

SETTING UP AN NTP SERVER AT THE ROYAL OBSERVATORY OF BELGIUM

Fabian Roosbeek, Pascale Defraigne, and André Somerhausen
Royal Observatory of Belgium

Abstract

This paper describes the setup of an NTP server for time synchronization via the Internet at the Royal Observatory of Belgium. The time server is realized by a “heartbeat” system of two PC servers connected to an NTS-3000 time server synchronized on three different sources: (1) UTC (ORB), our national time scale; (2) UTC broadcast by GPS satellite system; and (3) DCF-77 System.

The precision of the synchronization between our time server and the different time sources is a few microseconds. A remote synchronization on our server, using the NTP protocol, is directly traceable to UTC with an utmost precision of a few milliseconds when the synchronization is done under optimal conditions.

1. INTRODUCTION

As a governmental institution, one of the missions of the Royal Observatory of Belgium (ROB), and in particular of its time lab, is to integrate Belgium in international space-time reference systems and, hence, to be able to provide precise time to Belgian official institutions, Belgian industry, and private users. Even though Internet timekeeping is well established in the industrial countries all over the world, not so many time servers have a direct link to UTC. By setting up a facility for timekeeping via the Internet, a direct connection to our time scale UTC (ORB) and, by extension, to UTC is established.

NTP is a good choice for time synchronization in a variety of circumstances, specifically for Internet environments. Flexibility of the client/server relationship and security methods allow NTP to work well in almost any environment. NTP not only corrects the current time, but it can also keep track of consistent time variations and automatically adjust for time drift of the client. This allows for less network traffic and keeps client clocks more stable, even if the network is unavailable. In addition, the NTP daemon can automatically adjust the time by periodic increments. NTP can also operate through firewalls and has a number of security features.

NTP works on a hierarchical model in which a small number of servers give time to a large number of clients. The clients on each level, or stratum, are, in turn, potential servers to an even larger number of clients of a higher numbered stratum. Stratum numbers increase from the primary (stratum 1) servers to the low numbered strata at the leaves of the tree. Clients can use time information from multiple servers to automatically determine the best source of time and prevent bad time sources from corrupting their own time. Figure 1 illustrates the hierarchical strata model of servers used in NTP.

Under good conditions on a LAN (Local Area Network) without too many routers or other sources of network delay, synchronization to within a few milliseconds is normal. Anything that adds latency, such as hubs, switches, routers, or network traffic, will reduce this accuracy. The synchronization accuracy on a WAN (Wide Area Network) is typically within the range of 10-100 ms. For the Internet, synchronization accuracy is unpredictable, so special care is needed when configuring a client to use public NTP servers.

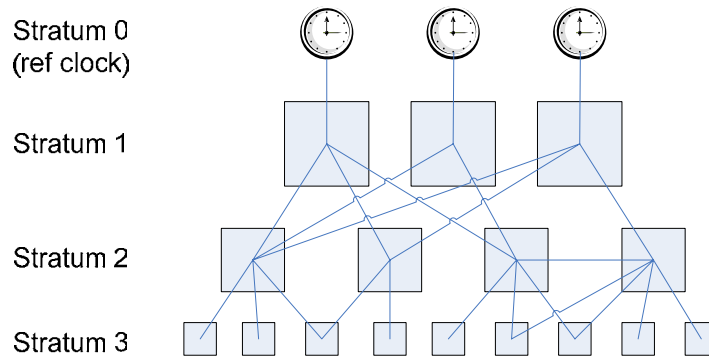


Figure 1. Hierarchical strata model of servers used in NTP.

2. IMPLEMENTATION OF THE TIME SERVER

The time lab of the Royal Observatory of Belgium is presently equipped with five clocks: three HP5071A Cesium clock and two H-Maser clocks (1 active CH1-75 and one passive CH1-76). The UTC realization UTC (ORB) is obtained from the 5 MHz frequency provided by the active H-Maser clock (CH1-75) in which the cavity auto-tuning is realized using the 5 MHz frequency of the passive H-Maser (CH1-76).

The time lab was installed in a new temperature-stabilized room in April 2002. Figure 2 gives a diagram of the time lab, emphasizing the implementation of the time server.

