

# Computer Time Synchronization

Michael Lombardi  
Time and Frequency Division  
National Institute of Standards and Technology

---

The personal computer revolution that began in the 1970's created a huge new group of time and frequency users, those people who need to keep computer clocks on time. As you probably know, computer clocks aren't particularly good at keeping time. Simple clocks like your wristwatch and most of the clocks in your home usually keep better time than a computer clock.

The poor performance of a computer clock can cause problems, since many computer applications require time kept to the nearest second or better. For example, the computers in a financial institution must keep very accurate records of when transactions were completed, for legal or other reasons. Computer systems that make physical measurements and acquire scientific data need to know precisely when the measurements were made. Software used on a manufacturing floor may need to turn a piece of equipment on or off at a specified time. Also, any system involved with synchronous communications must keep accurate time. For example, radio and TV stations may need computers that can switch feeds or link up with remotes at the right time.

This paper describes several methods to keep accurate time by computer. Before looking at these methods, let's look at how a PC-compatible computer keeps time.

## **Section 1.1 - How a Personal Computer keeps time**

Since the introduction of the IBM-AT personal computer in 1984, all PC-compatible computers have kept time the same way. Each PC contains two clocks, regardless of whether it uses the 286, 386, 486, or Pentium microprocessor (or a derivative). These clocks go by several different names, but for simplicity, we'll call them the software and hardware clocks. The software clock runs when the computer is turned on and stops when the computer is turned off. The hardware clock uses a battery and runs even while the computer is turned off.

The software clock is generated by an Intel 8254 timer-counter (or a functionally equivalent device). This timer-counter generates an interrupt every 54.936 milliseconds, or about 18.2 times per second. The computer's BIOS (Basic Input Output System) contains a software routine that counts the interrupt requests and generates a time-of-day clock that can

be read or set by other software programs. For example, the operating system might use the time-of-day information from the software clock to date and time stamp files.

The software clock is a poor timekeeper. Its timing uncertainty is limited by the stability of the interrupt requests. Any change in the interrupt request rate causes the clock to gain or lose time. If you leave your computer turned on for long periods, the software clock might be off by large amounts, perhaps a minute or more for every day that the computer was left turned on. It's also possible for an ill-behaved software program to use the timer-counter for another purpose and change its interrupt rate. This could cause the clock to rapidly gain or lose time.

The software clock also has limited resolution. It can only display values that are even multiples of the time interval between interrupts (55 milliseconds). For example, 00:00:01.00 could never be displayed by the software clock. The closest possible values it can display are 00:00:00.98 and 00:00:01.04.

The single biggest drawback of the software clock, however, is that when the computer is turned off, the clock stops running and loses all of its time-of-day information. For this reason, a hardware clock is also necessary. The hardware clock is based on the Motorola 146818 Real Time Clock Chip, or an equivalent device. When the computer is turned off, the hardware clock runs off batteries. When the computer is turned back on, the software clock starts running again and sets itself (within 1 second) to the hardware clock. Although the hardware and software clocks are synchronized at power-up, they run at different rates and will gain or lose time relative to each other while the computer is running.

The hardware clock is updated once per second and cannot display fractions of a second. Its timing uncertainty is determined by the quality of the crystal oscillator it uses as its time base. These crystals cost less than \$1 in single quantities and offer only marginal timekeeping performance. They are sensitive to temperature and other factors and their frequency uncertainty is not likely to be better than  $1 \times 10^{-5}$  (about 1 second per day). In actual operation, most hardware clocks gain or lose about 5 to 15 seconds per day, with 10 seconds per day being typical. Although the hardware clock usually outperforms the software clock, its performance pales in comparison to even a low-cost wristwatch.

As you can see, neither the software or hardware clock is suitable for accurate timekeeping. Fortunately, there are ways to solve the PC timekeeping problem. Let's start by looking at how to synchronize a computer clock using an Internet service.

## Section 1.2 - Internet Time Setting Services

If your computer is connected to the Internet, you can synchronize its clock to an Internet time server. The process is fast and easy. All you need is an active Internet connection and client software.

Internet time servers use several standard timing protocols defined in a series of RFC (Request for Comments) documents. The three major timing protocols are the Time Protocol, the Daytime Protocol, and the Network Time Protocol (NTP). The time servers are continually “listening” for timing requests sent using any of these three protocols. When the server receives a request, it sends the time to your computer in the appropriate format. The protocol that you use depends upon the type of client software that you have. Most client software requests that the time is sent using either the Daytime Protocol or NTP. Client software that uses the Simple Network Time Protocol (SNTP) makes the same timing request as an NTP client, but does less processing and provides less accuracy. Table 1.20 summarizes the protocols and their port assignments.

