
An Introduction to Frequency Calibration

Part II

Michael Lombardi

National Institute of Standards and Technology
Time and Frequency Division

This article is part II of a two-part series on frequency calibrations. Part I introduced the topic of frequency calibrations, discussed the specifications involved, and described several types of oscillators. Part II presents an overview of transfer standards, describes how frequency calibrations are made, and presents an automated frequency calibration system developed at NIST.

Transfer Standards

As discussed in part I of this article, a frequency calibration is a comparison between the *device under test* (DUT) and a reference. The DUT is usually a quartz, rubidium, or cesium oscillator. The reference is an oscillator of higher performance than the DUT or a *transfer standard* received by radio.

There are advantages to using a transfer standard for calibrations. First, all transfer standards have a cesium oscillator at their source. A cesium oscillator is an intrinsic frequency standard with a typical frequency uncertainty of about 1×10^{-13} . Cesium oscillators are expensive both to buy and maintain, and not all calibration laboratories can afford them. It is far less expensive to buy a radio receiver and "share" a cesium oscillator with many other users. Second, even if a calibration laboratory already has a cesium oscillator, they still need to check its performance. The only practical way to do this is by comparing the cesium to a transfer standard.

Another advantage of transfer standards is *traceability*. Most transfer standards are traceable to the national frequency standard maintained by the National Institute of Standards and Technology (NIST). Some, like the high frequency (HF) radio stations WWV and WWVH and the low frequency (LF) station WWVB, are traceable because they are directly controlled by NIST. Others, like LORAN-C and the Global Positioning System (GPS) satellite broadcasts, are traceable because their reference is regularly compared to NIST. Other transfer standards, like signals broadcast from outside the United States, are also considered to be traceable. This is because NIST compares its frequency standard to the standards maintained in other countries.

Of course, some compromises are made when you use a transfer standard. Even though the signal may be referenced to a nearly perfect source of frequency, its performance is degraded as it travels along its radio path. The more stable the path, the better the performance of the transfer standard. A stable radio path is one where the signals always take about the same amount of time to get from the transmitter to the receiver.

Since radio signals travel at the speed of light, they take about 3.3 microseconds to travel one kilometer. This means that a signal from a station 1000 kilometers (621 miles) away will already be 3.3 milliseconds late by the time you receive it. If you are using this signal for a timing application, you would want to compensate for this delay to recover the most accurate time possible. However, if you are using this signal for frequency calibrations, the amount of delay doesn't matter. The only thing that matters is the amount that the path delay varies over the course of a measurement. The transfer standards with the most stable paths provide traceability with the smallest amount of uncertainty.

To illustrate this, consider the signal broadcast from WWV, located in Ft. Collins, Colorado. WWV is a HF radio station (often called a *shortwave* station) that transmits on 2.5, 5, 10, 15, and 20 MHz. WWV is referenced to the national frequency standard at NIST, but by the time the signal gets to your receiver, much of its potential performance has been lost. This is because shortwave radio signals travel along different paths at different times of the day. Most shortwave users receive the *skywave* or the part of the signal that travels up to the ionosphere and bounces back down to earth. The skywave is responsible for the primary appeal of shortwave radio, which is its ability to travel great distances (many thousands of kilo-

meters) and provide a large coverage area for listeners. However, the same characteristics that allow WWV to travel great distances limit its usefulness as a transfer standard. Since the height of the ionosphere constantly changes, the path delay constantly changes, often by as much as 500 to 1000 microseconds. Therefore, even though WWV is traceable to NIST, its frequency uncertainty is limited to 1×10^{-7} .

Other transfer standards have more stable paths and much lower uncertainty values. LF and VLF radio stations (like NIST radio station WWVB) can provide traceability to NIST with a frequency uncertainty of about 1×10^{-11} . The LF and VLF radio paths are much more stable than an HF path, but will experience a path delay change when the height of the ionosphere changes at sunrise and sunset. The two finest transfer standards are LORAN-C and the Global Positioning System (GPS). LORAN-C is traceable to NIST with an frequency uncertainty of just 1×10^{-12} . The performance of GPS is even better, with a frequency uncertainty of about 5×10^{-13} .

