Synchronizing Stock Market Clocks to UTC(NIST)

Michael A. Lombardi
National Institute of Standards and Technology (NIST)
Boulder, Colorado, USA
lombardi@nist.gov

XXXIII General Assembly and Scientific Symposium (GASS) of the International Union of Radio Science
(Union Radio Scientifique Internationale - URSI)
August 2020

Introduction

- ✓ To reduce the possibility of fraudulent activity and market manipulation, the world's 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.

Why time synchronization is important to stock markets

Regulators are concerned with the fraudulent activity and market manipulation that can occur when stock market trades are not processed in the order that they are received.

- ✓ To prevent trades from being executed in an order that differs from when they were received, every part of a stock market transaction must be recorded and time stamped.
- ✓ To ensure that time stamps are accurate, every clock involved in a stock market transaction must be synchronized with a common reference clock.
- ✓ The common reference clock must be a trusted source of accurate and internationally traceable time. In the United States, the common reference clock is UTC(NIST).
- ✓ Operating a fair and equitable stock exchange requires all stock market clocks to be synchronized to within a known accuracy specification. This specification is provided by the regulatory bodies that oversee the stock markets.

High Frequency Trading (HFT)

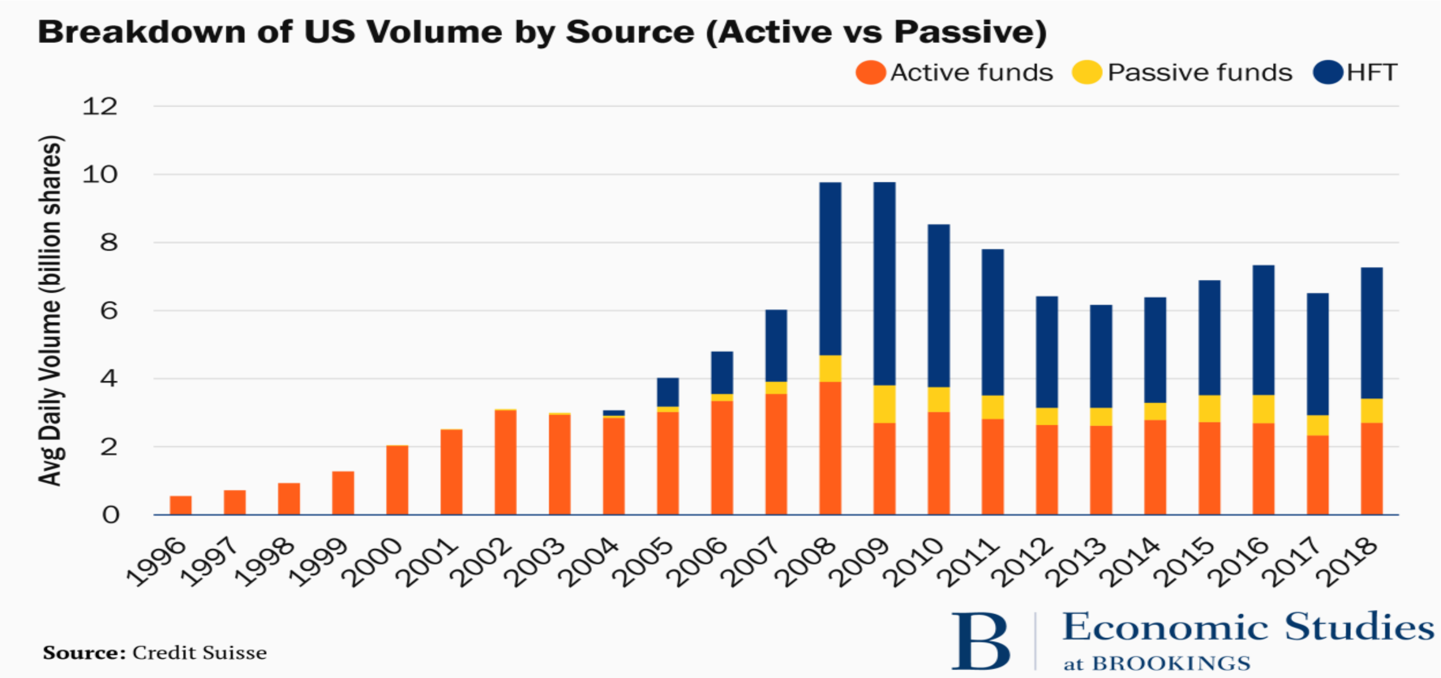

High frequency trading (HFT) is an offshoot of electronic trading that began in the late 1990s. It involves the use of computers running complex trading algorithms that automatically place and execute orders based on market conditions. It has made markets harder to regulate and accurate time synchronization more important.