*Table 1.20 - Internet Time Protocols*

<b>Name</b>	<b>Document</b>	<b>Format</b>	<b>Port Assignments</b>
Time Protocol	RFC-868	Unformatted 32-bit binary number contains time in UTC seconds since January 1, 1900.	Port 37 tcp/ip, udp/ip
Daytime Protocol	RFC-867	Exact format not specified in standard. The only requirement is that time code is sent as standard ASCII characters.	Port 13 tcp/ip, udp/ip
Network Time Protocol (NTP)	RFC-1305	The server provides a data packet that includes a 64-bit timestamp containing the time in UTC seconds since January 1, 1900 with a resolution of 200 picoseconds. NTP provides accuracy of 1 to 50 milliseconds. NTP client software normally runs continuously and gets periodic updates from the server.	Port 123 udp/ip
Simple Network Time Protocol (SNTP)	RFC-1769	The data packet sent by the server is the same as NTP, but the client software does less processing and provides less accuracy.	Port 123 udp/ip

NIST operates a Internet Time Service from Boulder, Colorado using multiple servers distributed around the country. The NIST servers distribute time using the Time, Daytime, and NTP formats and currently handle many millions of timing requests per day. For a current list of IP addresses for the NIST servers and sample client software (Daytime Protocol) that you can download, see:

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

NIST maintains a listing of other client software packages for a wide variety of computers. This list is available here:

<http://www.boulder.nist.gov/timefreq/general/softwarelist.htm>

### **Section 1.3 - Dial-Up Time Setting Service**

If your computer is not connected to the Internet, you can synchronize its clock to NIST time using a standard telephone line and a analog modem. The service which allows you to do this is called the Automated Computer Time Service (ACTS), which began in 1988.

ACTS requires only a computer, an analog modem and phone line, and some simple software. When a computer connects to ACTS by telephone, it receives an ASCII time code. The information in this time code is then used to set the computer clock to the correct time. ACTS is usable at modem speeds up to 9600 baud with 8 data bits, 1 stop bit, and no parity. To receive the full time code, you must connect at a speed of at least 1200 baud. The full time code is transmitted every second and contains more information than the 300 baud time code, which is transmitted once every 2 seconds. Table 1.30 (next page) describes the ACTS time code.

The last character in the time code is an asterisk (\*). The asterisk is called the on-time marker (OTM). The time values sent by the time code refer to the arrival time of the OTM. In other words, if the time code says it is 12:45:45, this means it is 12:45:45 when the OTM arrives.

To compensate for the time it takes for the OTM to travel from NIST to your computer, ACTS sends the OTM out 45 milliseconds early. This 45 milliseconds includes the 8 milliseconds that it takes to send the OTM at 1200 baud, 7 milliseconds transmission time to allow for travel from NIST to an average user in the United States, and 30 milliseconds to allow for the modem processing delay. The 45 millisecond advance was chosen based on experiments conducted at NIST using 1200 baud modems.

Advancing the OTM by 45 milliseconds always removes some of the delay. However, to get the least amount of timing uncertainty, the OTM should be advanced by the amount of the actual path delay. ACTS can do this by using a *loop-back* technique to calibrate the path. The loop-back technique works if the user's computer software returns the OTM to NIST after it is received. Each time the OTM is returned, ACTS measures the amount of time it took for the OTM to go from NIST to the user and back to NIST. This quantity is the round-trip path delay so it is divided by 2 to get the one-way path delay.

After a loop-back measurement is made, ACTS advances the time by the amount of the one-way path delay. For example, if the one-way path delay is 50.4 milliseconds, ACTS sends the OTM out 50.4 (instead of 45) milliseconds early. At this point, the path is calibrated, and OTM changes from an asterisk to a pound sign (#). If you calibrate the path, ACTS can set a computer clock with an uncertainty of less than 10 milliseconds.