Table 1 shows some of the transfer standards available and the NIST traceability that they provide when averaged for a measurement period of at least one day.

LORAN-C

LORAN is an acronym for Long Range Navigation. The LORAN-C navigation system consists of nearly 20 synchronized "chains" or networks of stations. These chains provide coverage for most of the Northern Hemisphere. Each chain consists of one master station and two to five secondary stations.

All LORAN-C stations use the same carrier frequency (100 kHz). Therefore, the receiver must distinguish between signals broadcast from many different stations on many different chains. Each chain is identified by a unique Group Repetition Interval (GRI). The length of the GRI is fixed, and each chain is named according to its GRI. For example, the 7980 chain has a GRI of 79.8 milliseconds. This means that every 79.8 milliseconds (about 12 times per second) each station in the chain transmits a group of pulses.

When the pulses leave the transmitter, they radiate in all directions. The groundwave travels parallel to the surface of the Earth. The skywave travels upward and is reflected off of the ionosphere. Receiving the skywave is less desirable than receiving the groundwave, because the skywave path is less stable. Fortunately, the groundwave is usually easy to receive in most areas in the Northern Hemisphere. You will usually receive the skywave

| Transfer Standard | Frequency Uncertainty over 24-hour measurement period (with respect to NIST) |
|--------------------------------------|--|
| HF Radio Signals (like WWV and WWVH) | 1×10^{-7} |
| LF and VLF Radio Signals (like WWVB) | 1×10^{-11} |
| LORAN-C | 1×10^{-12} |
| Global Positioning System (GPS) | 5×10^{-13} |

Table 1. Traceability levels provided by various transfer standards with level of NIST traceability when averaged for 1 day period.

only if the groundwave signal has traveled a long distance and is too weak and noisy for the receiver to track. If your receiver is within 1600 kilometers (1000 miles) of a LORAN-C transmitter, you should be able to receive the groundwave without any problems.

The LORAN-C pulse shape was designed so that the receiver can distinguish between groundwave and skywave. Most receivers are designed to track the third cycle of the pulse. A picture of a LORAN-C pulse with the third cycle identified is shown in Figure 1. The third cycle is a good cycle to track for two reasons. First, it arrives early in the pulse so we know that it is groundwave. Second, it has considerably more amplitude than the first and second cycles in the pulse. This makes it easier for the receiver to track. Of course, if the receiver moves to a different cycle, you'll get a 10 microsecond phase step in your data (the period of the 100 kHz carrier.) However, under normal conditions this seldom happens.

LORAN-C Receiving Equipment

There are two types of LORAN-C receivers: timing receivers and navigation receivers. Timing receivers are made specifically for time transfer and frequency calibrations. Since the market is small, only a few companies manufacture these receivers. They range in price from about \$3000 to over \$10,000. Timing receivers have a mul-




Figure 1. LORAN-C pulse with the third cycle identified.

titude of features. They provide several different frequency outputs (usually 10 MHz, 1 Hz, and the GRI pulse). Some models have a computer interface, so that the data can be sent to a computer and stored.

LORAN-C navigation receivers are different. They are mass produced for the marine market and are very inexpensive (some cost as little as \$200). However, navigation receivers were not designed for frequency calibrations and do not provide frequency outputs. However, you can often find the GRI pulse on the circuit board, amplify it if necessary, and use it as a frequency reference.

LORAN-C Performance

As with all transfer standards, the performance of LORAN-C is degraded by variations in the radio path. The size of these variations depends upon the signal strength, your distance from the transmitter, the weather and atmospheric conditions, and the quality of your receiver and antenna. Figure 2 shows the type of path variations you can typically expect. It shows a phase comparison between the LORAN-C 8970 chain as received in Boulder, Colorado, and the NIST national frequency standard. The 8970 signal is broadcast from the master station in Dana, Indiana, a site about 1512 kilometers (940 miles) from Boulder.

