Tick-to-Trade Latency

Myricom® ARC Network Adapters for market feeds - offering lowest latency and the highest functionality.

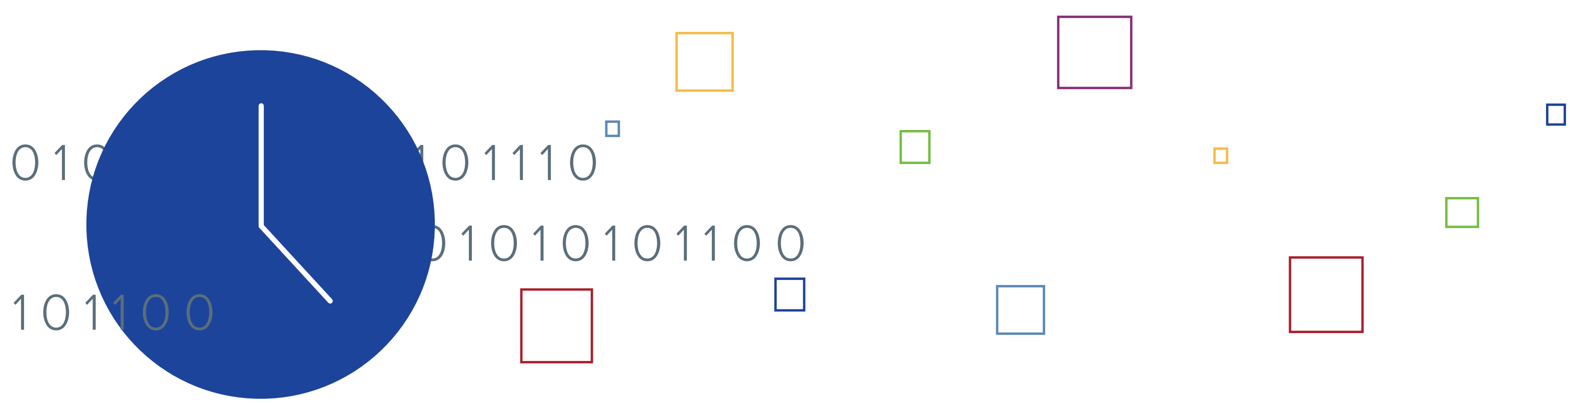

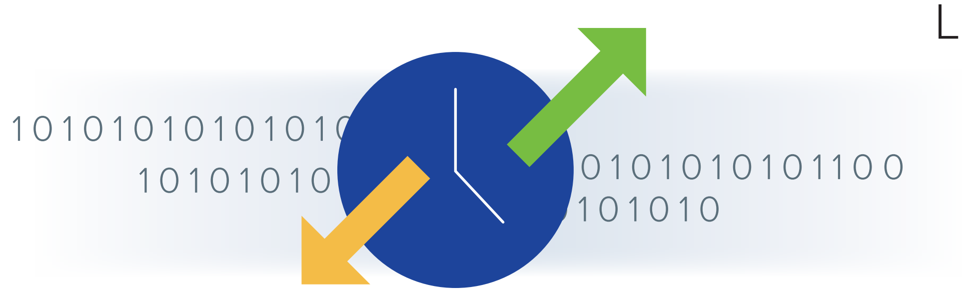

Tick-to-Trade Latency – Get Numbers That Mean Something

Why Latency Matters

Your trading operations must seize investment opportunities that last only for fractions of seconds. Competitors are also waiting to pounce. The key to who wins and who loses is the latency between digesting the market feed and placing the order; those who move fastest are the ones who will profit.

Tick-to-Trade – the latency you care about

Tick-to-Trade latency is the metric that matters, the time interval between receiving a market Tick showing opportunity to your algorithm and sending the Buy/Sell order.

Three Layers of Latency and Network Adapters

Let's analyze the latency in terms of 3 layers: (1) Network Adapter Receive Latency, (2) Application Latency and (3) Network Adapter Transmit Latency.

