

UTC TIME TRANSFER FOR HIGH FREQUENCY TRADING USING IS-95 CDMA BASE STATION TRANSMISSIONS AND IEEE-1588 PRECISION TIME PROTOCOL

Michael D. Korreng
EndRun Technologies
2270 Northpoint Parkway
Santa Rosa, CA 95407, USA
E-mail: mkorreng@endruntechnologies.com

Abstract

Today, High Frequency Trading (HFT) requires low latency and knowledge of UTC time in capital markets where a difference of better than one millisecond provides a competitive advantage to trading firms. HFTs co-locate high speed computer systems in rented space next to the market centers to reduce latency of raw data, giving them advanced knowledge of the market orders. Better than 10-microsecond accuracy to UTC can be achieved using Code Division Multiple Access (CDMA), in compliance with Telecommunications Industry Association (TIA/EIA) Standard IS-95 and IEEE-1588 Precision Time Protocol (PTP) without routing an external antenna.

This paper will present a method to provide UTC time using CDMA and PTP to these trading firms, why microsecond accuracy is required, and the performance of UTC time transfer as compared to GPS-derived UTC time.

INTRODUCTION

Providing precise UTC time indoors can be problematic in rented space. The most common method for precise UTC time transfer is accomplished worldwide by GPS. Mounting a GPS antenna in rented space may not be an option if access to the roof is not granted. Installing a window-mount GPS antenna is one option if there is a full view-of-the-sky, without reflections from other buildings or absorption by foliage. Using IS-95 Code Division Multiple Access (CDMA) is another method to provide precise UTC time transfer indoors. Since the CDMA base stations are each synchronized by GPS receivers, they provide an indirect link to GPS system time and UTC time.

The major stock exchanges are allowing traders to move their equipment into rented space. Since the rented space used by traders typically does not allow access to mount an external GPS antenna, a CDMA time server is the solution. These traders are commonly referred to as High Frequency Traders (HFTs). High Frequency Trading (HFT) is a term used for computer-algorithmic, high-volume, short-term trading. More recently HFT has become a popular subject, especially since the May 6, 2010 stock-market crash when the Dow plunged nearly 1000 points. It is argued that HFT creates liquidity and is viewed as both good and bad for the market.

More trading firms are using HFT technology in the major exchanges by co-locating their equipment in rented space next to the market center. By co-locating this equipment, HFTs reduce latency of market data giving them advance knowledge of the market just milliseconds before others. Reductions in latency and UTC timestamps are required at the market center where servers are co-located for HFT market orders. A CDMA receiver with an indoor antenna synchronizes the Local Area Network (LAN) to within 10 microseconds of UTC using the IEEE-1588 Precision Time Protocol (PTP). This is an ideal solution for HFTs.

Figure 1 provides an overview of UTC time transfer to HFTs using CDMA and PTP. Table 1 identifies the time and corrections in the chain. For the purpose of this paper, the relationship between