The graph in Figure 2 compares GRI pulses from the 8970 chain to the NIST national frequency standard. The measurement period was 100 seconds. During this period, the frequency offset between the NIST frequency standard and the cesium standard at the LORAN-C transmitter would produce only a tiny amount of phase shift (less than 1 nanosecond). Therefore, the graph simply shows the LORAN-C path noise or the difference in the arrival

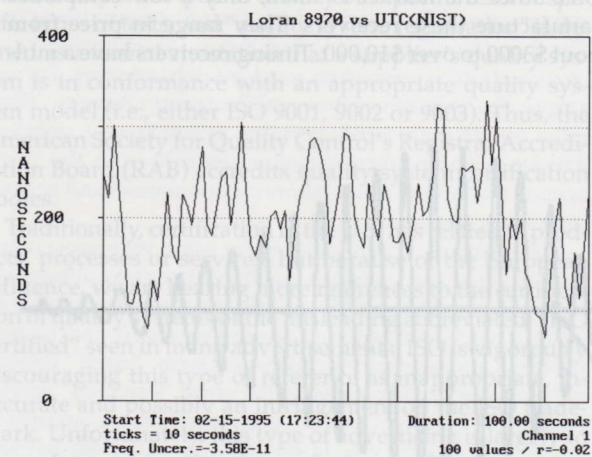


Figure 2. LORAN-C path variations over 100-seconds.




Figure 3. A 96-hour cesium calibration using LORAN-C.

times of the GRI pulses due to path variations. The range of this path noise is slightly more than 200 nanoseconds.

Although the path noise averages out over time, the short term stability of LORAN-C is poor. This influences the time required to make a calibration. The better the oscillator, the longer it takes to calibrate. For example, a typical rubidium oscillator with a frequency uncertainty of 1×10^{-11} will accumulate 1 microsecond of phase shift in 24 hours, or just slightly more than 1 nanosecond of phase shift in 100 seconds. Can we calibrate a rubidium oscillator using LORAN-C in 100 seconds? No. The 200 nanoseconds of path noise will completely hide the 1 nanosecond phase shift. We need to wait much longer than 100 seconds before we can determine the frequency uncertainty of a rubidium oscillator.

You can use LORAN-C to make a quick calibration if your DUT has a large frequency offset. For example, if your DUT is high in frequency by 1×10^{-7} , none of the path noise shown in Figure 2 would be visible. Even if a measurement period of just 100 seconds is used, the frequency uncertainty contributed by path noise will have no effect on the calibration.

For most types of calibrations, however, it's best to use a 24-hour measurement period when using LORAN-C. Over 24 hours the uncertainty contributed by path noise won't affect the calibration of a high quality quartz oscillator or a rubidium. For cesium oscillators a 24-hour calibration may not be long enough. However, it is still possible to calibrate a cesium using a longer measurement interval.

Figure 3 shows the results of a 96-hour calibration of a cesium oscillator using LORAN-C. The thick line is an estimate of the frequency uncertainty. Although the LORAN-C path noise is clearly visible, we can tell the cesium oscillator is low in frequency by about 3.4×10^{-12} .

Global Positioning System (GPS)

LORAN-C signals originate from earth and this limits their potential as a transfer standard. GPS signals originate from a satellite and are capable of better results. A satellite is an ideal vehicle for a transfer standard. This is because there is a clear path between the satellite and the receiver that is not influenced by earthbound noise sources.

Like LORAN-C, GPS is a radio navigation system. GPS is operated by the U.S. Department of Defense (DOD). The GPS system consists of a constellation of 24 earth orbiting satellites (21 primary satellites and 3 in-orbit spares). The 24 satellites orbit the Earth in 6 fixed planes that are inclined 55° from the equator. Each satellite is 20,200 kilometers above the Earth and has an 11 hour, 58 minute orbital period, which means a satellite will pass over the same place on earth 4 minutes earlier each day.

