

NISTIR 6610

NIST Frequency Measurement and Analysis System: Operator's Manual

Michael A. Lombardi



NIST
National Institute of Standards and Technology
Technology Administration, U.S. Department of Commerce

NISTIR 6610

NIST Frequency Measurement and Analysis System: Operator's Manual

Michael A. Lombardi
*Physics Laboratory
Time and Frequency Division
325 Broadway
Boulder, CO 80305*

August 2001



U.S. Department of Commerce
Donald Evans, Secretary

National Institute of Standards and Technology
Karen H. Brown, Acting Director

Table of Contents

INTRODUCTION	iii
CHAPTER 1	
INSTALLATION	1
List of Parts	1
Block Diagram of System	2
System Setup	3
GPS Antenna	3
Uninterruptible Power Supply (UPS)	4
Equipment Rack	4
Video Monitor	5
GPS Receiver	5
Printer	6
Telephone Modem	6
Connecting Oscillators to the System	7
Starting the System	9
CHAPTER 2	
LEARNING TO USE THE MENUS	13
Selecting a Menu Item	13
Exiting a Menu	14
Exiting the FMAS software	14
CHAPTER 3	
MAKING MEASUREMENTS	15
Automatic Plotting	17
Changing Oscillators	19
Quick Check of System	20
CHAPTER 4	
ANALYZING DATA	23
Loading the Demonstration Data from Disk	24
Viewing the Data on Screen	24
Printing the Data	26

Viewing the Data Statistics	26
Plotting the Data	27
Editing the Plot on Screen	29
Removing Data Points from the Plot	31
Plotting a Portion of the Plot	33
YAXIS, XAXIS, and MAGNIFY	34
Changing the Title of the Plot	35
CHAPTER 5	
RETRIEVING YOUR DATA	37
Retrieving All of the Data	37
Retrieving Some of the Data	38
CHAPTER 6	
UTILITIES	39
ShortTerm	39
Retrieving Short Term Measurement Data	43
Adjusting an Oscillator Using METER	44
History	44
GPS	46
CalibrateTIC	49
CHAPTER 7	
SETUP SOFTWARE	51
APPENDICES	
A - An Introduction to Frequency Calibrations	55
B - Calibration Reports	87
C - Specifications and Compliance Documents	91
D - Glossary	93
INDEX	99

Introduction to

NIST Frequency Measurement and Analysis System

The NIST Frequency Measurement and Analysis System (FMAS) measures and calibrates oscillators and helps you analyze their performance. The system can measure and calibrate quartz, rubidium, and cesium oscillators. Each day, it prints a calibration plot of the data collected in the last 24 hours. Up to five oscillators can be connected to the system and calibrated simultaneously.

The FMAS is housed in a metal equipment rack. Its components include a Global Positioning System (GPS) radio receiver, a time interval counter, a graphics printer, a computer, and a color video monitor. The FMAS also includes a backup power supply that can keep it running for about 30 minutes during power outages. The FMAS will stop operating during longer power outages, but will restart automatically when power is restored.

NIST provides all of the equipment and supplies you need to operate the FMAS. Your only requirements are to supply one or more oscillators and a dedicated telephone line (so that NIST can access the system by modem).

This manual tells you how to install and operate the system. The first four chapters are especially important. Chapter 1 describes how to install the system. Chapter 2 describes how to use the menus. Chapter 3 describes how to run the system in the measurement mode. Chapter 4 describes how to run the system in the analysis mode.

Please call NIST whenever you have a question or need assistance. You can reach us Monday through Friday from 8 a.m. to 5 p.m. Mountain Time by calling (303) 497-3212 or (303) 497-3378.

Chapter 1

Installation

The NIST Frequency Measurement and Analysis System consists of the following parts:

- metal equipment rack
- rack-mount computer
- rack-mount keyboard with sliding drawer
- rack-mount monitor shelf with color monitor
- telephone modem
- printer and cable
- multi-channel time interval counter
- Global Positioning System (GPS) receiver with antenna
- Uninterruptible Power Supply (UPS)

In addition to the parts listed above, the system is also shipped with all of the other items necessary for operation. These include cables, power cords, manuals, printer paper, and printer ribbons/ink cartridges. All parts should be unpacked and accounted for before installation. If any part is missing, call NIST at (303) 497-3212 or (303) 497-3378.

Figure 1.1 shows a block diagram of the NIST Frequency Measurement and Analysis System.




Figure 1.1. Block Diagram of the NIST Frequency Measurement and Analysis System.

The following sections tell how to install the NIST Frequency Measurement and Analysis System (FMAS). The installation is described in two parts. The first part, SYSTEM SETUP, tells how to set up the FMAS. The second part explains how to connect oscillators to the FMAS.

SYSTEM SETUP

This section describes how to set up the FMAS. We recommend that you install the parts in the order listed.

1. GPS ANTENNA

Install the GPS antenna as follows:

- a. Choose a location for the GPS antenna that is as high above the ground and as obstruction free as practical. GPS reception is line of sight and you will be receiving signals from orbiting satellites, so make sure the antenna has an unobstructed view of the sky on all sides.
- b. The antenna size and shape might vary, but most are small white plastic cones about 5 or 6 inches high (12 to 15 cm). The antenna has a threaded mounting base that attaches to an 18 inch metal pole (46 cm) included with the system. Mount the antenna by any convenient means, but make sure the pole is mounted vertically. If possible, you can affix the pole to an existing roof mount using the supplied clamps.
- c. Route the supplied cable down to your system location. The system is shipped with 50 feet (15 m) of RG-58 cable, unless previous arrangements were made with NIST. Make sure that all connections are made carefully, and do not kink or stress the cable. If you require a longer cable length, contact NIST and a custom lower loss cable will be sent to you. RG-58 will not work properly at the GPS carrier frequency (1575.42 MHz) when cables longer than 20 m are used.

2. UNINTERRUPTIBLE POWER SUPPLY (UPS)

Install the UPS as follows:

- a. Set the UPS on the table or desk top where the system will be located. Plug the UPS into a secure three-pronged wall outlet. Do not use an outlet that will be turned off or used by power drills, vacuum cleaners, motors, or other electrical devices.
- b. Turn on the UPS so that the batteries can begin charging.

3. EQUIPMENT RACK

The equipment rack is shipped with the computer, keyboard, GPS receiver, and monitor shelf already installed. The computer has been configured and tested at NIST and is ready to use in the system. You do not have to open the computer case or install any software.

- a. Remove the equipment rack from its shipping box and gently set it on the table or desk top where you will locate the system. The equipment rack is heavy, and you may want to have two people available to lift it. Handles are provided on both sides of the equipment rack.
- b. Note that the front panel of the equipment rack has two locks. One lock is for the monitor shelf. The second lock is for a door on the computer that covers a floppy drive and the on/off and reset buttons for the computer. You will receive a set of identical keys with the FMAS and the same key can be used to open either lock. Use the keys to open the computer door. Make sure the power switch is OFF for now.
- c. Locate the keyboard cable at the back of the equipment rack and plug it into the round keyboard jack on the back of the computer. The keyboard jack might not be marked, but it is the only jack where the keyboard connector will fit. Then pull out the keyboard drawer. Make sure that you can open and close the drawer without snagging the cable.

- d. Connect a power cord to the back of the computer, and plug the computer into the UPS. Leave the computer's power switch (behind the door) OFF for now.

4. VIDEO MONITOR

Install the monitor as follows:

- a. Place the video monitor on the monitor shelf through the large opening on the back of the equipment rack. Position the monitor so that the entire screen is visible even when the door to the monitor shelf is closed.
- b. Plug the monitor directly into a wall outlet, if available. Do not plug the monitor into the UPS unless no other outlets are available. The UPS will be able to endure longer power outages if it does not have to supply power to the monitor.
- c. Connect the monitor's signal cable to the video jack on the back of the computer. The video jack is the only jack where the monitor connector will fit.
- d. You may turn the monitor ON now. Keep in mind that the monitor will last longer if you turn it OFF when you leave for the night or when no one is using the system.

5. GPS RECEIVER

The FMAS uses a GPS receiver that is embedded inside the rack mount computer system. This receiver draws power from the computer's power supply, and is controlled by the computer's software. Its antenna connector is mounted on the computer's back panel.

To install the receiver, simply connect the cable from the GPS antenna to the TNC connector marked ANTENNA on the back of the computer.

6. PRINTER

Install the printer as follows:

- a. Remove the printer from the packaging and set it on the desk or table top.
- b. Load the printer with paper. It uses standard 8.5×11 inch cut sheet paper. A package of this paper is supplied with the system.
- c. Connect the printer cable (a round cable) to the jack on the back of the printer.
- d. Connect the other end of the printer cable to the jack on the back of the computer. The jack is not marked, but it will be the only jack that fits.
- e. Plug the power cord from the printer into the UPS.
- f. Turn the printer ON.

7. TELEPHONE MODEM

The modem is already installed inside the computer. You simply need to plug the modem into a telephone wall jack.

- a. Plug the supplied phone cable into the telephone wall jack.
- b. Plug the other end of this cable into the phone jack marked "WALL" on the back of the computer. Sometimes there will be two modem jacks and neither jack is labeled. In that case, either jack will usually work.

This completes the installation. The next section explains how to connect your oscillators to the measurement system.

CONNECTING OSCILLATORS TO THE SYSTEM

This section explains how to connect your oscillators to the FMAS. All connections are made using cables (typically RG-58, 50 Ω impedance) with standard BNC connectors. Several cables of this type were supplied with the system.

The first thing required is to select one of the oscillators in your laboratory as your primary oscillator. Your primary oscillator should meet the following criteria:

- *It should be the best oscillator available that is a permanent part of your laboratory. For example, select an oscillator that will remain permanently at your site and not one that will be sent to another location.*
- *Your oscillator must provide a frequency output of either 1, 5, or 10 MHz.*
- *Your oscillator should be connected to a stand-by power supply so that power outages do not interrupt the signals.*
- *Your oscillator should run 24 hours a day, 7 days a week, so that NIST can continually monitor its performance and send you a monthly calibration report, and so that you can continuously establish traceability to the United States national frequency standard.*
- *The signal from your primary oscillator at 50 Ω should be at least 200 mV peak-to-peak (about 70 mV RMS) and should not exceed 3.5 V peak-to-peak (about 1.2 V RMS). It is unlikely that the signal will be too small, but it could be too large. If the signal is too large, you need to use an attenuator to reduce it to an acceptable level.*

Oscillators are connected to the BNC connectors on the back panel of the measurement system. The time interval counter already installed inside the system can

accept either a 1, 5, or 10 MHz input signal from your primary oscillator as a time base.

Connect your oscillators to the system as follows:

Primary Oscillator (Channel 1) - Connect a cable from your primary oscillator to the BNC connector labeled PRIMARY. This signal serves as the counter's time base, and will also be measured via GPS.

Channel 1 measures the frequency uncertainty of your primary oscillator using GPS as a reference frequency. Since GPS signals are monitored by NIST and since the measurement system is linked by modem to NIST, your calibrations are traceable to the national frequency standard. You will receive a monthly calibration report that documents traceability (see Appendix B).

Secondary Oscillators (Channels 2, 3, 4, and 5) - Since your primary oscillator is traceable to NIST, you can use it as a reference frequency for calibrating other oscillators, thus completing their traceability chain. As many as four other oscillators may be calibrated simultaneously on channels 2, 3, 4, and 5. To connect additional oscillators to the system, simply connect their output to channel 2, 3, 4, or 5 on the back panel of the measurement system. The input signal can be any frequency from 1 Hz to 120 MHz in 1 Hz increments.

IMPORTANT NOTE: All signals that you connect to the system should not exceed 3.5 V peak-to-peak (about 1.2 V RMS). You should check the signal level of every oscillator before connecting it to the system. If the level exceeds 3.5 V peak-to-peak, use an attenuator to reduce the signal amplitude. Signals that are too large can cause interference with adjacent counter channels, or cause the counter to make trigger errors. In either case, the system will make measurement errors and reduce the quality of your data.

STARTING THE SYSTEM

To start the system, turn the computer's power switch ON. If the video monitor is not ON, turn it ON at this time. Before the system can make measurements, it has to make sure that the GPS signal is usable. After a few seconds, a screen will appear like the one shown below:

Configuring NIST-FMAS GPS Receiver			
Number of Satellites	6	Pulse (1 or 100 pps)	1 pps
Timing Stability (ns)	44	UTC Date (YYYY-MM-DD)	2001-06-19
Tracking Status Byte	32	UTC Time (HH:MM:SS)	16:59:46
Position Hold Mode	ON	Latitude	39° 59m 43.60s N
PDOP	2.8	Longitude	105° 15m 44.52s W

The integrity of the GPS Signal has been proven, and the system is now ready to use for frequency calibrations.

The system will be started in 8 second(s).

Figure 1.2. The GPS acquisition screen.

The screen shows the status of the GPS signal acquisition process. The receiver automatically acquires the signal. There are several conditions that must be met before the system decides that GPS is usable for frequency calibrations. The receiver must be tracking more than three satellites. The timing stability must be less than 80 ns (this internal check should not be confused with the true stability of the receiver which is generally much smaller). The value for PDOP must not exceed 3.0. PDOP is an acronym for *Position Dilution of Precision*, which is a measure of the quality of the position fix. The lower the PDOP, the more "spread out" the satellites are in the sky and the better the position fix. And finally, the *Position Hold Mode* must be turned on. This indicates that your antenna's position has been surveyed with an uncertainty of

less than 10 m. At this point, the system is ready to use for frequency calibrations. The system will wait 10 s and then load the measurement software.

The first time you install the system, it may take as long as 30 min to complete the process of GPS signal acquisition. The acquisition process will be repeated only if you restart the system, or if a power outage occurs that lasts longer than about 30 min (about what the UPS can handle). However, after the first time, the acquisition process should take a few minutes or less since the system has already stored most of the data that it needs. If the position of the antenna changes for any reason, you should restart the system. This will force the system to again go through the signal acquisition process and resurvey the new position of the antenna.

Load View Copy Plot Stats Measure Names Utilities DOS Quit Enter names for each device being calibrated			
	Channel Names	Nameplate (Hz)	Ref (P/G)
1			G
2			P
3			P
4			P
5			P
When you finish editing the names, press <ESCAPE> to save			

Figure 1.3. The channel name input screen.

Once the measurement software is loaded, the screen shown in Figure 1.3 should appear on the monitor. If this screen does not appear, go through the previous

steps in this chapter again and make sure the system has been properly installed. If you need help, call NIST at (303) 497-3212 or (303) 497-3378.

Once the screen shown in Figure 1.3 appears, enter a name and input frequency, and select a reference frequency (described at end of chapter) for each channel that has a signal connected. For example, if you connected signals to channels 1, 2, and 3 on the counter, enter a channel name and input frequency, and select a reference frequency for channels 1, 2, and 3. You can move around the input screen by pressing the <UP ARROW> key to go up, or the <DOWN ARROW> key to go down. Channel names, input frequencies, and the reference frequency selection are discussed below.

Channel Names - Give a clear, simple name to each channel. For example, if your primary oscillator is a Model X Rubidium and you are comparing it to GPS, then enter a name like:

Model X Rubidium vs GPS

On channel 2, you might be measuring a Model Z Quartz oscillator. You can use a name like:

Model Z Quartz (Serial Number 101)

Clear, simple titles are very helpful when you look at your data in the future. You can enter names up to 50 characters long, which allows you to be quite descriptive. You can include such information as the oscillator's serial number, property number, and so on. When you finish entering a channel name, press <ENTER>. If you make a mistake and want to erase the channel name, press the key to erase the channel name.

Input Frequencies - The input frequency is the nameplate or nominal frequency of the oscillator you will be measuring. In other words, this is the frequency that the oscillator is supposed to produce. For example, if you are measuring a quartz oscillator that is supposed to produce a 5 MHz output, the input frequency to the system is 5 MHz. The FMAS will then measure the difference, or frequency offset between the the nameplate frequency and the actual frequency produced by the oscillator.

The input frequency must be entered in hertz, so 5 MHz will be entered as 5000000 Hz. The FMAS is capable of measuring any frequency from 1 Hz to 120 MHz in increments of 1 Hz. This allows you to measure nearly any frequency that you are likely to encounter in the calibration laboratory.

Reference Frequency - The FMAS can compare each input frequency either to your primary oscillator (the oscillator connected to channel 1), or to GPS. The primary oscillator is designated as P and GPS is designated as G, so enter either a P or a G in this column. Keep in mind that the primary oscillator on channel 1 can be compared only to GPS, so you are not allowed to change the G in the channel 1 column.

It is conventional to compare channels 2 through 5 to the primary oscillator, if the primary oscillator is a significantly better performer than the other oscillators you are measuring. For example, if your primary oscillator is a cesium, and the oscillators on channels 2 through 5 are rubidium or quartz devices, you should select P on channels 2 through 5. However, if all of your oscillators are rubidiums, you should use GPS as the reference on all channels, since this makes more sense than comparing one rubidium oscillator to another.

After you have entered the necessary information for every channel in use, press the <ESC> key to save the channel names and descriptors to disk. You are now ready to go on to the next chapter.

Chapter 2

Learning to Use the Menus

This section tells how to use the FMAS menus. At this point, the MAIN MENU should be on screen as shown below:

```
Load View Copy Plot Stats Measure Names Utilities DOS Quit
Load data from disk
```

All FMAS menus work the same way. Each menu is located inside the box at the top of the screen. The top line of the menu lists the menu items. One menu item is always highlighted. The second line of the menu describes the function of the highlighted menu item. In the above menu, for example, the menu item LOAD is highlighted. The words "Load data from disk" are printed on the second line. This tells you that LOAD is used to load data from disk.