Table 1.30 - The ACTS Time Code

<b>JJJJ YR-MO-DA HH:MM:SS TT L UT1 msADV UTC(NIST) &lt;OTM&gt;</b>
JJJJJ is the Modified Julian Date (MJD). The MJD is the last five digits of the Julian Date, which is simply a count of the number of days since January 1, 4713 B.C. To get the Julian Date, add 2.4 million to the MJD.
YR-MO-DA is the date. It shows the last two digits of the year, the month, and the current day of month.
HH:MM:SS is the time in hours, minutes, and seconds. The time is always sent as Coordinated Universal Time (UTC). An offset needs to be applied to UTC to obtain local time. For example, Mountain Time in the U. S. is 7 hours behind UTC during Standard Time, and 6 hours behind UTC during Daylight Saving Time.
TT is a two-digit code (00 to 99) that indicates whether the United States is on Standard Time (ST) or Daylight Saving Time (DST). It also indicates when ST or DST is approaching. This code is set to 00 when ST is in effect, or to 50 when DST is in effect. During the month in which the time change actually occurs, this number will deincrement every day until the change occurs. For example, during the month of October, the U.S. changes from DST to ST. On October 1, the number will change from 50 to the actual number of days until the time change. It will deincrement by 1 every day, and reach 0 the day the change occurs.
L is a one-digit code that indicates whether a leap second will be added or subtracted at midnight on the last day of the current month. If the code is 0, no leap second will occur this month. If the code is 1, a positive leap second will be added at the end of the month. This means that the last minute of the month will contain 61 seconds instead of 60. If the code is 2, a second will be deleted on the last day of the month. Leap seconds occur at a rate of about one per year. They are used to correct for irregularity in the earth's rotation. UT1 is a correction factor for converting UTC to an older form of universal time that is still used in navigation. It is always a number ranging from -0.8 to +0.8 seconds. This number is added to UTC to obtain UT1.
msADV is a five-digit code that displays the number of milliseconds that NIST advances the time code. It is originally set to 45.0 milliseconds. If you return the on-time marker it will change to reflect the actual one way line delay.
The label UTC(NIST) is contained in every time code. It indicates that you are receiving Coordinated Universal Time (UTC) from the National Institute of Standards and Technology (NIST).
The on-time marker (OTM) is a single character sent at the end of each time code. The OTM is originally an asterisk (*) and changes to a pound sign (#) if ACTS has successfully calibrated the path.

For more information about ACTS and to obtain software, see:

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

## **Section 1.4 - Radio Clocks**

One shortcoming of both Internet and dial-up services is that they require you to connect to a time server each time you need to set your clock. If you need to keep your PC's clock on the right second all the time, this requires an Internet connection that is always on, or a substantial number of phone calls at substantial cost. You might find that your application demands continuous access to an accurate time code, without making phone calls or without having a network connection.

If this is the case, you can get accurate time, all the time, by using a radio clock. There are a number of different types of radio clocks that receive time codes transmitted by radio. The costs vary widely, from less than \$200 to \$20,000 or more. Radio clocks come in several different forms. Some are standalone devices with a digital time display. These are often interfaced to the PC (or to other types of computers) using an interface like the RS-232, RS-422, or IEEE-488. Others are available on plug-in expansion cards that work on the PC or AT bus. Your application can use a radio clock to constantly set the PC clock, or it can use a software driver to get all of its timing information from the radio clock and bypass the PC clock entirely.

Before purchasing a radio clock, make sure that the signal you choose is usable in your area. Also, remember that you'll need to be able to mount an outdoor antenna so you can receive the radio signal. Radio clocks are available that receive WWV/WWVH, WWVB, and GPS.

Radio clocks often provide the reference for Internet time servers, and many of the radio clocks listed in this table support one or more of the protocols listed in Table 1.20. They also often support protocols supported by other operating systems or local area networks. Be sure and check that the clock that you choose produces time in a format compatible with your application.

For a current list of radio clock manufacturers, see:

<http://www.boulder.nist.gov/timefreq/general/receiverlist.htm>

## Section 1.5 - Replacing the Hardware Clock

You can also improve the timekeeping performance of your PC by replacing its clock with a better clock. We mentioned that the hardware clock on a PC uses a very low cost (less than \$1) time base oscillator. If you use a clock with a better time base, you can obviously keep better time.

Precision clock boards are available that plug into the PC bus and effectively replace the hardware clock. These boards include a better time base oscillator than the one inside the PC and allow you to use an external oscillator to get even better results. This means that if you have access to a frequency standard (like a quartz, rubidium, or cesium oscillator) you can use it as the time base. With a good time base oscillator, these boards can keep the correct time for many years. Of course, you will still need to synchronize the clock and check it occasionally.

Another advantage of using a precision clock board is resolution. As we mentioned earlier, the software clock in a PC-compatible has a resolution of 55 milliseconds, and the hardware clock only provides 1-second resolution. There are many applications in science and metrology where more resolution is necessary. With the best precision clock boards, sub-microsecond resolution is attainable.

As we have seen, there are a number of ways to keep accurate computer time. A wide variety of computer timekeeping products and services exist. Services are available to synchronize computer clocks through dial-up, network, and radio links. Precision clock boards are available to replace the built-in hardware clock. In short, if you have a computer timekeeping problem, a solution is readily available.

*The mention of company or product names in this paper does not constitute any endorsement by the National Institute of Standards and Technology.*