GPS is an excellent transfer standard for two reasons. First, each satellite carries 4 atomic oscillators (two rubidiums and two cesiums). These oscillators are referenced to the United States Naval Observatory (USNO) and are traceable to NIST. Second, the coverage area of GPS is unmatched. With 24 satellites in orbit, GPS should be usable anywhere on the Earth's surface. No other transfer standard can lay claim to this distinction.

Each GPS satellite broadcasts two carrier frequencies, L1 at 1575.42 MHz and L2 at 1227.6 MHz. A spread spectrum waveform, called a pseudo random noise (PRN) code is transmitted on L1 and L2, and each satellite is identified by the PRN code it transmits. There are two types of PRN codes. The first type is a coarse acquisition code (called the C/A-code) with a 1.023 MHz chip rate and a period of 1 millisecond. The second is a precision code (called the P-code) with a 10.23 MHz chip rate. The C/A code is broadcast on L1, and the P code is broadcast on both L1 and L2.

Through the use of a technique called Selective Availability (SA), the DOD adds jitter to both the C/A-code and P-code to intentionally degrade its positioning accuracy. This is done for national security reasons. Special P-code receivers are available that employ cryptographic logic to remove the effects of SA from the P-code. However, these receivers may only be used by U. S. and allied government agencies and their military forces, and if in the national interest, selected civilian GPS users. Therefore, nearly all calibration laboratories use the C/A code. The jitter added to the C/A-code by SA adds to the frequency uncertainty of GPS, but averages out in much the same fashion as LORAN-C path noise.

GPS Receiving Equipment

GPS receivers can be divided into two categories: those intended for navigation and positioning applications, and

those designed for timing applications. As with LORAN-C, the market for navigation and positioning is many times larger than the market for timing receivers. This makes timing receivers more costly and less widely available than navigation receivers. However, since the overall GPS market is so large, GPS timing receivers are still common and cost less than their LORAN-C counterparts.

Most GPS timing receivers provide a 1 Hz, or 1 pulse per second (pps) output. This on-time pulse can provide time synchronization to better than one microsecond. Some receivers also provide a 1 kHz output (derived from the C/A code). Other receivers come equipped with an internal oscillator (usually a quartz but sometimes a rubidium), that phase locks the oscillator to GPS. These receivers usually produce at least one standard frequency output (1, 5, or 10 MHz), and sometimes provide other frequencies like 100 kHz. The cost of a GPS timing receiver ranges from about \$2,000 at the low end to \$20,000 or more for those receivers with built-in rubidium oscillators.

Before a GPS receiver can be used as a transfer standard, it has to determine its position on Earth. To get this information, you simply turn the receiver on. The receiver then performs a sky search, to find out which satellites are currently above the horizon and visible from the antenna site. The receiver then collects two blocks of data (called the *almanac* and *ephemeris*) from the satellites it finds. It can then compute a 3-dimensional coordinate (latitude, longitude, and altitude) as long as 4 satellites are in view. Once the receiver's location is known, it can provide a good output frequency.

GPS timing receivers often improve the quality of the received frequency by averaging data obtained from a number of different satellites. The averaging improves the receiver's performance by removing some of the intentionally introduced jitter. Most receivers have *multiple channels* and can track anywhere from about 4 to 12 satellites at once (although more than 8 will only be available in rare instances). If the number of satellites in view exceeds the number of channels, the receiver tracks the satellites that have the most favorable position in the sky. The receiver's software has some predefined criteria it uses to select the best satellites to track. For example, the criteria may be based on only selecting satellites whose elevation angles are higher than 20°. *All in view* receivers have enough capacity to assign a separate channel to every satellite in view.

GPS antennas are small, seldom more than 15 centimeters (6 inches) in diameter. The antenna must be mounted outdoors where it has a clear, unobstructed view of the sky. Signal loss can be a problem with GPS antennas. If long cable runs are used (more than 20 meters, for example) you may need to use a low-loss cable or a line amplifier to get the receiver to work properly.