You can move to a different menu item by using the right and left arrow keys. You should try this now. Each time you press either <RIGHT ARROW> or <LEFT ARROW>, a different menu item is highlighted. For example, if LOAD is highlighted, pushing <RIGHT ARROW> causes VIEW to be highlighted. Whenever you move to a different menu item, the text changes on the second line of the menu.

The menu is circular. For example, if you are on the left most menu item and press <LEFT ARROW>, you will move to the right most menu item. Likewise, if you are on the right most menu item and press <RIGHT ARROW>, you will move to the left most menu item. Try this by holding down either <RIGHT ARROW> or <LEFT ARROW> and watching the menu wrap around.

You can also move to the left most menu item by pressing <HOME>, or to the rightmost menu item by pressing <END>.

Selecting a Menu Item

There are two ways to select a menu item. The first way is to locate the item you want using the arrow keys, and to then press <ENTER>.

The second way is easier and faster. Notice that each menu item begins with a different letter. You can select a menu item by typing its first letter. For example, you could select LOAD by pressing the <L> key.

Exiting a Menu

Nearly every menu contains an item called QUIT. When you select QUIT, it returns you to the previous menu. If you make a mistake, and wind up with the wrong menu on screen, you can select QUIT and go back to where you were. Pressing the <ESC> key does the same thing as selecting QUIT.

Exiting the FMAS software

To exit the FMAS software, select QUIT from the MAIN MENU, and then select YES from the next menu. You are then returned to DOS. To run the FMAS software again, type FM2000 and press <ENTER>.

You may want to use DOS commands while you are running the FMAS. For example, you might want to copy a disk, format a disk, or view a directory. To do so, select DOS from the MAIN MENU. When you select DOS you will return to DOS. When you are ready to return to the FMAS, type EXIT and press <ENTER>.

Chapter 3

Making Measurements

To begin making measurements, select MEASURE from the MAIN MENU. In a few moments, the measurement screen will appear. The top part of the measurement screen is a bar graph that gives you a quick indication of how each oscillator is performing. The bars represent channels 1 through 5 (from top to bottom). Figure 3.1 shows a sample measurement screen:




Figure 3.1. A sample measurement screen.

The bars are color coded. Each bar corresponds to one of the channel names shown at the bottom of the measurement screen. The channel names represent channels 1 through 5 (from top to bottom). Each channel name is next to a colored box. The box colors match the bar colors. For example, the top bar is orange. It represents the performance of the oscillator connected to channel 1. The top box is also orange. It marks the channel name for channel 1.

When you first start the system, the measurement screen shows a bar and a channel name for every channel in use. However, the bars extend only to the number "5" on the far left side of the screen. After the system has been running for between 1 and 2 hours, the bars will extend further (assuming that the oscillator's performance is better than 1×10^{-5}) and show the true performance of the oscillator.

The bars show the estimated frequency offset of each oscillator. The numbers at the top of the screen (from 5 to 13) indicate frequency offset values ranging from 1×10^{-5} to 1×10^{-13} . For example, if the bar extends exactly to 10, the frequency offset is 1×10^{-10} . There are 10 divisions between the numbers. If the bar extends halfway between 10 and 11 (5 divisions past 10), the frequency offset is 5×10^{-11} . The longer the bar, the smaller the frequency offset of the device under test.

The range of the system is 1×10^{-5} to 1×10^{-13} . If the performance of an oscillator is worse than 1×10^{-5} , the bar will still extend to the "5". If the performance of an oscillator is better than 1×10^{-13} , the bar will extend only to the "13". Each bar also contains either a plus sign "+" or a minus sign "-". A plus sign means the oscillator is high in frequency. A minus sign means the oscillator is low in frequency.

When the system has been on for less than 24 hours, the bars are based on all of the available data. For example, if the system has been on for 12 hours, the bars show the oscillator's performance over the last 12 hours. However, when the system has collected data for more than 24 hours, only the most recent 24 hours are used.

The FMAS stores a new data point to disk every hour at 0 minutes and 0 seconds UTC. For example, the FMAS stores data at 17:00:00 UTC, 18:00:00 UTC, and so on. Whenever the system stores data, it also updates the bars to reflect the oscillator's performance over the last 24 hours. Stable oscillators will perform at the same level for long periods of time (many days), and so the length of the bars might not change very often. Even so, you should glance at the bars periodically to determine whether your oscillators are performing as expected.

Figure 3.1 shows the measurement screen of an FMAS with five channels connected. The top bar (for channel 1) extends all the way to the right. It shows that the frequency offset of UTC(NIST) is 1×10^{-13} or less. Channel 2 shows a quartz oscillator with a frequency offset of about 1×10^{-8} . Channels 3 through 5 show devices whose frequency offset is measured in parts in 10^{10} , 10^{11} , and 10^{12} , respectively.

The screen area below the bar graph contains information about the measurements being made. The small red box contains the current date and the current time (UTC). This information is received by satellite from the GPS broadcast. As discussed earlier, the channel names are located inside the blue box (below the date and time values).

The "Difference" column shows the difference between the last two counter readings. This difference is expressed in nanoseconds. For example, if the last two counter readings are 11736.32 and 11735.34, the difference is 0.98 ns (11736.32 minus 11735.34).

The "Count" column shows the most recent counter reading for each channel. The FMAS reads each counter channel in numerical order. The arrow moves each time a measurement is made, and points to the channel currently being read.

The FMAS has a number of powerful data analysis features that are described in Chapter 4. To use these features, you need to stop the measurements. However, you should always leave the system running in the measurement mode whenever you are not using the analysis software. The FMAS was designed to run 24 hours per day, 7 days per week. Every 24 hours the FMAS stops taking readings and plots the data collected in the past 24 hours. It then starts taking readings again. The only time you are required to stop the measurements is when you change the frequency inputs connected to the system. Automatic plotting and changing the frequency inputs are discussed in the next two sections.

AUTOMATIC PLOTTING

The FMAS produces a daily calibration plot for every oscillator connected to the system. If two oscillators are connected there are two daily plots; if four are connected there are four daily plots, and so on. The plots are normally printed every day at 1600 UTC (to change the plot time, see Chapter 7). After the plots are printed, the FMAS starts to take readings again. Each plot shows the data recorded from the oscillator during the past 24 hours.

You do not have to wait for the daily plot time to print a plot of the last 24 hours. You can print plots at any time by pressing <CONTROL-P>. If you have more than one channel, but want to print only the plot for Channel 1, press <CONTROL-P>,

then press <ESC> while the plot is being printed. The FMAS will then go back to taking readings without printing the other plots.

A typical daily calibration plot is shown in Figure 3.2. The plot's title is centered above the plot. The plot's title is the channel name. The y-axis of the plot is labeled MICROSECONDS. The x-axis intercepts the y-axis at 0 μ s. The scale of the plot is indicated by the numbers at the top and bottom of the y-axis. The plot has the number 100 printed at the top and bottom of the y-axis. This means that the areas above and below the x-axis each represent 100 μ s. The total range of the y-axis is 200 μ s (from -100 to +100).




Figure 3.2. A sample daily calibration plot.

There are three lines of text below the plot. The top line shows the time and date when the first and last data points were recorded. Since the plot normally covers a period of 24 hours, the time of day should be the same for the first and last data points, but the date will be advanced by one day.

The second line shows the tick mark interval and the channel from which the

data was recorded. For a 24 hour plot, the tick marks are always spaced 1 hour apart. For example, the left edge of Figure 3.2 represents 1900 UTC, the first tick mark is 2000 UTC, the second tick mark is 2100 UTC, and so on.

The bottom line shows the number of data points (usually 25 for a 24 hour plot), the estimated frequency offset of the oscillator, and the correlation coefficient (r). For more information about these values, please see Appendix A.

Each calibration plot is printed on a separate piece of paper. You may want to keep the plots in a notebook for future reference.

Please note that if the printer is turned off, is off-line, or is out of paper, the FMAS will not be able to print the plots. It will sound an alarm for 10 s when it tries to print each plot. If you do not fix the printer (turn it on or put it on-line) the FMAS will go back to making measurements without printing any plots. If you do not want or need the daily plots, you can leave the printer turned off. However, you should always keep the printer connected to the FMAS, or the system may not work properly.

CHANGING OSCILLATORS

From time to time, you'll need to change, remove, or add oscillators to the FMAS. To do so, stop the system by pressing <CONTROL-X>. You are then returned to the MAIN MENU. Then complete the following steps:

- Disconnect the oscillators that you want to remove and connect the new oscillators using the procedures described on pages 7 and 8.
- Select NAMES from the MAIN MENU. Type in the new channel names and input frequencies as described on pages 10 through 12 and press <ESC> to save. If a channel is no longer used, delete the channel name by pressing the key.
- When you finish, press <ESC> to save the new channel names to disk.
- Restart the measurements by selecting MEASURE.

When you restart the measurements, the bar for each new oscillator will be all the way to the left since the system does not yet have enough data to show the new

oscillator's performance. After 1 to 2 hours, the bars will extend to show the performance of the oscillator. If you deleted a channel name, the bar for that channel will no longer appear on the measurement screen.

Whenever the system prints plots of the last 24 hours the plots will include only data from the new oscillator. This feature was included so that the 24 hour plots will never "mix" data from two different oscillators. For example, if you change oscillators at 1800 UTC and press <CONTROL-P> at 2300 UTC, the plots will include data from the last 5 hours only. However, the data from the oscillator that you replaced is still on the disk and can be retrieved at any time. For information about retrieving data, please see Chapter 5.

QUICK CHECK OF SYSTEM

When the system is taking measurements it requires only a small amount of operator attention. However, make a habit of regularly visually checking the system to make sure that everything is working properly. Once you know what to look for, the visual check takes just a few seconds. The things to look for are summarized below:

- *Is the printer ON, ON-LINE, and loaded with paper?*

The printer must be on, on-line, and loaded with paper in order for the system to plot automatically.

- *Are the counter readings (displayed on screen) stable?*

If the counter readings are not stable, they are printed in red. If this only happens on channel 1, it could mean that the GPS receiver has failed or temporarily lost the signal. If it happens on all channels, it could indicate a problem with your primary oscillator. If it happens on an individual channel other than channel 1, it could mean that the oscillator you are calibrating is defective or is producing a signal with a frequency offset larger than 1×10^{-5} with respect to its nameplate frequency. If you have just turned the oscillator on, this situation may occur for a few minutes while the oscillator warms up.

- *Is any counter channel showing NO SIGNAL?*

If you see the NO SIGNAL message printed in place of the counter readings, it means that the system cannot detect whether an oscillator is connected to the channel. Check the cable between the oscillator and the counter (it could be loose or disconnected). Or, the signal may be too weak to drive the counter. If you think this is the case, please contact NIST for assistance.

- *Are the oscillators performing as well as usual?*

A quick glance at the bar graph will indicate whether your oscillators are working the way they should. For example, if a bar is shorter than usual, you may have a problem with one of your oscillators or your GPS reception (if GPS is used as the reference on that channel). Contact NIST if you need assistance.

Chapter 4

Analyzing Data

This chapter describes the powerful data analysis features of the FMAS. These features are used to analyze the data collected by the system. Some of the features are listed below:

- Up to 3600 data points can be included on one graph (150 days of data).
- Full-page, high-quality plots are printed.
- Data from all five channels can be loaded at one time, so that you can quickly switch back and forth between channels.
- Outlying data points can be removed.
- Phase steps can be removed.
- Any segment of data can be plotted.
- The scale of the y-axis can be increased.
- The plot title can be changed.
- The position of the x-axis can be changed.
- All data can be viewed on screen or printed.
- Data statistics can be viewed or printed.

In this chapter, we'll illustrate each analysis feature using demonstration data that is provided with the FMAS. The next section tells how to load the demonstration data from disk.

Before beginning, you need to have the MAIN MENU on your monitor. If the system is currently taking readings, press <CONTROL-X> to get the MAIN MENU to appear.

Loading the Demonstration Data from Disk

To load the demonstration data, select LOAD from the MAIN MENU. The LOAD MENU appears as shown below:

```
ALL Select Demo Quit
Load all data
```

The LOAD MENU lists four items: ALL, SELECT, DEMO, and QUIT. ALL loads all the data on the hard disk. SELECT loads selected data from the hard disk. QUIT returns you to the MAIN MENU. Pressing <ESC> does the same thing as selecting QUIT. You may try this now. Select QUIT (or press <ESC>) to return to the MAIN MENU. Then return to the LOAD MENU by selecting LOAD again.

ALL and SELECT are very useful, and are described in detail in Chapter 5. They are used when you use your own data with the system. Right now, however, select DEMO to load the demonstration data. The demonstration data are included with the system and will be used throughout this chapter.

When you select DEMO, the counter in the center of the screen counts each data point as it is loaded. There are 96 demonstration data points. When all 96 points are read from the disk, the system returns to the MAIN MENU.

The next section tells how to view the data on screen.

Viewing the Data on Screen

Select VIEW from the MAIN MENU. The VIEW MENU appears as shown below:

```
1[One] 2[Two] 3[Three] 4[Four] 5[Five] Copy Quit
View channel one data
```

QUIT returns you to the MAIN MENU. Pressing <ESC> does the same thing as selecting QUIT.

1[ONE], 2[TWO], 3[THREE], 4[FOUR], and 5[FIVE] refer to the time interval counter channels. 1[ONE] refers to channel 1, 2[TWO] to channel 2, and so on. When you select a channel, data from that channel appear on screen as shown in Figure 4.1. The measurement channel title is listed above the data. The column headings are:

Number is the number of the data point.

MJD is the Modified Julian Date when the data point was recorded.

DATE is the date when the data point was recorded.

UTC is the Coordinated Universal Time when the data point was recorded.

Phase Shift is the oscillator phase shift (in nanoseconds) since the last data point was recorded.

1[One] 2[Two] 3[Three] 4[Four] 5[Five] Copy Quit				
View channel one data				
CHANNEL 1 - Cesium vs GPS				
Number	MJD	Date	UTC	Phase Shift (ns)
1	49613	1994-09-18	17:00	-54.54
2	49613	1994-09-18	18:00	-4.55
3	49613	1994-09-18	19:00	-46.24
4	49613	1994-09-18	20:00	+69.90
5	49613	1994-09-18	21:00	-125.37
6	49613	1994-09-18	22:00	+11.40
7	49613	1994-09-18	23:00	+19.71
8	49614	1994-09-19	00:00	-5.03
9	49614	1994-09-19	01:00	+62.48
10	49614	1994-09-19	02:00	+33.73
11	49614	1994-09-19	03:00	-16.61
12	49614	1994-09-19	04:00	-46.78
13	49614	1994-09-19	05:00	+22.74
14	49614	1994-09-19	06:00	+20.84
15	49614	1994-09-19	07:00	-88.59

Press <PgUp> to move up 15 values, or <PgDn> to move down 15 values.
 Press <HOME> to go to the top, or <END> to go to the end of the list.
 Press <ESC> to return to the menu.

Figure 4.1. The view data screen.

One screen can hold up to 15 data points. To move to the next screen, press <PgDn>. To move to the previous screen, press <PgUp>. To move to the first data point, press <HOME>. To move to the last data point, press <END>. Notice that you cannot move forward beyond the last data point, or backwards before the first data point.

At any time, you can print the data points on screen by pressing <ESC> and then selecting COPY. When you are finished using VIEW, select QUIT to return to the MAIN MENU.

Printing the Data

COPY is used to print out all of the data for any channel. A total of 50 data points are printed per page. If you want this printout, select COPY and then select a channel number (1, 2, 3, 4, or 5). You can calculate the length of the printout by rounding ahead to the nearest number that is evenly divisible by 50, and then dividing by 50. In this case, we have 96 data points. The first number greater than 96 that is evenly divisible by 50 is 100. Since 100 divided by 50 is 2, the printout will be 2 pages long.

When you finish using COPY, select QUIT to return to the MAIN MENU.

Viewing the Data Statistics

STATISTICS is used to view or print statistics for any channel. When you select a channel, you'll see a screen similar to Figure 4.2.

The statistics for the selected oscillator (in this case, a cesium) are printed inside the box. By looking at the box, we can see that 96 data points were recorded. The first data point was recorded on September 18, 1994 at 1700 UTC, and the last data point was recorded on September 22, 1994 at 1600 UTC, a period of 3 days, 23 hours, and 0 minutes. During this time, the data range was 382.27 ns. The frequency offset of the cesium was $+3.47 \times 10^{-13}$, and the correlation coefficient (r) is +0.45 (see Appendix A for more information about these values).

By selecting 2[TWO], 3[THREE], 4[FOUR], or 5[FIVE] you can look at the statistics for the other channels. Whenever you want to print the statistics shown on screen, select COPY.


```

      2[Two] 3[Three] 4[Four] 5[Five] Filter Copy Quit
Show statistics for channel one

                                CHANNEL 1 - Cesium vs GPS

Number of Data Points:  96                                Data Filter:  OFF
First data point:      1994-09-18 (MJD 49613) at 17:00 UTC.
Last data point:      1994-09-22 (MJD 49617) at 16:00 UTC.

Measurement Period =           3 Day(s) 23 Hour(s)

Data Range (nanoseconds) is:           382.27

Frequency Offset is:                   +3.47E-13
Correlation Coefficient (r) is:         +0.45

```

Figure 4.2. The view statistics screen.

The menu includes another item called FILTER that is used to smooth the data by removing phase steps and bad data points. When you select FILTER, the data are filtered and the statistics will change. If the filter is in use, the data filter indicator in the statistics box will be set to ON. FILTER will not do anything to data that is already smooth, but you may find it useful in certain instances. In the next section, we'll plot the data and look at some filtered plots. You will then have a better understanding of what FILTER does.

When you finish using STATISTICS, select QUIT to return to the MAIN MENU.

