Delivering NIST Time to Financial Markets Via Common-View GPS Measurements

Michael Lombardi
NIST Time and Frequency Division
lombardi@nist.gov

55th CGSIC Meeting – Timing Subcommittee
Tampa, Florida
September 14, 2015

Outline

The basics of stock market trading (the bid, the ask, and the spread)

High frequency trading and its effect on the stock market

FINRA OATS Rule 7430 (including a little history)

How a NIST disciplined oscillator works

NIST disciplined oscillator performance

Using NTP/PTP for time stamping financial transactions

Summary

The basics of stock market trading (the bid, the ask, and the spread)

The bid, the ask, and the spread

- Every stock has a "bid" and an "ask" price.
  - ✓ Buyers of a stock pay the "ask" price
  - ✓ Sellers of a stock receive the "bid" price
  - ✓ The difference between the bid and the ask is the "spread," or the profit made by the broker or market maker who handles the transaction.

- For example, let's say stock ABC has a bid of $20.25 and an ask of $20.26. The broker buys shares from me at $20.25 and sells them to you at $20.26, thus making a profit of 1 cent per share on the transaction.

Accurate, high resolution time stamps are needed to avoid stock fraud

- Consider a simple example:
  - ✓ Dozens of small investors get in line to place market orders for ABC when the current ask is $20.26. No one else is in line when these small investors places their orders and they assume $20.26 is what they will pay.
  - ✓ Mister Big, who works for a large hedge fund, enters the line 0.4 seconds later with an order to buy one million shares of ABC. The broker knows the small investors arrived first, but waited for Mister Big to get there and executes his trade first.
  - ✓ Mister Big's huge buy raises the ask by 15 cents, to $20.41. The little guy's orders have already been placed, so they all have to pay the higher price. Their trades raise the ask even higher and also raises the bid much higher than the price Mister Big just paid. Mister Big happily decides to sell his million shares at the new bid, taking an instant profit and sticking the small investors with an instant loss.
  - ✓ Can the broker get away with this? Sure, if the trades were time stamped by a clock with a resolution of 1 second. Mister Big's order arrived 0.4 seconds late, but according to the time stamps, it arrived at the **exact same time** as the small orders. The small investors can't prove that they were in line first.

High frequency trading and its effect on the stock market

High Frequency Trading (HFT)

- High frequency trading (HFT) involves the use of high speed computers running complex trading algorithms that automatically place and execute orders based on market conditions. Most stock market volume now comes from HFT orders.

  - ✓ HFT platforms rapidly scan the market to give their traders the most favorable terms possible. Fast traders (investment banks, hedge funds, institutional investors) are usually more profitable than slow traders (retail investors).

- HFT became commonplace after the financial collapse of 2008, because liquidity was a major concern for investors who worried about getting out of the market.

  - ✓ To increase liquidity, the NYSE and other major exchanges began paying smaller exchanges a small fee (or “rebate”) for each transaction they handled. The small fees can turn into substantial amounts when the volume gets high. This created incentive for the exchanges to drive up the volume as high as possible, by enticing traders to buy and sell multiple times per second.

HFT has led to more stock exchanges

- As a result of HFT, the number of stock exchanges has gone up dramatically. There were once only a few exchanges (NYSE, NASDAQ, AMEX) but stocks now trade at over 70 venues in the United States. Much has changed during the past decade:

  - ✓ Most stock exchanges have no people present, the exchange is simply a rack of computers in a data center. The old image that people have of the market is gone (there are no more people on phones, or traders walking around carrying slips of paper). Everything is automated.

  - ✓ Stocks are still "listed" on a specific exchange like the NYSE or NASDAQ, but they can trade anywhere. For example when the NYSE crashed for three hours on July 8, 2015, the other exchanges picked up nearly all of the trades so it was barely noticeable.

  - ✓ As noted, the exchanges make more when they execute more trades, so they try hard to get the HFT business. One way to attract HFT business is to rapidly look at who is in line to buy or sell, and to then offer them a slightly more favorable bid or ask price than the other exchanges.

  - ✓ **HFT makes it more important than ever for all exchanges and trading platforms to have accurate, high resolution time stamps that originate from a common clock, to avoid stock market fraud and manipulation.**

FINRA OATS Rule 7430

A little history

- Most stock market clocks once had 0.1 minute resolution (6 seconds). This meant that trades placed as much as three seconds apart would have the exact same time tag, making it very easy to reorder trades and defraud investors.

- To reduce fraud, the Securities and Exchange Commission (SEC) forced the National Association of Securities Dealers (NASD) to implement an Order Audit Trail System (OATS) in August 1996. The OATS system needed a time requirement for all stock market clocks.

- To ensure getting rid of all of the old clocks with 6 second resolution, the NASD required all clocks to be within 3 seconds of a common reference (half of the previous resolution). The common reference was chosen to be NIST time, because the NASDAQ stock exchange was then using the GOES satellite time code service operated by NIST as its reference time source.

- NASD OATS Rule 6953, which became effective in August 1998, thus required all stock market clocks to be synchronized to within 3 seconds of NIST time. This rule made NIST the official time provider for the stock market.

FINRA OATS Rule 7430

- In 2008, FINRA OATS Rule 7430 superseded NASD OATS Rule 6953. The synchronization requirement was reduced by a factor of three, from 3 seconds to within 1 second of NIST time.
  - ✓ FINRA is the Financial Industry Regulatory Authority, an organization dedicated to investor protection and market integrity.
- Reducing the requirement to 1 second was in response to HFT, but it is still too coarse a requirement to prevent HFT fraud.
  - ✓ FINRA Regulatory Notice 14-47, which went out for comments in November 2014, proposed tightening the time requirement by a factor of 20, to within 0.05 s (50 milliseconds) of NIST time.
  - ✓ The new rule has yet to be adopted, so the requirement remains 1 s as of September 2015. However, even if the 50 millisecond requirement is not adopted, the notice indicates that a reduction to either 200 milliseconds (0.2 s) or 100 milliseconds (0.1 s) will occur soon.