After some investigations, we have chosen to buy a modified version of the NTS-3000 server from the company Elproma (www.ntp-servers.com) for our primary server (see Figure 3). The standard NTS-3000 network time server synchronizes its clock via the GPS satellite system and has also a fully redundant second source of time via DCF-77. We have asked Elproma to add a 1 PPS input to our NTS-3000 server. This 1 PPS input is used for connecting UTC (ORB). The advantage of adding a 1 PPS input to the time server, beside the better precision and the link to a national time scale, is that, contrary to our reference clock, both GPS and DCF-77 are outside the control of the users, and may be subject to changes without notice.

We have chosen not to allow a direct synchronization on the NTS-3000 from our LAN or from the Internet. We have rather setup a “heartbeat” system with two servers (the Proliant family from Compaq) running under Linux OS. The “heartbeat” system provides a full redundancy both from the software and the hardware point of view. If anything fails on one server, the second one takes the relay transparently. For the clients, this heartbeat system is seen as a single system with one address IP, available both from our LAN and from the Internet (see Figure 4). Strictly speaking, the time server is a stratum 2 time server available for time synchronization over the Internet; it synchronizes itself on our stratum 1 time server

(NTS-3000), which in turn has 3 UTC sources: UTC (ORB), GPS, and DCF-77.

The NTP daemon computes, about every 60 seconds, a mean and a standard deviation for each of the different sources available (UTC (ORB), GPS, DCF-77, etc.). The server oscillator is then steered to a value based on an internal algorithm in order to minimize the time offset between UTC and the server clock. The NTS-3000 server clock is then used as the source for the synchronization of the heartbeat system.

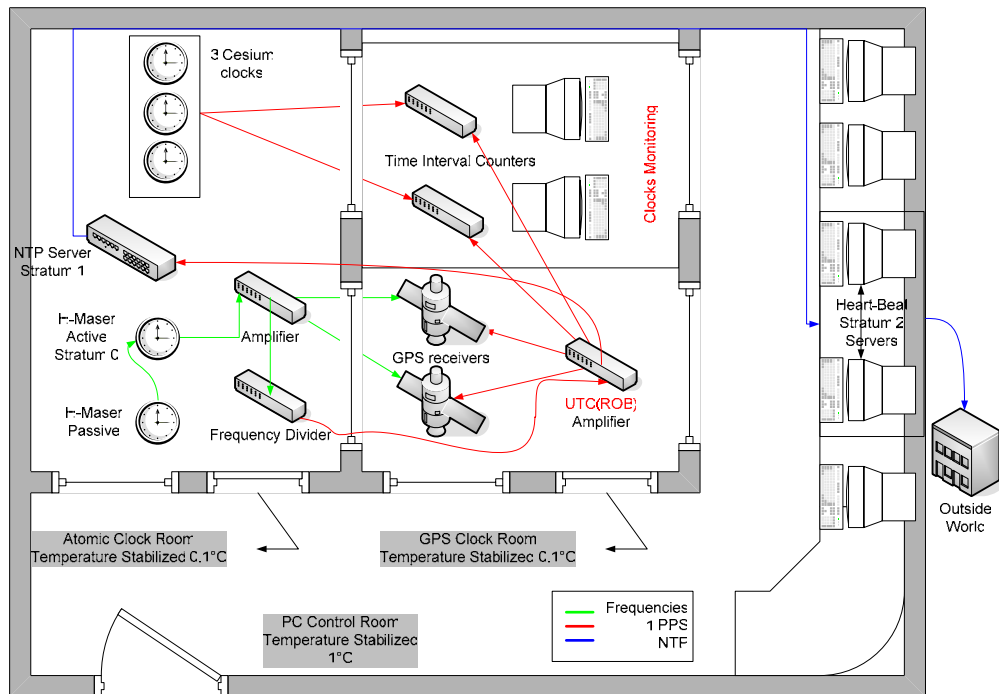


Figure 2. Diagram of the time lab of the ROB, emphasizing the implementation of the time server.



Figure 3. Elproma NTS-3000 Time Server.

