Accurate, Traceable, and Verifiable Time Synchronization for World Financial Markets

Executive Summary

Since the inception of the direct electronic trading of financial instruments in the 1990s, the speed of financial market transactions has increased at an exponential rate, leading to the present situation where trading decisions are made and trades are executed in microseconds. The world's financial markets now operate as high performance distributed systems where the timestamp of any particular trade can have a huge influence on the financial fortunes of investors. Recording each transaction with an accurate time stamp is an essential part of operating a fair and equitable financial market, not to mention satisfying the regulatory requirements for precise timekeeping by both the SEC in the United States and ESMA in the European Union. The most recent set of regulations, the Market in Financial Instruments Directive (MiFID II), will require time synchronization as low as 100µs for trades that are deemed to be part of High Frequency Trading (HFT). Whereas the SEC mandates synchronization to the NIST clock, MiFID II allows timekeeping to be synchronized to any timing laboratory that contributes to UTC, which includes NIST.

Transferring Time to Financial Markets with NIST Disciplined Clocks

Time transfer, the science of transferring time at high accuracies from one location to another, is an important and well researched area of time and frequency metrology. All time transfer systems have a reference clock at their source (point A). Information from the reference clock is encoded on a signal that is transmitted to its destination (point B), where the remote clock is located. In the simplest form of time transfer, the "one-way" technique (Fig. 3), the remote clock is then synchronized with the time from the reference clock, which is corrected to include path delay. Even if the reference clock is a nearly perfect source of time, the accuracy of the time transferred to the remote clock can be no better than the uncertainty of the path delay measurement, which is the first principle of time transfer.

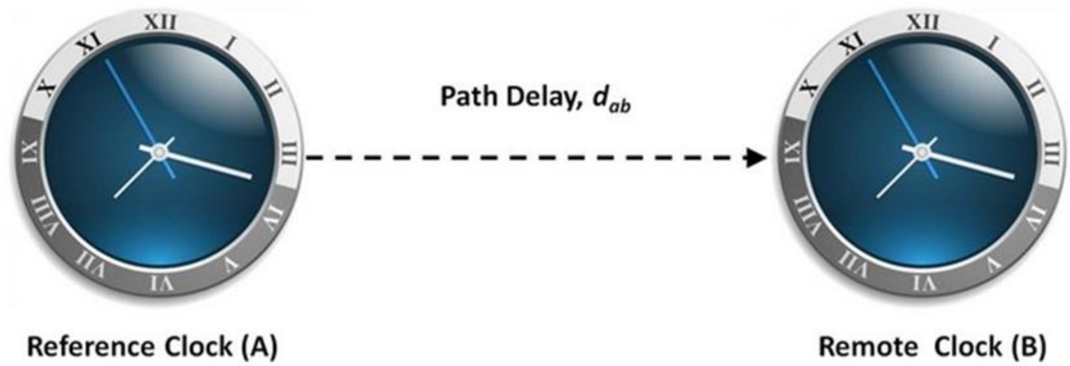

Fig. 1. A simple "one-way" time transfer system.

An effective method for reducing time transfer uncertainties is the common-view measurement technique. The technique compares two clocks located at different sites to a common-view signal, CVS, broadcast from an independent transmitter. The measurement at site A produces the time difference Clock A – CVS, and the measurement at site B produces the time difference Clock B – CVS. When the two measurements are subtracted from each other, the result is the time difference between the two clocks.

To better understand how the common-view technique works, imagine two people living at opposite ends of town who want to compare the time displayed by the clocks in their houses. Each person agrees to write down the time displayed by their clock, when a siren, located midway between them, is heard in their town. They then call or text each other to exchange the time readings. If Clock A displayed 08:01:07 and Clock B displayed 08:01:22 when the siren is first heard, then simple subtraction tells them that Clock A is 15 seconds behind Clock B. The time when the siren sounds is unimportant, it only matters that it was simultaneously heard at both clock locations. As this simplified example illustrates, the CVS doesn't have to be accurate because it does not supply the reference time, it is simply a vehicle used to transfer time from the reference to the remote clock.