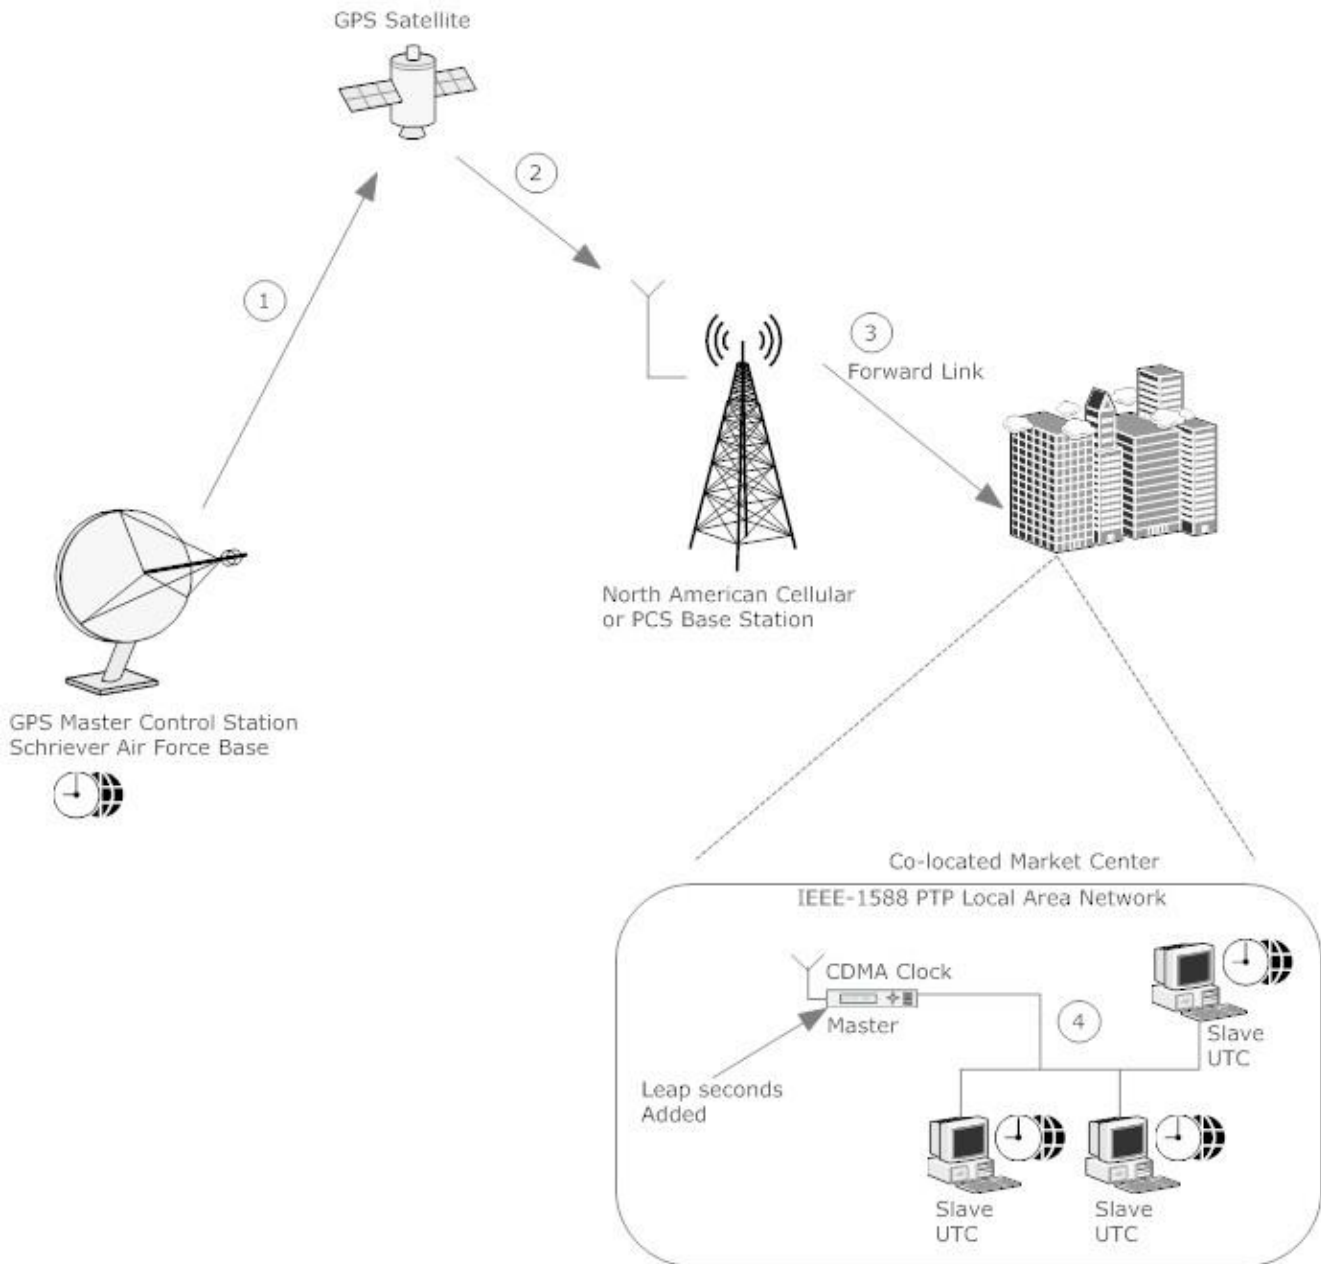


Figure 1. UTC time transfer using CDMA and PTP.

UTC (USNO), UTC (NIST), GMT, UTC, UT1, and TAI are not illustrated.

Table 1. Time and corrections in time transfer chain.

Link	Time and Corrections
1	GPS Master Control Time + UTC(USNO) Time Offset + Leap Seconds
2	GPS Satellite Time + UTC(USNO) Time Offset + Leap Seconds
3	CDMA System Time (GPS Time Scale)
4	UTC Time (CDMA System Time Corrected With Current Leap Seconds)

HIGH FREQUENCY TRADING

HFTs install high-speed hardware and proprietary software with high-bandwidth connections at the market centers. The co-located space provides an advantage over traders that are not co-located of several milliseconds, allowing them to identify trading opportunities, even at a profit of only \$0.01 per share. The HFTs trade at a rate of more than 1000 per day and it is estimated that up to 61% of daily US equity share volume and 70% of total daily trades are from HFTs [1]. The Commodity Futures Trading Commission (CFTC) and the Securities Exchange Commission (SEC) reported on the events of May 6, 2010 and identified that starting at 2:45:13 and lasting only 14 seconds, there were 27,000 trades attributed to HFTs [2]. That is nearly 2000 trades per second or one trade every 500 microseconds.

The major exchanges are adding more capacity for this technology by adding more square feet and bandwidth to the centers. They are also going to great lengths to make the access latency equal for all the various HFTs by matching network bandwidth, network speed, and even network cable lengths. One of the market centers specifies the market center gateway as ultra low-latency (<100 microseconds).