Plotting the Data

Select PLOT when you want to plot the data. The PLOT MENU appears listing eight items. The first five items are 1[ONE], 2[TWO], 3[THREE], 4[FOUR], and 5[FIVE]. These items work just as they did in the VIEW, COPY, and STATISTICS menus, except now they plot the data instead of performing the other

functions. 1[ONE] plots the data for channel 1, 2[TWO] plots the data for channel 2, and so on. QUIT also works like before. Selecting QUIT returns you to the MAIN MENU.

The other menu items, EDIT and COPY, won't work unless a plot is on screen. If you try using them without a plot on screen, you'll receive a message on the second line of the menu telling you that there is no plot to EDIT or no plot to COPY.

We'll discuss EDIT a little later. As you can probably guess, COPY lets you print out a copy of the plot currently on screen. In order for COPY to work, the printer must be connected to the computer, turned on, on-line, and loaded with paper. If you cannot get a plot to print out, please check the status of your printer. If the printer is working properly, the menu disappears when you select COPY and a larger version of the plot appears on screen before it is copied to the printer. After the plot is printed, the menu and the original plot return to the screen.

For now, let's get started by selecting 1[ONE] to plot the demonstration data for channel 1. The plot is shown in Figure 4.3.




Figure 4.3. Demonstration data for channel 1.

Notice that the plot in Figure 4.3 contains a thick line showing the slope of the data obtained by doing a linear least squares curve fit. This line is printed on channel 1 only.

The format of the plot is identical to the daily plots printed by the system. The title of the plot is the title that you entered using NAMES from the MAIN MENU. You may change the title if you wish (we'll discuss that feature later on).

The y-axis of the plot is labeled MICROSECONDS. The above plot has a y-axis scaled for 0.4 μs (from plus to minus 0.2 μs).

There are three lines of text below the plot. The top line shows the time and date when the first and last data points were recorded. The second line shows the tick mark interval and the channel from which the data was recorded. Figure 4.3 uses data recorded from Channel 1 of the time interval counter. The spacing between any two ticks on the x-axis represents 2 hours. The tick mark interval can range from 1 to 96 hours, depending upon the number of data points. The third line below the x-axis shows the number of data points plotted (96), the frequency offset of the oscillator ($+3.47 \times 10^{-13}$), and the correlation coefficient (r) for the data (+0.45). This is the same information printed on the statistics screen in Figure 4.2.

Editing the Plot on Screen

The plot we just looked at is a good example of the results you can obtain when everything is working properly. The plot compared a GPS signal to a high-quality cesium oscillator. The fluctuations are quite small (tens of nanoseconds). Most of the time, the FMAS will produce good looking plots like this one. However, in some cases you may see a plot with a phase step or some outlying data points. To see an example of this type of plot, select 4[FOUR] to plot the demonstration data for channel 4. The plot shown in Figure 4.4 will appear on your monitor. You can see that most of the data in Figure 4.4 are probably good. However, the bad data prevent us from seeing the good data.

We can solve this problem by using FILTER. To do so, select EDIT to get to the EDIT MENU and then select FILTER. When FILTER is selected, the data are replotted and the outlying data points and the phase step are automatically removed. The filtered plot for channel 4 is shown in Figure 4.5.




Figure 4.4. A plot with a phase step and some outliers.




Figure 4.5. The Figure 4.4 graph after selecting FILTER.

With the bad data removed, the filtered plot shows the true performance of the rubidium oscillator. Notice that the filtered plot has only 91 points, as opposed to 96 on the original plot. Five data points were filtered out. Also, the new plot shows a frequency offset of $+1.04 \times 10^{-10}$. This is the true frequency offset of the oscillator (it formerly was $+6.47 \times 10^{-10}$). Also, removing the bad data reduced the scale of the plot from 4000 to 40 μ s.

FILTER works well with most types of data. However, if the plot has so much bad data that it is hard to separate bad data from good data, FILTER may not help. Also, FILTER won't do anything to data that have a range of less than 1 μ s, or data that are perfectly linear (r equals 1 or -1). If this is the case, you'll get the following message on the second line of the menu, "The system cannot FILTER this data".

FILTER works best with plots where most of the data are good (like the plot we just looked at), and where only a few points need to be filtered. If this is the case, FILTER removes the bad data and shows you the true performance of your oscillator. If the plot is filtered, you can tell by looking at the channel label below the x-axis. For example, you'll see the words "Channel 4" for a regular plot, and "Channel 4 Filtered" for a filtered plot.

As you probably have noticed, there are a number of other editing features on the EDIT menu (a total of eleven). We've already covered four of them. To review, COPY lets you make copies of the plot on screen. FILTER removes the bad data from the plot on screen. RESET restores the plot on screen to its original status. QUIT returns you to the PLOT MENU. The other seven menu items let you change, or edit, the plot on screen. These items are explained on the following pages. We'll begin with a discussion of LO and HI. Before trying LO and HI, select RESET to get the original plot back on screen.

Removing Data Points from the Plot

LO and HI remove data points from the plot. LO removes the lowest (or minimum) value from the plot. HI removes the highest (or maximum) value from the plot. To see how LO and HI work, select LO. A new plot appears on screen with the minimum data point removed. However, the plot still contains another "spike" caused by the highest valued data point. Select HI to remove this data point. You will now see the plot shown in Figure 4.6.




Figure 4.6. The figure 4.4 plot with the outliers removed.

Notice that the plot labels now show that the plot has 94 data points (two data points were removed).

LO and HI are different from FILTER. FILTER adjusts outlying data points and phase steps by comparing good data to bad data, and can remove multiple data points. LO and HI can remove only one data point at a time, and this point has to be either the minimum or maximum point. However, as with FILTER, the data points removed by LO and HI are not deleted from the disk or from memory, they are just no longer plotted. If you select RESET the original plot returns to the screen. Also, if you try using LO and HI on a filtered plot, you'll get an error message on the second line of the menu.

When you are finished using LO and HI, select RESET to restore the plot on screen to its original status. The next section tells how to plot a selected portion of the plot on screen.

Plotting a Portion of the Plot

PORTION lets you plot a selected portion of the plot on screen. The following example tells how to plot just the first 10 (out of 96) data points, using the data from Channel 1.

Select PORTION. You'll be prompted (on the second line of the menu) to enter the numbers for the first and last data points that you want to plot. Enter 1 for the first data point (and press <ENTER>), and enter 10 for the last data point (and press <ENTER>). A new plot appears containing only the first 10 data points, as shown in Figure 4.7. Notice that the data point label below the plot has changed from 96 to 10.




Figure 4.7. Plotting a portion of the channel 1 demonstration data.

PORTION is useful for "blowing up" a selected segment of the plot. Try plotting other segments of the data. You have 96 data points in memory. The first data point can be any number from 1 to 95. The second data point can be any number from 2 to 96, as long as it is greater than the number for the first data point. The system won't accept numbers outside these ranges.

You can use FILTER to adjust any data portion that is plotted on screen. You can also get a copy of any portion by selecting COPY. You can't, however, use PORTION to plot a portion of a filtered plot. If you try this, you'll receive an error message on the second line of the menu. When you finish using PORTION, select RESET to restore the plot to its original status. You may now go on to the next section.

YAXIS, XAXIS, and MAGNIFY

YAXIS lets you change the scale of the y-axis. The current scale of the y-axis is 0.4 μs (from plus to minus 0.2 μs). To change the scale of the y-axis, select YAXIS. A prompt appears on the second line of the menu asking you to enter the new scale (in microseconds). Type 4 and press <ENTER>. The data are replotted, and the scale is increased from 0.4 to 4 μs . You can see that the variations in the Channel 1 data are very small. When a 4 μs scale is used, the data is almost a straight line.

You may find changing the scale to be useful. You can even set all your plots to the same scale so you can lay them side by side for comparison. The scale can only be increased, however. If you enter a scale smaller than the original scale, the data are still plotted with the original scale. At any point, you can return to the original plot by selecting RESET.

One problem with using YAXIS is that the data are often plotted on top of the x-axis when you increase the scale. This makes it difficult to see the data. XAXIS lets you change the position of the x-axis so that you can see the data better. To illustrate this, select XAXIS now. The data are now plotted with the x-axis on the bottom of the plot instead of in the middle, and are now easier to see. This plot is shown in Figure 4.8. You can use COPY to make copies of plots with the x-axis at the bottom position.

XAXIS is a toggle function. If you select it again, the x-axis returns to the middle position.

MAGNIFY is another useful feature. When you select MAGNIFY, the data are plotted using the range of the data as the y-axis scale. This allows you to "magnify" the data and show as much detail as possible. Like XAXIS, MAGNIFY is a toggle function. If you select it again, the y-axis scale returns to its original value.




Figure 4.8. A plot with the x-axis moved to the bottom position.

Keep in mind that YAXIS, XAXIS, and MAGNIFY change only the appearance of the data as plotted on screen. The actual data (and the frequency offset) remain the same.

Changing the Title of the Plot

The last item on the EDIT MENU that we need to discuss is TITLE. TITLE lets you change the title of the plot. The new title appears on the plots on screen, and also on the plots that are printed.

To illustrate this, select TITLE now. A prompt appears on the second line of the menu asking you to enter a new plot title. Type in a new title and press <ENTER>. The new title replaces the old one in the plot on screen, and in any copies of the plot that you make. If you select RESET, the original plot is put back on screen along with its original title.

When you finish editing plots, select QUIT to go back to the PLOT MENU. Select QUIT again to go back to the MAIN MENU.

Chapter 5

Retrieving Your Data

In Chapter 4, we explained how to use the data analysis features of the FMAS. We used the demonstration data in all of our examples. Of course, you can use the same features on your own data. In this chapter, we'll explain how to retrieve the data collected by your system.

To retrieve your data, stop the system from making measurements (if necessary) by pressing <CONTROL-X>. Then, select LOAD from the MAIN MENU. The LOAD MENU appears, listing four items:

As usual, QUIT returns you to the MAIN MENU. Pressing <ESC> does the

```
ALL Select Demo Quit
Load all data
```

same thing as selecting QUIT. DEMO loads the demonstration data that we used in Chapter 4. The other items (ALL and SELECT) are explained below.

Retrieving All of the Data

ALL loads all the data that the system has available. When you select ALL, this message appears on the second line of the menu:

Are you sure?

If you decide **not** to use ALL, press <ENTER>. Otherwise, enter a "Y" or "y" and press <ENTER>. The counter in the center of the screen counts each data point as it is loaded. When all the data points are loaded, the system returns to the MAIN MENU.

The FMAS has a maximum storage capacity of 3600 data points. If more than 3600 data points are available, ALL stops loading data when it gets to 3600.




Figure 5.1. The file load window.

Retrieving Some of the Data

SELECT loads selected data from the disk. When you choose SELECT, a box appears that allows you to "tune in" the data that you want to see. You can load any of the data listed in the box. For example, if data are available from 09/18/96 to 10/07/96, you can load data recorded within that time period. A picture of the "tuning box" is shown in Figure 5.1.

Use the <UP ARROW> and <DOWN ARROW> keys to scroll through the lists of available dates one day at a time. Use <PgUp> to move back 30 days. Use <PgDn> to move ahead 30 days. If you decide to exit the SELECT function and go back to the LOAD MENU, press <ESC>.

When you scroll through the dates in the tuning box, notice that three dates are always shown. The currently selected date is highlighted. The two dates adjacent to the currently selected date are shaded. Locate the start date for the data you want to load, and press <ENTER>. Then, locate the end date for the data you want to load and press <ENTER> again. After a few seconds, the system will load the data you selected. A counter will briefly appear in the center of the screen as the system loads each data point. The system then returns to the MAIN MENU.

Chapter 6

Utilities

The FMAS has several utility features that you should know about. To use these features, select UTILITIES from the MAIN MENU. You'll see the menu shown below:

```
ShortTerm History GPS CalibrateTIC Quit
Plot the short term stability of any channel in use
```

As always, select QUIT or press <ESC> to leave this menu and go back to the MAIN MENU. The other four menu items, SHORTTERM, HISTORY, GPS, and CALIBRATETIC are discussed below.

ShortTerm

This feature allows you to look at the short-term performance of an oscillator by taking readings for a period from 2 to 3600 s. You can then produce graphs that show the frequency offset and stability of the oscillator. When you select SHORTTERM, the menu shown below appears on your monitor:

```
PickChannel Length Go Offset Stability File Meter Copy Quit
Begin measurements
```

Three items on the SHORTTERM MENU are used to set up a short term measurement run. These items are:

PICKCHANNEL - This item is used to select the channel that you want to use for the measurement. The default channel number is 1. However, you can take measurements from any of the five channels.

LENGTH - This item allows you to select the length of the measurement run in seconds. The default length is 10 s, but you can make measurements for any period from 2 to 3600 s.

GO - After you have selected a channel and the length of the measurement run, select GO to start the measurements.

When the system begins taking measurements, the message "Measurement Run in Progress" will appear. The system will complete a new data point every second. When the measurements are finished, this message will change to "Measurement Run Completed". The screen will then look similar to the screen shown in Figure 6.1.

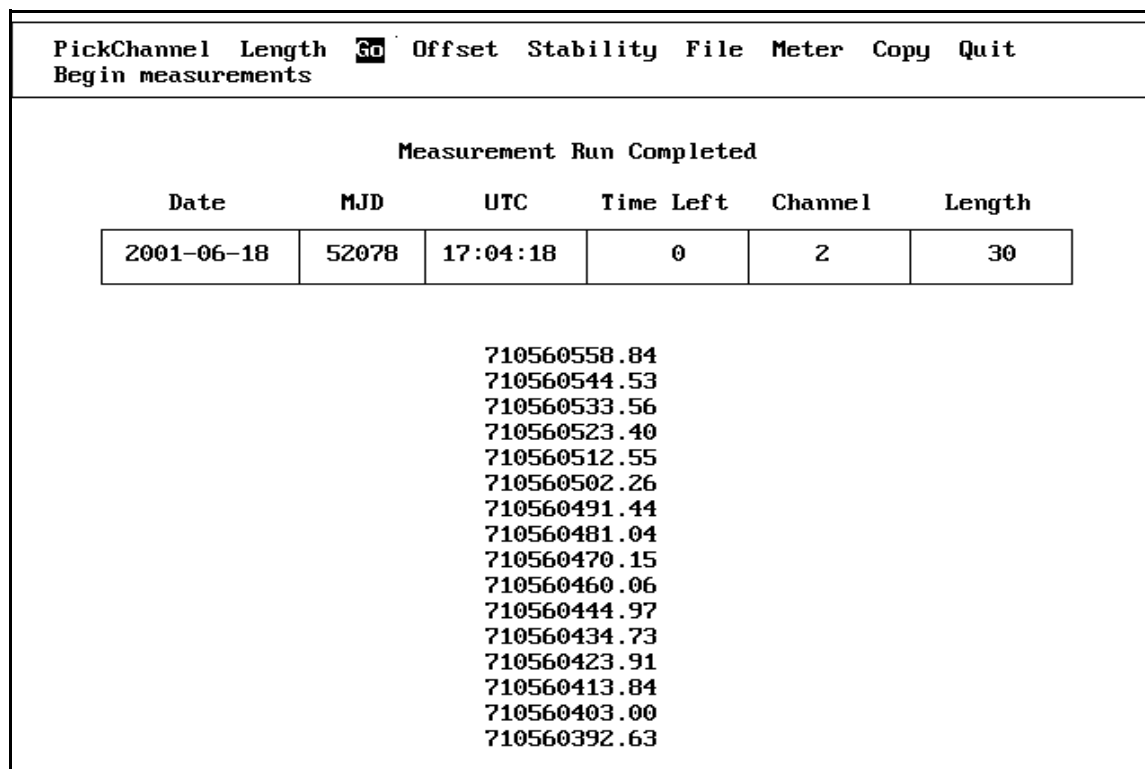


Figure 6.1. The short term measurement screen.

At this point you can graph the short term frequency offset of the oscillator under test by selecting OFFSET. A sample graph is shown in Figure 6.2. The plot shows each individual data point recorded during the measurement run. It shows the start time for the measurement run, the duration of the run, the number of data points

plotted, and the frequency offset and correlation coefficient (r) of the data. You can print a copy of the plot on screen by selecting COPY from the menu.

You can also plot the stability of the oscillator under test by selecting STABILITY. A stability plot is also shown in Figure 6.3. The technique used to estimate stability is the Allan deviation. Unlike an uncertainty plot, the stability plot doesn't show how close the oscillator is to the correct frequency. Instead, it shows how closely the oscillator produces the same frequency over a time interval. By doing so, it can tell us a lot about the quality of the oscillator under test.




Figure 6.2. A short term frequency offset graph.

To illustrate the difference between frequency offset and stability, consider a 10 MHz oscillator that is high in frequency by 1 part in 10 million (1×10^{-7}). This means that the oscillator is actually high in frequency by 1 Hz, and is producing 10 000 001 Hz instead of 10 000 000 Hz. If we measure this same frequency offset day after day, we know that the oscillator is stable enough to produce frequency closer to the correct value. In this case, we should set the oscillator closer to its correct frequency to take advantage of its full potential. A stability plot provides a quick estimate of the inherent stability of an oscillator, even if the oscillator has a large frequency offset.




Figure 6.3. A short term stability graph.

Stability graphs use a logarithmic scale. The values along the x-axis (called tau values) represent the length of the averaging period in seconds. Each division represents an averaging period 10 times longer than the previous division. For example, a tau value of 2 represents an averaging period of 100 s (10^2). A tau value of 3 represents an averaging period of 1000 s (10^3). The data points (rectangles) represent measurement periods whose length is a power of 2. For example:

2^0	=	1 s
2^1	=	2 s
2^2	=	4 s
2^3	=	8 s
2^4	=	16 s
2^5	=	32 s
2^6	=	64 s
2^7	=	128 s
2^8	=	256 s
2^9	=	512 s
2^{10}	=	1024 s

We need at least two samples to estimate stability. This means that to estimate stability at 4 s, we need to measure for at least 8 s (two 4 s samples). To estimate stability at 1024 s we must measure for at least 2048 s (two 1024 s samples). Keep this in mind when creating stability graphs.