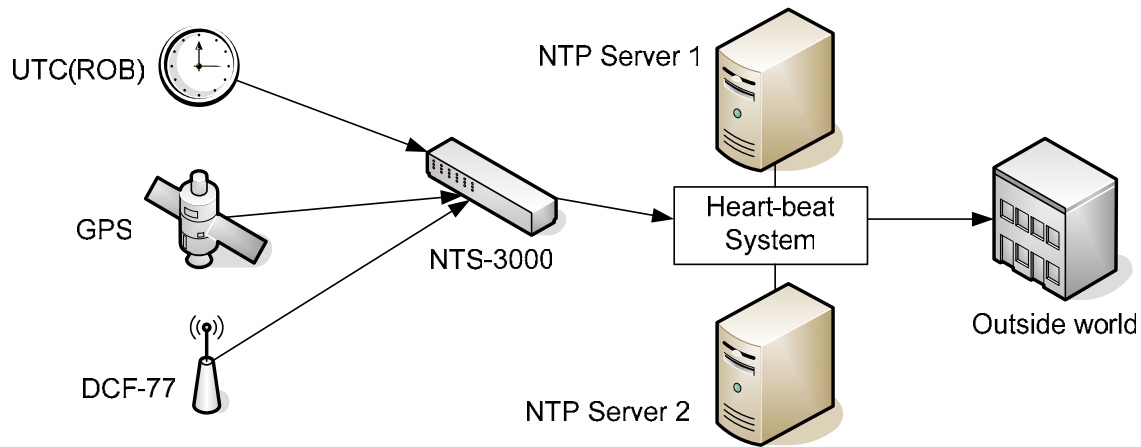


Figure 4. NTP Servers at ROB.

3. PERFORMANCE TESTS

In order to evaluate the performance of the NTS-3000 server, we compare its internal clock to the different UTC sources available (UTC (ORB), GPS, and DCF-77).

Figure 5 shows the time offset and jitter (due to interrupt latency, processing delays, and similar effects) between the NTS-3000 server and the 1 PPS UTC source: UTC (ORB), over a period of 23 days. It can be seen from Figure 5 that the NTP daemon is able to synchronize the clock in the NT-3000 server with a mean of 0.6 microsecond (μs) and with a standard deviation of 5.3 μs . We see also that there is a clear correlation between peaks in the time offset and jitter.

Figure 6 shows the time offset and jitter between the NTS-3000 server and the GPS UTC source over the same period. Due to the fact that our UTC realization is kept very close to UTC and, by extension, to GPS time (maximum difference of 0.1 μs), the results are very similar (mean of 0.8 μs and standard deviation of 5.2 μs).

Figure 7 shows the time offset and jitter between the NTS-3000 server and the DCF-77 UTC source over a few hours. As we can see, the DCF-77 signal is much less precise than the two other sources. But, at the level of precision generally needed for time synchronization (millisecond), the DCF-77 can act as a valuable backup in case of failure of the two main UTC sources.




Figure 5. Time offset and jitter between the NTS-3000 server and UTC (ORB).




Figure 6. Time offset and jitter between the NTS-3000 server and GPS.




Figure 7. Time offset and jitter between the NTS-3000 server and DCF-77.

4. TIME SYNCHRONIZATION BY LAN

For testing purposes, we have set up a PC that synchronizes its clock by NTP protocol to our stratum 2 NTP server system.

Figure 8 shows the delay, time offset, and jitter between this client and the NTP server over a period of 10 days. It can be seen that the NTP daemon is able to synchronize the clock in the client PC with a mean of 0.038 milliseconds (ms) and with a standard deviation of 2.12 ms. Like for the NTS-3000 performance tests, there is a clear correlation between peaks in the time offset and jitter. The delay remains very short, with a mean of 0.29 ms and a standard deviation of 0.21 ms. Only two peaks of a few ms have appeared during the 10-day period.




Figure 8. Time offset and between the ROB time server and a client PC over our LAN.

5. TIME SYNCHRONIZATION BY INTERNET

In order to emphasize the effect of the distance between the client and the PC, we have synchronized the clock of our test PC with the stratum 2 time server of our Internet Provider (BELNET), located in Brussels at a few kilometers from the Observatory.