Co-located HFTs require precise UTC time to track orders in the market by placing timestamps in their buy and sell messages. Most HFTs have many computer systems running on dedicated LANs – each LAN using one PTP Grandmaster Clock and each computer on the LAN synchronized to that Grandmaster. These parallel, high-speed computer systems must be synchronized for the algorithms to process market buy-side and sell-side data. The timestamps and synchronization of computers are also used in network analysis of log files and all trading activities.

Understanding and measuring latency on the LAN is key to the success of HFTs. Without precise time synchronization below 1 millisecond, HFTs would find it difficult to optimize both hardware and software used. Real-time market data travels over cables, switches, and routers with latency averaging from 1 millisecond to 5 milliseconds for non-co-located traders [3]. Co-located HFTs reduce this latency to below 1 millisecond, giving them a significant amount of time to process market data as compared to traders that are not co-located. HFTs also measure the algorithmic latency to identify hardware or software improvements. Their goal is to approach near-zero latency.

IS-95 CDMA

Code Division Multiple Access (CDMA) is a wireless spread-spectrum communication technology used in North America by Sprint and Verizon. Other countries, including Brazil, China, Korea, India, Japan, New Zealand, Hong Kong, and others, also use CDMA. Referred to as IS-95 CDMA, it competes with technologies such as GSM provided by AT&T. However, worldwide GSM has more coverage, and in Europe, it is the only option. Most smart phones offered in the market today have a GSM and CDMA version available, with some providing dual CDMA/GSM hybrid models for more coverage. Currently, GSM and other technologies do not provide precise timing capabilities.