- ✓ HFT allows trades to execute in intervals measured in microseconds.
- ✓ Less than one billion shares per day were traded before electronic trading began in the late 1990s. Trading volume increased to nearly 10 billion shares per day in 2008 and 2009 when HFT became commonplace.
- ✓ HFT did not significantly impact volume until 2004, but has accounted for at least 50% of volume since 2008.

Synchronization Requirements (US and EU)

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

- ✓ 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.
- ✓ 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.
- ✓ The accuracy requirements are applicable to all clocks involved in a transaction, including the clocks in server and client computers. Ensuring computer clock accuracy is the most difficult part of meeting regulatory requirements.

NIST disciplined clocks (NISTDC) provide UTC(NIST) synchronization to numerous stock market locations

- ✓ A NISTDC is a rack mount instrument (PC-based) that is co-located in the same data center as the computers that record stock exchange transactions.
- ✓ All NISTDCs are calibrated at NIST prior to shipment to compensate for cable and hardware delays.
- ✓ NISTDC units currently 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.

Block diagram of NISTDC

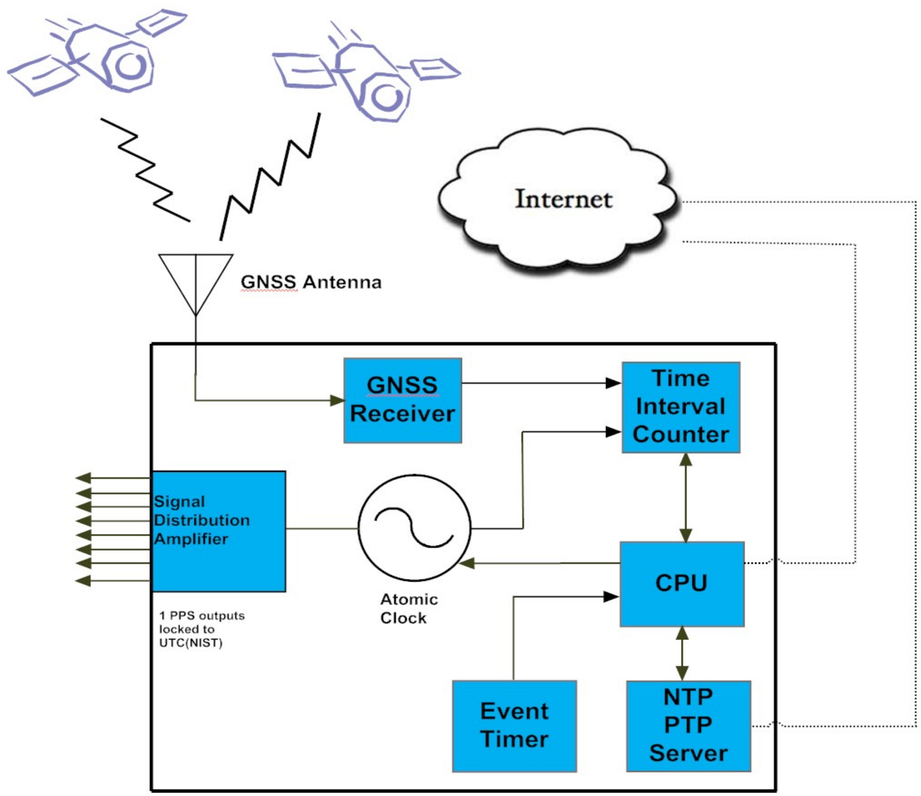

✓ When fully configured the NISTDC chassis includes a network connection, a GNSS receiver and associated time interval measurement hardware, either a rubidium or cesium atomic clock (rubidium clocks are integrated inside the NISTDC chassis but cesium clocks are rack mounted separately), a computer time 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.

