URSI GASS 2020, Rome, Italy, 29 August - 5 September 2020

Synchronizing Stock Market Clocks to UTC(NIST)

Michael A. Lombardi
Time and Frequency Division, National Institute of Standards and Technology (NIST), Boulder, CO, US, 80305
lombardi@nist.gov

Abstract

To reduce the possibility of fraudulent activity and market manipulation, the world's major stock exchanges require every clock involved in a stock market transaction to be synchronized to agree with a common reference clock that keeps accurate and internationally traceable time. This paper describes a distributed system that synchronizes stock market clocks to UTC(NIST), the national time standard for the United States. This system is currently installed at stock market data centers in the United States, the European Union, and Asia, and meets the time synchronization regulatory requirements for all three regions.

1 Time Accuracy Requirements

The accuracy and reference clock requirements for stock market clocks differ in the United States (US) and the European Union (EU), as summarized in Table 1. The time accuracy requirement is defined as the maximum allowable time offset with respect to the reference clock.

Table 1. Summary of stock market synchronization requirements.

Region	Reference Clock	Time accuracy requirement	
US [1, 2]	UTC(NIST)	Automated orders	50 ms
		Manual orders	1 s
EU [3]	Any time scale that contributes to UTC	Manual orders	1 s
		High frequency trading	100 µs
		All other trading	1 ms

Three points should be noted when looking at Table 1. First, the time scale of the National Institute of Standards and Technology (NIST) is the reference clock for US stock markets, whereas any time scale (including NIST) that contributes to Coordinated Universal Time (UTC) can serve as the reference clock in the EU. Second, the 100 µs accuracy requirement for high frequency trading (HFT) in the EU is 500× more stringent than the 50 ms requirement in the US for automated orders. However, many US stock exchanges already maintain synchronization to within 100 µs of NIST [2]. Finally, note that the accuracy requirements are applicable to all clocks involved in a transaction, including the clocks in server and client computers. This is important, because whereas accuracy to within tens of nanoseconds is fairly easy to obtain, for example, with a global navigation satellite system (GNSS) clock installed at a stock market data center; the delay asymmetries found in computer networks, in networking hardware, and even in application software and operating systems, typically limit the accuracy of computer clocks to tens of microseconds, or 1000× worse. Thus, for the computer clocks that time stamp transactions, the 100 µs accuracy requirement can be difficult to achieve and verify.

The NIST system not only provides a reference clock at stock market sites that can be synchronized to ~10 ns (0.01 µs) with respect to UTC(NIST), but also continuously monitors and verifies that the time kept by stock market computer clocks is within 100 µs. By doing so it allows stock exchanges to demonstrate to the US Securities and Exchange Commission (SEC) and other regulatory agencies that they comply with all time synchronization requirements. It also establishes a continuous traceability chain that originates with UTC and UTC(NIST) and extends to all server and client computer clocks.

2 NIST Disciplined Clock (NISTDC)

The heart of the distributed synchronization system is a NIST disciplined clock (NISTDC), which is typically co-located in the same data center as the computer systems that record the stock exchange's transactions. The NISTDC is a rack mounted instrument (Figure 1) with a variety of configuration options.

Figure 1. Front panel of NIST disciplined clock.

When fully configured the NISTDC chassis includes a network connection, a GNSS receiver and associated time interval measurement hardware, either a rubidium or cesium atomic clock (rubidium clocks are integrated inside the NISTDC chassis but cesium clocks are rack mounted separately), a computer time server supplying both the network time protocol (NTP) and precision time protocol (PTP), an event timing board used to measure the accuracy of packets sent by stock exchange time servers, and an amplifier that distributes the atomic clock signals to time servers and other stock market clocks [4]. The distribution amplifier provides two 10 MHz frequency outputs and eight 1 pulse per second (pps) time outputs that are synchronized to UTC(NIST), as shown in the block diagram in Figure 2.