The values along the y-axis represent the result of the measurement. A value of -10 means that the oscillator has a stability of 1×10^{-10} . This stability value should never be confused with the frequency offset value. For example, an oscillator with a frequency offset of 1×10^{-8} might be stable to 1×10^{-12} at 1000 s. This means that even though the oscillator is producing a frequency with a large offset, the frequency offset isn't changing very much.

A typical stability graph (such as Figure 6.3) shows the stability values improving as the measurement period increases in length. During this period, phase noise from the oscillator, the reference, and the measurement system is being removed by averaging. This improvement continues until the oscillator noise floor is reached, and no further gains are made by averaging. When the noise floor is reached, the stability values level out, and then eventually start to get worse. Most precision oscillators reach their noise floor at an averaging period of about 1000 s. Obviously, we want the measurement system and the reference to have lower noise floors than the device under test for the selected averaging period. For this reason, stability measurements should be made only on channels 2 through 5 using your primary oscillator as the reference. Although GPS has excellent long term stability, short term stability graphs tend to be dominated by GPS phase noise, since it might take hours or days of averaging for GPS to reach its noise floor.

Keep in mind that the FMAS is designed for long term measurements. Its noise floor is too high to accurately measure the stability of the best oscillators at short averaging times (such as 1 s). However, at longer averaging times it can measure the stability of nearly any oscillator, since its noise floor is near 1×10^{-13} at 100 s, and near 1×10^{-14} at 1000 s.

Retrieving Short Term Measurement Data

Each time a short term measurement run is completed, the data are automatically saved. You can retrieve these files and review the graphs at a later time. To do so, select FILE. A menu appears listing the available file names. After you select the file you want, it is automatically loaded into memory. You can then use OFFSET and STABILITY to plot the data, and COPY to copy the plots to the printer.

The first five digits of the short-term filenames represent the Modified Julian Date (MJD). This is followed by a hyphen. The next two digits represent the UTC hour when the data were recorded. This is followed by a period. The next two digits represent the UTC minute when the data were recorded. And finally, the last digit in the filename is the number of the channel from which the data were taken.

To illustrate this, consider the filename "49629-16.293". This file was recorded on MJD 49629 at 16 hours and 29 minutes UTC. The data were taken from channel 3.

When you are finished making short-term measurements, you can return to the UTILITIES MENU by selecting QUIT or by pressing <ESC>.

Adjusting an Oscillator Using METER

The FMAS also includes a "frequency meter" that is a useful tool for adjusting oscillators. To use this feature, select METER from the UTILITIES MENU. The frequency meter will appear on screen.

The frequency meter estimates and displays the current frequency offset of the oscillator. The possible frequency offset values range from 1×10^{-7} to 1×10^{-10} . The yellow line on the meter indicates the current value. By adjusting the oscillator, you can try to get the yellow line as close to the center of the meter (1×10^{-10}) as possible.

METER is mostly useful for adjusting quartz oscillators, or rubidium oscillators that are well off in frequency. Keep in mind that adjustments must be made slowly and carefully, and oscillators will not always "stay" at the frequency where you set them. Also, oscillators should have gone through their full warm-up period before you even try to make adjustments.

When you finish using METER, press any key to return to the menu.

History

HISTORY is a powerful feature of the FMAS that allows you to look at the performance history of your primary oscillator. HISTORY produces plots similar to those that you receive as part of your monthly calibration report (see Appendix B).

Each day, the FMAS records the frequency offset of the primary oscillator over the last 24 hours. This information is stored, and one file is kept for each month that the system has been running. For example, if your system has been running for three months, you will have three files. By using HISTORY, you can plot the information in these files, and see how your primary oscillator has performed over past months.

When you select HISTORY, a menu appears with three items; PLOT, COPY, and QUIT. QUIT and COPY work as explained earlier. QUIT returns you to the previous menu (the UTILITIES MENU in this case), and COPY prints a copy of the plot on screen if one is available. PLOT produces a list of all of the available frequency files. The files are named using the month and year when they were created. For example, "10-1996" is the file for October 1996.

The menu sorts the files so that the most recent file is at the top of the list. The most recent file is already highlighted. To select this file, simply press <ENTER>. To select another file, use <DOWN ARROW> to move down, <UP ARROW> to move up, and so on. When the file you want is highlighted, press <ENTER>. You will then see a plot similar to Figure 6.4:




Figure 6.4. A sample HISTORY graph.

Figure 6.4 shows the performance of a primary frequency standard (cesium oscillator) during the month of October 1996. Each marker represents the performance of the cesium oscillator for one day. The x-axis shows the days of the month. Each tick mark represents one day. If no data are recorded on a given day (for example, if the system is turned off), no marker will be present.

The y-axis shows the frequency offset of the primary oscillator. The middle of the y-axis represents a frequency offset of $\pm 1 \times 10^{-13}$ or better. In this example, the frequency offset of the primary oscillator was better than $\pm 1 \times 10^{-12}$ for all but two days of the month.

The history plot can accommodate just about any kind of oscillator. If an oscillator is high in frequency, the markers will appear in the top half of the plot. Conversely, if an oscillator is low in frequency, the markers will appear in the bottom half of the plot. The plot shows values as large as $\pm 1 \times 10^{-8}$. If an oscillator is high or low in frequency by more than that, the markers will appear along the top or bottom line of the plot.

The important thing to look for on the history plots is stability. For example, if your primary oscillator is a low cost quartz unit, it might not be capable of performance much better than 1×10^{-9} . However, if your system and oscillator are operated properly, your measurements should be able to reproduce this value day after day, and your history plot should show a nearly straight line of markers.

GPS

The GPS menu item shows the current status of GPS reception. When you select this option you'll see a screen similar to the one shown in Figure 6.5. This screen was included as a diagnostic tool that can assist in troubleshooting GPS reception problems. NIST will routinely check your GPS reception by modem. If you think you have a problem with your data (or if you are interested), you might want to check it also. Please contact NIST when you have questions, need assistance, or notice any condition that indicates a problem. The text below Figure 6.5 annotates each value shown on the GPS status screen.

ShortTerm History GPS CalibrateTIC Quit Show status of Global Positioning System (GPS) receiver					
GPS Tracking Status - Press <ESC> to Exit					
Tracking Status Byte		8	UTC Date (YYYY-MM-DD)		2001-06-18
Satellites (Visible/Track)		8/8	UTC Time (HH:MM:SS)		17:55:03
Timing Stability (ns)		38	Pulse Status		ON
PRN Number	Health (S/T)	Elevation	Offset (ns)	Carrier to Noise	
18	H/ON	25	5	47	
15	H/ON	27	-4	46	
30	H/ON	17	-14	49	
17	H/ON	46	1	49	
6	H/ON	80	-1	48	
10	H/ON	37	-1	49	
26	H/ON	38	4	49	
23	H/ON	41	10	48	

Figure 6.5. The GPS status screen.

Tracking Status Byte - This byte provides diagnostic information about the receiver. It will normally have a value of 8, which means that more than three satellites are being successfully received, the almanac is good, and the receiver is in position hold mode. For brief periods of time, a value other than 8 might be shown (usually a 32), but it should normally be at 8 if the receiver is working properly.

Satellites (Visible/Tracked) - These two values show the number of satellites within view of the receiver and the number of satellites actually being tracked. Often, the two values will be equal. However, as many as 12 or 13 GPS satellites could theoretically be in view, and the receiver can track a maximum of 8 satellites at one time.

Timing Stability (ns) - This is an internal estimate made by the receiver of the stability (standard deviation) of the 1 Hz output signal. It will normally be a value smaller than 50. Keep in mind that this is not a measured value, but simply a rough estimate provided for diagnostic purposes. The actual stability of the receiver should be much better than this estimate.

UTC Date and Time - These values should be correct, but the time might be off by a fraction of second due to delays in the screen updates. However, if you see something unusual (such as the wrong date) it indicates a receiver problem. If this is the case, please call NIST for assistance.

Pulse Status - This indicates whether the 1 Hz output on the receiver is turned on or off. It should always be set to ON if the receiver has completed the acquisition process and is working properly.

The bottom part of the GPS status screen shows information about the satellites being tracked (the receiver tracks as many as eight satellites). New satellites are constantly coming into view and this information is always being updated. Five columns of information are provided for each satellite:

PRN Number - The pseudo random noise (PRN) number of the satellite being tracked. Each satellite in the GPS constellation is identified by a unique PRN number. The GPS constellation was designed to use 24 satellites, but more than 24 may be in orbit at a given time. For this reason, the possible PRN numbers range from 1 to 32.

Health (S/T) - This box indicates the “health” of the satellite and whether or not it is being used by the receiver to produce the timing pulse. An H indicates that the satellite is healthy, a U indicates that it is not healthy. An ON indicates that the satellite is being used in the timing solution, OFF indicates that it is not. The box will normally contain this notation, H/ON, which means that the satellite is healthy and is being used in the timing solution. On rare occasions, satellites do broadcast bad data and are marked as unhealthy. However, the receiver is designed to automatically remove bad satellites from the timing solution. Therefore, the probability of a U/ON message being shown is extremely low.

Elevation - The vertical angle measured from the horizon up to the satellite. The receiver automatically selects those satellites that have the most favorable elevation angle. Once a satellite drops below the horizon (below 0° elevation), the receiver begins tracking another satellite in its place. GPS was designed so that at least four satellites are above the horizon (in view) at all times. If fewer than four satellites are in view, it could mean that your antenna has an obstructed view of the sky.

Offset (ns) - This value (in nanoseconds) represents the time offset of the satellite with respect to the average time offset of all of the satellites being tracked. Since all satellites are usually in very close agreement with each other, this value typically ranges from about -20 to +20 ns.

Carrier to Noise - The signal-to-noise ratio for the signal being tracked. A number of 45 or higher is considered to be very good, with 51 the maximum value. A number less than 40 indicates a possible reception problem, but a 0 could simply mean that the receiver is still processing data for the specified satellite and has yet to produce a value. If everything is working properly, you should see numbers of 40 or higher for all satellites. Lower numbers could mean that you are using a cable with excessive signal loss. In this case, either a lower loss cable should be used, or the length of the cable should be shortened. Please contact NIST if you need assistance.

When you finish viewing the GPS status box, press <ESC> to return to the UTILITIES MENU.

CalibrateTIC

The time interval counter inside the FMAS is automatically calibrated each time the system starts and is recalibrated every 24 hours while the system is running. If you select the CalibrateTIC menu item, it automatically calibrates the counter and displays the calibration values on screen. Although there is usually no reason to do this, it was included as a diagnostic tool that NIST can use to see whether the counter is working properly.

Chapter 7

Setup Software

The software used by the FMAS was installed at NIST before it was shipped to your site. We have included a setup program with the software called FMSETUP. You may never need to use FMSETUP, but we'll discuss it here in case you are interested.

To use the setup software, type FMSETUP from the DOS prompt and <ENTER>. You will see a screen similar to the one in Figure 7.1.

Pathname for Software	C:\NISTFMAN	
Pathname for Phase Data	C:\NISTFMAN\PHASDATA\	
Pathname for Demo Data	C:\NISTFMAN\DEMODATA\	
Pathname for Frequency Data	C:\NISTFMAN\FREQDATA\	
Pathname for Short Term Data	C:\NISTFMAN\SHORTERM\	
Plot Time (UTC hour, 0-23)	16	FREQUENCY MEASUREMENT SETUP Use the arrow keys to select items that you want to enter or change. When you finish, press <ESCAPE> to save, or <END> to exit without saving.
Is Tape Drive Present? (Y or N)	Y	
Is Counter Present? (Y or N)	Y	
Is Receiver Present? (Y or N)	Y	
Use Channel 1 Filter? (Y or N)	Y	
Use GPS on channels 2-5? (Y or N)	N	
Printer: C=Cut L=Laser R=Roll N=None	L	

Figure 7.1. The FMSETUP screen.

The screen contains 12 boxes of information. You can use the arrow keys to highlight any of these boxes and then change the information in the box. Be sure to **consult with NIST before changing any setting**. The 12 boxes of information are:

Pathname for Software - This is where the file FM2000.EXE (or a file with a similar name) is located. This executable file controls the entire system.

Pathname for Phase Data - This is where the system stores the data it records when you use MEASURE to make measurements. When the system saves data to disk each hour, the data are stored in this directory.

Pathname for Demo Data - This is where the system stores the demonstration data used in the examples in Chapter 4.

Pathname for Frequency Data - This is where the system stores the daily frequency offset values for the primary oscillator. These data can be retrieved with HISTORY (see Chapter 6) and are used by NIST to produce the monthly calibration reports.

Pathname for Short Term Data - This is where the system stores the data it records when you use SHORTTERM to make measurements.

Plot Time (UTC hour, 0-23) - This setting determines the UTC hour when the system produces the daily phase plot. This is normally set to 16. This means that the system will plot automatically at 1600 UTC every day. If you change this number to 17, the system will plot at 1700 UTC every day, and so on. You may set the plot time to any hour that you choose. However, you should not make a habit of changing the plot time. Once you find a setting that is convenient for you, we recommend that you leave the plot time at that setting.

Is Tape Drive Present? (Y or N) - This setting remains from earlier versions of the FMAS which used a hard drive/tape drive combination, and periodically backed up the contents of the hard drive to the tape drive. Since the FMAS now uses a solid state flash disk to store data, this value should always be set to "N".

Is Counter Present? (Y or N) - This setting specifies whether the system has a time interval counter present. Of course, the FMAS includes a time interval counter inside the computer system, so the normal setting is "Y". However, by changing the setting to "N", the software will run on a computer without a time interval counter installed. Obviously, measurements can't be made on a system without

a counter, but the plotting and data analysis features of the software will still be usable.

Is Receiver Present? (Y or N) - This setting specifies whether the system has a GPS receiver present. Since a GPS receiver is part of the standard FMAS equipment, the normal setting is "Y". If the setting is "N", most of the features of the system will still work. However, the time on screen will no longer be synchronized to GPS (it will come from the computer's clock), and the GPS option on the UTILITIES menu will be unusable. Of course, oscillator to oscillator comparisons will still be usable.

Use Channel 1 Filter? (Y or N) - The FMAS has a FILTER function (discussed in Chapter 4) that can remove phase steps and outliers from the data. This option allows you to automatically use this feature on the channel 1 data. This means that when the system plots, it will automatically produce graphs that have the filter turned on. The bars on the measurement screen will still indicate the performance of your primary oscillator with the filter turned off, so you can be alerted to possible problems. This feature was designed to only remove phase steps caused by a disturbance of the measurement system, and not phase changes that indicate a frequency change in your primary oscillator.

Use GPS on Channels 2-5? (Y or N) - This setting pertained to an earlier version of the FMAS software and is now ignored. GPS can now be used on channels 2 through 5 without changing the value here.

Printer - This setting tells the system what kind of printer is connected. The standard printer supplied with the FMAS is either a dot matrix with a cut sheet feeder (use the C setting), or an inkjet printer with a cut sheet feeder (use the L setting). Both types of printer allow each plot to be printed on an individual sheet. However, the system also supports a dot matrix using roll paper (R setting), or a laser printer (L setting). There is also a memory-resident printer driver that might need to be changed, so simply changing this setting might not be all that is required when changing printers. Contact NIST if you need assistance.

When you finish using FMSETUP, press <ESC> to save the changes you have just made and go back to DOS. To go back to DOS without saving any changes, press <END>.

Appendix A

An Introduction to Frequency Calibrations

Frequency is the rate of occurrence of a repetitive event. If T is the period of a repetitive event, then the frequency $f = 1/T$. The International System of Units (SI) states that the period should always be expressed in units of seconds (s), and the frequency should always be expressed in hertz (Hz). The frequency of electrical signals often is measured in units of kilohertz (kHz) or megahertz (MHz), where 1 kHz equals one thousand (10^3) cycles per second and 1 MHz equals one million (10^6) cycles per second.

Average frequency over a time interval can be measured very precisely. Time interval is one of the four basic standards of measurement (the others are length, mass, and temperature). Of these four basic standards, time interval (and frequency) can be measured with the most resolution and the least uncertainty. In some fields of metrology, one part per million (1×10^{-6}) is considered quite an accomplishment. In frequency metrology, measurements of one part per billion (1×10^{-9}) are routine, and even one part per trillion (1×10^{-12}) is commonplace.

Devices that produce a known frequency are called *frequency standards*. These devices must be calibrated so that they remain within the tolerance required by the user's application. Let's begin our discussion with an overview of frequency calibrations.

Overview of Frequency Measurements and Calibration

Frequency calibrations measure the performance of frequency standards. The frequency standard being calibrated is called the *device under test* (DUT). In most cases, the DUT is a *quartz*, *rubidium*, or *cesium* oscillator. In order to perform the calibration, the DUT must be compared to a *standard* or *reference*. The standard should outperform the DUT by a specified ratio in order for the calibration to be valid. This ratio is called the *test uncertainty ratio* (TUR). A TUR of 10:1 is preferred, but not always possible. If a smaller TUR is used (5:1, for example) then the calibration will take longer to perform.

Once the calibration is completed, the metrologist should be able to state how close the DUT's output is to its *nameplate frequency*. Often called the nominal

frequency, the nameplate frequency is labeled on the oscillator's output. For example, a DUT with an output labeled "5 MHz" is supposed to produce a 5 MHz frequency. The calibration measures the difference between the actual frequency and the nameplate frequency. This difference is called the *frequency offset*. There is a high probability that the frequency offset will stay within a certain range of values, called the *frequency uncertainty*. The user specifies an uncertainty requirement for the frequency offset that the DUT must meet or exceed. In many cases, users base their requirements on the specifications published by the manufacturer. In other cases, they may "relax" the requirements and use a less demanding specification. Once the DUT meets specifications, it has been successfully calibrated. If the DUT cannot meet specifications, it fails calibration and is repaired or removed from service.

