



Low Latency Trading

Three Hurdles in the Race to Zero

May 2009

by

Scott Parsons

Chief Scientist / Chief Architect

Exegy, Inc

In today's high frequency trading community where everything is carried out electronically there is much talk about latency. But few truly understand what latency is, where it comes from or how it manifests itself. The first step to eliminate latency is to understand it.

In the race to zero latency there are three major hurdles:

- Propagation latency
- Transmission latency
- Processing latency

Of these three, one has all but been eliminated by the world's leading high frequency traders. A second has been tamed, and the third is in the cross hairs for demolition by innovative new computing techniques.

Latency defines how quickly a trading system can react to changing conditions. In any system of networked computers, overall latency can be broken down into the three fundamental categories of propagation, transmission and processing.

Other forms of latency, such as queuing latency or network latency can be described in terms of these three fundamental types.

Understanding the relative magnitudes of latency in each category can yield valuable insight for optimizing the performance of a high-frequency trading system.

Propagation latency is the time required for a signal to move from one end of a communication link to another. This quantity of latency is a function of the speed of light, the physical nature of the communication medium and the distance between the endpoints. For an optical signal traveling through a fiber optic cable, the rate of signal propagation is roughly two-thirds the speed of light in a vacuum, or about 200 million meters per second. Another way of stating this is that every 20 kilometers between two endpoints equates to about 100 microseconds of propagation latency.

Since we can't change the speed of light, the only significant way to reduce propagation latency is to move the endpoints closer together. This fact gives rise to the increasing popularity of proximity solutions, where electronic trading systems are placed in the same geographic area as the trading venues themselves. If a trading system is within 20 kilometers of an electronic venue, the propagation latency can be as small as 100 microseconds. If the trading system

is co-located in the same data center and within 200 meters, the propagation latency can be smaller than a single microsecond.

Transmission latency is often confused with propagation latency, but it is a distinct phenomenon. Unlike propagation latency, transmission latency is not dependent on distance. It is directly related to how quickly a data packet can be converted into a series of signals to be transmitted on a communication link.

To send a data packet from one computer to another, it must first be turned into a stream of individual bits. Each bit is then transmitted in a serial fashion, one after another. The rate at which individual data bits are converted to signals is termed the bit rate of the communication link, or its link speed.

Many of today's high-speed networks utilize link speeds of 1Gbps, or one billion bits per second. A data packet consisting of 1,000 bits (or 125 bytes) requires one millionth of a second to transmit on such a link, or exactly one microsecond. More and more trading systems today are moving to even higher speed interconnect technologies, such as 10 Gigabit Ethernet or InfiniBand, which can transmit data up to 40Gbps. At these link speeds, a 1,000 byte data packet (8,000 bits) can be serialized and transmitted in a fraction of a microsecond.

It's clear that two of the three major component categories of overall latency in a high-frequency trading system can easily be reduced to less than a microsecond by moving close to trading venues and utilizing modern high-speed communication links.

Processing latency then becomes the category that dwarfs all others. It's the time required to act upon a data packet once it has been received. In a high-frequency trading system, this can involve transforming the data into alternate representations, updating the current view of the market, deciding whether to act upon this new view and issuing buy, sell or cancel orders into the market.

General purpose computer systems are typically used to host the processing logic of trading systems in the form of software. As their name implies, these systems are designed to be very general, and can be programmed through software to solve a wide variety of computational chores. These systems are a marvel of modern technology, but their very nature often leads to increased latency for specific applications.

The fundamental operation of all general-purpose processors today, regardless of the physical form they take, is to execute a sequence of stored instructions in a serial fashion, one instruction after another. Each instruction takes a fixed, finite

amount of time to complete that is related to the CPU's clock speed. The more instructions needed to implement a given piece of functionality, the longer the entire sequence takes to complete and the greater the processing latency.

A modern day microprocessor with a clock speed of 2.0 GHz can, optimistically, execute a single machine instruction every clock cycle. This equates to two billion instructions per second, or 2,000 machine instructions per microsecond. To put this into perspective, a single line of a high-level programming language will typically yield anywhere from 3 to 8 machine instructions. That means on average a single microprocessor can execute at most a few hundred lines of software code in a single microsecond.