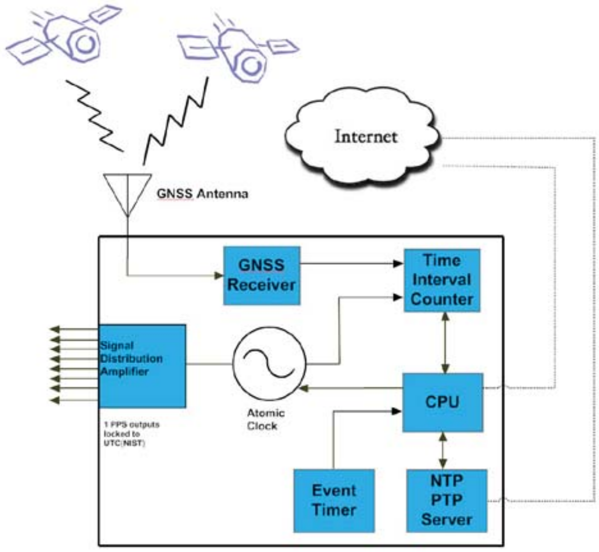

Figure 2. Block diagram of NIST Disciplined Clock.

Time from the UTC(NIST) time scale in Boulder, Colorado is transferred to the NISTDC by use of the common-view disciplining method first described in [5]. The NISTDC continuously adjusts its local rubidium or cesium atomic clock by applying frequency and/or time corrections obtained through common-view observations of GNSS satellites. As illustrated in Figure 3, a system at NIST measures the time difference UTC(NIST) – GNSS, and measurements performed at the NISTDC site produce NISTDC – GNSS. Every 10 minutes, each NISTDC and the NIST system simultaneously send their measurement data to an Internet cloud server where the NISTDC data are subtracted from the NIST data. This removes the contribution of GNSS time and results in an estimate of the UTC(NIST) – NISTDC time difference. The NISTDC converts this time difference to a frequency correction by use of an adaptive proportional-integral-derivative (PID) controller and then applies the correction to its local atomic clock. The process is continuously repeated to keep the NISTDC locked to UTC(NIST).

The NISTDC is considered locked when it is accurate to within 50 ns (0.05 µs) of UTC(NIST) and stable to within 5 ns (0.005 µs). However, it internally distinguishes between a soft lock based on the 50/5 criteria, and a hard lock that requires accuracy to within 10 ns (0.01 µs) of UTC(NIST) and stability to within 2 ns (0.002 µs). The hard lock condition is always maintained during normal operating conditions.

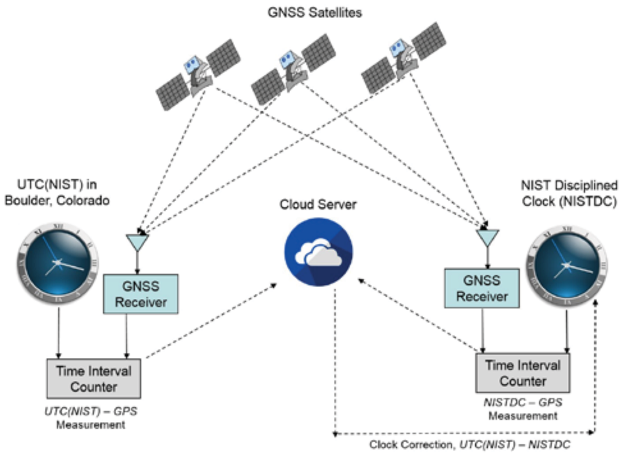

Figure 3. NISTDC common-view link to UTC(NIST).

At this writing (January 2020), NISTDC units synchronize some of the world's largest stock exchanges and are installed at data centers near New York City and Chicago in the US, as well as in London and Frankfurt in the EU, and in Tokyo, Japan. Stock exchange clients who utilize the NISTDC have full 24/7 access to their clock data via a web portal.