GPS satellites have served as the CVS source for several decades. Common-view GPS measurements involve a GPS satellite (S), and two receiving sites (A and B), each containing a GPS receiver and a local clock. The satellite transmits a signal that is received at both A and B, and A and B each compare the received signal to their local clock. The measurements at sites A and B then compare the GPS signal received over the path to the local clock. The difference between the two measurements is an estimate of Clock A – Clock B. Delays that are common to both paths cancel even if they are unknown, but uncorrected delay differences between the two paths add uncertainty to the measurement result.

Fig. 2. Common-view time transfer via satellite.

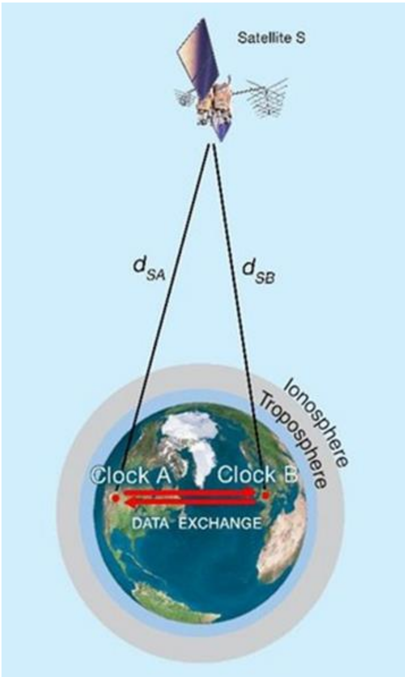

Time transfer techniques were first used to measure the frequency and/or time offset of a remote clock. If a remote clock's frequency and time could be measured, it could also be controlled. Accurate disciplined oscillators and clocks have existed for many decades, based on one-way time transfer obtained via LF radio or satellite signals. The concept and demonstration of a common-view disciplined clock, however, is relatively new and was first implemented in 2010. This service is used to replicate UTC(NIST) near the site of financial market exchanges.

Delivering NIST time to financial markets at high levels of accuracy requires a NIST disciplined clock (NISTDC) to be installed as close to the electronic trading platform as possible, which means installation in a data center where the servers that handle stock market transactions are co-located.

Fig. 3. Common-clock calibration of a NISTDC prior to shipment.