CDMA IS-95 does provide precise timing capabilities. CDMA base stations transmit timing signals that are sourced from GPS timing receivers in order to keep base stations synchronized to GPS time within 10 microseconds, even during periods of GPS satellite unavailability lasting up to eight hours [4]. A CDMA timing receiver provides precise time when GPS is available, and for critical timing requirements such as HFT, CDMA provides a backup to GPS. If GPS were unavailable, the CDMA base station continues to provide time and the CDMA receiver maintains the time for HFTs.

The CDMA receiver is a miniaturized module similar to an OEM GPS receiver and is embedded in the rackmount PTP Grandmaster Clock with an indoor antenna, as shown in Figure 2. The receiver is designed to demodulate and recover time in the forward link transmissions of an IS-95 CDMA network.



Figure 2. PTP Grandmaster Clock.

There are two types of IS-95 CDMA systems: the 869 to 894 MHz band, referred to as the Cellular band provided by Verizon in North America, and the 1930 to 1990 MHz band, referred to as the Personal Communications Systems (PCS) band provided by Sprint in North America.

The CDMA receiver is passive, that is, it only receives the transmission on the forward link from the base station. There is no need for a cell phone subscription, since there is no reverse link as required for a mobile cell phone. Figure 3 shows the overall architecture of the CDMA receiver [5]. The receiver preselects the Cellular or PCS band; the signal is amplified and down-converted in quadrature to an intermediate frequency (IF) that is then digitized by the analog-to-digital converters. The digital signal processor (DSP) and microcontroller receive the baseband data and phase-lock to the carrier and pilot code of the forward link. After the data are decoded, the DSP and microcontroller provide phase alignment of the 1 Pulse Per Second (1PPS) and Time of Day.

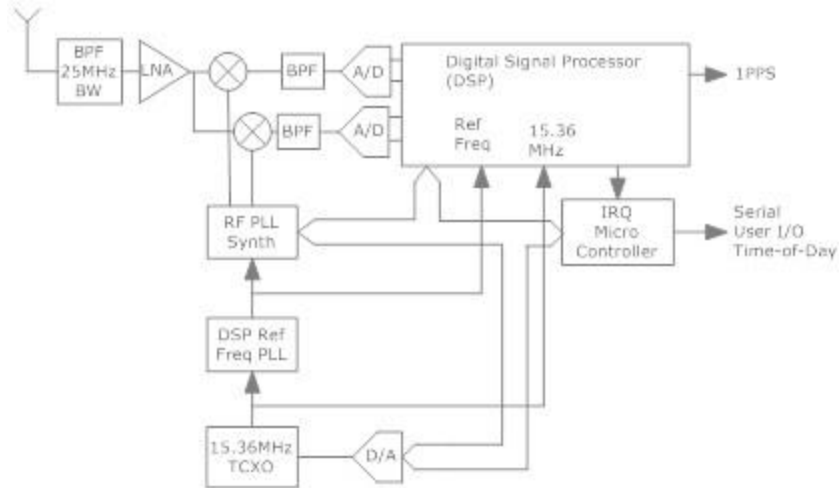


Figure 3. CDMA receiver architecture.

IEEE-1588 (PTP)

PTP is an IEEE-1588 defined protocol for synchronizing real-time clocks over a LAN to UTC with sub-microsecond accuracy [6]. Clock synchronization on the LAN requires at least one master and one slave. Multiple slaves can synchronize to a single master. The master clock provides synchronization messages that the slaves use to correct their local clocks. Precise timestamps are captured at the master and slave clocks. These timestamps are used to determine the network latency which is required to synchronize the slave to the master. There is a synchronization message typically transmitted every 2 seconds from the master, and a delay request message from the slave less frequently. UTC time synchronization is provided at each high-speed computer using PTP, as shown in Figure 1 (Link 4).

In Figure 4, the PTP Grandmaster Clock is a software PTP daemon coupled with a hardware clock and CDMA receiver. The CDMA clock providing PTP is referred to as the master and the HFT high-speed computers are referred to as the slaves.