3 Accuracy compared to UTC(NIST)

Common-view observations routinely show that a locked NISTDC seldom deviates by more than ±10 ns (±0.01 µs) from UTC(NIST) and that its average time offset is near 0. To illustrate this, Figure 4 shows a 6-month (July to December 2019) comparison of a NISTDC, located at a major US stock exchange, to UTC(NIST). The peak-to-peak variation over the 6-month interval is ~25 ns, but most data points fall within ±5 ns and the average time offset is less than 0.1 ns, or essentially 0.

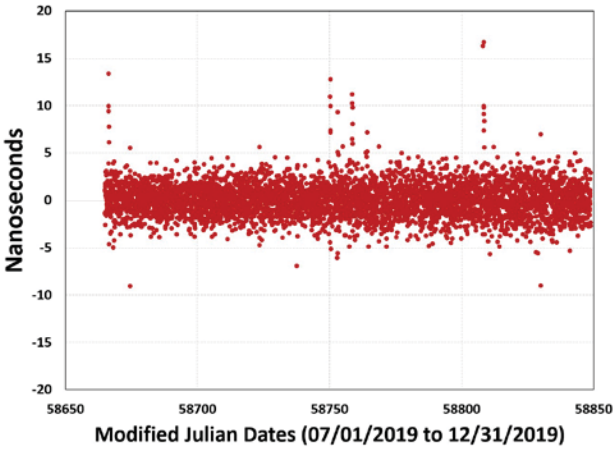

Figure 4. NISTDC – UTC(NIST) time differences.

Even when common-view data show that the time offset of a NISTDC is essentially 0, the actual time offset may be larger due to uncertainties in the common-view method. These uncertainties (k = 2, or 2σ) typically range from ~10 ns (0.01 µs) in the best case to ~50 ns (0.05 µs) in the worst case. A complete NISTDC uncertainty evaluation is provided in [4].

Figure 5 shows the time deviation (stability) of this data set for averaging periods ranging from one hour to about one month. After averaging for one hour, the stability is about 1.5 ns, dropping below 0.4 ns after one day and below 0.2 ns after one week. This high level of stability is possible because the time differences between UTC(NIST) and the NISTDC are always compensated for by the common-view corrections.

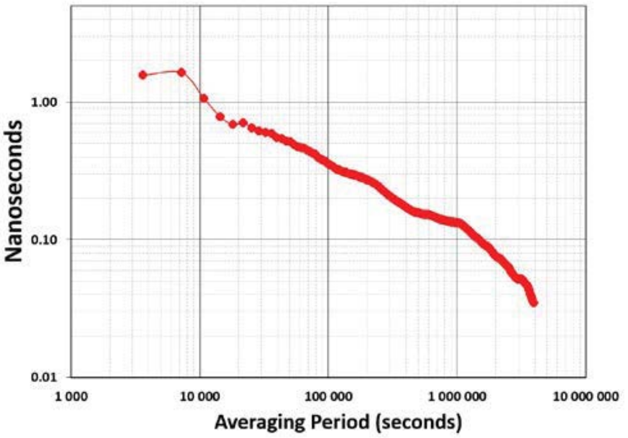

Figure 5. NISTDC stability with respect to UTC(NIST).

4 Reliability and holdover performance

Stock exchange activities have a huge impact on the global economy, and thus the reliability of trading platforms, including synchronization systems, is of critical importance. To provide the highest level of synchronization reliability, the NISTDC software was recently enhanced to implement a patented technique [6] known as a multi-source common-view disciplined clock (MSCVDC) with fail-safe redundancy. The MSCVDC method protects against UTC(NIST) failures, GNSS reception failures, and network failures. For example, if the UTC(NIST) time scale in Boulder, Colorado is unavailable or malfunctioning, the NISTDC automatically locks to the NIST back up time scale located ~80 km away in Fort Collins, Colorado. Other time scales can be added to provide even more redundancy. If all reference time scales are unavailable due to a network outage, the NISTDC automatically locks to the GNSS time signals. When the NISTDC is forced to switch from the primary to the backup NIST time scale, or from either NIST time scale to GNSS, it typically results in an insignificant time step of < 20 ns (0.02 µs), such as the small outliers shown in Figure 4. If the GNSS constellation providing the common-view signals is unavailable, for example if the Global Positioning System (GPS) is being jammed, the MSCVDC technique potentially allows switching to another satellite constellation such as Galileo or GLONASS, providing the NISTDC with common-view signal redundancy. In extreme situations where all reference time scales and all GNSS signals are unavailable, the NISTDC enters holdover mode, where its performance then depends on the accuracy and stability of its free running atomic clock.