How a NISTDC works – Part I

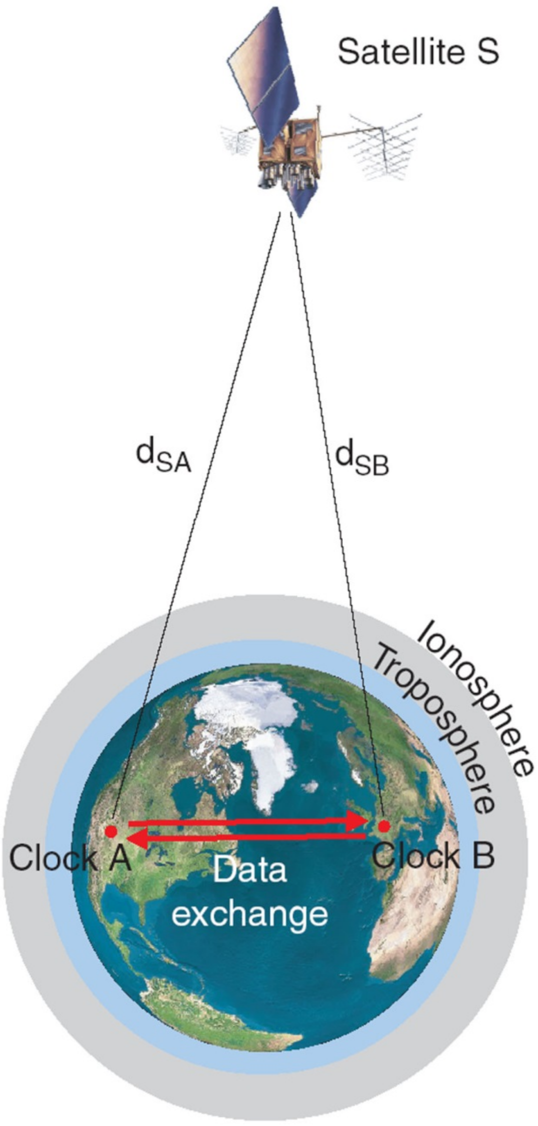

- ✓ Common-view GNSS signals are used to “relay” time from the NIST time scale to the stock market clocks.
- ✓ Common-view simply means that GNSS signals can be received nearly simultaneously at NIST and at each stock exchange.

How a NISTDC works - Part II

A measurement system is installed at the NIST time scale site in Boulder, Colorado. This system continuously measures the time difference between UTC(NIST) and Global Navigation System Satellites (GNSS).

Two time difference measurements are made:

- ✓ *NIST – GNSS* (recorded in Boulder, Colorado)
- ✓ *NISTDC – GNSS* (recorded at each stock market site)

Both measurement results are uploaded to a cloud server every 10 minutes. A clock correction for the NISTDC is obtained by simply subtracting the measurement made at NIST from the measurement made by the NISTDC. For example:

- ✓ *NISTDC Clock Correction* = *(NIST – GNSS) – (NISTDC – GNSS)*

A control loop continuously applies clock corrections to the NISTDC to adjust its frequency and time so that the device stays "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.