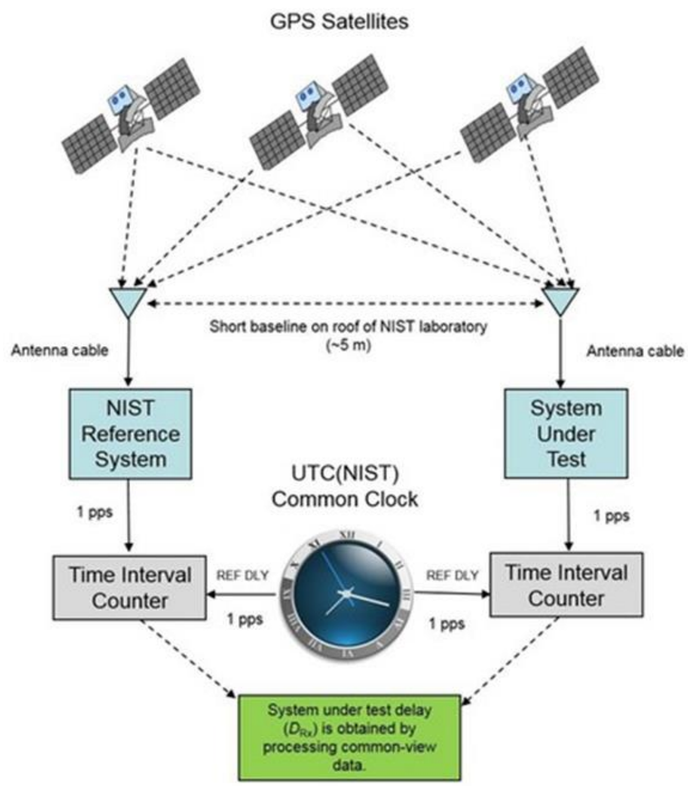

The NISTDC consists of a rack mounted instrument with an integrated computer that runs software developed at NIST; a network interface card, a 12-channel L1 band GPS receiver, a time interval counter with sub-nanosecond resolution, an atomic clock, and a distribution amplifier for the atomic clock signals. Time (1 pulse per second, pps) signals are distributed and serve as the reference for the time servers. The NISTDC is connected to a "pinwheel" type GPS antenna that is mounted on the roof of the data center. The NISTDC is also connected to the Internet. The primary purpose of the Internet connection is to transmit data files to servers operated by NIST every 10 minutes, allowing common-view data subtraction to be performed in real-time.

The NISTDC continuously adjusts its local atomic clock. The local atomic clock is either a rubidium or a cesium clock, but rubidium clocks are typically used at financial market sites. The adjustments are made by applying frequency and time corrections obtained through common-view GPS measurements. The measurements performed at NIST produce the time difference UTC (NIST) - GPS, and the measurements performed at the remote site produce NISTDC-GPS.
The data are collected by having the measurement systems at both sites average time interval counter readings for 10 minutes. At the end of each 10-minute segment the reference NIST system and the NISTDC systems simultaneously upload their measurements to a server located at NIST.

After the NISTDC is calibrated, but prior to it being shipped, a direct comparison between NISTDC time and UTC(NIST) time is performed in the same laboratory to verify consistency between the two systems. During this period, each NISTDC is tested and locked and its 1 pps output is compared to the 1 pps output of UTC(NIST).

Without requiring any measurement uncertainty analysis, the direct comparison results verify that a NISTDC can replicate the accuracy of UTC(NIST) to within a few nanoseconds, mostly limited by the uncertainty of the delay calibration. However, once a NISTDC is delivered to a financial market data center and installed in a non-controlled environment, it becomes necessary to mathematically estimate the measurement uncertainty of the common-view comparison method. This is done routinely, and monthly reports of measurement uncertainty are prepared at NIST for each NISTDC.

