time and Frequency Stability", Radio Science, 14, 521, 1979.

(A review with much the same material on time and frequency domains, Allan variances and their relationships. This source may be available where some others are not. This paper is also in NIST 1337)

Allan, D. W., "Time and Frequency (Time-Domain) Characterization, Estimation, and Predictions of Precision Clocks and Oscillators", IEEE Trans. Ultrasonics, Ferroelectrics, and Frequency Control, 34, 647, 1987

Allan, D. W. et al, "Standard Terminology for Fundamental Frequency and Time Metrology", Proc. of 42 nd Frequency and Control Symposium, p 419, 1988.
(This has a nice table of frequency and time domain characteristics of clock noise types. The regions of an Allan variance curve are defined and discussed. This paper is also in NIST 1337. The same Allan variance tables are in the Allan-Ashby-Hodge articles.)

Allan, D.W., Ashby, N., and Hodge, C. "Fine Tuning Time in the Space Age", IEEE Spectrum, 35, 42, 1999.

Allan, D.W., Ashby, N., and Hodge, "A Brief History of Precision Time and GPS", Supplement to GPS World, December 1998.

Hewlett Packard Application Notes which are available online.

Allan, D.W., N. Ashby, and C. Hodge, "The Science of Timekeeping", App. Note 1289.
------------, "Fundamentals of Time Interval Measurements", App. Note 200-3
------------, "GPS and Precision Timing Applications", App. Note 1272