How a NIST disciplined oscillator works

How a NISTDO works – Part I

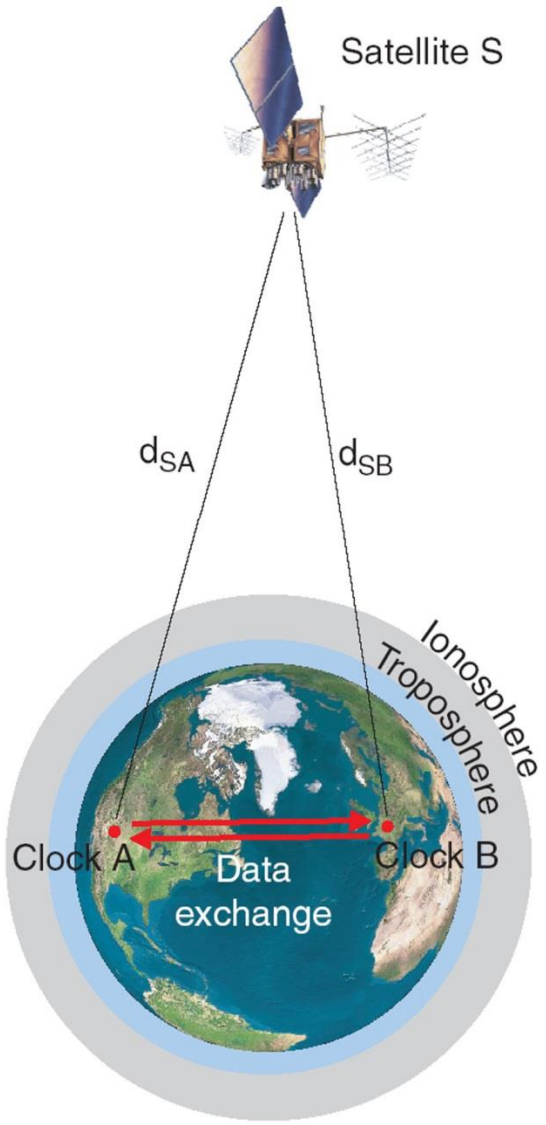

- Common-view GPS signals are used to “relay” time from the NIST time scale to the stock market clocks.
- Common-view simply means that GPS signals can be received nearly simultaneously at NIST and at each stock exchange.

How a NISTDO 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 GPS.

Two time difference measurements are made:

- ✓ *NIST – GPS* (recorded in Boulder, Colorado)
- ✓ *NISTDO – GPS* (recorded at each financial market site)

Both measurement results are sent via the Internet to a NIST server. A clock correction for the NISTDO is obtained by simply subtracting the measurement made at NIST from the measurement made by the NISTDO. For example:

- ✓ *NISTDO Clock Correction = (NIST – GPS) – (NISTDO – GPS)*

The clock correction adjusts the frequency and time of NISTDO so that it stays "locked" to UTC(NIST). This allows NIST time to be replicated at stock exchange sites with an uncertainty of about 10 nanoseconds.