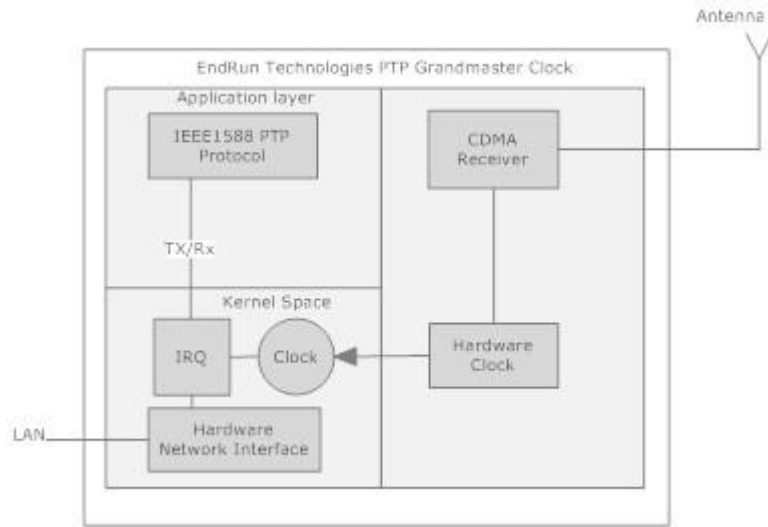


Figure 4. PTP/IEEE-1588 Grandmaster Clock

UTC TIME TRANSFER PERFORMANCE

Synchronization of clocks on a HFTs network is embedded in servers in the co-located market centers. Since the servers and trading algorithms are highly secretive, the following analysis of UTC time transfer data is representative of performance the HFTs may experience. It is measured lab data, indoors on a LAN at EndRun Technologies in Santa Rosa, California, with a CDMA Grandmaster Clock and PTP slave as compared to GPS. Tables 2 and 3 show the CDMA receiver status and the PTP configuration.

Table 2. CDMA receiver status.

CDMA Channel	North American Cellular
TFOM	6 (time error < 100 microsecond)
AGC	134 (average)
SNR	2.8(average)
FER	0.001 (average)
PNO	237

Table 3. PTP configuration.

Preferred Master	TRUE
PTP Identifier	GPS (CDMA is indirect GPS)
PTP Stratum level	3 (>100 nanoseconds)
PTP Subdomain	_DFLT
PTP Sync Interval	2 seconds

Measuring network-synchronized clocks on computers and servers is not possible using traditional methods. They do not have clock outputs such as 10 MHz or 1PPS that can be measured with a scope or analyzer for accuracy or stability. Because of this, most PTP systems rely on statistics logged at each slave. For the purpose of this paper, a PTP hardware measurement clock was configured. This clock contains the same PTP slave daemon that many HFTs are using, with a GPS-referenced clock providing hardware timestamping.

Statistics gathering at a slave clock is a simple method of understanding synchronization capability; however, it does not directly relate to the actual time of the slave clock. Hardware measurements of a slave to master using lab equipment such as a time-interval counter or oscilloscope are valid techniques only if both the master and slave clocks provide 1PPS. This method can measure the difference between the phases of the 1PPS pulses. While this is valid, it lacks the major time of day (Days, Hours, Minutes, Seconds), and can have an offset of several seconds and not be correlated with the collected data.

For the measurements presented here, a device was configured with a GPS receiver and a hardware real-time clock. A time-compare register was used to capture the minor (sub-seconds) and major times of the hardware clock and the PTP slave clock. The time compare was captured at 16-second intervals with the data collected in a data file on the system, as shown in Figure 5, and the UTC time transfer test configured as shown in Figure 6. Collected data were then removed from the system for analysis. Data were collected for 6.5 hours, the same time period that the market is open daily.