The NISTDC locks to UTC(NIST) via common-view GNSS

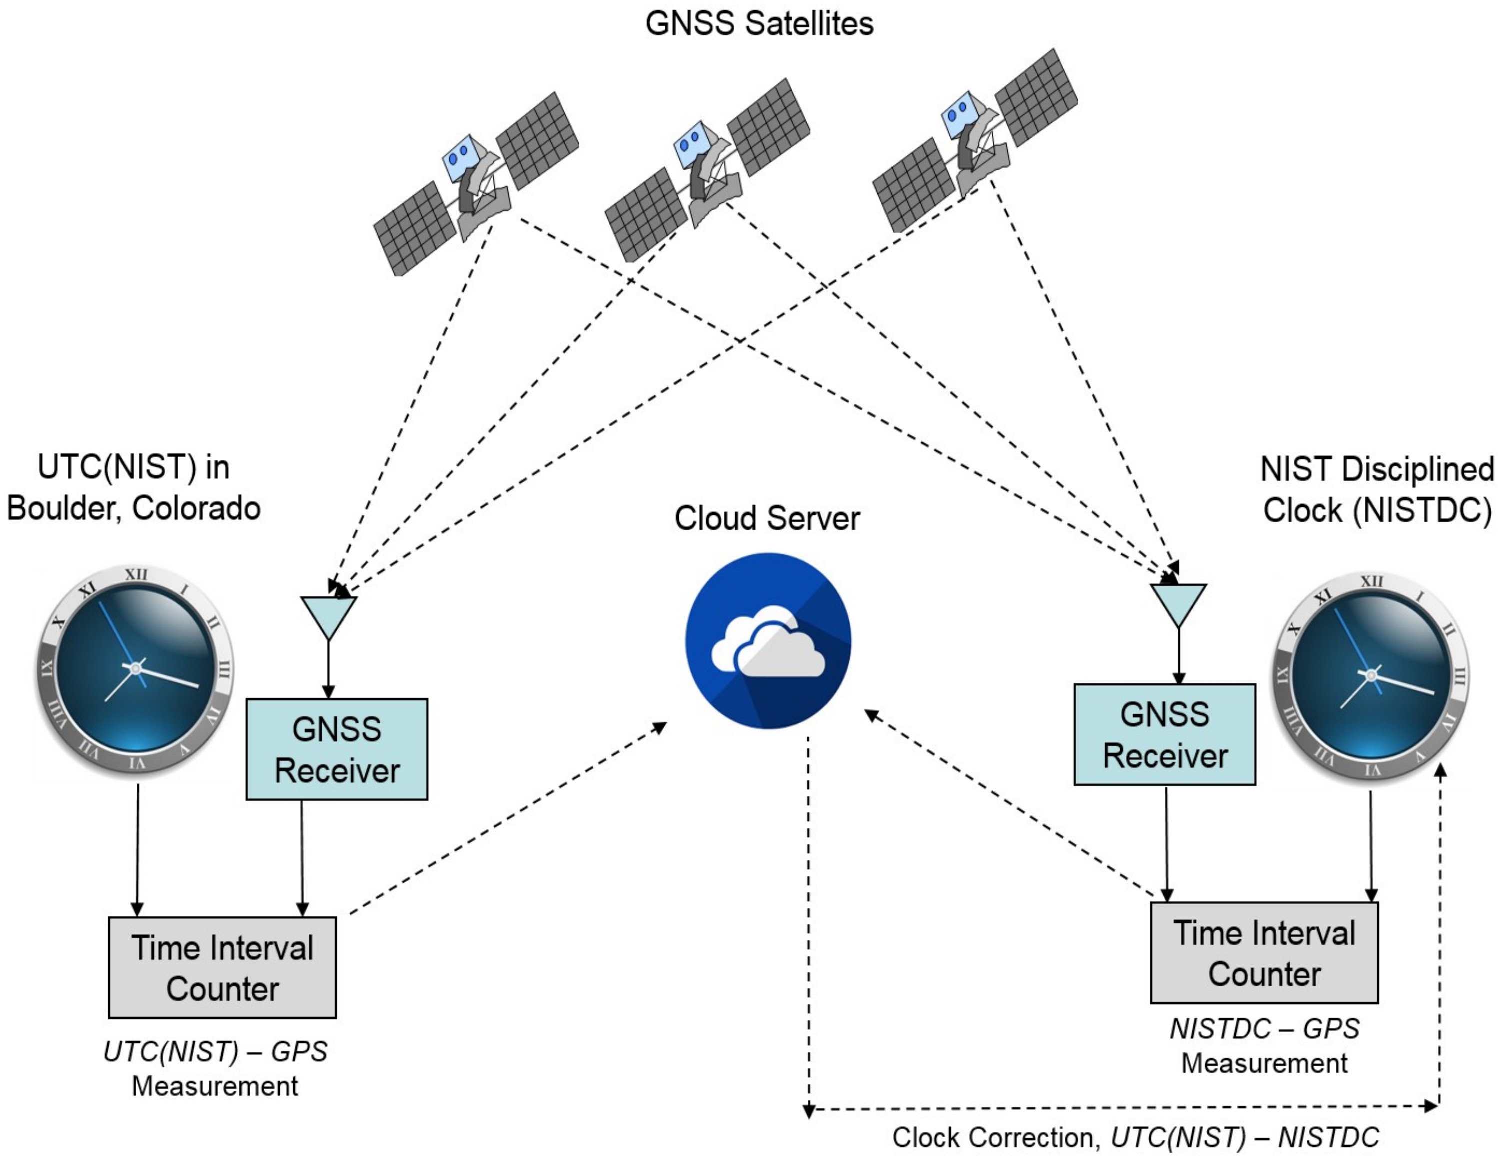

Accuracy of NISTDC, at major US stock exchange, compared to UTC(NIST)