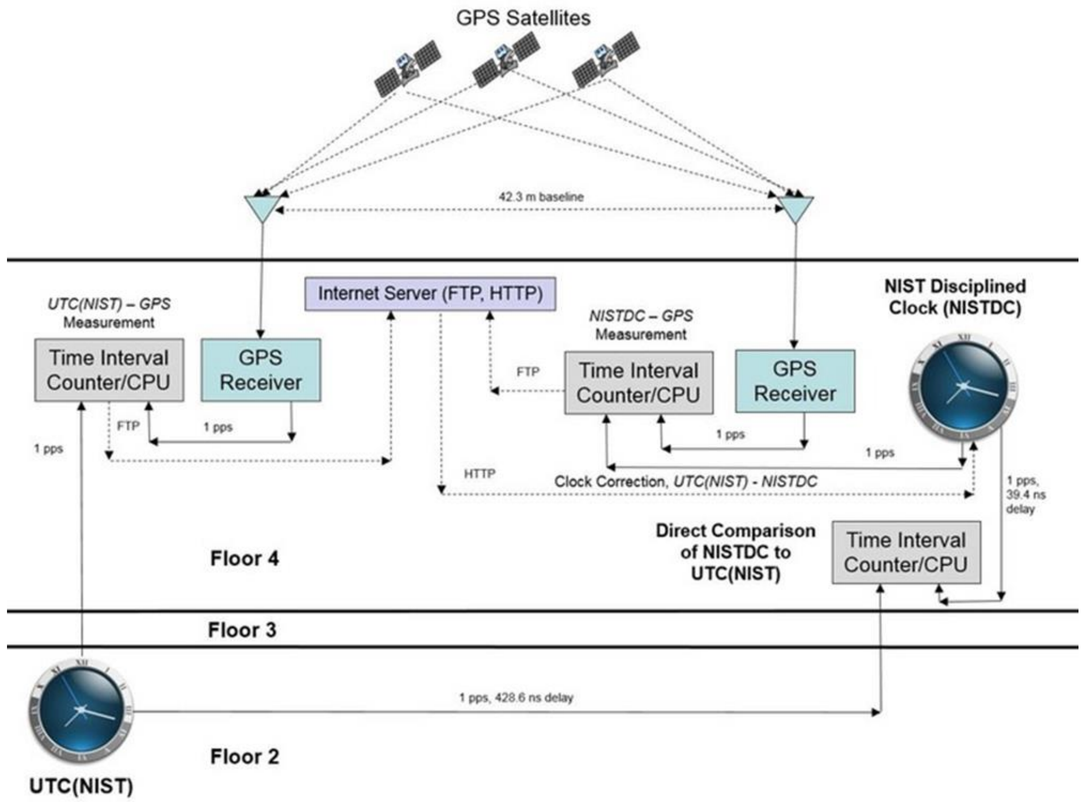

Fig. 4. Measurement configuration for direct comparison of a NISTDC to UTC(NIST).

Failure Modes and Holdover Capability

The primary cause of failure for a NISTDC, other than hardware failures or power outages at the data center sites, is the loss of data from the UTC(NIST) time scale. This can occur for a variety of reasons, including an Internet outage at either the NIST laboratories or the NISTDC site, a NIST server failure, a GPS reception problem that prevents the common-view comparison measurements from occurring, or if no data is being transmitted from UTC(NIST). If signal loss occurs, the rubidium clocks at each site go into holdover mode where they run without being disciplined.

A NISTDC can maintain 1µs synchronization in holdover mode for one day in the worst case, and for several days in the best case. This was evident on October 17, 2015, when a hardware failure prevented UTC(NIST) data from being transmitted for 9 hours and 18 minutes, from 0634 to 1552 UTC causing NISTDCs to exceed their normal 50ns lock threshold. During the UTC(NIST) outage the maximum time differences with respect to UTC(NIST) ranged from 69ns to 394ns.

Delivering Time Stamps to Financial Market Computers

A NISTDC installed in a data center near a stock exchange can routinely keep time that is accurate to within nanoseconds of UTC. To meet the synchronization requirements, however, all of the computers involved in the operations of the stock exchange must also keep accurate time. Dedicated time servers, located adjacent to a NISTDC and referenced to a 1 pps signal, are utilized to transfer time to the computers that perform and record the financial transactions.

Two industry standard time code protocols are used to transfer accurate time to stock market computers: the Network Time Protocol (NTP) and the Precision Time Protocol (PTP). PTP is currently the most commonly utilized time synchronization protocol in financial markets.

Both protocols are designed to transfer time information via data packets containing timestamps, and both transmit these packets over Ethernet connections. The timing accuracy of NTP and PTP does not come from the protocols themselves but from other factors; including the accuracy of the source clock, the frequency of the synchronization requests, the holdover capability of the client, software stack delays in both the server and the client, and the asymmetry of the network. The first factor, the accuracy of the source clock, is equivalent for both NTP and PTP because both protocols are commonly supported by the same server and referenced to the same clock. However, PTP is typically more accurate than NTP due to advantages in the other mentioned areas. PTP issues more frequent synchronization updates (unlike NTP, the server initiates the synchronization request rather than the client). The use of more frequent packet exchanges means that PTP collects more information about network delays and provides quicker, and smaller, corrections to the client clock. PTP is often implemented in the network interface card (NIC) installed inside the client computer. NICs that support PTP may include more accurate clock crystals further improving the quality of the received time signal. In contrast, NTP, which was designed for a wide area network such as the Internet, is generally implemented entirely in software on client computers foregoing many of the advantages that are present in PTP enabled network hardware.

To illustrate how this works at a financial exchange location, consider the following example. An NTP server, operated by GTT, is referenced to a 1 pps signal from a NISTDC that is kept within nanoseconds of NIST time. The client is a computer system that is part of an automated trading platform. The client computer initiates a synchronization request by sending a UDP packet to port 123 of the NTP server. The packet includes the time of the request, T1, as obtained from the client computer’s clock.