The NISTDO links to NIST through common-view GPS

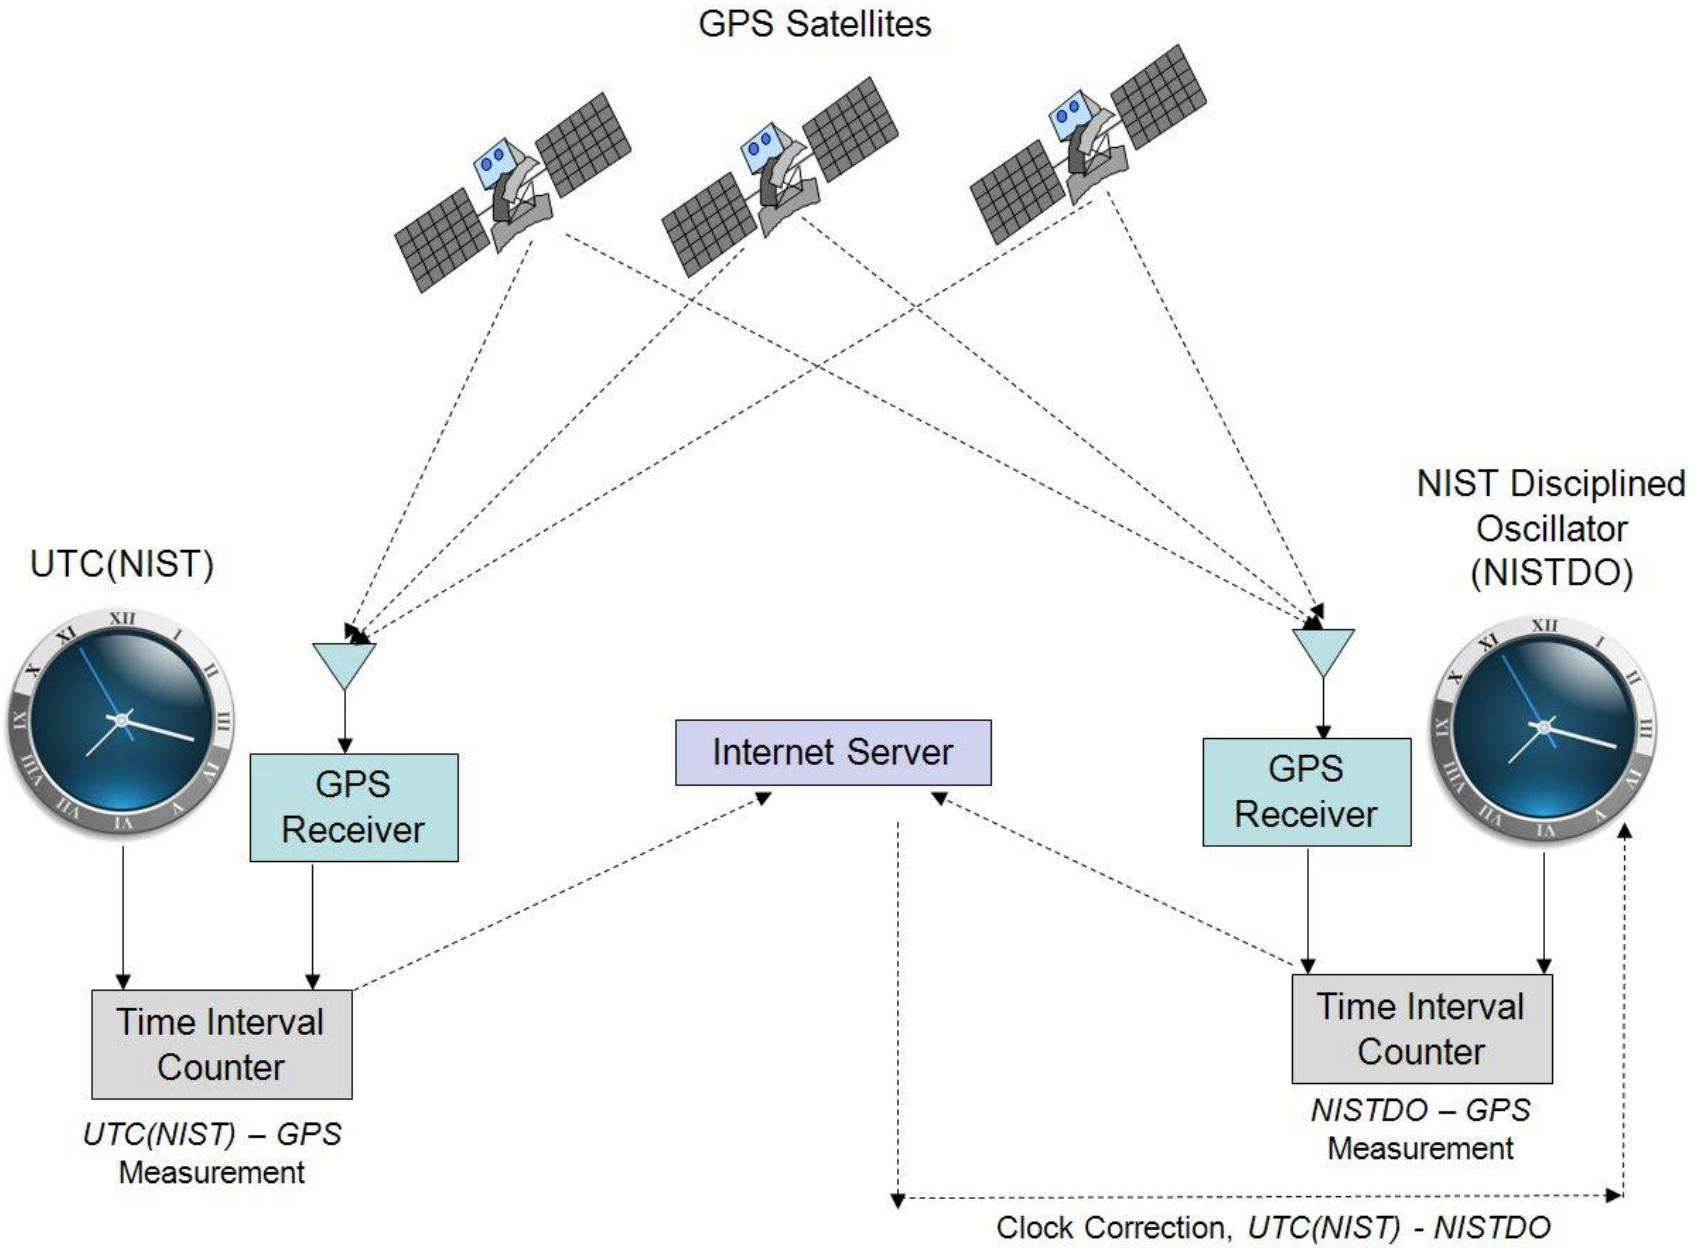

NISTDO technical characteristics

- A NISTDO is a rack mount instrument (PC-based) with a touch screen display
- Local oscillator is a rubidium with a built-in distribution amplifier (multiple 1 pps and 10 MHz outputs)
- Receiver is a 12-channel L1 band GPS receiver
- Phase comparator is a time interval counter (30 ps resolution)
- Correction method is a PID controller implemented in software, the correction interval is 10 minutes
- Utilizes "classic" common-view method over baselines shorter than 5000 km, utilizes "all-in-view" method over longer baselines
- Sends files to NIST via Internet (FTP)
- All NISTDOs are calibrated at NIST prior to shipment to compensate for cable and hardware delays

NIST disciplined oscillator performance

(accuracy and stability)

Intercomparison grid of NISTDOs located near stock exchanges

NIST 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)		-5.5	-9.1	-3.0	-6.4	-1.1	-8.7	-9.1	-3.2
Chicago (Equinix)	5.5		-3.6	2.5	-0.9	4.3	-3.2	-3.6	2.3
Secaucus (NY4)	9.1	3.6		6.1	2.7	7.9	0.4	-0.1	5.9
London (LD4)	3.0	-2.5	-6.1		-3.4	1.9	-5.7	-6.1	-0.2
Aurora	6.4	0.9	-2.7	3.4		5.2	-2.3	-2.7	3.2
London (LHC)	1.1	-4.3	-7.9	-1.9	-5.2		-7.6	-8.0	-2.0
Tokyo (Equinix)	8.7	3.2	-0.4	5.7	2.3	7.6		-0.4	5.5
NYC (1400 Fed)	9.1	3.6	0.1	6.1	2.7	8.0	0.4		6.0
UTC(NIST)	3.2	-2.3	-5.9	0.2	-3.2	2.0	-5.5	-6.0	
Last Update (HHMM UTC)	1850	1850	1850	1850	1850	1850	1850	1850	1850

This table was created at 09-11-2015 (MJD 57276) 18:52:46 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.

Demonstrating how NIST Time can be replicated at remote sites - comparing UTC(NIST) and four NISTDOs to UTC(CNM)