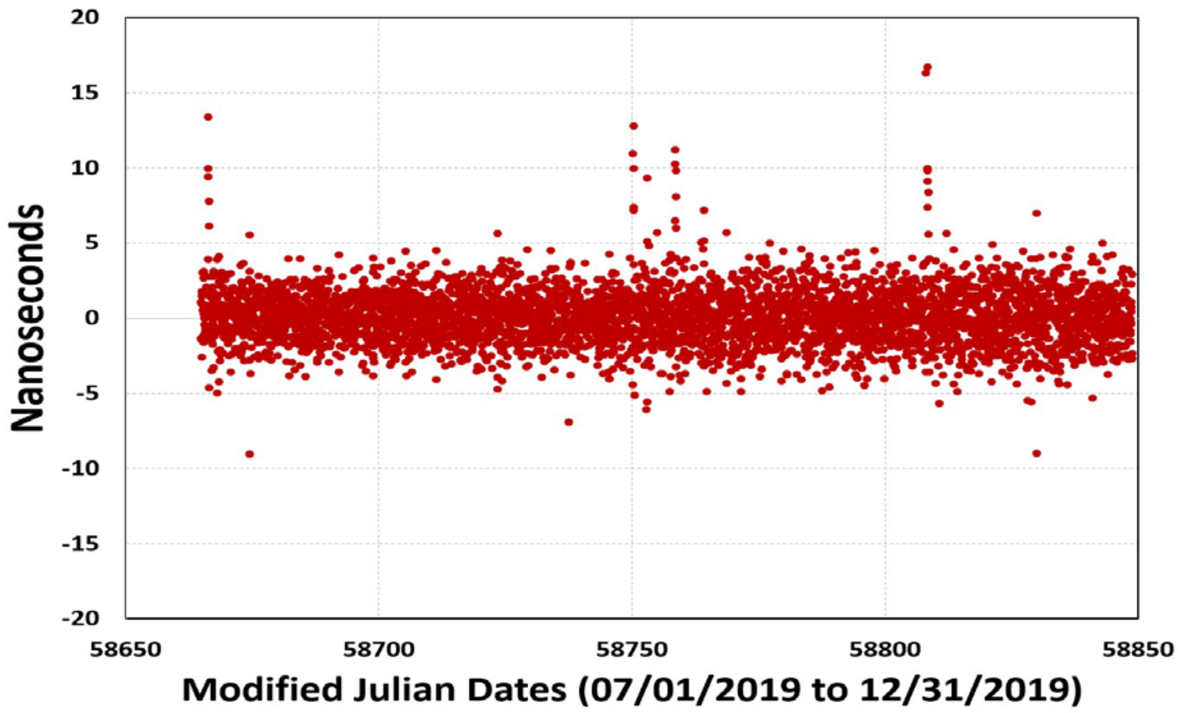

- ✓ A locked NISTDC seldom deviates by more than ±10 ns (±0.01 µs) from UTC(NIST).
- ✓ The peak-to-peak variation over the 6-month interval shown in the graph is ~25 ns, but most data points fall within ±5 ns and the average time offset is less than 0.1 ns, or essentially 0.
- ✓ The reported time differences between the NISTDC and UTC(NIST) may actually 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. These uncertainties are estimated and reported to stock market clients.