A NISTDC in holdover mode can stay within 1 µs of UTC(NIST) for at least 48 hours with a rubidium clock, and for at least several months with a cesium clock. To illustrate this, Figure 6 shows the performance of a rubidium NISTDC in holdover mode. This device maintained 1 µs synchronization for ~73 hours after its GNSS antenna was disconnected. The error increased to ~5 µs after ~110 hours when GNSS reception was restored and it began to relock to UTC(NIST).

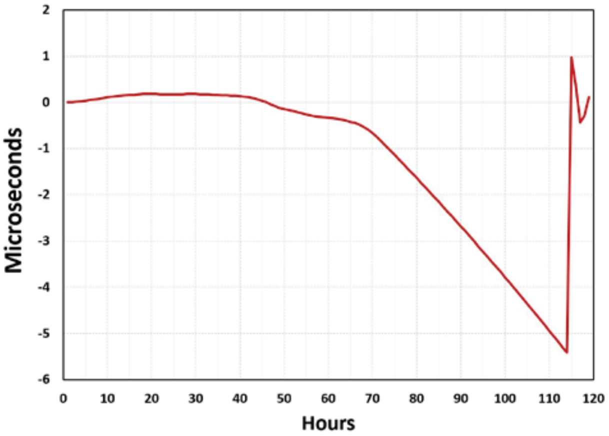

Figure 6. A rubidium-based NISTDC in holdover mode.

Figure 7 graphs the performance of a cesium NISTDC in holdover mode that remained within 300 ns (0.3 µs) of UTC(NIST) after free running for about eight months. Because the cesium clock frequency had been optimally adjusted while it was locked to UTC(NIST), the time offset increased at a rate of just 1.2 ns per day (frequency offset of ~1×10^{-14}) when in holdover mode.

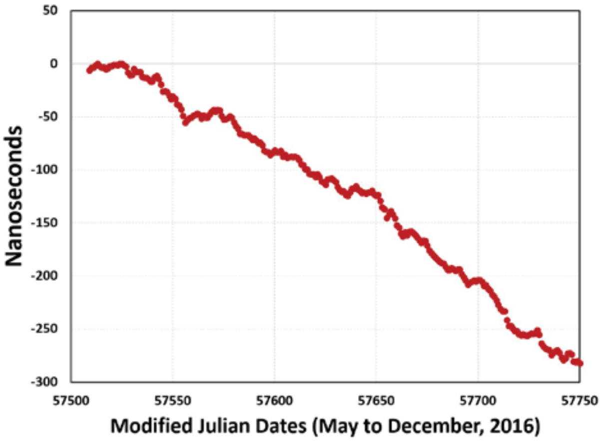

Figure 7. A cesium-based NISTDC in holdover mode.

The combination of the MSCVDC method and rubidium clock holdover usually provides enough time to troubleshoot and correct most failure modes. When coupled with cesium clock holdover, the MSCVDC method assuages most stock market concerns related to reliability, vulnerability, and GNSS dependency.

5 Time server synchronization

Time servers co-located in the same data center as a NISTDC are often referenced to GNSS signals. Preferably, however, they should utilize the 1 pps signals distributed from the NISTDC via coaxial cable as their reference for frequency and time interval. If the NISTDC has been configured to include the time code output option, the servers can also periodically obtain time-of-day information from the NISTDC's integrated NTP/PTP server, and thus be entirely controlled by NIST time signals, without dependence on GNSS. The integrated NTP/PTP server automatically adjusts when necessary to account for leap seconds.