The reference used for the calibration must be *traceable*. The International Organization for Standardization (ISO) definition for *traceability* is:

The property of the result of a measurement or the value of a standard whereby it can be related to stated references, usually national or international standards, through an unbroken chain of comparisons all having stated uncertainties [1].

In the United States, the "unbroken chain of comparisons" should trace back to the National Institute of Standards and Technology (NIST). In some fields of calibration, traceability is established by sending the standard to NIST (or to a NIST-traceable laboratory) for calibration, or by sending a set of reference materials (such as a set of artifact standards used for mass calibrations) to the user. Neither method is practical when making frequency calibrations. Oscillators are sensitive to changing environmental conditions and especially to being turned on and off. If an oscillator is calibrated and then turned off, the calibration could be invalid when the oscillator is turned back on. In addition, the vibrations and temperature changes encountered during shipment can also change the results. For these reasons, laboratories should always make their calibrations on-site.

Fortunately, we can use *transfer standards* to deliver a frequency reference from the national standard to the calibration laboratory. Transfer standards are devices that receive and process radio signals that provide frequency traceable to NIST. The radio signal is a link back to the national standard. Several signals are available, including NIST radio stations WWV, WWVH, and WWVB, and radionavigation signals from LORAN-C and GPS. Each signal delivers NIST traceability at a known level of uncertainty. The ability to use transfer standards is a tremendous advantage. It allows traceable calibrations to be made simultaneously at a number of sites as long as

each site is equipped with a radio receiver. It also eliminates the difficult and undesirable practice of moving frequency standards from one place to another.

Once a traceable transfer standard is in place, the next step is developing the technical procedure used to make the calibration. This procedure is called the *calibration method*. The method should be defined and documented by the laboratory, and ideally a measurement system that automates the procedure should be built. ISO/IEC Guide 17025, *General Requirements for the Competence of Testing and Calibration Laboratories*, states:

The laboratory shall use appropriate methods and procedures for all tests and/or calibrations within its scope. These include sampling, handling, transport, storage and preparation of items to be tested and/or calibrated, and, where appropriate, an estimation of the measurement uncertainty as well as statistical techniques for analysis of test and/or calibration data.

In addition, Guide 17025 states:

The laboratory shall use test and/or calibration methods, including methods for sampling, which meet the needs of the client and which are appropriate for the test and/or calibrations it undertakes When the client does not specify the method to be used, the laboratory shall select appropriate methods that have been published either in international, regional, or national standards, or by reputable technical organizations, or in relevant scientific texts or journals, or as specified by the manufacturer of the equipment [2, 3].

Calibration laboratories, therefore, should automate the frequency calibration process using a well documented and established method. This helps guarantee that each calibration will be of consistently high quality, and is essential if the laboratory is seeking ISO registration or laboratory accreditation.

Now that we've provided an overview of frequency calibrations, we'll take a more detailed look at the topics introduced. We'll begin by looking at the specifications used to describe a frequency calibration. Then, we'll discuss the various types of frequency standards and transfer standards. We'll conclude with a discussion of how the NIST Frequency Measurement and Analysis Service provides a complete solution to the frequency calibration problem.

The Specifications: Frequency Offset and Stability

In this section, we'll look at the two main specifications of a frequency calibration, frequency offset and stability. We'll define frequency offset and stability and show how they are measured. Keep in mind during this discussion that frequency offset is often referred to simply as *accuracy* (or *frequency accuracy*), and that stability is nearly the same thing as frequency uncertainty.

Frequency Offset

When we make a frequency calibration, our measurand is a DUT that is supposed to produce a specific frequency. For example, a DUT with an output labeled 5 MHz is supposed to produce a signal at a frequency of 5 MHz. Of course, the DUT will actually produce a frequency that isn't exactly 5 MHz. After we calibrate the DUT, we can state its frequency offset and the associated uncertainty.

Measuring the frequency offset of a DUT requires comparing it to a reference. This is normally done by making a *phase comparison* between the frequency produced by the DUT and the frequency produced by the reference. There are several calibration methods (described later) that allow us to do this. Once we know the amount of *phase deviation* and the *measurement period*, we can estimate the frequency offset of the DUT. The measurement period is the length of time over which phase comparisons are made. Frequency offset is estimated as follows, where Δt is the amount of phase deviation, and T is the measurement period:

$$f(\text{offset}) = \frac{-\Delta t}{T}$$

To illustrate, let's say that we measure +1 μs (microsecond) of phase deviation over a measurement period of 24 hours (h). The unit used for measurement period (h) must be converted to the unit used for phase deviation (s). The equation then becomes:

$$f(\text{offset}) = \frac{-\Delta t}{T} = \frac{1\mu\text{s}}{86,400,000,000\mu\text{s}} = -1.16 \times 10^{-11}$$

The smaller the frequency offset, the closer the DUT is to producing the same frequency as the reference. An oscillator that accumulates +1 μs of phase deviation/day has a frequency offset of about -1×10^{-11} with respect to the reference. Table A.1 lists the approximate offset values for some standard units of phase deviation and some standard measurement periods.

Table A.1. Frequency offset values for given amounts of phase deviation.

Measurement Period	Phase Deviation	Frequency Offset
1 s	1 ms	1.00×10^{-3}
1 s	1 μ s	1.00×10^{-6}
1 s	1 ns	1.00×10^{-9}
1 h	1 ms	2.78×10^{-7}
1 h	1 μ s	2.78×10^{-10}
1 h	1 ns	2.78×10^{-13}
1 day	1 ms	1.16×10^{-8}
1 day	1 μ s	1.16×10^{-11}
1 day	1 ns	1.16×10^{-14}

The frequency offset values in Table A.1 can be converted to units of frequency (Hz) if the nameplate frequency is known. To illustrate this, consider an oscillator with a nameplate frequency of 5 MHz that is high in frequency by 1.16×10^{-11} . To find the frequency offset in hertz, multiply the nameplate frequency by the dimensionless offset value:

$$(5 \times 10^6) (+1.16 \times 10^{-11}) = 5.80 \times 10^{-5} = +0.000\ 058\ 0\ \text{Hz}$$

The nameplate frequency is 5 MHz, or 5 000 000 Hz. Therefore, the actual frequency being produced by the frequency standard is:

$$5\ 000\ 000\ \text{Hz} + 0.000\ 058\ 0\ \text{Hz} = 5\ 000\ 000.000\ 058\ 0\ \text{Hz}$$

To do a complete uncertainty analysis, the measurement period must be long enough to insure that we are measuring the frequency offset of the DUT, and that other sources are not contributing a significant amount of uncertainty to the measurement. In other words, we must be sure that Δt is really a measure of only the DUT's phase deviation from the reference and is not being contaminated by noise from the reference or the measurement system. This is why a 10:1 TUR is desirable. If a 10:1 TUR is

maintained, many frequency calibration systems are capable of measuring a 1×10^{-10} frequency offset in 1 s [4].

Of course, a 10:1 TUR is not always possible, and the simple equation we gave for frequency offset is often too simple. When transfer standards such as LORAN-C or GPS receivers are used (discussed later), radio path noise contributes to the phase deviation. To get around this problem, a measurement period of at least 24 hours is normally used when calibrating frequency standards using a transfer standard. This

period is selected because changes in path delay between the source and receiver often have a cyclical variation that averages out over 24 hours. In addition to averaging, curve-fitting algorithms and other statistical methods are often used to improve the uncertainty estimate and to show the *confidence level* of the measurement [5].




Figure A.1. Simple phase comparison graphs.

better is being maintained. The bottom graph shows a small amount of phase noise, which could mean that a TUR of less than 10:1 exists and that some uncertainty is being contributed by the reference.

To summarize, frequency offset is the quantity of greatest interest to a calibration laboratory because it tells us how close a DUT is to its nameplate frequency. You will probably notice that the term *frequency accuracy* (or just

accuracy) often appears on oscillator specification sheets instead of the term frequency offset. Frequency accuracy and frequency offset are equivalent terms that refer to the result of a measurement at a given time. Frequency uncertainty indicates the limits (upper and lower) of the frequency offset. ISO defines uncertainty as a:

Parameter, associated with the result of a measurement, that characterizes the dispersion of values that could reasonably be attributed to the measurand [1].

In other words, the frequency uncertainty shows us the possible range of values (or limits) for the frequency offset. It is now standard practice to use a 2σ uncertainty test. This means that there is a 95.4 % probability that the frequency offset will stay within the stated range during the measurement period. The range of values is obtained by both adding the frequency uncertainty to and subtracting it from the average (or mean) frequency offset. Therefore, frequency uncertainty is sometimes stated with a “plus or minus” sign ($\pm 1 \times 10^{-12}$) to show the upper and lower limits of the offset. If the “ \pm ” symbol is omitted, it is still implied. The largest contributor to the frequency uncertainty is usually the stability of the device under test. Stability is the topic of the next section.

Stability

Before beginning our discussion of *stability*, we should mention an important distinction between frequency offset and stability. Frequency offset is a measure of how closely an oscillator produces its nameplate frequency, or how well an oscillator is adjusted. It doesn't tell us about the quality of an oscillator. For example, a stable oscillator that

needs adjustment might produce a frequency with a large offset. An unstable oscillator that is well adjusted might temporarily produce a frequency very close to its nameplate value.




Figure A.2. Comparison of unstable and stable frequencies.

Stability indicates how well an oscillator can produce the same frequency over a given period of time. It doesn't indicate whether the frequency is "right" or "wrong," only whether it stays the same. Also, the stability doesn't necessarily change when the frequency offset changes. You can adjust an oscillator and move its frequency either further away from or closer to its nameplate frequency without changing its stability at all. Figure A.2 illustrates this by displaying two oscillating signals that are of the same frequency between t_1 and t_2 . However, it's clear that signal 1 is unstable and is fluctuating in frequency between t_2 and t_3 .

Stability is defined as the statistical estimate of the frequency fluctuations of a signal over a given time interval. *Short-term* stability usually refers to fluctuations over intervals less than 100 s. *Long-term* stability can refer to measurement intervals greater than 100 s, but usually refers to periods longer than 1 day. A typical oscillator specification sheet might list stability estimates for intervals of 1, 10, 100, and 1000 s [6, 7].

Stability estimates can be made in the *frequency domain* or *time domain*, and statistical tools exist to easily convert from one domain to the other. Time domain estimates are more widely used, since time interval counters are often used to measure frequency. To estimate stability in the time domain, we must start with a set of frequency offset measurements y_i that consists of individual measurements, y_1, y_2, y_3 , and so on. Once this data set is obtained, we need to determine the dispersion or scatter of the y_i as a measure of oscillator noise. The larger the dispersion, or scatter, of the y_i , the greater the instability of the output signal of the oscillator.

Normally, classical statistics such as *standard deviation* (or *variance*, the square of the standard deviation) are used to measure dispersion. Variance is a measure of the numerical spread of a data set with respect to the average or mean value of the data set. However, variance works only with stationary data, where the results must be time-independent. This assumes the noise is *white*, meaning that its power is evenly distributed across the frequency band of the measurement. Oscillator data is usually nonstationary, since it contains time-dependent noise contributed by the frequency offset. For stationary data, the mean and variance will converge to particular values as the number of measurements increases. With nonstationary data, the mean and variance never converge to any particular values. Instead, we have a moving mean that changes each time we add a new measurement.

For these reasons, a non classical statistic is used to estimate stability in the time domain. This statistic is often called the *Allan variance*, but since it is actually the square root of the variance, its proper name is the *Allan deviation*. By recommendation of the Institute of Electrical and Electronics Engineers (IEEE), the Allan deviation is used by manufacturers of frequency standards as a standard specification for stability. The equation for the Allan deviation is:

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(M-1)} \sum_{i=1}^{M-1} (\bar{y}_{i+1} - \bar{y}_i)^2}$$

where M is the number of values in the y_i series, and the data are equally spaced in segments τ seconds long. Note that while classical deviation subtracts the mean from each measurement before squaring their summation, the Allan deviation subtracts the previous data point. Since stability is a measure of frequency fluctuations and not of frequency offset, successive data points are differenced to remove the time-dependent noise contributed by the frequency offset.

Table A.2. Using phase measurements to estimate stability (1 s data intervals).

Phase Measurements (ns)	Phase Deviation (ns), Δt	Frequency Offset $\Delta t/1$ s (y_i)	First Differences ($y_{i+1} - y_i$)	First Difference Squared ($y_{i+1} - y_i$) ²
3321.44	(-----)	(-----)	(-----)	(-----)
3325.51	4.07	4.07×10^{-9}	(-----)	(-----)
3329.55	4.04	4.04×10^{-9}	-3×10^{-11}	9×10^{-22}
3333.60	4.05	4.05×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3337.67	4.07	4.07×10^{-9}	$+2 \times 10^{-11}$	4×10^{-22}
3341.72	4.05	4.05×10^{-9}	-2×10^{-11}	4×10^{-22}
3345.78	4.06	4.06×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3349.85	4.07	4.07×10^{-9}	$+1 \times 10^{-11}$	1×10^{-22}
3353.91	4.06	4.06×10^{-9}	-1×10^{-11}	1×10^{-22}
3357.96	4.05	4.05×10^{-9}	-1×10^{-11}	1×10^{-22}

Table A.2 shows how stability is estimated. The first column is a series of phase measurements recorded at 1 s intervals. Each measurement is larger than the previous measurement. This indicates that the DUT is offset in frequency from the reference and this offset causes a phase deviation. By differencing the raw phase measurements, we obtain the phase deviations shown in the second column. The third column divides the phase deviation (Δt) by the 1 s measurement period to get the frequency offset. Since the phase deviation is about 4 ns/s, it indicates that the DUT has a frequency offset of about 4×10^{-9} . The frequency offset values in the third column form the y_i data series. The last two columns show the first differences of the y_i and the squares of the first differences. Since the sum of the squares equals 2.2×10^{-21} , the equation (where $\tau = 1$ s) becomes:

$$\sigma_y(\tau) = \sqrt{\frac{2.2 \times 10^{-21}}{2(9-1)}} = 1.17 \times 10^{-11}$$

Using the same data, the Allan deviation for $\tau = 2$ s can be computed by averaging pairs of adjacent values and using these new averages as data values. For $\tau = 4$ s, take the average of each set of four adjacent values and use these new averages as data values. More data must be acquired for longer averaging times. Keep in mind that the confidence level of a stability estimate improves as the averaging time increases. In the above example, we have eight samples for our $\tau = 1$ s estimate. However, we have only two samples for an estimate of $\tau = 4$ s.




Figure A.3. A sample Allan deviation graph.

More data must be acquired for longer averaging times. Keep in mind that the confidence level of a stability estimate improves as the averaging time increases. In the above example, we have eight

The confidence level of our estimate (1σ) can be roughly estimated as:

$$\frac{1}{\sqrt{M}} \times 100\%$$

In our example, where $M = 9$, our error margin is 33 %. With just 2 samples, our estimate may be in error by as much as 70 %. With 10^4 samples, the error margin is reduced to 1 %.

A sample Allan deviation graph is shown in Figure A.3. It shows the stability improving as the measurement period gets longer. Part of this improvement is because measurement system noise becomes less of a factor as the measurement period gets longer. At some point, however, the oscillator will reach its *flicker floor*, and from a practical point of view, no further gains will be made by averaging. The flicker floor is the point where the white noise processes begin to be dominated by nonstationary processes such as frequency drift. Most quartz and rubidium oscillators reach their flicker floor at a measurement period of 10^3 s or less, but cesium oscillators may not reach their flicker floor for 10^5 s (more than 1 day). Figure A.3 shows a sample Allan deviation graph of a quartz oscillator that is stable to about 5×10^{-12} at 100 s and is approaching its flicker floor [8, 9, 10].

Be sure not to confuse stability with frequency offset when you read a specifications sheet. For example, a DUT with a frequency offset of 1×10^{-8} might still reach a stability of 1×10^{-12} in 1000 s. This means that the output frequency of the DUT is stable, even though it is not particularly close to its nameplate frequency. To help clarify this point, Figure A.4 is a graphical representation of the relationship between frequency offset (accuracy) and stability.




Figure A.4. The relationship between frequency offset (accuracy) and stability.

Frequency Standards

Frequency standards all have an internal device that produces a periodic, repetitive event. This device is called the *resonator*. Of course, the resonator must be driven by an energy source. Taken together, the energy source and the resonator form an *oscillator*. Two main types of oscillators are used as frequency standards: quartz oscillators and atomic oscillators.

Quartz Oscillators

Quartz crystal oscillators first appeared during the 1920s and quickly replaced pendulum devices as laboratory standards for time and frequency [11]. More than 10^9 quartz oscillators are produced annually for applications ranging from wristwatches and clocks to communications networks and space tracking systems [12]. However, calibration and standards laboratories usually calibrate only the more expensive varieties of quartz oscillators, such as those found inside electronic instruments (such as frequency counters) or those designed as stand-alone units. The cost of a high-quality quartz oscillator ranges from a few hundred to a few thousand dollars.




Figure A.5. Block diagram of quartz oscillator.

The quartz crystal inside the oscillator can be made of natural or synthetic quartz, but all modern devices are made of synthetic material. The crystal serves as a mechanical resonator which creates an oscillating voltage due to the *piezoelectric effect*. This effect causes the

crystal to expand or contract as voltages are applied. The crystal has a *resonance frequency* that is determined by its physical dimensions and the type of crystal used. No two crystals can be exactly alike or produce exactly the same frequency. The output frequency of a quartz oscillator is either the fundamental resonance frequency

or a multiple of that frequency. Figure A.5 is a simplified circuit diagram that shows the basic elements of a quartz oscillator. The amplifier provides the energy needed to sustain oscillation.