Figure 1

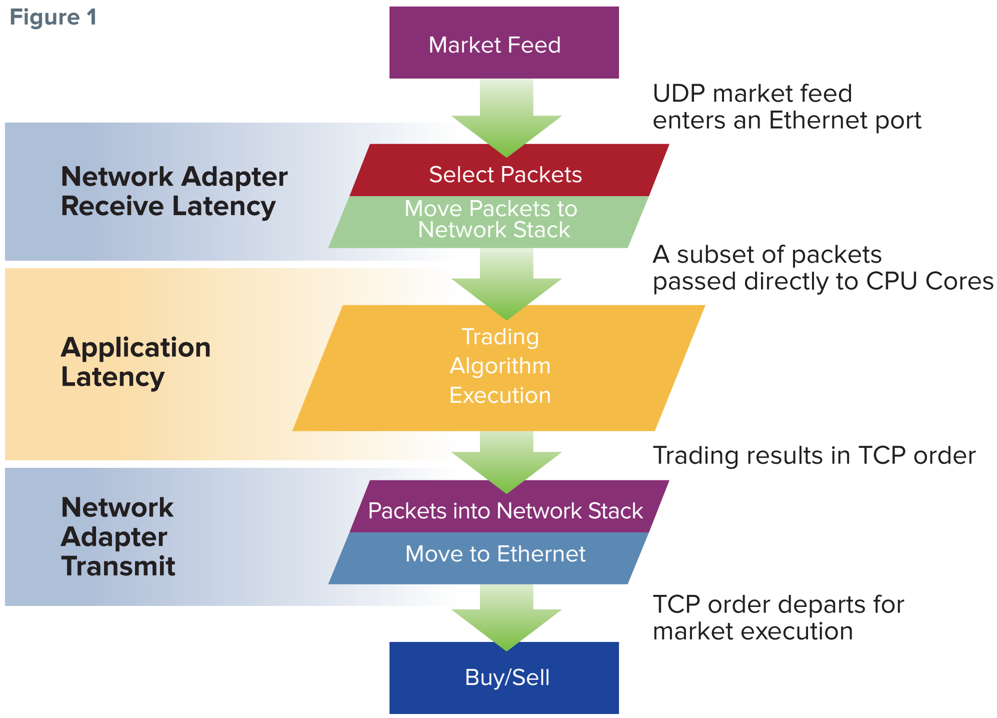

A network adapter should (a) drive the Receive and Transmit layers close to zero and (b) help your application compress the middle layer. Both are critical. This document will focus on a meaningful measurement for Receive and Transmit latency. To learn more about how an adapter can reduce your latency for all three layers, see the ARC Network Adapters data sheet and the Myricom DBL Integrated FPGA Firmware and Software technology brief, both available on our website at www.cspi.com.

Measuring Transmit and Receive Latency

This is a hardware performance number, so it should be a fairly straightforward measurement. The issue, however, is how do you set up the test?

The traditional and still most common measuring stick is the ping-pong test. This test is based on moving packets back and forth between two servers utilizing the Linux netperf command, first with a UDP protocol roundtrip and then with a TCP/IP protocol roundtrip. These roundtrip numbers are typically divided by two in an attempt to approximate a one-way trip, since trading applications take UDP in and send TCP/IP out.

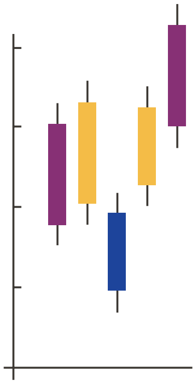

This is a fairly simple test, which is good, but not a clean capture of transmit and receive latency. The ping-pong test gives an average of Send and Receive for both protocols, when the functions of interest are just Receive for UDP and Send for TCP/IP. See the ping-pong sidebar for a complete description of the test, its limitations and some popular benchmarking tricks.

A more valid test is to measure a network adapter's tick-to-trade latency using a real round trip, with an absolutely minimal amount of time in the trading application server. This direct measurement of the Receive and Transmit layers is possible with current network adapters. We will illustrate this with CSPi's Myricom hardware but it can be implemented with any adapter that can accurately timestamp both received and transmitted packets.

Figure 2

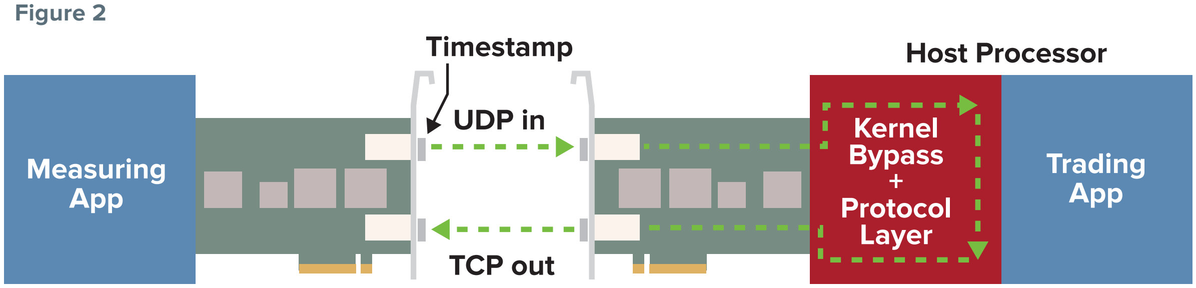

$$\text{Tick-To-Trade} = T_{\text{TCPout}} - T_{\text{UDPin}}$$

The test harness in Figure 2 directly measures the Tick-To-Trade latency of a network adapter connected to a host processor but bypassing application processing in that host. The receive and send messaging does go through Kernel Bypass and Protocol Layer functions, just as it would in real operations. Accurate measurements are achieved by using another Myricom ARC Series network adapter to initiate and terminate the communication; Myricom ARC adapters are equipped with both egress and ingress timestamping, supported by an extremely precise onboard clock.