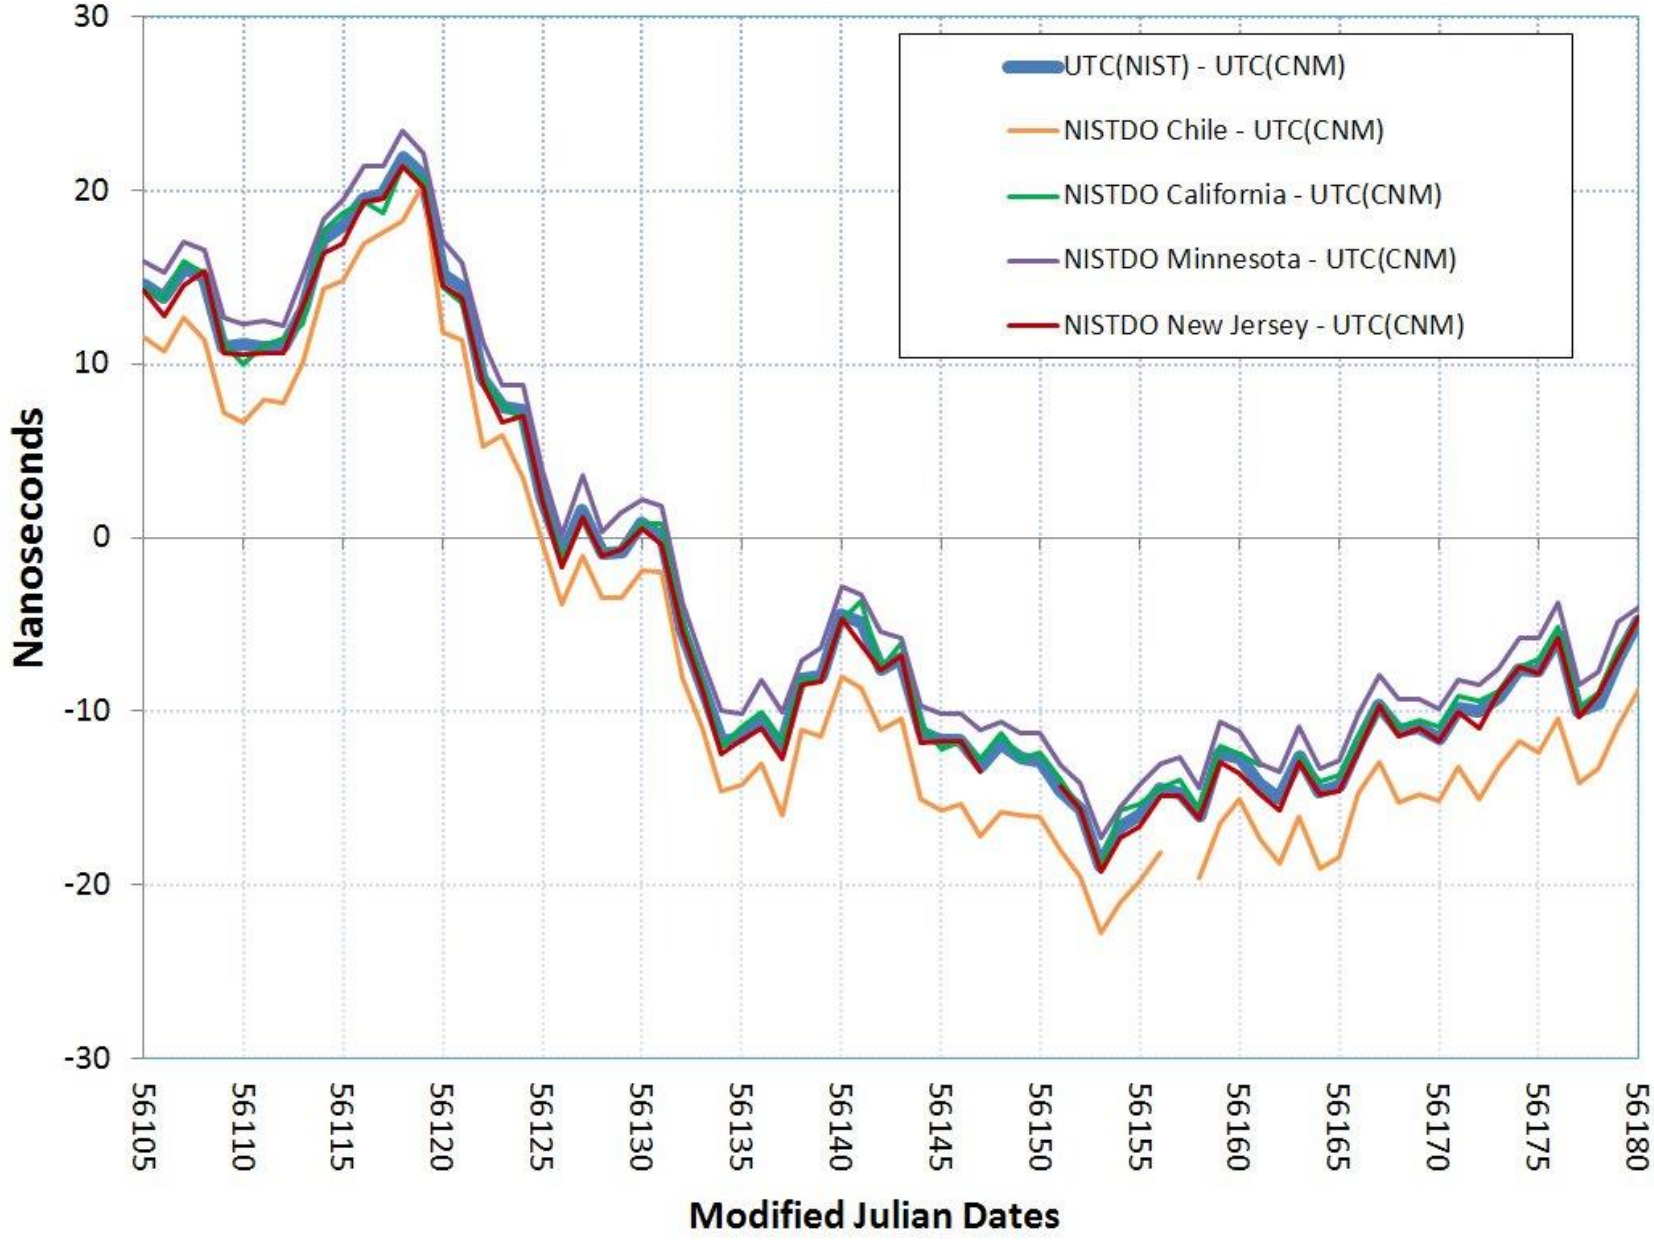

Direct comparisons of NISTDOs to UTC(NIST)

- Before shipping a NISTDO, we routinely do a direct comparison of the instrument to UTC(NIST) to verify that it works properly and that it has been properly calibrated.

- When using one-day averages, the average time offset of the NISTDO is typically near 1 nanosecond with respect to UTC(NIST). The range (peak-to-peak) is typically 2 or 3 nanoseconds.

- The small differences are due to calibrations biases, uncertainties in cable delay calibrations, and environmental effects. Because the NISTDO corrections are based on past measurements made at 10 minute intervals, they can lag behind environmental effects such as temperature changes that disturb the rubidium oscillator.