When you consider the hundreds of lines of code that are executed by an Operating System simply to receive or transmit a data packet over a network, coupled with the thousands of lines of code needed to process a market data feed or implement a trading strategy or matching engine, it's easy to see how even simple software applications can quickly dwarf other sources of latency in a trading system. Even the most highly optimized software-based systems can easily take several thousand microseconds (milliseconds) to complete their chores.

Improving the processing latency on general-purpose computers means reducing the number of software instructions, increasing the processor clock speed, or introducing parallel threads of execution by way of multiple processors. But these approaches offer diminishing returns today. The increasing sophistication of trading algorithms demands more instructions, not fewer.

Meanwhile microprocessor speeds seem to have reached a plateau. Until recently, increases in CPU performance have been able to ride the coattails of Moore's Law, and their advances have been sufficient to meet the demands of trading systems. But that free ride has come to an end.

Moore's Law originally stated that "The complexity of an integrated circuit, with respect to minimum component cost, doubles every 24 months." This complexity is often measured in the number of transistors per unit area that can be implemented on a silicon chip. Increasing transistor density has led to greater capacities in memory devices and until recently, in greater clock speeds for microprocessors.

While Moore's law still appears to apply for increasing complexities, other economic factors mean that these advances do not translate into clock frequency speed increases for microprocessors. The increased clock frequencies made

possible by greater transistor densities have come at a cost. It takes more energy to switch a transistor on or off at higher frequencies, regardless of its size. The greater the clock frequency, the more electrical power is consumed and the more heat is generated by individual transistors. The increased power and cooling requirements are creating prohibitive operational costs.

So rather than increase clock speeds, microprocessor producers are incorporating multiple instruction execution cores to boost performance. This parallelism increases the number of instructions executed per unit of time, thereby helping to reduce processing latency. But the degree of parallelism achieved with this approach is limited to the number of instruction cores. More importantly, there is an upper limit on the speedup achievable with parallel instruction processors.

According to Amdahl's Law, the speedup of a program using multiple processors in parallel is limited by the time needed for the sequential fraction of the program – that portion of the program that cannot be parallelized.

Every software-based program has some inherent sequential nature. Indeed, the very structure of the instruction execution cores dictates sequential processing at the individual machine instruction level. There is nothing a software system can do to increase the parallelism of a single sequential instruction.

For this reason, parallel computing with multiple instruction execution cores is only useful for either small numbers of cores, or for so-called “embarrassingly parallel” problems. Relying on this approach alone is not sufficient to reduce total processing latency of trading systems to microsecond levels at a reasonable cost.

A new approach that avoids this limitation uses reconfigurable hardware to increase the degree of parallelism by several orders of magnitude. Reconfigurable hardware can host complex computational circuits whose logic is a direct implementation of a specific algorithm, but without the time-sequential nature seen in software. This “instruction-less” computing is able to utilize fine-grained, multi-layer parallelism in a way that surpasses the computational speed of coarse-grained, instruction-level and thread-level parallelism obtained with multi-core processors and software.

When this type of hardware acceleration is coupled with traditional software systems, the overall processing latency of the trading system can finally be reduced to microsecond scale, and no longer dwarfs other forms of latency. Microseconds can now rule the world of high-frequency trading and the high frequency trader can leap over the last significant hurdle in the race to zero.



www.exegy.com

Headquarters

349 Marshall Ave.
Suite 100
St. Louis, MO 63119
314-218-3600
877-247-5534 (sales)

Washington D.C

6701 Democracy Plaza
Suite 555
Bethesda, MD 20817
877-247-5534

For sales:

sales@exegy.com

For support:

support@exegy.com

For media

media@exegy.com

For all other inquiries

info@exegy.com

Copyright © 2009 Exegy Incorporated. All Rights Reserved. This document is a proprietary work of Exegy Incorporated and is distributed under an Exegy Incorporated license and is subject to its limitations. Except as set forth in the license, this document may not be copied, distributed, sold, bought, traded or posted on the Internet. All license restrictions, warranty limitations and limitations on liability apply.

Patent Statement:

The Exegy Ticker Plant is protected in part by the following patents: US 6,711,558, US 7,139,743 and US 7,181,437. Other patents pending.