Time deviation (stability) of NISTDC

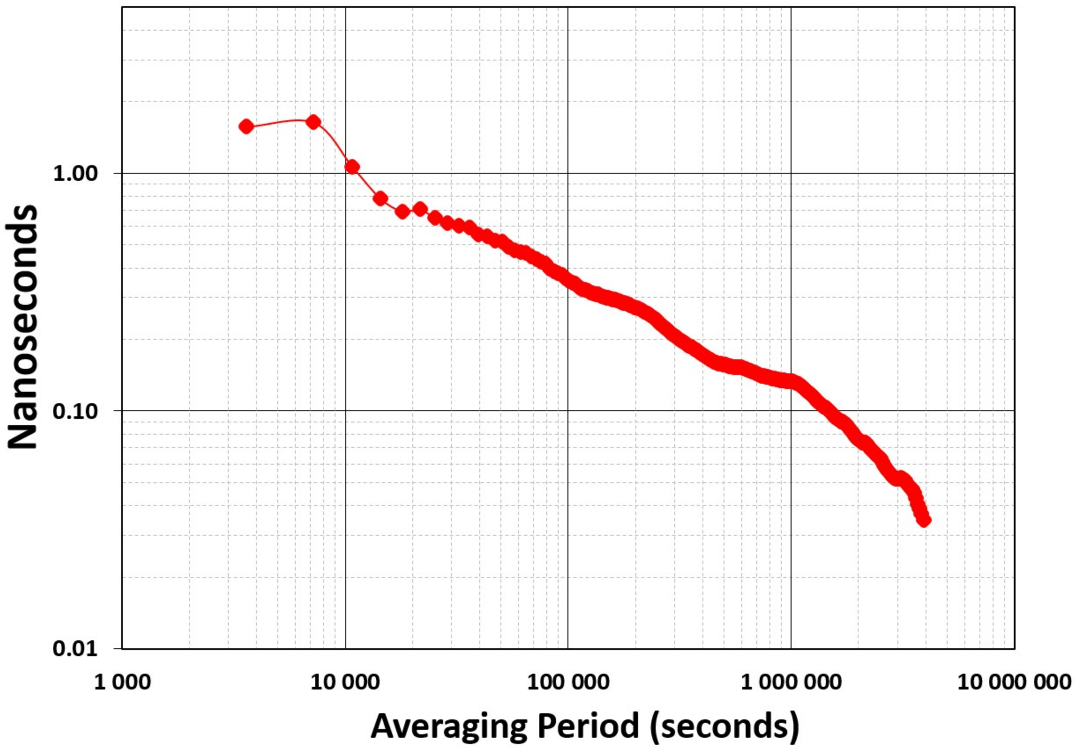

- ✓ The graph shows the time deviation (stability) of a NISTDC for averaging periods ranging from one hour to about one month.
- ✓ After averaging for one hour, the stability is about 1.5 ns, dropping below 0.4 ns after one day and below 0.2 ns after one week.
- ✓ This high level of stability is possible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

Multi-Source Common-View Disciplined Clock Technique Improves Reliability

- 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 implements a patented technique 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).

  - If the GNSS constellation providing the common-view signals is unavailable, for example if the satellite signals are 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 accuracy then depends on the accuracy and stability of its free running atomic clock.

Holdover Performance (NISTDC with rubidium clock)

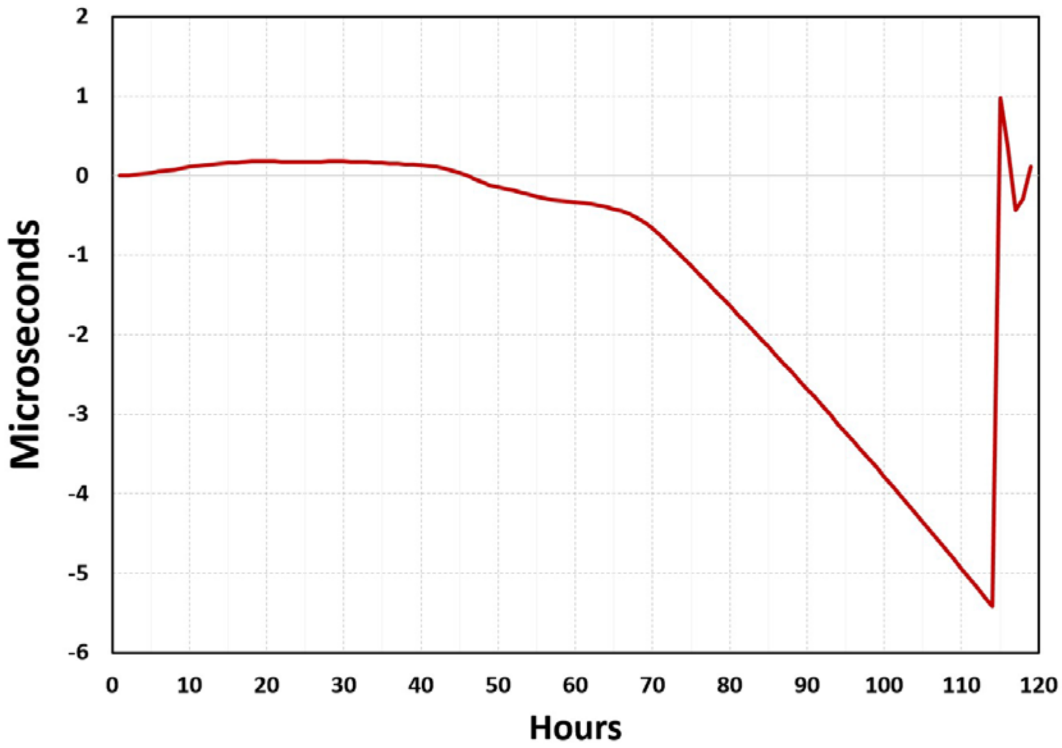

- ✓ A NISTDC in holdover mode can stay within 1 µs of UTC(NIST) for at least 48 hours with a rubidium clock.
- ✓ To illustrate this, the graph shows the performance of a rubidium NISTDC in holdover mode. This device maintained 1 µs synchronization for ~73 hours after its GNSS antenna was disconnected. The error increased to ~5 µs after ~110 hours when GNSS reception was restored and it began to relock to UTC(NIST).
- ✓ The combination of the MSCVDC method and rubidium clock holdover usually provides enough time to troubleshoot and correct most failure modes

Holdover Performance (NISTDC with cesium clock)