The server responds to the timing request by returning a data packet to the client. The NTP client software decodes this data packet, which includes three time stamps. One of the time stamps returned by the server simply echoes back T1, the time when the client made the request (measured by the client).

Two other time stamps contain the time, T2, when the request was received by the server, and the time, T3, when the server transmitted its response. When the client receives the packet, it again queries its own clock and records a fourth-time stamp, T4, the time of the packet's arrival. Using these same four time stamps, the round trip delay between the client and server is computed by the client.

This computation assumes that the delay from the server to the client is equal to one half of the round trip delay. If this assumption were true, the path delays to and from the server would be equal and dividing by two would fully compensate for all delays. In practice networks are asymmetric, which means that the incoming and outgoing delays are not equal, and this asymmetry adds error to the time received by the client. To reduce the maximum possible error, the round trip delay should be made as small as possible by locating the time servers very close to the stock exchange; ideally in an adjacent equipment rack in the same data center.

By use of NTP/PTP grand masters that are synchronized to NIST via NISTDCs and located in data centers adjacent to stock exchanges, GTT is able to offer its customers computer clock synchronization ranging from 25 µs to 50 µs for NTP and from 1 µs to 50 µs for PTP. Actual performance can vary based on the client software and hardware platforms used by the customer.

Verification of Time Accuracy

To ensure traceability, the time information sent to financial exchanges should be measured and verified at every link in the chain that connects NIST to the financial market customer. To verify the currently transmitted time, UTC(NIST) is continuously compared to the national time scales of nations in North, Central, and South America, via the Interamerican Metrology System Time Network (SIMTN).

Each of the NISTDCs located at financial system data centers are compared to NIST and to each other via near real-time common-view measurements. Every 10 minutes, each NISTDC uploads its latest measurement results to a NIST server. The results are made available to GTT in the form of a grid that updates every 10 minutes.