6 Time server monitoring

The NISTDC can monitor the performance of up to 12 time servers. It does so by simulating a client computer and requesting NTP packets from each server every 10 seconds. It then uses an event timing board with 0.1 µs resolution to compare the time stamps in the received packets to the time kept by the NISTDC [4]. When the NISTDC resides on the same local area network (LAN) as the time server, the measured time offset of a properly synchronized server typically ranges from a few microseconds to about 50 µs when compared to the NISTDC. Figure 8 shows the time offset of an NTP server, operated by a major US stock exchange, when compared to UTC(NIST) during the last 100 days of 2019 (one data point per hour). The average time offset of the NTP server clock is 20 µs and its time deviation (stability) is < 1 µs at an averaging period of 1 day.

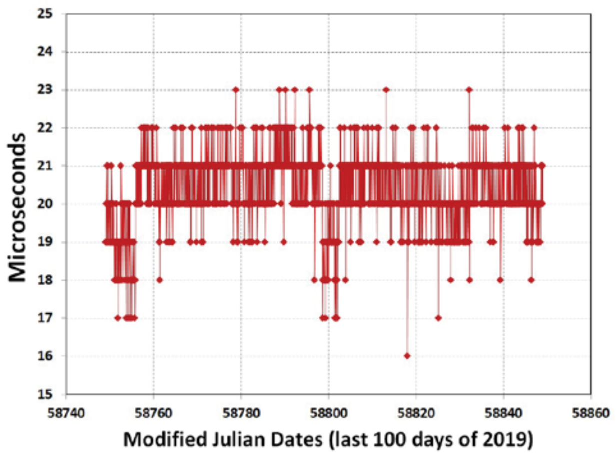

Figure 8. Accuracy of NTP time server operated by a major US stock exchange, with respect to UTC(NIST).

The server measurements have an uncertainty that equals one half of the asymmetry in network delays, and thus is likely to be small if the round trip delay between the server and client is also small. In a typical LAN configuration, such as the one utilized for the Figure 8 measurement, the round trip delay between the server and the NISTDC client is small, typically ~300 µs. In this case, if we conservatively estimate that the network asymmetry is 1.05:1, meaning that the client to server path delay is 5% larger than the server to client path delay, the uncertainty is still < 5 µs.

7 Summary

To help stock exchanges keep accurate time, NIST has implemented a system that synchronizes stock market clocks to UTC(NIST), the national time standard for the United States. By placing a NISTDC inside their data center and then using it to synchronize all other clocks, stock exchanges can demonstrate to regulatory agencies that they meet all synchronization requirements, including the 100 µs requirement for HFT in the EU.

This paper is a contribution of the United States government and is not subject to copyright.

8 References

1. Financial Industry Regulatory Authority (FINRA), "Rule 6820. Clock Synchronization," Adopted March 15, 2017.

2. United States Securities and Exchange Commission (SEC), "Joint Industry Plan; Order Approving the National Market System Plan Governing the Consolidated Audit," Release No. 34-79318, File No. 4-698, November 15, 2016.

3. European Union, "Commission Delegated Regulation (EU) 2017/574 of 7 June 2016 supplementing Directive 2014/65/EU of the European Parliament and of the Council with regard to regulatory technical standards for the level of accuracy of business clocks," *Official Journal of the European Union*, pp. L 87/148 – L 87/151, March 31, 2017.

4. M. Lombardi, A. Novick, B. Cooke, and G. Neville-Neil, "Accurate, Traceable, and Verifiable Time Synchronization for World Financial Markets," *Journal of Research of the National Institute of Standards and Technology*, vol. 121, pp. 436–463, 2016. http://dx.doi.org/10.6028/jres.121.023

5. M. Lombardi and A. Dahlen, "A common-view disciplined oscillator," *Rev. Sci. Instrum.*, vol. 81, no. 5, pp. 055110–055111/6, 2010. DOI: 10.1063/1.3430071.

6. M. Lombardi, "Disciplined clock for providing a disciplined time and a disciplined frequency synchronous with a reference clock," *United States Patent 10,466,363 B2*, granted November 5, 2019.