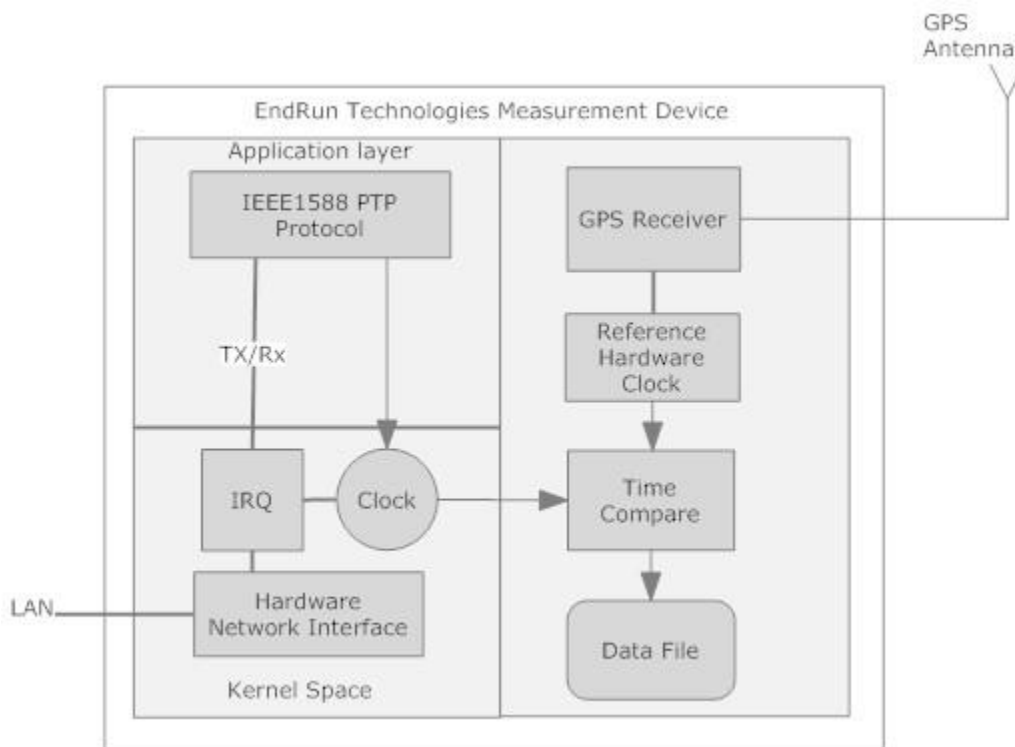


Figure 5. PTP measurement device.

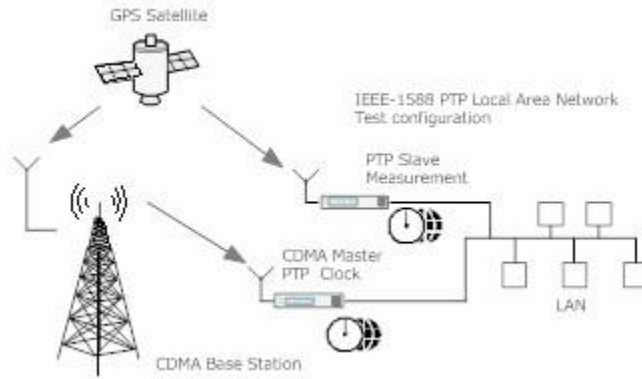


Figure 6. UTC time transfer test.

The analysis of the time synchronization is a combination of the CDMA master clock reference source, the network element that affects Packet Delay Variation (PDV), and the slave clock implementation.

The following plots in Figure 7 and Figure 8 include the PTP slave offset from the CDMA Grandmaster and the Histogram of the PTP slave locked to the CDMA Grandmaster.




Figure 7. PTP offset.