GPS Performance

GPS has many technical advantages over LORAN-C. The signals are usually easier to receive, the equipment is less expensive, and the coverage area is much larger. In addition, the performance of GPS is superior to that of LORAN-C. However, like all transfer standards, the short term stability of GPS is not particularly good (made worse by SA), and it influences the time required to make a calibration. As with LORAN-C, a 24-hour measurement period is recommended for frequency calibrations using GPS.

To illustrate this, Figure 4 shows a 100-second comparison between GPS and a cesium oscillator. The cesium oscillator has a frequency uncertainty of 1×10^{-13} and its accumulated phase shift during the 100-second measurement period is much less than 1 nanosecond. Therefore, all of the noise on the graph can be attributed to GPS path variations. Most of the noise can be attributed to SA.

Figure 5 shows a 1-week comparison between GPS and the same cesium oscillator used in Figure 4. The scale of this graph is 550 nanoseconds. The thick line is an estimate of the frequency uncertainty (based on a *linear least squares curve fit*). Although the GPS path noise is still clearly visible, we can see the trend contributed by the frequency offset of the cesium. This trend tells us that the cesium oscillator is high in frequency by about 5×10^{-13} .

Comparing Two Frequencies: The Time Interval Method

In part I of this article, we introduced the topic of *phase comparisons*. In order to complete a calibration, we must perform a phase comparison and measure the difference




Figure 4. GPS compared to cesium over 100-seconds.

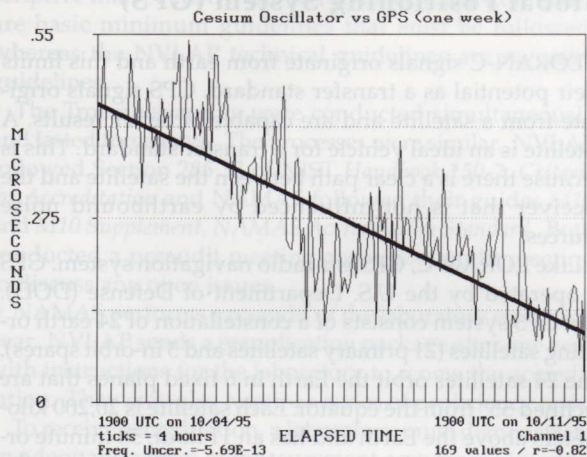


Figure 5. GPS compared to cesium for 1-week period.

between the DUT and the reference. Figures 2 through 5 all show the results of phase comparisons made between a transfer standard and a cesium oscillator. In this section, we'll discuss how phase comparisons are made using the *time interval method*.

The *time interval method* involves using a device called a *time interval counter (TIC)* to measure the time interval between two signals. If the two signals have the same frequency, the time interval will not change. If the two signals have different frequencies, the time interval will change, although usually very slowly. By looking at the *rate of change*, you can calibrate the device. It is exactly as if you had two clocks. By reading them each day, you could determine the amount of time one clock gained or lost relative to the other clock.

It takes two time interval measurements to produce any useful information. By subtracting the first measurement from the second, we can tell whether time was gained or lost. Normally, a computer is used to read the TIC and display, process, and store each measurement. The computer's software can then go through the appropriate steps to compute the frequency uncertainty and stability of the DUT.

A TIC has inputs for two electrical signals. One signal serves as a start pulse to the TIC, and the other serves as a stop pulse. The TIC starts measuring the time interval when the start pulse arrives, and stops measuring when the stop pulse arrives.

TIC's differ in specification and design details, but they all contain several basic parts. These parts are known as the *timebase*, the *main gate*, and the *counting assembly*. The timebase provides evenly spaced pulses used to measure time interval. The timebase is usually an internal quartz oscillator. It must be stable because timebase errors will directly affect the measurements. The main gate controls




Figure 6. A time interval measurement.

the time at which the count begins and ends. Pulses which pass through the gate are routed at the counting assembly where they are displayed on the TIC's front panel, or read by computer. The counter can then be reset to begin another measurement.