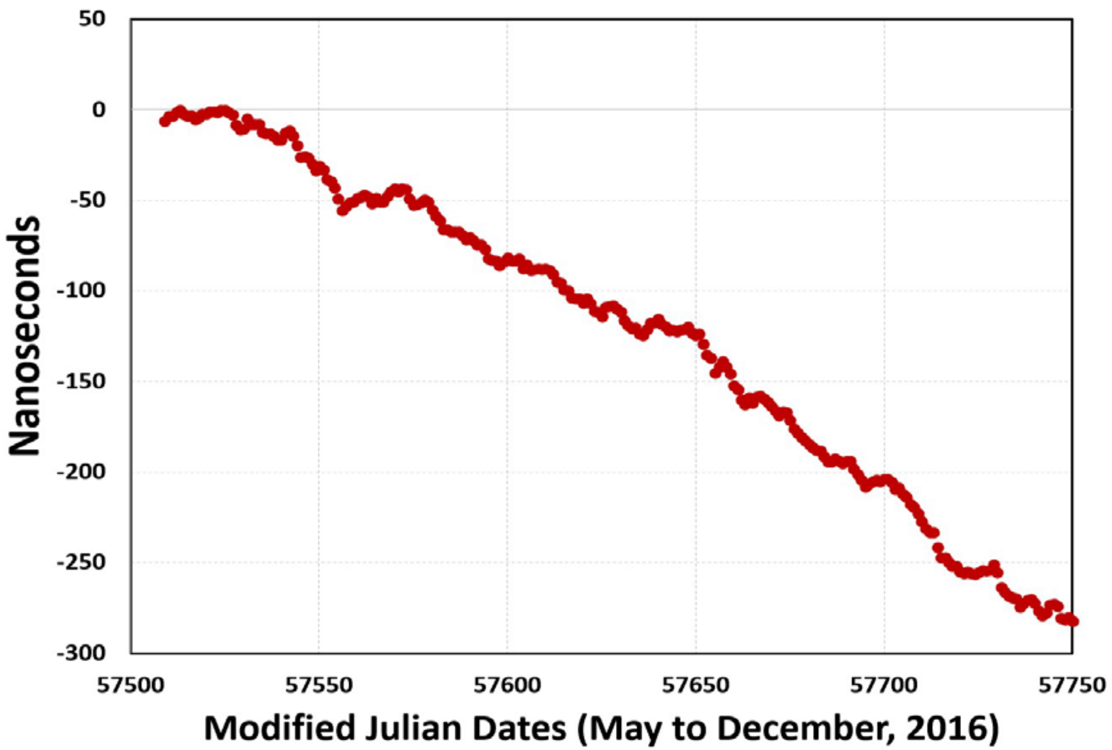

- ✓ A NISTDC in holdover mode can stay within 1 µs of UTC(NIST) for at least several months with a cesium clock.
- ✓ The illustrate this, the graph shows the performance of a cesium NISTDC in holdover mode that remained within 300 ns (0.3 µs) of UTC(NIST) after free running for about eight months. Because 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.
- ✓ When coupled with cesium clock holdover, the MSCVDC method assuages most stock market concerns related to reliability, vulnerability, and GNSS dependency.

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.

- ✓ The time servers referenced to the NISTDC 1 pps 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.

Time Server Monitoring

Server Information		Measurement Results and Uncertainty, *U*, contributed by Network Asymmetry (values are in milliseconds)						
Name	**Address**	**Server Time - UTC(NIST)**	**Round Trip Delay**	***U*, Worst Case**	***U*, 3:1**	***U*, 1.5:1**	***U*, 1.1:1**	***U*, 1.05:1**
US02-01		-0.021	0.334	±0.167	±0.084	±0.033	±0.008	±0.004
US02-03		-0.021	0.327	±0.164	±0.082	±0.033	±0.008	±0.004
US02-06		-0.017	0.314	±0.157	±0.079	±0.031	±0.007	±0.004
US02-07		-0.018	0.316	±0.158	±0.079	±0.032	±0.008	±0.004
US16-01		0.004	1.051	±0.525	±0.263	±0.105	±0.025	±0.013
US01-01		1.338	11.830	±5.915	±2.958	±1.183	±0.282	±0.144
US18-01		-0.020	16.156	±8.078	±4.039	±1.616	±0.385	±0.197
US18-02		-0.020	16.141	±8.070	±4.035	±1.614	±0.384	±0.197
CA07-01		-0.007	11.598	±5.799	±2.900	±1.160	±0.276	±0.141
CA07-02		0.003	11.616	±5.808	±2.904	±1.162	±0.277	±0.142
ISEAD01		0.450	1.387	±0.694	±0.347	±0.139	±0.033	±0.017
-----	-----							

This table was created at 05-28-2019 (MJD 58631) 22:59:30 UTC and will refresh every five minutes. Click on a time difference number to graph *Server Time - UTC(NIST)*.