Measurement configuration for direct comparison

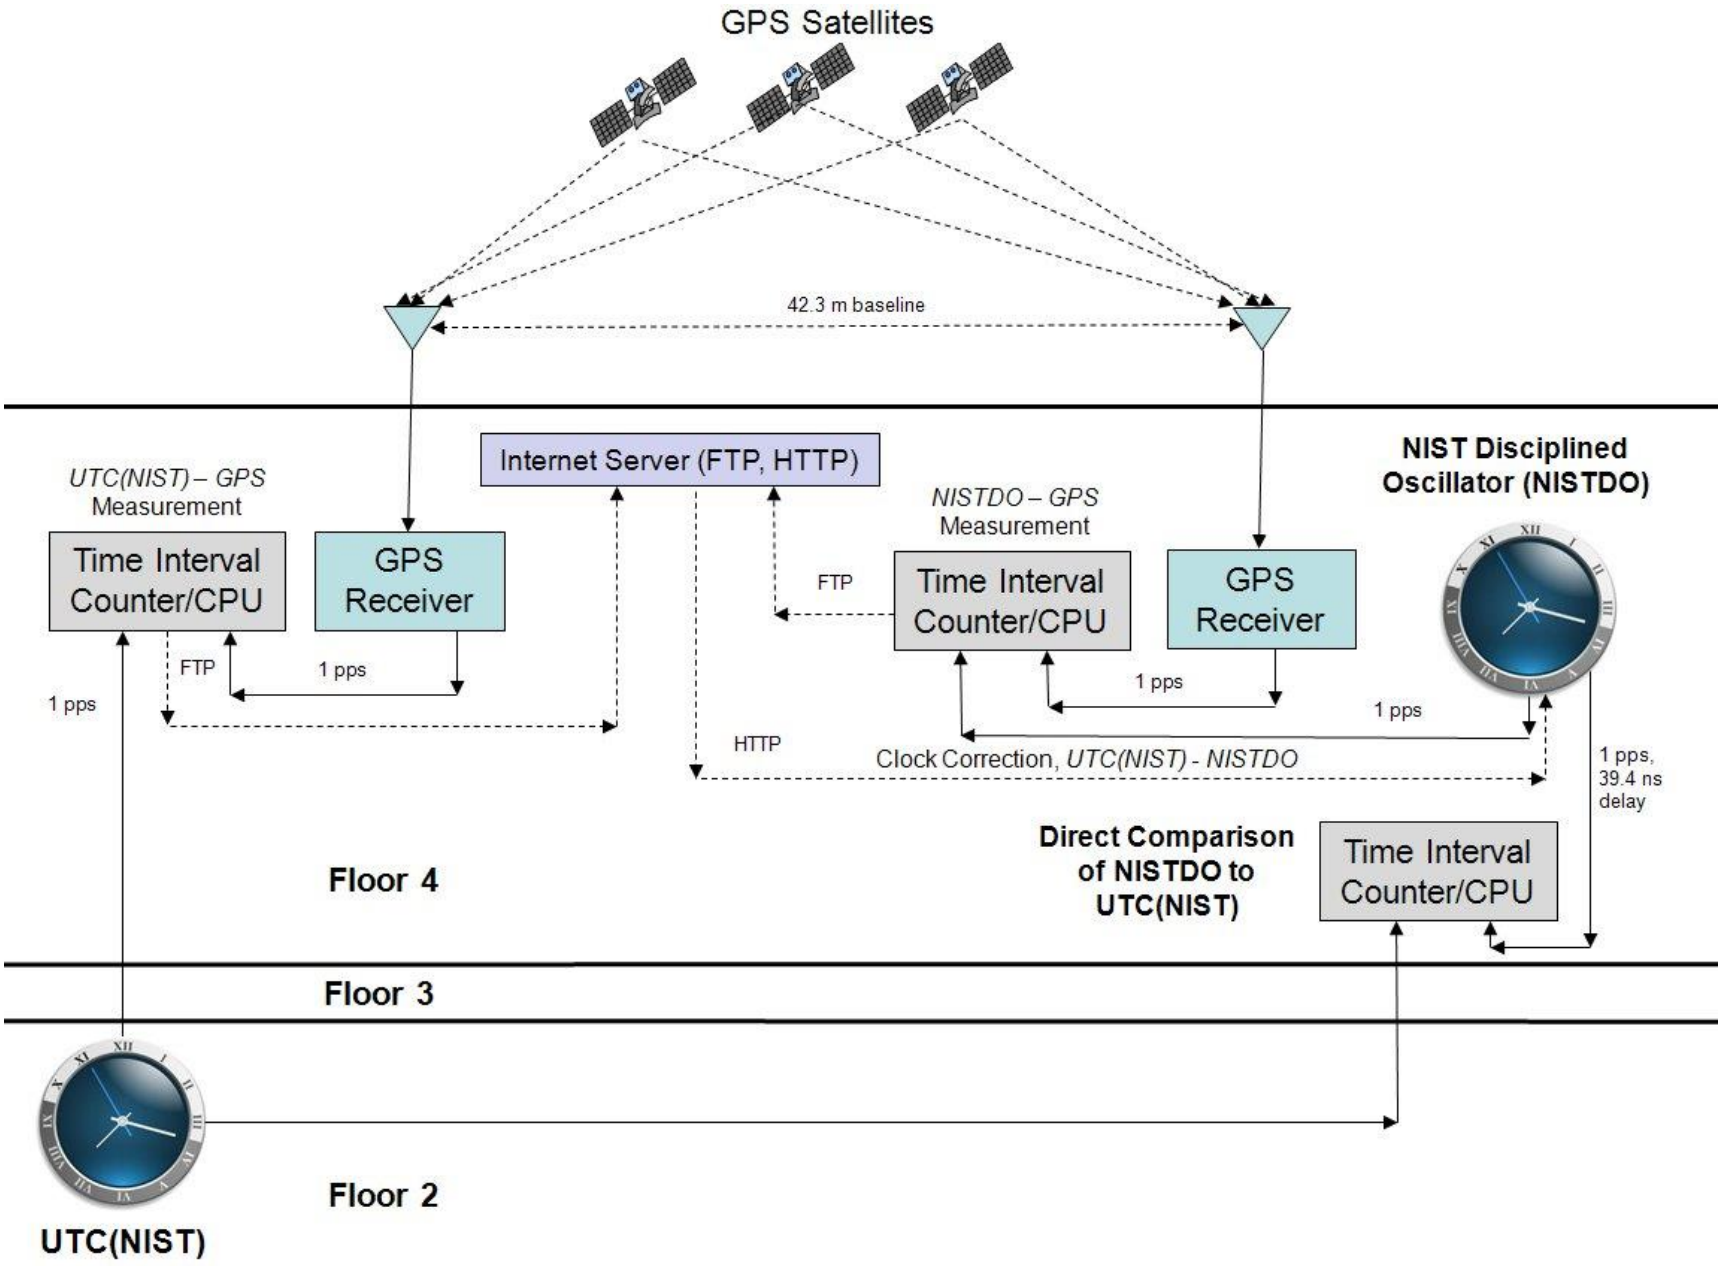

Direct comparison results

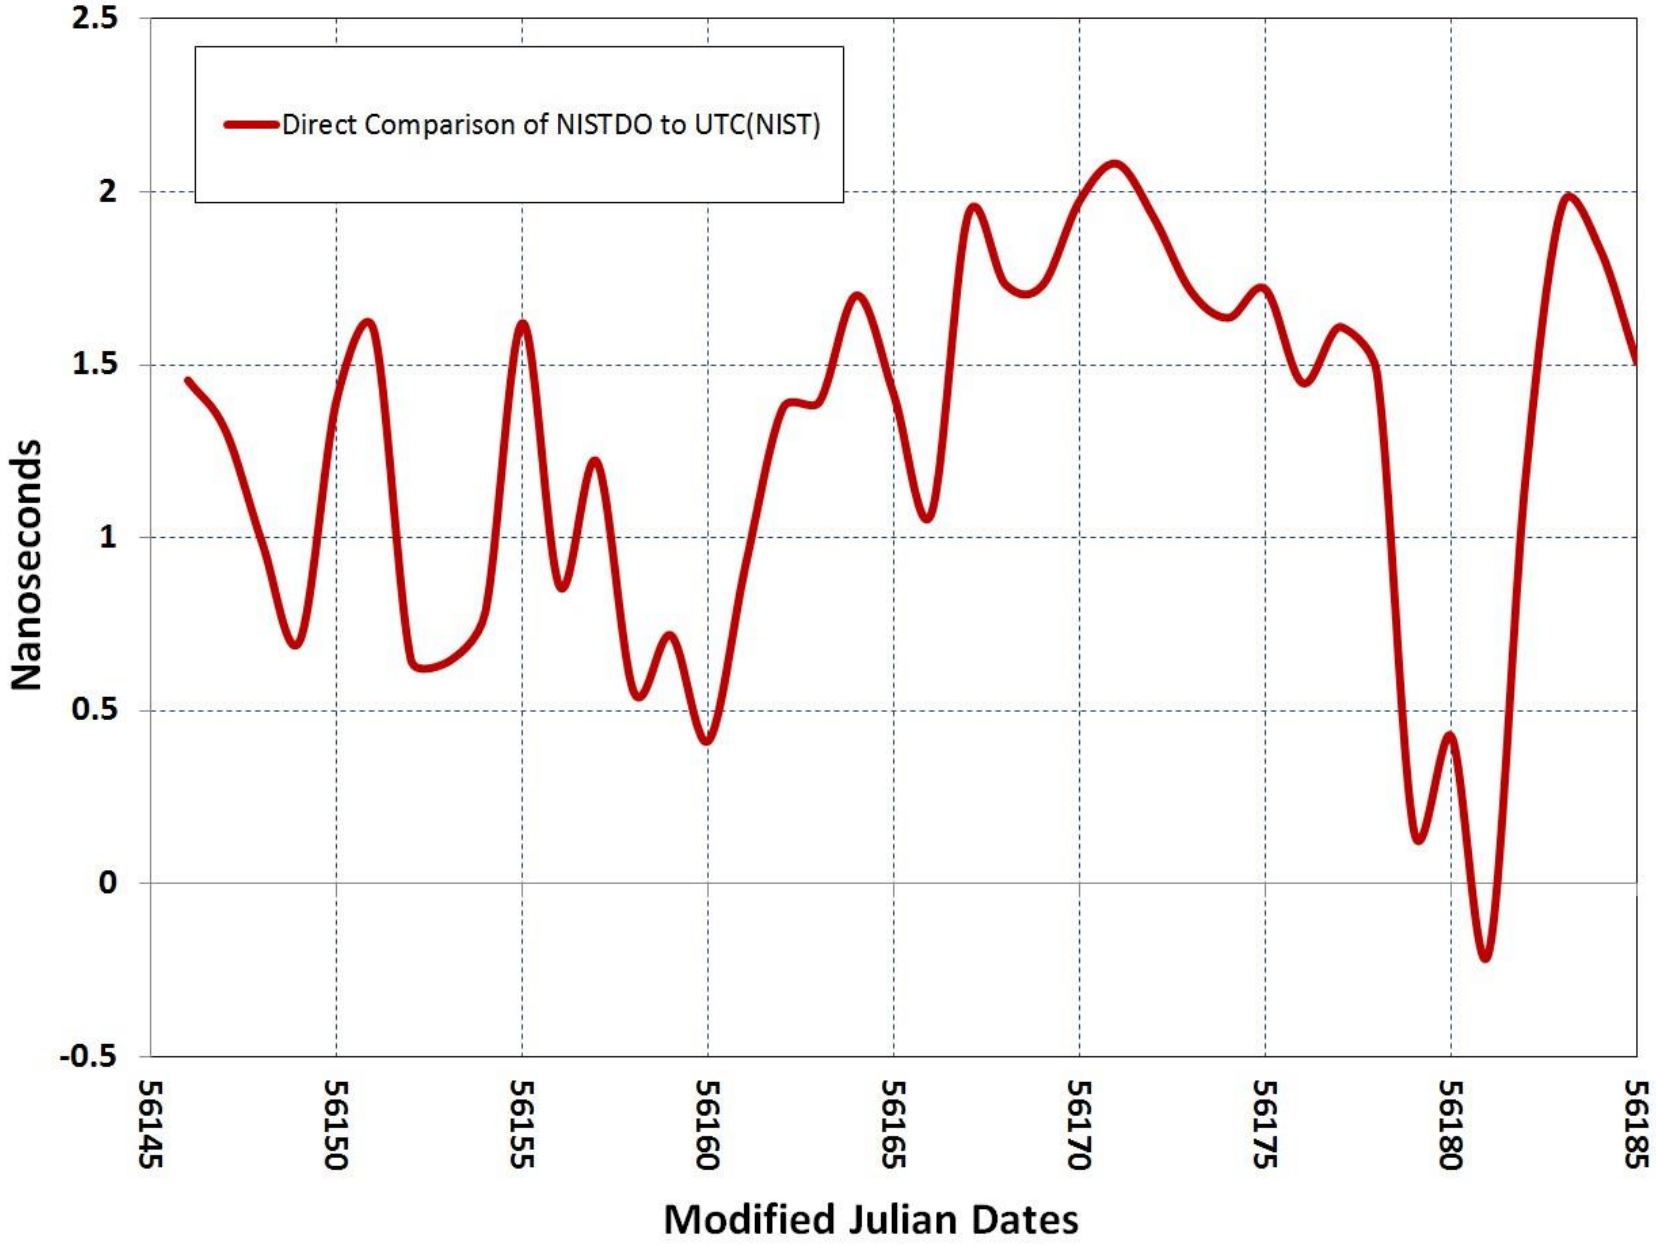

Comparing NISTDO stability to GPSDOs

- About half of the GPSDOs we have tested at NIST reach a frequency stability of 1×10^{-13} or less after one day of averaging.
- The worst GPSDOs are stable to about 4×10^{-13} at one day.
- The best GPSDOs are about one order of magnitude better than the worst devices, stable to about 4×10^{-14} at one day.
- We compared a NISTDO operating in Minnesota to the “best” and “worst” GPSDOs over a 125-day interval, using UTC(NIST) as our reference.