Figure 9 shows the delay, time offset, and jitter between this client and the NTP server of BELNET over a period of 7 days. It can be seen that the NTP daemon is able to synchronize the clock in the client PC with a mean of 0.69 ms and with a standard deviation of 6.66 ms. The delay is of course more important than over our LAN, with a mean of 1.56 ms and a standard deviation of 0.89 ms. A lot of peaks of a few ms in the delay appear now during the 7-day period.

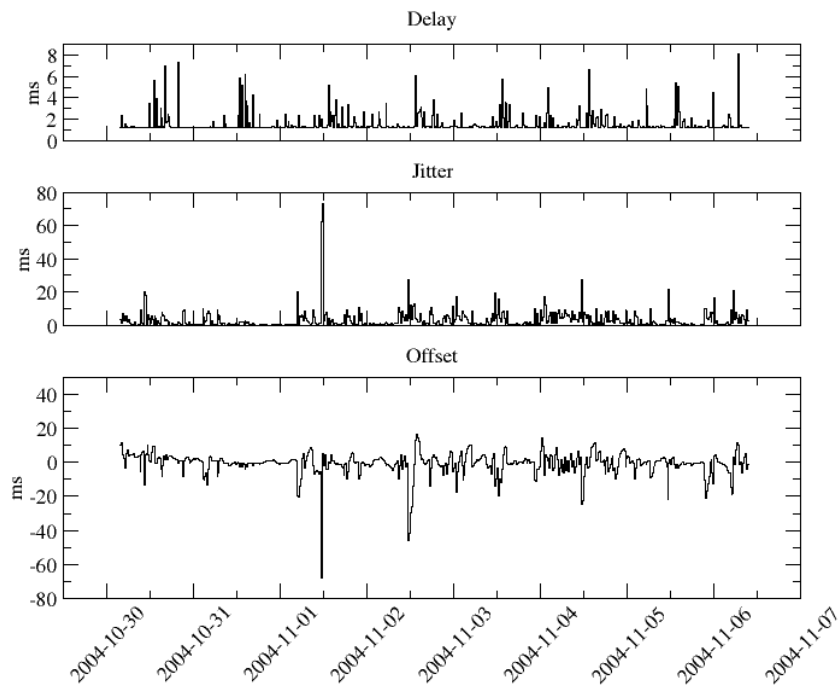


Figure 9. Time offset and between the BELNET time server and a client PC over the Internet.

6. HIGH TRAFFIC LOAD

In order to see if our time server is able to handle many requests at the same time, we simulated high traffic load condition by stressing the CPU with benchmark programs, as well as by flooding the time server with millions of ping requests. After one hour of tests, no impact was seen on the offset or on the delay between a client and the NTP server. Of course, this could not be the same under real high traffic load coming from many NTP requests, which could not be simulated.

7. CONCLUSION

We have described the implementation of a time server system at the Royal Observatory of Belgium for time synchronization via the Internet.

One advantage of this system is its direct traceability to UTC via our local realization UTC (ORB). Another advantage is the reliability; the time server is composed of two time servers linked together with a heartbeat system, which is in turn synchronized to an NTS-3000 time server connected to three different UTC sources: the primary UTC source is UTC (ORB), the second UTC source is a GPS receiver and the third is a DCF-77 receiver. All these components are located in a temperature-stabilized room.

Performance tests show that the NTS-3000 server is synchronized to UTC (ORB) or GPS within a precision of less than 6 μ s, while a synchronization based on DCF-77 has a precision limited to a few ms. Time synchronization on our time server over our LAN from a client PC gives a precision of less than 3

ms with very small delay or jitter. Synchronization over the Internet increases the delay as well as the jitter, resulting in a precision of about 7 ms in the time synchronization on a baseline of a few kilometers.

NOTE: The Internet address of the time server at the Royal Observatory of Belgium is *ntp.oma.be*. The time server is not official as of November 2004 and a full service is not guaranteed until the server is declared officially operational.

REFERENCES

- [1] D. Deeths and G. Brunette, 2001, “Using NTP to Control and Synchronize System Clocks,” Sun BluePrints™ OnLine (<http://www.sun.com/blueprints>).
- [2] D. L. Mills, 1992, “Network Working Group Request for Comments: 1305,” University of Delaware (<http://www.ietf.org/rfc/rfc1305.txt>).