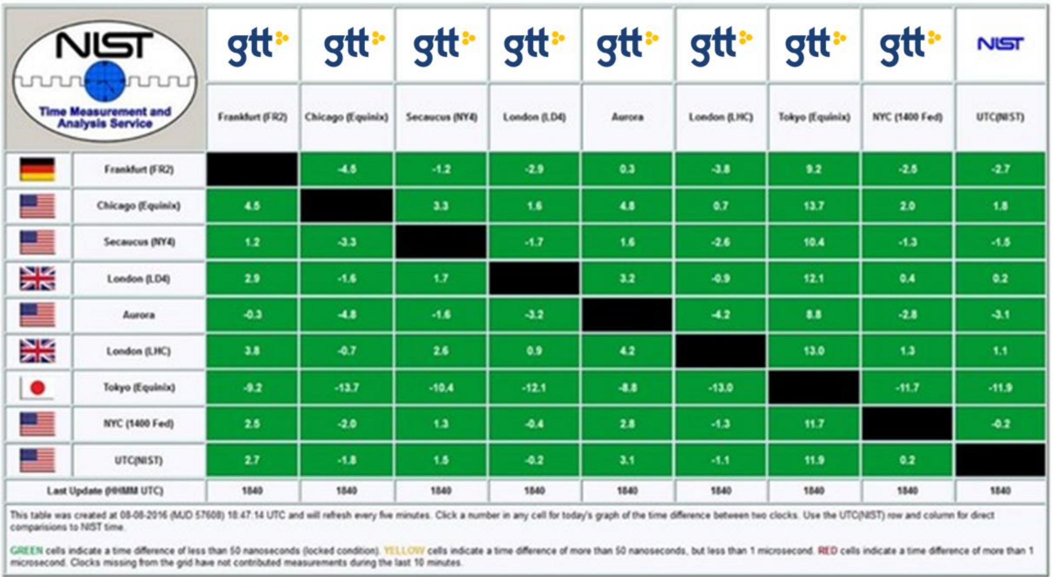

Time Measurement and Analysis Service	Frankfurt (FR2)	Chicago (Equinix)	Secaucus (NY4)	London (LD4)	Aurora	London (LHC)	Tokyo (Equinix)	NYC (1400 Fed)	UTC(NIST)
Frankfurt (FR2)		-4.5	-1.2	-2.9	0.3	-3.8	9.2	-2.5	-2.7
Chicago (Equinix)	4.5		3.3	1.6	4.8	0.7	13.7	2.0	1.8
Secaucus (NY4)	1.2	-3.3		-1.7	1.6	-2.6	10.4	-1.3	-1.5
London (LD4)	2.9	-1.6	1.7		3.2	-0.9	12.1	0.4	0.2
Aurora	-0.3	-4.8	-1.6	-3.2		-4.2	8.8	-2.8	-3.1
London (LHC)	3.8	-0.7	2.6	0.9	4.2		13.0	1.3	1.1
Tokyo (Equinix)	-9.2	-13.7	-10.4	-12.1	-8.8	-13.0		-11.7	-11.9
NYC (1400 Fed)	2.5	-2.0	1.3	-0.4	2.8	-1.3	11.7		-0.2
UTC(NIST)	2.7	-1.8	1.5	-0.2	3.1	-1.1	11.9	0.2	
Last Update (HHMM UTC)	1840	1840	1840	1840	1840	1840	1840	1840	1840

This table was created at 08-08-2016 (MJD 57608) 18:47:14 UTC and will refresh every five minutes. Click a number in any cell for today's graph of the time difference between two clocks. Use the UTC(NIST) row and column for direct comparisions to NIST time.

GREEN cells indicate a time difference of less than 50 nanoseconds (locked condition). YELLOW cells indicate a time difference of more than 50 nanoseconds, but less than 1 microsecond. RED cells indicate a time difference of more than 1 microsecond. Clocks missing from the grid have not contributed measurements during the last 10 minutes.

Fig. 5. Grid displaying the most recently measured time differences between NISTDCs and UTC(NIST).

The grid allows both NIST and GTT to check the current status of the system at a glance. The values in the grid are in units of nanoseconds. Under normal operation, the cells in the grid will have a green background indicated that the NISTDC is locked and the time difference displayed in that cell is less than 50 ns. A yellow cell indicates that the current time difference for a particular comparison is between 50ns and 1000ns (1µs), and a red cell indicates a time difference that exceeds 1 µs. A yellow cell is usually not serious and can be caused by a brief loss of lock due a short network outage or other factors. However, a red cell generally indicates a failure condition that requires immediate attention by NIST and/or GTT.

The figure below displays a graph of the eight NISTDCs located at GTT data center sites compared to UTC(NIST) for the month of July 2016. The measured average daily time offset of each NISTDC was always within ±1ns and the average time offset for the month was near 0 at all locations.