Quartz oscillators are sensitive to environmental parameters such as temperature, humidity, pressure, and vibration [12, 13]. When environmental parameters change, the fundamental resonance frequency also changes. There are several types of quartz oscillator designs that attempt to reduce the environmental problems. The *oven-controlled crystal oscillator* (OCXO) encloses the crystal in a temperature-controlled chamber called an oven. When an OCXO is first turned on, it goes through a "warm-up" period while the temperatures of the crystal resonator and its oven stabilize. During this time, the performance of the oscillator continuously changes until it reaches its normal operating temperature. The temperature within the oven then remains constant, even when the outside temperature varies. An alternate solution to the temperature problem is the *temperature-compensated crystal oscillator* (TCXO). In a TCXO, the output signal from a special temperature sensor (called a thermistor) generates a correction voltage that is applied to a voltage-variable reactance (called a varactor). The varactor then produces a frequency change equal and opposite to the frequency change produced by temperature. This technique does not work as well as oven control, but is much less expensive. Therefore, TCXOs are normally used in small, usually portable units when high performance over a wide temperature range is not required. A third type of quartz oscillator is the *microcomputer-compensated crystal oscillator* (MCXO). The MCXO uses a microprocessor and compensates for temperature using digital techniques. The MCXO falls between a TCXO and an OCXO in both price and performance.

All quartz oscillators are subject to *aging*, which is defined as "a systematic change in frequency with time due to internal changes in the oscillator." Aging is usually observed as a nearly linear change over time in the resonance frequency. Aging can be positive or negative, and occasionally, a reversal in aging direction is observed. Often, the resonance frequency decreases, which might indicate that the crystal is getting larger. Aging has many possible causes including: contamination of the crystal due to deposits of foreign material, changes in the oscillator circuitry, or changes in the quartz material or crystal structure. The vibrating motion of the crystal can also contribute to aging. High quality quartz oscillators age at a rate of no more than 5×10^{-9} per year.

The best quartz oscillators have excellent short-term stability. An OCXO might be stable to 1×10^{-13} at 1 s. The limiting factor in the short-term stability is often noise from electronic components in the oscillator circuits. However, due to aging and other factors, the long term stability of a quartz oscillator is poor. Even the best OCXOs require periodic adjustments to stay within 1×10^{-10} of their nameplate

frequency. The TCXOs found in test equipment like counters and signal generators typically have frequency offsets ranging from 1×10^{-7} to 1×10^{-9} . Quartz oscillators without temperature control (such as those found in wristwatches, computers, radios, and so on) typically have frequency offsets measured in parts per million.

Atomic Oscillators

Atomic oscillators use the quantized energy levels in atoms and molecules as the source of their resonance frequency. The laws of quantum mechanics dictate that the energies of a bound system, such as an atom, have certain discrete values. An electromagnetic field can boost an atom from one energy level to a higher one. Or, an atom at a high energy level can drop to a lower energy level by emitting electromagnetic energy. The resonance frequency (f) of an atomic oscillator is the difference between the two energy levels divided by Planck's constant (h) [14]:

$$f = \frac{E_2 - E_1}{h}$$

All atomic oscillators are *intrinsic standards*, since their frequency is inherently derived from a fundamental natural phenomenon. There are three main types of atomic oscillators: rubidium standards, cesium standards, and hydrogen masers (discussed individually in the following sections). All three types contain an internal quartz oscillator that is locked to a resonance frequency generated by the atom of interest. Locking the quartz oscillator to the atomic frequency is advantageous. Most of the factors that degrade the long-term performance of a quartz oscillator disappear, since the atomic resonance frequency is much less sensitive to environmental conditions than the quartz resonance frequency. As a result, the long-term stability and uncertainty of an atomic oscillator are much better than those of a quartz oscillator, but the short-term stability is unchanged.

Rubidium Oscillators

Rubidium oscillators are the lowest priced members of the *atomic* oscillator group. They offer perhaps the best price/performance ratio of any oscillator. They perform much better than a quartz oscillator and cost much less than a cesium oscillator.

A rubidium oscillator operates at the resonance frequency of the rubidium atom (^{87}Rb), which is 6 834 682 608 Hz. This frequency is synthesized from a lower quartz frequency (typically 5 MHz) and the quartz frequency is steered by the rubidium resonance. The result is a very stable frequency source with the short-term stability of quartz, but much better long-term stability. Since rubidium oscillators are more stable

than quartz oscillators, they can be kept within tolerance with fewer adjustments. The initial price of a rubidium oscillator (typically from \$3000 to \$8000) is higher than that of a quartz oscillator, but since fewer adjustments are needed, labor costs are reduced. As a result, a rubidium oscillator might actually be less expensive to own than a quartz oscillator when used as a frequency standard.

The typical frequency offset of a rubidium oscillator ranges from 5×10^{-10} to 5×10^{-12} . Stability is typically about 1×10^{-12} at one day. Maintaining frequency within 1×10^{-11} can be done routinely with a rubidium oscillator but is nearly impossible with even the best quartz oscillators. The performance of a well-maintained rubidium oscillator can approach the performance of a cesium oscillator, and a rubidium oscillator is much smaller, more reliable, and less expensive.

Cesium Oscillators

Cesium oscillators are *primary frequency standards* because the SI second is based on the resonance frequency of the cesium atom (^{133}Cs), which is 9 192 631 770 Hz. This means that a cesium oscillator that is working properly should be very close to its nameplate frequency without any adjustment, and there should be no change in frequency due to aging. The time scale followed by all major countries, Coordinated Universal Time (UTC), is derived primarily from averaging the performance of a large ensemble of cesium oscillators, although some hydrogen masers also contribute to UTC.

Cesium oscillators are the workhorses in most modern time and frequency distribution systems. The primary frequency standard for the United States is a large cesium fountain oscillator named NIST-F1 that has a frequency uncertainty of about 1.7×10^{-15} . Commercially available cesium oscillators use cesium beam technology, and are small enough to fit in a standard equipment rack. They differ in quality, but their frequency offset should still be no greater than $\pm 5 \times 10^{-12}$ after a brief warmup period, and is typically parts in 10^{13} . Stability generally reaches parts in 10^{14} at one day, and it might take days or weeks before a cesium oscillator reaches its noise floor.

Reliability and cost are both issues to consider when purchasing a cesium oscillator. The main component of a cesium oscillator, called the *beam tube*, has a life expectancy of about 3 to 25 years, depending upon the type of tube and the amount of beam current used. The beam tube is needed to produce the resonance frequency of the cesium atom, and this frequency is then used to discipline a quartz oscillator. When the beam tube fails, the cesium oscillator performs like an undisciplined quartz oscillator. For this reason, a cesium oscillator needs to be constantly monitored to make sure that it is still delivering a cesium-derived frequency. Cost is also a major issue. The initial purchase price of a cesium oscillator ranges from \$30,000 to

\$80,000, and the cost of a replacement beam tube is a substantial fraction of the cost of the entire oscillator. Laboratories that use cesium oscillators need to budget not only for their initial purchase, but for the cost of maintaining them afterwards.

Hydrogen Masers

The *hydrogen maser* is the most elaborate and most expensive commercially available frequency standard. Few masers are built and most are owned by observatories or national standards laboratories. "Maser" is an acronym that stands for "Microwave Amplification by Stimulated Emission of Radiation." Masers use the resonance frequency of the hydrogen atom, which is 1 420 405 752 Hz.

There are two types of hydrogen masers. The first type, called an *active maser*, oscillates spontaneously and a quartz oscillator is phase-locked to this active oscillation. The second type, called a *passive maser*, frequency-locks a quartz oscillator to the atomic reference. The "passive" technique is also used by rubidium and cesium oscillators. Since active masers derive the output frequency more directly from the atomic resonance, they have better short-term stability than passive masers. Both types of maser have better short-term stability than a cesium oscillator. However, since the maser's performance depends upon a complex set of environmental conditions, its frequency uncertainty is greater than that of a cesium oscillator. And due to their complexity and low volume of production, masers typically cost \$200,000 or more [15,16].

Table A.3 summarizes the characteristics of the different types of oscillators.

Transfer Standards

To briefly review, a frequency calibration compares the *device under test* (DUT) to 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* that receives a radio signal. All transfer standards receive a signal that has a cesium oscillator at its source, and this signal delivers a cesium derived frequency to the user. This benefits many users, since not all calibration laboratories can afford to buy and maintain a cesium. Even if a laboratory has a cesium oscillator, it still needs to check its performance, and the only practical way to do this is by using a transfer standard.

Table A.3. Summary of oscillator types.

Oscillator Type	Quartz (TCXO)	Quartz (MCXO)	Quartz (OCXO)	Rubidium	Cesium	Hydrogen Maser
Primary Standard	No	No	No	No	Yes	No
Intrinsic Standard	No	No	No	Yes	Yes	Yes
Resonance Frequency	Mechanical (varies)	Mechanical (varies)	Mechanical (varies)	6.834682608 GHz	9.19263177 GHz	1.42040575 GHz
Leading Cause of Failure	None	None	None	Rubidium Lamp (15 years or more)	Cesium Beam Tube (3 to 25 years)	Hydrogen Depletion (7 years or more)
Stability, σ_y (τ), $\tau=1s$	1×10^{-9}	1×10^{-10}	1×10^{-12}	5×10^{-11} to 5×10^{-12}	5×10^{-11} to 5×10^{-12}	1×10^{-12}
Noise Floor, $\sigma_y(\tau)$	1×10^{-9} ($\tau = 1$ to 10^2 s)	1×10^{-10} ($\tau = 1$ to 10^2 s)	1×10^{-12} ($\tau = 1$ to 10^2 s)	1×10^{-12} ($\tau = 10^3$ to 10^5 s)	1×10^{-14} ($\tau = 10^5$ to 10^7 s)	1×10^{-15} ($\tau = 10^3$ to 10^5 s)
Aging/year	5×10^{-7}	5×10^{-8}	5×10^{-9}	2×10^{-10}	None	$\sim 1 \times 10^{-13}$
Frequency Offset after warm up	1×10^{-6}	1×10^{-7} to 1×10^{-8}	1×10^{-8} to 1×10^{-10}	5×10^{-10} to 5×10^{-12}	5×10^{-12} to 1×10^{-14}	1×10^{-12} to 1×10^{-13}
Warm-Up Time	< 10 s to 1×10^{-6}	< 10 s to 1×10^{-8}	< 5 min to 1×10^{-8}	< 5 min to 5×10^{-10}	30 min to 5×10^{-12}	24 hours to 1×10^{-12}
Cost	\$100	\$1000	\$2000	\$3000 to \$8000	\$30,000 to \$80,000	\$200,000 to \$300,000

Transfer standards provide *traceability*. Most transfer standards receive signals traceable to the national frequency standard maintained by NIST. Some signals, such as those transmitted by HF (high frequency) radio stations WWV and WWVH and the LF (low frequency) station WWVB, are traceable because they are directly controlled by NIST. Other signals, such as the LORAN-C and Global Positioning System (GPS) broadcasts, are traceable because their received signals are regularly compared to the NIST standard. Some signals broadcast from outside the United States are also traceable. This is because NIST compares its frequency standard to the standards maintained in other countries.

Some compromises are made when using a transfer standard. Even if the radio signal is referenced to a nearly perfect frequency, its performance is degraded as it travels along the radio path between the transmitter and receiver. To illustrate, consider a laboratory that has a rack-mounted frequency standard that produces a 5 MHz signal. Metrologists need to use this signal on their work bench, so they run a length of coaxial cable from the frequency standard to their bench. The signal is delayed as it travels from the standard to the bench, but since the cable is of fixed length the delay is constant. Constant delays don't change the frequency. The frequency that goes into one end of the cable is essentially the same frequency that comes out the other end. However, what if the cable length were constantly changing? This would generally cause the frequency to fluctuate. Over long periods of time, these fluctuations will average out, but the short-term frequency stability would still be very poor. This is exactly what happens when you use a transfer standard. The "cable" is actually a radio path that might be thousands of kilometers in length. The length of the radio path is constantly changing and appears to introduce frequency fluctuations, even though the source of the frequency (a cesium oscillator) is not changing. This makes transfer standards unsuitable as a reference when making short-term stability measurements. However, transfer standards are well suited for long term measurements, since you can minimize these frequency fluctuations if you average for a long enough measurement interval. Eventually, you will recover the performance of a cesium oscillator.

Some radio signals have path variations that are so pronounced that they are not well suited for high level frequency calibrations. To illustrate this, consider the signal broadcast from WWV, located in Fort 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. Most shortwave users receive the *skywave*, or the part of the signal that travels up to the ionosphere and is reflected back to earth. Since the height of the ionosphere constantly changes, the path delay constantly changes, often by as much as 500 to 1000 μ s. Since there is so

much variability in the path, averaging leads to only limited improvement. Therefore, although WWV is traceable to NIST, its frequency uncertainty is limited to parts in 10^9 when averaged for one day.

Other radio signals have more stable paths and much lower uncertainty values. Low frequency (LF) radio stations (such as WWVB and LORAN-C) can provide traceability to NIST with a frequency uncertainty of 1×10^{-12} /day. An LF path is

Table A.4. Typical frequency uncertainties of various transfer standards.

Transfer Standard	Frequency Uncertainty over 24 h measurement period (with respect to NIST)
HF receiver (WWV and WWVH)	$\pm 5 \times 10^{-9}$
LF receiver (LORAN-C and WWVB)	$\pm 1 \times 10^{-12}$
Global Positioning System receiver (GPS)	$\pm 2 \times 10^{-13}$

much more stable than an HF path, but still experiences a path delay change when the height of the ionosphere changes at sunrise and sunset. Currently, the most widely used signals originate from the Global Positioning System (GPS) satellites. GPS signals have the advantage of an unobstructed path between the transmitter and receiver. The frequency uncertainty of GPS is about 2×10^{-13} /day. WWVB, LORAN-C and GPS are described in the next three sections.

Table A.4 shows some of the transfer standards available, as well as their frequency uncertainty with respect to NIST when averaged for at least 24 h [17].

WWVB

Many countries broadcast time and frequency signals in the LF band from 30 to 300 kHz, as well as in the VLF (very low frequency) band from 3 to 30 kHz. Since part of the LF signal is *groundwave* and follows the curvature of the earth, the path stability of these signals can be quite good. One such station is NIST's WWVB, which transmits on 60 kHz from the same site as WWV in Fort Collins, Colorado and provides coverage to most of North America.

Although far more stable than an HF path, the WWVB path length is still influenced by environmental conditions along the path and by daily and seasonal

changes. Path length is important because part of the signal travels along the ground (*groundwave*) and another part is reflected from the ionosphere (*skywave*). The groundwave path is far more stable than the skywave path. If the path is relatively short (less than 1000 km), the receiver might continuously track the groundwave signal since it always arrives first. For longer paths, a mixture of groundwave and skywave is received. And over a very long path, the groundwave can become so weak that it will be possible to receive only the skywave. In this instance, the path becomes much less stable.

The characteristics of an LF path are different at different times of day. During the daylight and nighttime hours, the receiver might be able to distinguish between groundwave and skywave, and path stability might vary by only a few hundred nanoseconds. However, if some skywave is being received, *diurnal* phase shifts occur at sunrise and sunset. For instance, as the path changes from all darkness to all daylight, the ionosphere lowers and the path gets shorter. The path length then stabilizes until either the transmitter or receiver enters darkness. At this point, the ionosphere rises and the path gets longer.

WWVB receivers have several advantages when used as a transfer standard. They are low cost and easy to use, and the received signals are directly traceable to NIST. With a good receiver and antenna system, you can achieve a frequency uncertainty of 1×10^{-12} by averaging for one day [18].

LORAN-C

LORAN-C is a radionavigation system that operates in the LF band. Most of the system is operated by the U. S. Department of Transportation (DOT), but some stations are operated by other governments. The system consists of groups of stations called *chains*. Each chain has one master station and from two to five secondary stations. The stations operate at high power, typically 275 to 800 kW, and broadcast on a carrier frequency of 100 kHz using the 90 to 110 kHz band.

Since all LORAN-C chains use the same carrier frequency, the chains transmit pulses so that individual stations can be identified. Each chain transmits a pulse group that includes pulses from all of the individual stations. The pulse group is sent at a unique Group Repetition Interval (GRI). For example, the 7980 chain transmits pulses every 79.8 ms. 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. The pulse shape was designed so that the receiver can distinguish between groundwave and skywave and lock on to the more stable groundwave signal. Most receivers stay locked to the groundwave by tracking

the third cycle of the pulse. The third cycle was chosen for two reasons. First, it arrives early in the pulse so we know that it is groundwave. Second, it has more amplitude than the first and second cycles in the pulse, which makes it easier for the receiver to stay locked. Generally, a receiver within 1500 km of the transmitter can track the same groundwave cycle indefinitely and avoid skywave reception. The variations in groundwave path delay are typically quite small (<500 ns/day). However, if the path length exceeds 1500 km, the receiver might lose lock, and jump to another cycle of the carrier. Each cycle jump introduces a 10 μ s phase step, equal to the period of 100 kHz.

LORAN-C Performance

The frequency uncertainty 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. Path variations cause the short-term stability of LORAN-C to be poor. However, since the path variations average out over time, the long-term stability is quite good. This means that we can use LORAN-C to calibrate nearly any frequency




Figure A.6. LORAN-C compared to cesium oscillator over 96 h interval.

standard if we can average for a long enough interval. For this reason, a measurement period of at least 24 h is recommended when using LORAN-C to calibrate atomic oscillators. Figure A.6 shows the results of a 96 h calibration of a cesium oscillator using LORAN-C. The thick line is a least squares estimate of the frequency offset. Although the path noise is clearly visible, the slope of the line indicates that the cesium oscillator is low in frequency by 3.4×10^{-12} .