NISTDO frequency stability w/respect to UTC(NIST), compared to GPSDOs

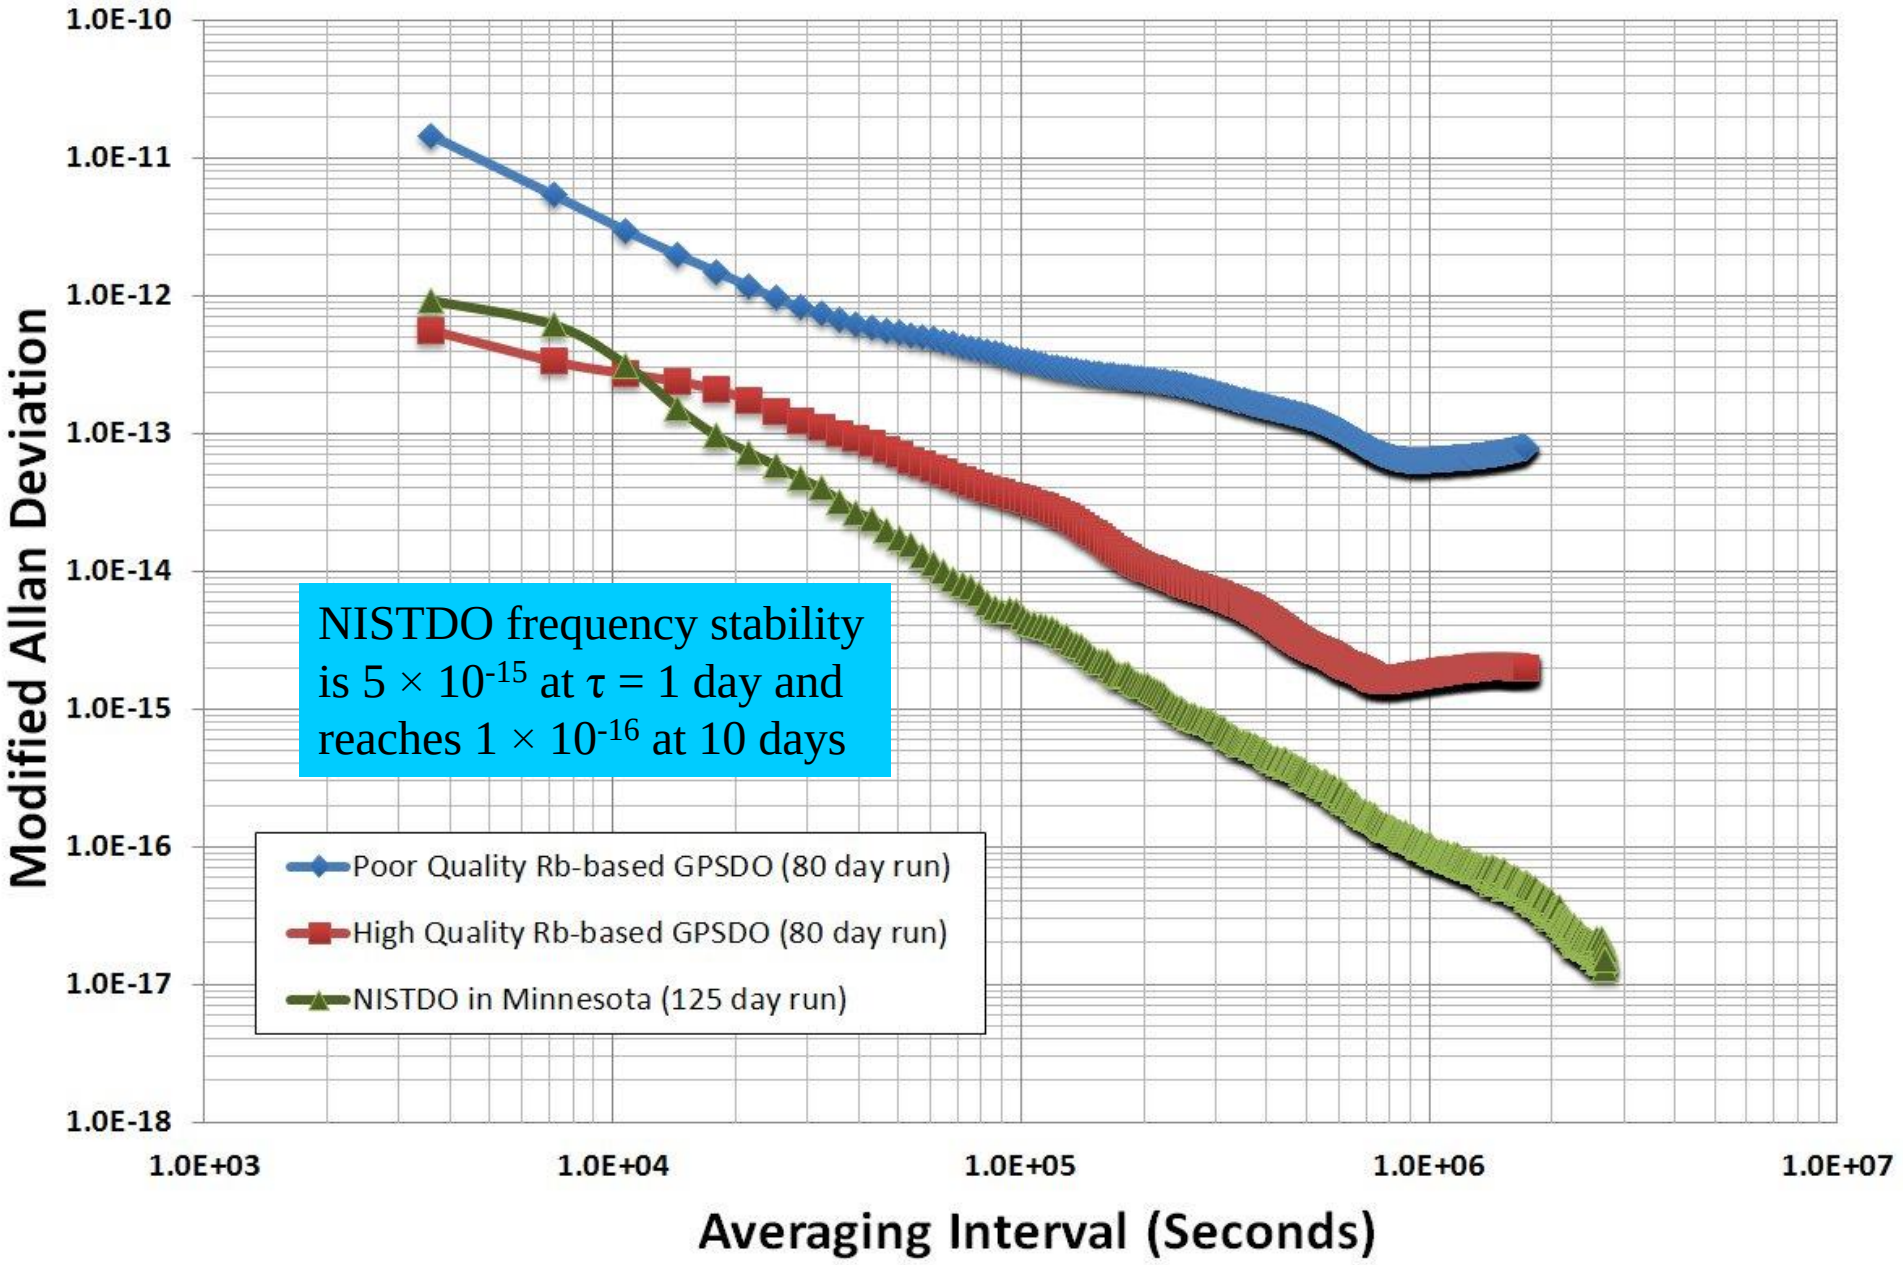

NISTDO time stability compared to UTC(NIST)