The **BLUE** column estimates the uncertainty for typical network conditions where asymmetry equals a small percentage of the round trip delay.

If *Server Time - UTC(NIST)* is:

- **GREEN** then the server clock appears to be properly synchronized because the time difference is small enough to be reasonably attributed to network conditions.
- **YELLOW** then the server clock appears to not be properly synchronized because a time difference this large would require > 3:1 network asymmetry.
- **RED** then the time difference is too large to be attributed to even worst case network conditions and the server clock should be resynchronized.

✓ The NISTDC can monitor the performance of up to 12 time servers. It does so by simulating a client computer and requesting 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.

✓ 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.

✓ The stock exchange client has a private web portal (pictured) showing the performance of its time servers.

Time Server Monitoring

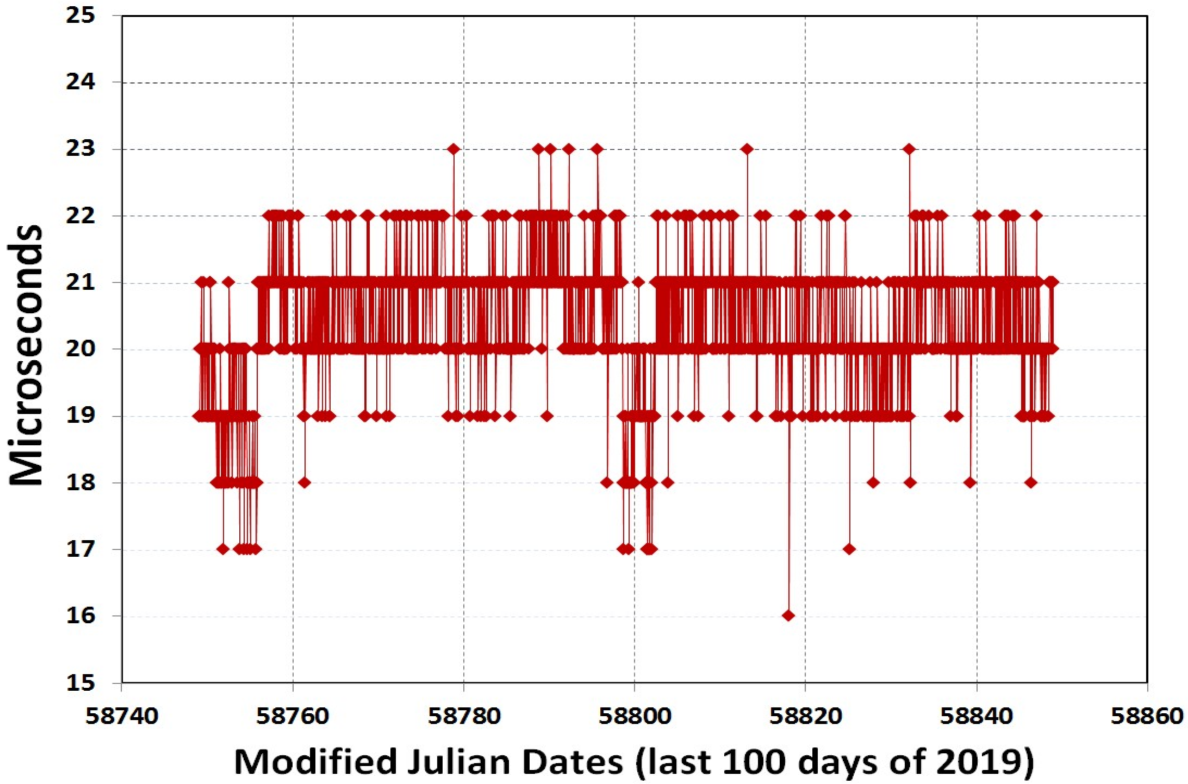

- ✓ The graph shows the time offset of an NTP server, operated by a major US stock exchange, when compared to UTC(NIST) during the last 100 days of 2019 (one data point per hour). The average time offset of the NTP server clock is 20 µs and its time deviation (stability) is < 1 µs at an averaging period of 1 day.

- ✓ 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 this 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.

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.
- ✓ Because NISTDCs are measured and compared in real-time, NIST customers are able to provide continuous evidence to auditors that they are in compliance with all current stock market synchronization requirements.