Global Positioning System (GPS)

GPS is a radionavigation system developed and operated by the U.S. Department of Defense (DOD). The system consists of a constellation of at least 24 earth-orbiting satellites (21 primary satellites and 3 in-orbit spares). Each satellite carries its own atomic frequency standards (cesium and/or rubidium oscillators) that are referenced to the United States Naval Observatory (USNO) and traceable to NIST. The 24 satellites orbit the Earth in six fixed planes inclined 55° from the equator. Each satellite is 20 200 km above the earth's surface and has an orbital period of 11 h, 58 min, which means a satellite will pass over the same place on earth 4 min earlier each day. Since the satellites continually orbit the earth, GPS should be usable anywhere on the earth's surface.

The GPS satellites broadcast on two carrier frequencies: L1 at 1575.42 MHz, and L2 at 1227.6 MHz. Each satellite broadcasts a spread-spectrum waveform, called a pseudo-random noise (PRN) code 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 (C/A) code with a chip rate of 1023 chips per millisecond. The second type is a precision (P) code with a chip rate of 10230 chips per millisecond. The C/A code is broadcast on L1, and the P code is broadcast on both L1 and L2. The C/A code repeats every millisecond. The P code repeats only every 267 days, but for practical reasons is reset every week. The C/A code is broadcast on L1, and the P code is broadcast on both L1 and L2 [19, 20]. GPS reception is line-of-sight, which means that the antenna must have a clear view of the sky. If a clear sky view is available, the signals can be received nearly anywhere on earth.

Each satellite carries either rubidium or cesium oscillators, or a combination of both. The on-board oscillators provide the reference for both the carrier and code broadcasts. They are steered from USDOD ground stations and are referenced to Coordinated Universal Time (UTC) maintained by the USNO. By mutual agreement UTC(USNO) and UTC(NIST) are maintained within 100 ns of each other, and the frequency offset between the two time scales is always $<1 \times 10^{-13}$.

GPS Receiving Equipment

Most GPS receivers provide a 1 *pulse per second* (pps) output. Some receivers also provide a standard frequency output (1, 5, or 10 MHz). To use these receivers, you simply mount the antenna, connect the antenna to the receiver, and turn the receiver on. The antenna is often a small cone or disk (normally about 15 cm in diameter) that must be mounted outdoors where it has a clear, unobstructed view of the sky. Once the receiver is turned on, it performs a sky search to find out which satellites are currently above the horizon and visible from the antenna site. It then computes its three-dimensional coordinate (latitude, longitude, and altitude as long as four satellites are in view) and begins producing a frequency signal. The simplest GPS receivers have just one channel and look at multiple satellites using a sequencing scheme that rapidly switches between satellites. More sophisticated models have parallel tracking capability and can assign a separate channel to each satellite in view. These receivers typically track from 5 to 12 satellites at once (although more than 8 will be available only in rare instances). By averaging data from multiple satellites, a receiver can reduce the frequency uncertainty [20].

GPS Performance

GPS has many technical advantages over LF radio signals. The signals are usually easier to receive, the equipment is often less expensive, the coverage area is much larger (worldwide), and the performance is better. However, as with all transfer




Figure A.7. GPS compared to cesium oscillator over 100 s interval.

standards, the short-term stability of a GPS receiver is not particularly good and this lengthens the time required to make a calibration. As with LORAN-C, a measurement period of at least 24 h is recommended when calibrating atomic frequency standards using GPS.




Figure A.8. GPS compared to cesium oscillator over 1 week interval.

To illustrate this, Figure A.7 shows a 100 s 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 s measurement period is <1 ns. Therefore, most of the phase noise on the graph can be attributed to GPS path variations.

Figure A.8 shows a one-week comparison between a GPS receiver and the same cesium oscillator used in Figure A.7. The range of the data is well under 100 ns. The thick line is the result of a linear least squares fit. Although phase noise caused by GPS path variations is visible, we can clearly see the linear trend contributed by the small frequency offset of the cesium ($<1 \times 10^{-13}$).

Calibration Methods

To review, frequency standards are normally calibrated by comparing them to a traceable reference frequency. In this section, we'll discuss how this comparison is made. To begin with, let's look at the electrical signal produced by a frequency standard. This signal can take several forms, as illustrated in Figure A.9.

The dashed lines represent the supply voltage inputs (ranging from 5 to 15 V for CMOS), and the bold solid lines represent the output voltage.




Figure A.9. Oscillator outputs.

If the output frequency is an oscillating sine wave, it might look like the one shown in Figure A.10. This signal produces one cycle (2π radians of phase) in one period. Frequency calibration systems compare a signal like the one shown in Figure A.10 to a reference signal of higher quality, and measure and record the change in phase between the two signals. If the two frequencies were exactly the same, their phase relationship would not change. Since the two frequencies are not exactly the same, their phase relationship does change, and the rate of change determines the frequency offset of the DUT. Under normal circumstances, the phase changes in an orderly, predictable fashion. However, external factors such as power outages, component failures, or human errors can cause a sudden phase change, or *phase step*. A calibration system measures the total amount of phase shift (caused either by the frequency offset of the DUT or a phase step) over a given measurement period.




Figure A.10. An oscillating signal.

Figure A.11 shows a phase comparison between two sinusoidal frequencies. The top sine wave represents a signal from the DUT, and the bottom sine wave represents a signal from the reference. Vertical lines have been drawn through the points where each sine wave crosses zero. The bottom of the figure shows “bars” that indicate the size of the phase difference between the two signals. If the phase relationship between the signals is changing, the phase difference will either increase or decrease to indicate that the DUT has a frequency offset (high or low) with respect




Figure A.11. Two signals with a changing phase relationship.

to the reference. Earlier, we introduced this simple equation to estimate frequency offset:

$$f(\text{offset}) = \frac{-\Delta t}{T}$$

In Figure A.11, both Δt and T are increasing at the same rate and the phase difference “bars” are getting uniformly wider. This indicates that the frequency from the DUT is stable, but is offset with respect to the reference.

Several types of calibration systems can be used to compare phase. The simplest type of system involves directly counting and displaying the frequency output of the DUT with a device called a *frequency counter*. This method has many applications but is unsuitable for measuring high performance devices. The DUT is compared to the counter's time base (typically a OCXO), and the uncertainty of the system is limited by the uncertainty of the time base, typically $\pm 1 \times 10^{-8}$. Some counters allow use of an external time base, which can improve the results. The biggest limitation is that frequency counters display frequency in hertz and have a limited amount of resolution. Detecting small changes in frequency may take many days or weeks, which might make it difficult or impossible to use this method to adjust a precision oscillator or to measure stability. For this reason, most high performance calibration systems collect time series data that can be used to estimate both frequency uncertainty and stability. We'll discuss how phase comparisons are made using the *time interval method* [9, 10].




Figure A.12. Measuring time interval.

The time interval method uses a device called a *time interval counter* (TIC) to measure the time interval between two signals. A TIC has inputs for two electrical signals. One signal starts the counter and the other signal stops it. 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 useful information. By subtracting the first measurement from the second, we can tell whether time was gained or lost.

TIC's differ in specification and design details, but they all contain several basic parts known as the *time base*, the *main gate*, and the *counting assembly*. The time base provides evenly spaced pulses used to measure time interval. The time base is usually an internal quartz oscillator that can often be phase locked to an external reference. It must be stable because time base errors will directly affect the measurements. The main gate controls the time at which the count begins and ends. Pulses passing through

the gate are routed to the counting assembly where they are displayed on the TIC's front panel or read by computer. The counter can then be reset (or *armed*) 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 might stop or start when it detects noise or other unwanted signals and produce invalid measurements.

Figure A.12 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 signal reaches its trigger level and stops measuring when the stop signal reaches its trigger level. The time interval between the start and stop signals is measured by counting cycles from the time base. 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 ns, it could produce a reading of 3340 ns or 3350 ns but not a reading of 3345 ns. This is because 10 ns 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 time base frequency. For example, a TIC with a 10 MHz time base would be limited to a resolution of 100 ns. This is because traditional TIC designs count whole time base cycles to measure time interval and cannot resolve time intervals smaller than the period of one cycle. To improve this

situation, some TIC designers have multiplied the time base frequency to get more cycles and thus more resolution. For example, multiplying the time base frequency to 100 MHz makes 10 ns resolution possible, and 1 ns counters have even been built using a 1 GHz time base. However, a more common way to increase resolution is to detect parts of a time base cycle through interpolation and not be limited by the number of whole cycles. Interpolation has made 1 ns TICs commonplace, and even 20 ps TICs are available [21, 22].




Figure A.13. Time interval system.

A time interval system is shown in Figure A.13. This system uses a TIC to measure and record the difference between two signals. Low frequency start and stop signals must be used (typically 1 Hz). Since oscillators typically produce frequencies such as 1, 5, and 10 MHz 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 circuits divide by multiples of 10, so it is common to find circuits that divide by one thousand, one million, and so on. For example, dividing a 1 MHz signal by 10^6 produces a 1 Hz signal. Using low frequency signals reduces the problem of counter overflows and underflows and helps prevent errors that could be made if the start and stop signals are too close together. For example, a TIC might make errors when attempting to measure a time interval of <100 ns.

The time interval method is probably the most common method in use today. It has many advantages, including low cost, simple design, and excellent performance when measuring long term frequency offset or stability.

The NIST Frequency Measurement and Analysis System

We have now completed our discussion of how frequency calibrations are made. We've seen that a frequency calibration system must include several basic components: an oscillator to calibrate, a reference (preferably a NIST-traceable transfer standard), and a phase comparison device. We'll conclude our tutorial with a brief look at how the NIST Frequency Measurement and Analysis System (FMAS) implements the concepts we've discussed.

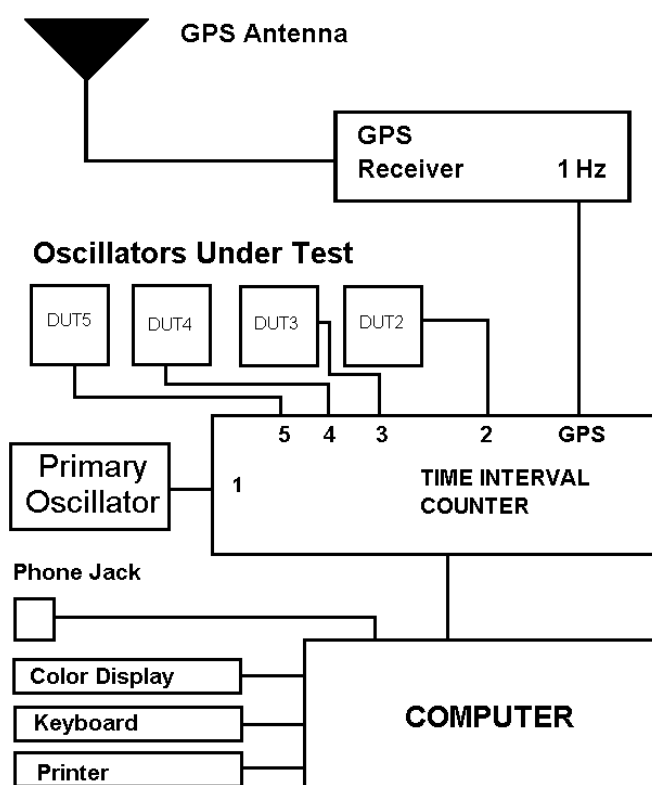


Figure A.14. The NIST Frequency Measurement and Analysis System.

A block diagram of the FMAS is shown in Figure A.14. The FMAS fits in a metal equipment rack (75 cm tall), and is computer controlled. Software developed at NIST controls all aspects of the calibration process. It makes measurements and stores and graphs them automatically. Up to five oscillators can be calibrated at once. The system produces calibration graphs that show the performance of each DUT for intervals ranging from 2 s to 150 days.

By looking at Figure A.14, you can see all the various components of a frequency calibration system. The FMAS uses a GPS receiver as a transfer standard. The GPS receiver produces a 1 kHz output frequency that provides traceability to NIST with an uncertainty of 2×10^{-13} .

The FMAS makes phase comparisons using the time interval method. It includes a TIC with a single-shot resolution of less than 40 ps. 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 an input frequency ranging from 1 Hz to 120 MHz on each of the five channels. The TIC is software controlled, so users don't have to set 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.

As you can see, the FMAS implements many of the concepts discussed in this overview. It provides a calibration laboratory with a well-defined and documented calibration method, and is a major asset to laboratories seeking ISO compliance or laboratory accreditation.

References

- [1] International Organization for Standardization (ISO). International Vocabulary of Basic and General Terms in Metrology. Switzerland: ISO; 1993. 59 p.
- [2] International Organization for Standardization (ISO). ISO/IEC Guide 17025, General Requirements for the Competence of Testing and Calibration Laboratories. Switzerland: ISO; 1999. 26 p.
- [3] American National Standards Institute (ANSI). ANSI/ISO/IEC Guide 17025:2000, General Requirements for the Competence of Testing and Calibration Laboratories. Boulder, Colorado: NCSL International; 2001. 26 p.
- [4] Lombardi, Michael A. An Introduction to Frequency Calibration: Part I. Cal Lab Int. J. Metrology: 17-28. 1996 January-February.
- [5] Taylor, B. N.; Kuyatt, C. E. Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results. Nat. Inst. Stand. Technol. Tech. Note 1297; 1994 September. 24 p.
- [6] IEEE Standards Coordinating Committee 27 on Time and Frequency. IEEE Standard 1139-1999, IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities. New York, NY: IEEE, 1999. 31 p.
- [7] Sullivan, D. B.; Allan, D. W.; Howe, D. A.; Walls, F. L.; ed. Characterization of Clocks and Oscillators. Nat. Inst. Stand. Technol. Tech Note 1337; 1990 March. 352 p.
- [8] Jespersen, James. Introduction to the Time Domain Characterization of Frequency Standards. Proc. 23th Annu. Precise Time and Time Interval (PTTI) Meeting; 1991 December; Pasadena, CA. 83-102.
- [9] Stein, S. R. Frequency and Time - Their Measurement and Characterization, chapter 12 in Precision Frequency Control, Vol. 2. Gerber, E. A.; Ballato, A.; ed. New York, NY: Academic Press; 1985. 191-232. Reprinted in [7].
- [10] Howe, D. A.; Allan, D. W.; Barnes, J. A. Properties of Signal Sources and Measurement Methods. Proc. 35th Annu. Sym. Freq. Cont; 1981. 47 p. Reprinted in [7].
- [11] Marrison, W. A. The Evolution of the Quartz Crystal Clock. Bell Systems Tech. 27(3): 510-588; 1948.

- [12] Vig, J. R. Introduction to Quartz Frequency Standards. Army Research and Development Technical Report: SLCET-TR-92-1; 1992. 49 p.
- [13] Walls, F. L.; Gagnepain, J. Environmental Sensitivities of Quartz Oscillators. IEEE Trans. Ultrasonics, Ferroelectrics, Frequency Control 39(2): 241-249; 1992 March.
- [14] Itano, W. M.; Ramsey, N. F. Accurate Measurement of Time. Sci. Am. 269(1): 56-65; 1993 July.
- [15] Lewis, L. An Introduction to Frequency Standards. Proc. IEEE 79(7): 927-935; 1991 July.
- [16] Major, F. G. The Quantum Beat: The Physical Principles of Atomic Clocks. New York: Springer Verlag; 1998, 475 p.
- [17] Lombardi, Michael A. An Introduction to Frequency Calibration: Part II. Cal Lab Int. J. Metrology: 28-34. 1996 March-April.
- [18] Beehler, R.; Lombardi, M. A. NIST Time and Frequency Services. *Nat. Inst. Stand. Technol. Spec. Publ.* 432; 1991 June. 30 p.
- [19] Hoffmann-Wellenhof, B.; Lichtenegger, H.; Collins, J. Global Positioning System: Theory and Practice. New York: Springer-Verlag; 1997. 370 p.
- [20] Lombardi, Michael A; Nelson, Lisa. M.; Novick, Andrew N.; Zhang, Victor S. Time and Frequency Measurements Using the Global Positioning System. Cal Lab Int. J. Metrology: 26-33. 2001 July-September.
- [21] Zhang, Victor S., Davis, Dick D.; Lombardi, Michael A. High Resolution Time Interval Counter. Proc. 26th Annu. Precise Time and Time Interval (PTTI) Meeting; 1994 December; Reston, VA. 191-200.
- [22] Novick, Andrew N., Lombardi, Michael A.; Zhang, Victor S.; Carpentier, Anthony. High Performance Multi-Channel Time Interval Counter with Integrated GPS Receiver. Proc. 31st Annu. Precise Time and Time Interval (PTTI) Meeting; 1999 December; Dana Point, CA. 561-568.

Appendix B

Calibration Reports

Calibrations made using the FMAS are traceable to the National Institute of Standards and Technology (NIST) under NIST Test Number 76100S. If your system is recording data and is connected to a phone line, you will receive a monthly calibration report from NIST that documents measurement traceability. The calibration report that you receive is in full compliance with ISO Guide 17025.

GPS is officially referenced to the United States Naval Observatory (USNO), but the time offset between UTC(NIST) and UTC(USNO) is kept within 100 ns by mutual agreement, and the frequency offset is $<1 \times 10^{-13}$. Therefore, GPS can be used to establish traceability to either agency with very small uncertainties. The current time offset between UTC(USNO) and UTC(GPS) is available on the NIST web site:

<http://www.boulder.nist.gov/timefreq/pubs/bulletin/nistusno.htm>

NIST also establishes traceability through GPS by monitoring each of the individual satellites as they fly over the Boulder, Colorado laboratories. This information is published daily at:

<http://www.boulder.nist.gov/timefreq/service/gpstrace.htm>

Each calibration report is three pages long. The content of each page is described in the following sections. To obtain the lowest possible uncertainties on your calibration report, you must operate the FMAS as described in the manual and follow good laboratory practices. You must keep the FMAS running 24 hours a day, whenever possible. Your primary frequency standard should have an uninterruptible power supply and should also run 24 hours a day.