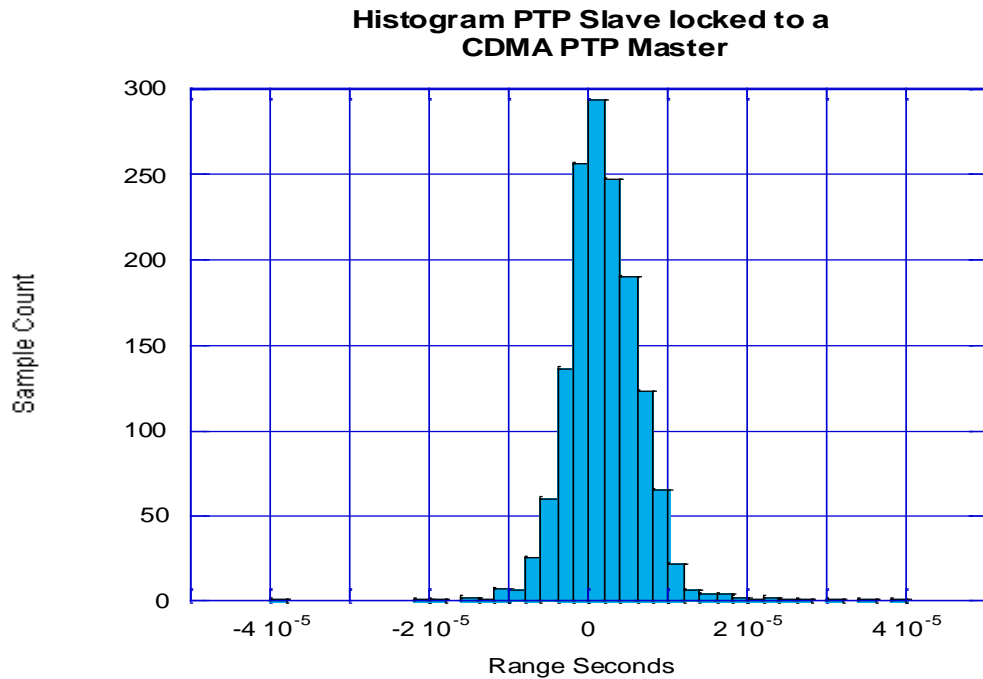


Figure 8. PTP histogram.

Taking a closer look at the statistics in Table 4, we can see that there is 5.04 microseconds RMS synchronization to the PTP Grandmaster and a mean offset of 1.39 microseconds.

Table 4. PTP measurement statistics.

Samples	1500 (6.5 hours)
Standard Deviation	4.84E-06
Mean	1.39E-06
RMS	5.04E-06
Min	-4.00E-05
Max	3.80E-05

CONCLUSION

The transfer of UTC time to HFT computer networks is currently provided using the combined technology of IS-95 CDMA and IEEE-1588 PTP. The need for indirect, GPS-derived UTC time is increasing as more traders move to co-located services provided by the market centers. Other applications can benefit from this combined technology, where better than 10 microseconds to UTC is

required and the ability to mount a GPS antenna on a roof top is either costly or not available.

REFERENCES

- [1] L. Tabb, 2009, “*TABB GROUP Response to Securities and Exchange Commission File No. S7-27-09*,” Release No. 34-60997 (New York).
- [2] US Commodity Futures Trading Commission and U.S. Securities & Exchange Commission, “*Findings Regarding The Market Events of May 6, 2010*,” 30 September 2010. Report of the staffs of the CFTC and SEC to the joint advisory committee on the emerging regulatory issues (Washington, D.C.).
- [3] NYSE Technologies Corvil, 2010, “*LatencyStats.com Latency Transparency for Market Data*,” White Paper (New York).
- [4] CDMA Development Group, 2010, “*System Time*,” http://www.cdg.org/technology/cdma_technology/a_ross/systemtime.asp
- [5] B. Penrod, 2002, “*A New Class of Precision UTC and Frequency Reference Using IS-95 CDMA Base Station Transmissions*,” in Proceedings of the 33rd Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 27-29 November 2001, Long Beach, California (U.S. Naval Observatory, Washington, D.C.), pp. 279-292.
- [6] J. Eidson, 2006, **Measurement, Control, and Communication Using IEEE 1588** (Springer, New York).