The stop and start inputs are usually provided with level controls that set the amplitude limit (or *trigger level*) at which the counter responds to input signals. If the trigger levels are set improperly, a TIC may stop or start when it detects noise or other unwanted signals and produce invalid measurements.

Figure 6 illustrates how a TIC measures the interval between two signals. Input A is the start pulse, and Input B is the stop pulse. The TIC begins measuring a time interval when the start pulse reaches its trigger level, and stops measuring when the stop pulse reaches its trigger level. The time interval between the start and stop pulses is measured by counting cycles from the timebase. The measurements produced by a TIC are in time units, milliseconds, microseconds, nanoseconds, and so on. These measurements assign a value to the phase difference between the reference and the DUT.

The most important specification of a TIC is *resolution*, which is the degree to which a measurement can be determined. For example, if a TIC has a resolution of 10 nanoseconds, it could produce a reading of 3340 nanoseconds, but not a reading of 3345. This is because 10 nanoseconds is the smallest significant difference the TIC can measure. Any finer measurement would require more resolution.

In traditional TIC designs, the resolution is limited to the period of the TIC's timebase frequency. For example, if a TIC uses a 10 MHz timebase, its resolution is limited to 100 nanoseconds. This is because a complete cycle of the 10 MHz signal only occurs every 100 nanoseconds. Traditional TIC designs counted cycles to measure time interval, and could not resolve time intervals smaller than the period of one cycle. To get around this problem, some TIC designers multiplied the timebase to get more cycles and thus more resolution.

For example, multiplying the timebase up to 100 MHz

makes 10 nanosecond resolution possible. Going beyond 10 nanoseconds is difficult, but 1 nanosecond counters have been built using a 1 GHz timebase. However, these types of TIC designs are expensive to build and design, and the use of high frequency components can make them unreliable.

A better way to get more resolution is to be able to detect parts of a timebase

cycle through *interpolation* and not be limited by the number of whole cycles. All modern, high performance TIC's use this method. Interpolation has made counters with 1 nanosecond resolution commonplace. Some TICs have a resolution of less than 20 picoseconds.

When we use a TIC for calibrations, we can't just plug any two signals in and expect to get useful measurements. To illustrate this, consider a case where we want to calibrate a 5 MHz DUT versus a 5-MHz reference. If we plug both 5-MHz signals into a TIC, the largest reading we can get is 200 nanoseconds (the period of 5 MHz). This is because both signals produce a complete cycle every 200 nanoseconds. This leads to two problems. The first problem is that when the DUT shifts its phase by more than 200 nanoseconds, the counter will either *overflow* (go from 200 back to 0) or *underflow* (go from 0 back to 200). A DUT with a frequency uncertainty of 1×10^{-10} would cause an over or underflow every 20 seconds! This makes the measurements *ambiguous* to 200 nanoseconds. If the oscillator has a large enough frequency offset, we could get overflows and not even know it.

If we record a reading of 150 nanoseconds, for example, and record another one of 155 nanoseconds, we can't be sure if the phase changed by 5 nanoseconds or 205 nanoseconds. The second problem is that when the time interval gets too small, the TIC makes errors. Many TIC's aren't able to measure a time interval of 100 nanoseconds. If the start and stop signals are too close together, the TIC returns an invalid reading.

If the signals are too high in frequency, the solution is to use a *frequency divider* to convert them to a lower frequency. A frequency divider could be a stand-alone instrument, a small circuit, or just a single chip. Most divider chips divide by multiples of 10, so it is fairly common to find circuits that divide by a thousand, a million, and so on. A common practice is to divide everything down to 1 Hz, then plug 1 Hz signals into the TIC. If you use 1 Hz signals you will seldom get overflows or underflows. You don't have to divide all the way down to 1 Hz, as long as your signals are related in frequency. For example, you can use 1 kHz or 10 kHz signals, but you may get frequent overflows and underflows that the system designer will have to take care of in software.