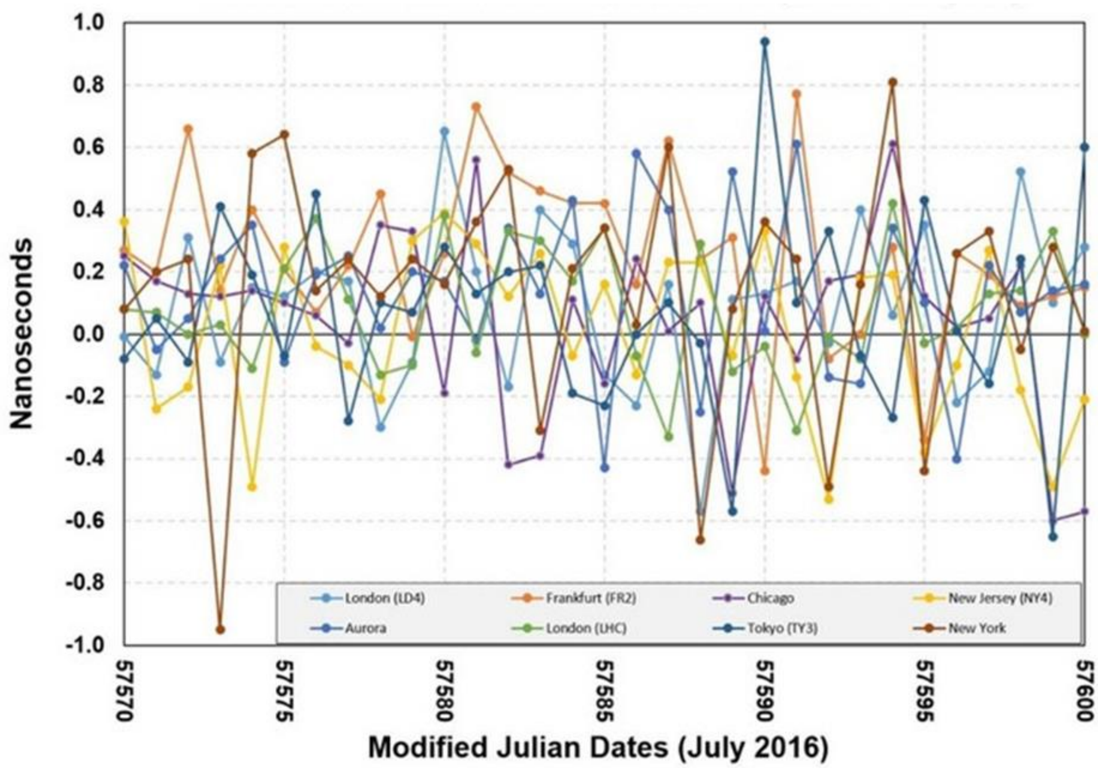

Fig. 6. The average daily time offsets between eight NISTDCs at GTT data centers and UTC(NIST).

Conclusion

In collaboration with GTT, NIST provides accurate, traceable, and verifiable time synchronization to world financial markets. This is accomplished by installing a replica of the UTC(NIST) time scale, in the form of a NISTDC, as close as possible to the stock exchange (often in the same data center). The frequency uncertainty of a NISTDC is less than 1 × 10-14 after one day of averaging, and the time uncertainty is less than 15ns, with respect to UTC(NIST) and UTC. Because NISTDCs are continuously measured and compared to national and international time standards, clients who utilize this system can provide verification to auditors and regulators that they are in full compliance with all requirements for financial market time synchronization.

The preceding report contains excerpts from: Accurate, Traceable, and Verifiable Time Synchronization for World Financial Markets, Journal of Research of the National Institute of Standards and Technology, Vol 121 (2016)
http://nvlpubs.nist.gov/nistpubs/jres/121/jres.121.023.pdf

About GTT

GTT provides multinationals with a better way to reach the cloud through its suite of cloud networking services, including optical transport, wide area networking, internet, managed services, voice and video services. The company's Tier 1 IP network, ranked in the top five worldwide, connects clients to any location in the world and any application in the cloud. GTT delivers an outstanding client experience by living its core values of simplicity, speed and agility. For more information on how GTT is redefining global communications, please visit http://www.gtt.net

Worldwide Sales Office Locations

North America

Washington, DC (HQ)

Atlanta, GA

Austin, TX

Boston, MA

Chicago, IL

Dallas, TX

Denver, CO

Halifax, NS

Los Angeles, CA

New York, NY

Philadelphia, PA

Phoenix, AZ

Pittsburgh, PA

Rochester, NY

San Francisco, CA

San Jose, CA

Seattle, WA

South America

São Paulo, Brazil

Europe

Amsterdam, The Netherlands

Belfast, Ireland

Cagliari, Italy

Dublin, Ireland

Frankfurt, Germany

London, UK

Southport, UK

Asia

Hong Kong

Pune, India

Contact Numbers

North America

+1 703 442 5500

EMEA

+44 (0)20 7489 7200

South America

+55 11 3443 1547

APAC

+852 8107 1088