Page 1

The top portion of page 1 contains some general information about each customer including: the customer's address, the main point of contact, and the manufacturer, model number, and type of the customer's primary frequency standard (called the device under test). If this information is incorrect or if it changes, please contact NIST.

The rest of page 1 contains a brief description of the calibration method, and some general information about the calibration service. This text may change slightly from time to time if changes are made to the calibration service.

The bottom of page 1 contains a footer (also found on pages 2 and 3) that lists the issue date of the report, the report number, the NIST Test number, and the page number. The report number is a six-digit number that uniquely identifies the report.

Page 2

Page 2 contains a table that lists the daily frequency offset values for the device under test. Each offset value is the result of a 24-hour calibration. The table also lists a confidence level for each measurement, and a status code for each measurement. The status codes are assigned by NIST. The possible value for each status code is listed below.

Code	Description
0	<i>Valid Calibration</i> - NIST has determined that all components of the measurement system were working properly during the calibration and that the calibration results are valid. This usually means that the calibration was within two standard deviations (2σ) of your average frequency. However, in some cases a valid calibration falls outside of the 2σ range by a small amount.
1	<i>No Data</i> - No data was recorded. The measurement system was either turned off or malfunctioning.
2	<i>GPS Reception Error</i> - NIST has determined that you had a problem with GPS reception during the calibration. The cause could be a malfunctioning receiver, antenna, cable, or an interference problem.
3	<i>GPS Broadcast Error</i> - NIST has determined that a GPS satellite in view of your location was broadcasting invalid data that adversely influenced the calibration. This situation occurs only in rare instances.
4	<i>Measurement System Error</i> - NIST has determined that the measurement system was malfunctioning due to a hardware or software failure.
5	<i>DUT error</i> - Your primary frequency standard was out of normal tolerance during the calibration.
6	<i>DUT change</i> - Your primary frequency standard was replaced by another standard.

If a status code other than 0 is assigned, NIST will provide a brief comment that explains the problem. If a status code of 0 is assigned it means that NIST has reviewed the calibration, found it to be valid, and no further comment is necessary.

Page 3

Page 3 contains a plot showing the daily frequency offset values for the device under test. The frequency offset values represent the performance of your primary standard during each 24 hour calibration. The x-axis shows the days of the month. Each tick mark represents one day. The y-axis shows the frequency offset of the primary oscillator. The middle of the y-axis represents a frequency uncertainty of 1×10^{-12} or better. The plot accommodates oscillators that are high in frequency (the top half of the plot), or low in frequency (the bottom half). The plot shows frequency offset values as large as 1×10^{-8} . If the frequency offset is $>1 \times 10^{-8}$, it is plotted along the top or bottom line of the plot.

Page 3 also contains an uncertainty statement. This statement lists the number of days in the month when data was recorded, and the number of days when your primary standard was in tolerance. It also lists the following values:

Mean Frequency Offset (MFO) - This is your primary standard's mean (or average) frequency offset for the entire month.

Uncertainty (U), 2σ - This is the measurement uncertainty assigned to the calibration. Measurement uncertainty is contributed by the GPS receiver, by DUT aging and frequency drift, and by measurement system noise. NIST monitoring indicates that the received GPS signals have an uncertainty of 2×10^{-13} with respect to UTC(NIST). Measurement system noise contributes an uncertainty of 2×10^{-15} .

FO_{lower} and FO_{upper} - These numbers represent the upper and lower bands of the performance of your frequency standard. There is a 95.4 % probability that the frequency offset of your standard was within the tolerance established by these bands at any given instant during the month covered by the report. For example, if the lower band is -1.52×10^{-12} and the upper band is $+9.44 \times 10^{-13}$ there is a 95.4 % probability that the frequency offset remained within these limits.

Page 3 also lists the signatures and telephone numbers of the preparer of the calibration report and the manager of time and frequency services at NIST. Please contact these individuals if you have questions about the report.

Appendix C

Specifications and Compliance Documents

This section lists the specifications for the NIST Frequency Measurement and Analysis System (FMAS). It also lists the compliance documents that define the terminology and techniques used by the FMAS hardware, software, and documentation.

Specifications

Number of measurement channels:	5
Frequencies directly accepted by the system:	1 Hz to 120 MHz (1 Hz increments)
Maximum number of data points per graph:	3600
Sampling rate (all measurements):	1 Hz
Averaging period (long-term measurements):	1 h
Averaging period (short-term measurements):	1 s
Single shot measurement resolution:	< 30 ps
Frequency uncertainty using GPS (24 h):	2×10^{-13}
Frequency uncertainty without GPS (24 h): (oscillator to oscillator comparisons)	2×10^{-15}
Stability during oscillator self test (1000 s):	1×10^{-14}
Auto-trigger sensitivity range (50 Ω):	200 mV to 3.5 V (peak to peak)
	70 mV to 1.25 V (rms)

Compliance Documents

ANSI/ISO/IEC Guide 17025, *General Requirements for the Competence of Testing and Calibration Laboratories*, International Organization for Standardization (ISO), 2000.

ANSI NCITS 310-1998, *Representation of Time for Information Interchange*, American National Standards Institute (ANSI), 1998.

ISO/IEC Guide 17025, *General Requirements for the Competence of Testing and Calibration Laboratories*, International Organization for Standardization (ISO), 1999.

International Organization for Standardization (ISO), *International Vocabulary of Basic and General Terms in Metrology (VIM)*, 1993.

NVLAP Calibration Laboratories Technical Guide, NIST Handbook 150-2, 2001.

Taylor, Barry N.; Kuyatt, Chris E., *Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results*, NIST Technical Note 1297, 1993.

Appendix D

Glossary

The definitions in this glossary have been obtained from several sources, including ISO (International Organization for Standardization) and ITU (International Telecommunications Union) documents.

ACCURACY - The degree of closeness to which a measured value agrees with its definition at a given point in time.

AGING - A systematic change in frequency with time due to internal changes in the oscillator.

ALLAN DEVIATION - A widely-used statistical test for estimating the stability of an oscillator. Also referred to as the Allan variance.

ANSI - American National Standards Institute.

ATOMIC OSCILLATORS - Any oscillator whose resonance frequency is derived from an atomic frequency. The most common atomic oscillators are rubidium and cesium oscillators.

ATOMIC TIME (TA) - A time scale based on the resonance frequency of the cesium atom.

BIPM - International Bureau of Weights and Measures. The BIPM is located near Paris, France and maintains the International System (SI) of units, and Coordinated Universal Time (UTC).

CALIBRATION - A set of operations that establish, under specified conditions, the relationship between values of quantities indicated by a measuring instrument or measuring system, and the corresponding standard or known values derived from the standard. Practically speaking, a calibration is a comparison between an unknown quantity (produced by the device under test) and a known quantity (produced by a reference).

CALIBRATION METHOD - The technical procedure used to make a calibration.

CESIUM OSCILLATOR - An atomic oscillator based on the resonance frequency of the cesium atom. Cesium oscillators provide the basis for the international definition of the second.

CHARACTERIZATION - The process of measuring both the frequency uncertainty and stability of an oscillator, and stating both characteristics in a quantitative way.

CLOCK - A device that counts a specific number of oscillations from a frequency source, and then displays or records the amount of time that has elapsed. For example, a quartz wristwatch counts 32 768 oscillations of a quartz crystal and records that one second has elapsed.

COARSE ACQUISITION CODE (C/A CODE) - A 1023 bit pseudo random noise code that is transmitted by the GPS satellites every millisecond.

COORDINATED UNIVERSAL TIME (UTC) - An international atomic time scale governed by the BIPM. UTC is computed by averaging data from the world's major metrology laboratories.

DEVICE UNDER TEST (DUT) - The device being calibrated. Sometimes referred to as the unit under test (UUT).

DIURNAL PHASE SHIFT - A phase shift that occurs at about the same time each day. A common example is a shift in the path delay of a radio signal caused by a change in the height of the ionosphere at sunrise and sunset. Diurnal phase shifts are common with LF and HF transfer standards.

EPHEMERIS - A table that predicts the positions of extraterrestrial objects (such as the Sun, Moon, planets, or artificial satellites).

FLICKER FLOOR - The averaging period that produces the best possible stability value for a frequency standard. Once a standard reaches its flicker floor, averaging for a longer period will not improve the stability estimate. Often called the flicker noise floor, or simply the noise floor.

FREQUENCY COUNTER - An electronic device that counts the number of cycles produced by a signal over a given time interval and, by doing so, determines the frequency of the signal.

FREQUENCY DIVIDER - A device that divides a frequency to a lower frequency while preserving its phase.

FREQUENCY OFFSET - The difference between the nameplate frequency and the reference frequency. Usually expressed as a unitless value, for example, -1×10^{-9} , but it can be expressed in units of hertz. Sometimes called frequency accuracy, frequency difference, frequency error, or relative frequency.

FREQUENCY UNCERTAINTY - The limits (upper and lower) of the measured frequency offset. Typically, frequency uncertainty is expressed using a 2σ test, meaning that there is a 95.4 % probability that the frequency offset of the device under test will remain within the stated limits during the calibration.

GLOBAL POSITIONING SYSTEM (GPS) - A radionavigation system maintained by the United States Department of Defense. GPS consists of a constellation of earth-orbiting satellites that carry atomic oscillators, and is widely used as a reference for frequency calibrations.

HYDROGEN MASER - An advanced type of atomic oscillator based on the resonance frequency of hydrogen.

INTRINSIC STANDARD - A standard used to define a unit of measurement that is based on a fundamental constant in nature. Atomic oscillators are intrinsic standards.

ISO - International Organization of Standardization.

LORAN-C - A ground-based radionavigation system operated by the United States Coast Guard and often used as a frequency reference. LORAN-C signals are broadcast at 100 kHz.

LOW FREQUENCY (LF) - The portion of the radio spectrum from 30 to 300 kHz.

MEASURAND - The quantity being measured during a calibration.

MICROSECOND - One millionth of a second (10^{-6} s).

MILLISECOND - One thousandth of a second (10^{-3} s).

MODIFIED JULIAN DATE (MJD) - A way of telling what day it is with the least possible ambiguity. The Julian date number is obtained by counting days with a starting point of noon on January 1, 4713 B.C., which was defined as Julian day zero. The Modified Julian date is equal to the Julian date shifted so that its origin occurs at midnight on November 17, 1858. The MJD differs from the Julian date by exactly 2 400 000.5 days.

NAMEPLATE FREQUENCY - The frequency labeled on an oscillator's output, or the frequency that an oscillator is supposed to produce. The nameplate frequency is an ideal frequency with zero uncertainty. Often called the nominal frequency.

NANOSECOND - One billionth of a second (10^{-9} s).

NCSL - National Conference of Standards Laboratories.

NIST - National Institute of Standards and Technology.

NVLAP - National Voluntary Laboratory Accreditation Program, a division of NIST.

OSCILLATOR - An electronic device that produces signals of a specified frequency.

OVEN - A temperature-controlled chamber used to enclose a quartz crystal. By providing a stable environment for the crystal, an oven improves the oscillator's performance.

PATH DELAY - The amount of time that elapses as a time or frequency signal travels from the transmitter to the receiver.

PHASE COMPARISON - A measurement technique used to measure the difference between two frequencies.

PHASE STEP - A sudden change in phase.

PICOSECOND - One trillionth of a second (10^{-12} s).

PIEZOELECTRIC EFFECT - The oscillation of a quartz crystal that takes place when a voltage is applied.

PSEUDO RANDOM NOISE CODE (PRN) - A spread spectrum waveform that delivers encoded information but appears to be noise.

PULSE PER SECOND (PPS) - An electrical pulse that is generated once per second. If the pulse is synchronized to UTC, it is said to be on-time.

QUARTZ OSCILLATOR - A common type of oscillator that produces a frequency derived from the mechanical resonance of a quartz crystal.

REFERENCE - A device with a known uncertainty relative to national or international standards that is used to calibrate another device. Also called a standard.

RESOLUTION - The degree to which a measurement can be determined.

RESONANCE FREQUENCY - The fundamental frequency of an oscillator. Atomic oscillators derive their resonance frequency from the the physical properties of an atom. Quartz oscillators derive their resonance frequency from the mechanical properties of a quartz crystal.

RUBIDIUM OSCILLATOR - An atomic oscillator which derives its frequency from the resonance frequency of the rubidium atom.

SECOND - The basic unit of time interval, defined as the duration of 9 192 631 770 energy transitions of a cesium atom. All units of frequency and time interval are defined based on the second.

STABILITY - The statistical estimate of the frequency fluctuations of a signal over a given time interval.

STANDARD - A device with a known level of uncertainty that is used to calibrate another device. Also called a reference.

TEST UNCERTAINTY RATIO (TUR) - The performance ratio between the reference and the device under test during a calibration. A TUR of 10:1 is preferred, but a lower ratio can be used if more averaging is done.

TIME BASE - The component inside an instrument (such as a counter) that provides the reference for all time and frequency related functions.

TIME CODE - A digital data stream that contains date and time-of-day information.

TIME INTERVAL COUNTER - A device that measures the amount of time that elapses between the arrival of two electrical signals.

TIME INTERVAL METHOD - A phase comparison technique that uses a time interval counter.

TIME SCALE - An unambiguous ordering of events, or a system for keeping time.

TRACEABILITY - The property of the result of a measurement or the value of a standard whereby it can be related to stated references, usually national or international standards, through an unbroken chain of comparisons, all having stated uncertainties.

TRANSFER STANDARD - A radio receiver that receives, processes, and distributes a signal that serves as a traceable reference for calibrations. Transfer standards used for frequency typically receive GPS, LORAN-C, or WWVB signals.

UNCERTAINTY - Parameter, associated with the result of a measurement, that characterizes the dispersion of values that could reasonably be attributed to the measurand.

USNO - The United States Naval Observatory, located in Washington, DC.

WWV - A high frequency (HF) radio station operated by NIST in Fort Collins, Colorado. WWV broadcasts standard time and frequency information on 2.5, 5, 10, 15, and 20 MHz. NIST operates a similar station named WWVH from Hawaii.

WWVB - A low frequency (LF) radio station operated by NIST in Fort Collins, Colorado. WWVB broadcasts standard time and frequency information on a carrier frequency of 60 kHz.

Index

accuracy 61, 93
Allan deviation (or Allan variance) 41-43, 62-65, 93
antenna, GPS 3, 6
automatic plotting 17-19

bar graph 15-17, 19, 21
block diagram of system 2, 83

calibration 55-57, 94
changing oscillators 19
channel names 12
characterization 94
color monitor 5
commands accepted by the system
 ALL 37
 CalibrateTIC 49
 CONTROL-P 17
 CONTROL-X 20, 23
 COPY 26-28, 34, 43, 45
 DEMO 24, 37
 DOS 14
 EDIT 29, 31
 EXIT 14
 FILE 43
 FILTER 31-32, 34, 53
 GO 40
 GPS 46-48
 HI 31-32
 HISTORY 44-46
 LENGTH 40
 LO 31-32
 LOAD 24, 37
 MAGNIFY 34

MEASURE 15, 19
METER 44
NAMES 12, 19, 29
OFFSET 40-41
PICKCHANNEL 39
PLOT 27
PORTION 33-34
QUIT 14, 24-28, 31, 35, 37, 39, 45
RESET 31-32, 34
SELECT 38
SHORTTERM 39
STABILITY 41-43
STATISTICS 26-27
TITLE 35
UTILITIES 39
VIEW 24-26
XAXIS 34-35
YAXIS 34-35
YES 14

connecting oscillators 8-9
Coordinated Universal Time (UTC) 17, 94

data
 printing 26
 retrieving 37-38, 43-44
 viewing 24-26
demonstration data 24

exiting the FMAS software 14

frequency files 45-46, 52
frequency offset 16, 19, 26, 29, 31, 40-41, 43-46, 56, 58-61, 87, 95
frequency uncertainty 56, 58-61, 89, 91, 95

glossary 93-98
Global Positioning System (GPS) 1-3, 6, 8-12, 46-48, 53, 76-78, 87-89, 91, 95

installing the system 1-7

measurement screen 15-17
menus
 exiting a menu 14
 learning to use them 13-14
 selecting a menu item 13-14
microsecond 18, 100
modem 7
Modified Julian Date (MJD) 25, 44, 96
monitor, color 5

nanosecond 17, 100
NO SIGNAL message 20

oscillators
 changing 19-20
 connecting to system 8-9
 primary 8

parts included in system
 list of 1
plotting a portion of the plot on screen 33-34
power outages iii, 2
primary oscillator
 performance history 44-46
 selecting 8
 timebase, used as 8
 traceability 9, 87
printer
 cable 7
 changing types 53
 checking status 20
 paper 1, 9
printing the data 26

quick check of system 20-21

r (correlation coefficient) 26, 29, 41
range of data 26
reference frequency 9, 89

removing data points 31, 32
retrieving your data 37-38, 43-44

setup software
 pathname for software 52
 pathname for phase data 52
 pathname for demo data 52
 pathname for frequency data 52
 pathname for short term data 52
 select a printer 53
 changing the plot time 52
 setting the file capacity 52

short-term oscillator performance 39-44
specifications 95
stability 39-43, 61-65, 97
starting the system 10-12
supplies iii, 1

tape backup 3-4
telephone line iii, 7
timebase jumper 8
time interval counter 1, 8-9, 49, 81-84, 98
title of plot 29
traceability
 definition 56, 98
 monthly reports 9, 55-57, 87-89

uninterruptible power supply (UPS) 1, 4-7

viewing data 24-26
viewing the data statistics 26-27

xaxis
 labels 18-19
 position 34-35
 tick mark interval 29

yaxis
 labels 18-19
 scale 34-35