The NIST Frequency Measurement and Analysis System

We'll conclude our tutorial with a brief look at an actual frequency calibration system. The system we'll look at is the NIST Frequency Measurement and Analysis System (FMAS). The FMAS was designed at NIST and uses the concepts discussed in this tutorial.

A block diagram of the FMAS is shown in Figure 7. The entire FMAS fits in a metal equipment rack (75 centimeters tall) and is controlled by an industry standard 486 computer system. Software developed at NIST controls all aspects of the measurement process. It makes measurements and stores and graphs them automatically. It backs up the data automatically on tape every 10 days. Up to 5 oscillators can be calibrated at once. The system produces calibration graphs that show the performance of each DUT for measurement periods ranging from 2 seconds to 150 days.

Looking at Figure 7, you can see the various components of a frequency calibration system. The FMAS uses a GPS receiver as a transfer standard. The receiver is connected to the computer using a standard RS-232 interface. The GPS receiver produces a 1 kHz output frequency that provides traceability to NIST with an uncertainty of 5×10^{-13} .

The FMAS makes phase comparisons using the time interval method. It includes a TIC with a single-shot resolution of less than 40 picoseconds. The TIC achieves its high resolution through an interpolation scheme and includes a built-in multiplexer that enables it to switch between five inputs. This allows the FMAS to calibrate up to five oscillators simultaneously. The TIC also includes built-in divider circuitry and can directly accept either a 1, 5, or 10 MHz input on each of the five channels. The TIC is software controlled, so users don't have to worry about setting trigger levels. The system is connected to NIST through a telephone modem so that NIST personnel can troubleshoot the system remotely and analyze the calibration data.

A system like the FMAS provides a calibration laboratory with a well-defined and documented calibration method. Such a system can be a major asset to a laboratory seeking ISO registration or compliance with a laboratory accreditation program.

References

1. George Kamas and Michael A. Lombardi, *Time and Frequency User's Manual*, NIST Special Publication 559, 1990.

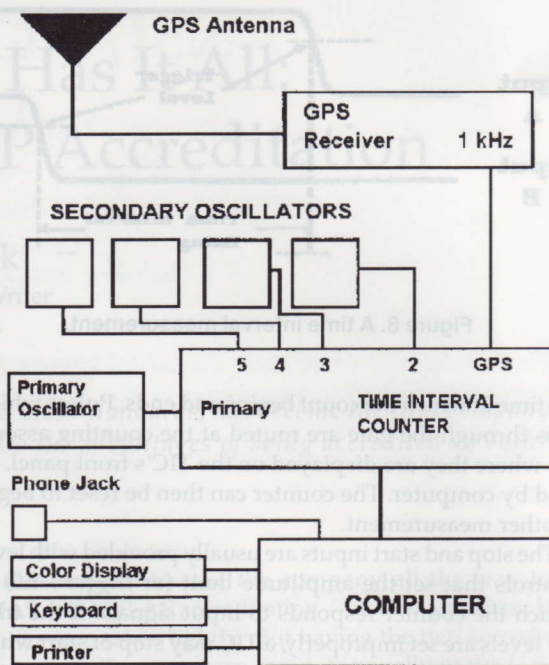


Figure 7 - Block Diagram of the NIST Frequency Measurement and Analysis System.

2. Victor S. Zhang, Dick D. Davis, and Michael A. Lombardi, "High Resolution Time Interval Counter," *Precise Time and Time Interval Conf. (PTTI)*, pp. 191-200, 1994.
3. B. Hoffmann-Wellenhof, H. Lichtenegger, and J. Collins, *GPS Theory and Practice*, Springer-Verlag, 2nd edition, 1993.
4. D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, editors, *Characterization of Clocks and Oscillators*, NIST Technical Note 1337, 1990.
5. Michael A. Lombardi, "The NIST Frequency Measurement Service," *Cal Lab Magazine*, Vol. 2, No. 3, May-June 1995.

Michael Lombardi, NIST Time and Frequency Division, 325 Broadway, Boulder, CO 80303. 303-497-3212, fax 303-497-3228, E-mail lombardi@bldrdoc.gov