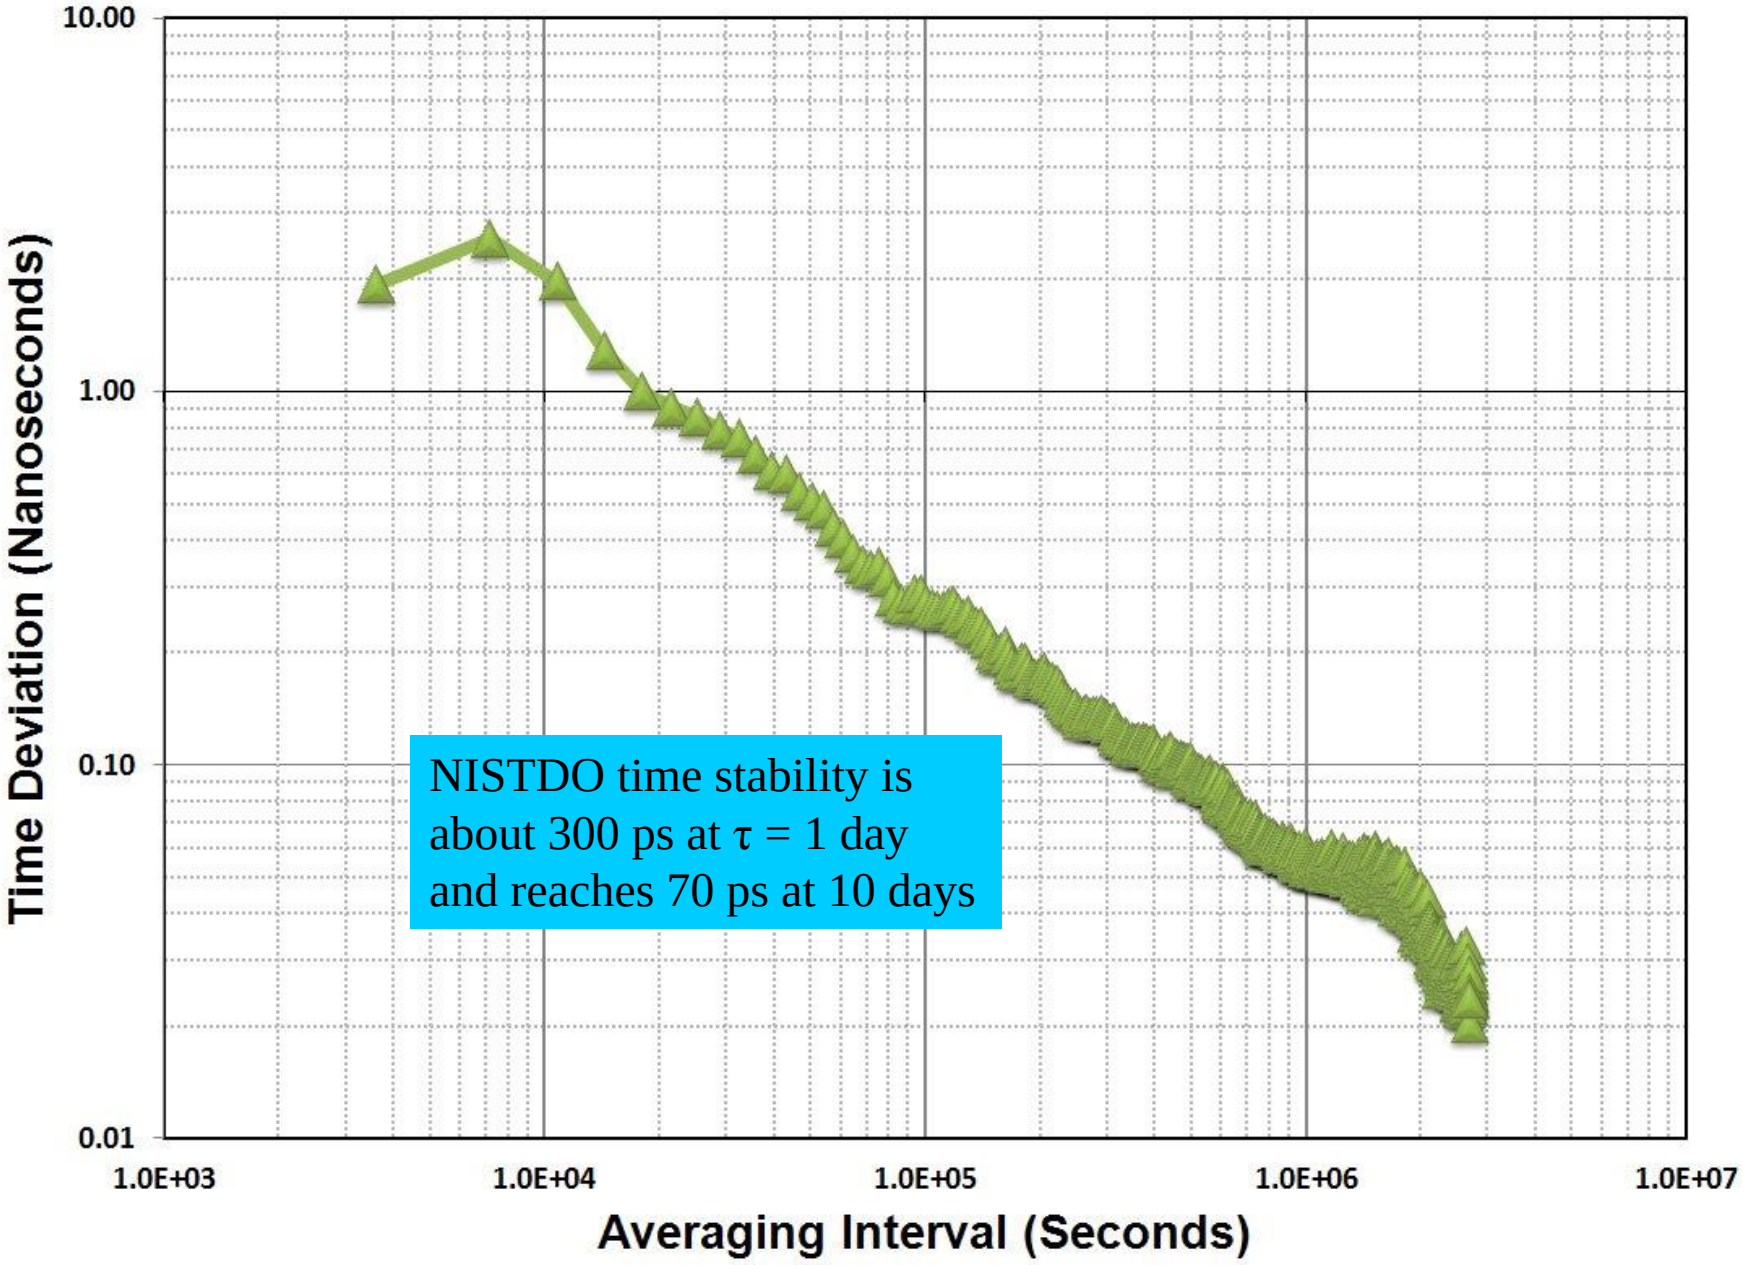

Using NTP/PTP for time stamping financial transactions

NTP and/or PTP packets synchronize the computers that time stamp stock transactions

- Network Time Protocol (NTP) and/or Precise Time Protocol (PTP) servers typically provide time for financial market computers.
  - ✓ For NIST customers, the time servers are referenced to a 1 pulse per second (pps) signal from a NISTDO.
  - ✓ The servers can only be accessed by exchanges and traders who pay to use them.
  - ✓ The time servers are located as close to the stock exchange as possible, sometimes in an adjacent equipment rack in the same data center. The short round trip delays over a local area network provide much better accuracy than, for example, accessing a NIST NTP server from across the country. Even so, the accuracy is degraded by about three orders of magnitude with respect to a NISTDO, going from nanoseconds to microseconds. Here are the accuracy numbers quoted by one provider:
    - ✓ Standard NTP/PTP service – 50 microseconds
      (1000 times better than proposed FINRA OATS rule change)
    - ✓ High Precision NTP – 25 microseconds
    - ✓ High Precision PTP – 1 microsecond

Summary

- NIST provides financial market customers with a replica of the UTC(NIST) time scale in the form of a NISTDO. The NISTDO is placed as close as possible to the stock exchange (often in the same data center).
- The frequency uncertainty of a NISTDO is less than 1×10^{-14} after one day of averaging, and the time uncertainty is ~10 ns, with respect to UTC(NIST).
- This uncertainty is increased by about a factor of 1000 by the NTP and PTP methods of transferring time to computers, but still at least 1000 times lower than the most recently proposed FINRA OATS requirement.
- Because NISTDOs are measured and compared in real-time, NIST customers are able to provide continuous evidence to auditors that they are in compliance with FINRA OATS requirements.