Benchmark Results

The CSPi Tick-To-Trade benchmark tests use payload sizes selected in consultation with our customers as representative of a real-world trading environment, specifically a UDP payload size of 64 bytes for the packets flowing into the system under test and a TCP payload size of 400 bytes for the packets flowing out.

The results shown in Table 1 were all generated using a 4 Ghz Intel Haswell i7 “Devil’s Canyon”, single-socket workstation running Linux; several CSPi customers use an identical workstation configuration in their R&D labs. To configure the SolarFlare adapter we first used it in with SolarFlare’s proprietary ping-pong test to verify that we achieved the ping-pong results published by SolarFlare.

Measured Tick-To-Trade Latency in Microseconds, 25,000 Runs						
Network Adapter	Minimum	Mean	Median	99% Less Than	Maximum	Std Deviation
CSPi ARC Series E-Class*	1.675	1.770	1.757	2.004	2.757	.052
SolarFlare Flareon**	2.150	2.245	2.234	2.457	3.297	.049

Table 1

* The ARC E-Class adapter running firmware 2.0.6 and plugged into a Centos 7.2.1151 Devil’s Canyon 4.0 GHz host running DBL 5.1.0.

** Flareon® Ultra SFN7122F Dual-Port 10GbE adapter running and using OpenOnload®-201606 (as opposed to their low functionality, “raw” ef_vi API).

The ping-pong test: what it measures and how it can be gamed

Because the ping-pong test measures two roundtrips (one for UDP and one for TCP/IP) it includes timings for UDP Send and TCP/IP Receive, which are not relevant. Dividing the roundtrip number by 2 is meant to yield a valid one-way number but it really includes an average of Send and Receive. And because it measures the speed of transactions between two servers, it includes the average send time, average receive time, and average overhead contributed by each server. In addition to its questionable metrics, the ping-pong test is inconvenient to setup, requiring not just two adapters but also two identical servers.

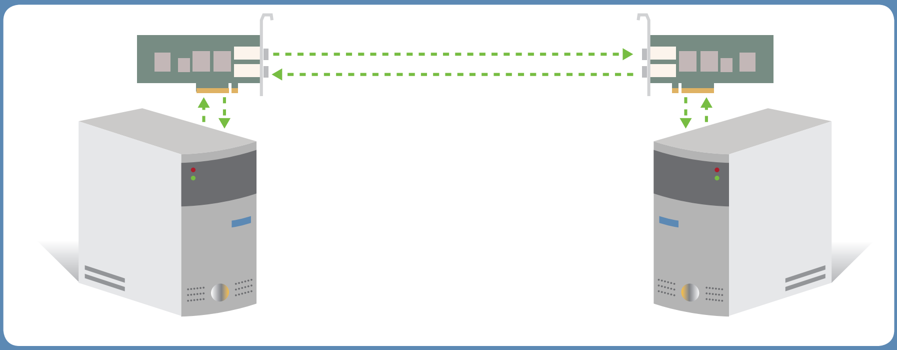

To improve test results, network card vendors use various benchmarking tricks which have nothing to do with card performance. Common techniques include:

- Creating a test harness with the fastest available servers, sometimes overclocked.
- Minimizing the packet size by using a zero payload option.
- Replacing the netperf function with a proprietary implementations of a similar request-response algorithm optimized for a specific adapter, sometimes even written to a “raw” hardware API (for example, SolarFlare’s 8000-series announcement)
- Reporting the lowest result number from 100,000 (or more) test instances. The 99% number is a much more valid indicator of true performance in any timing test.

All of these tricks demonstrate a level of technical cleverness but they obscure even further the value of ping-pong results.

Ping-pong test results

We argue in this paper against the usefulness of the ping-pong benchmark. Nevertheless, we are happy to present our results. On the hardware used for the Tick-To-Trade testing (4 GHz Haswell Devil’s Canyon and ARC E-Class network adapters) our minimum latency number for a useless, 1-byte-payload UDP ½ round trip is 1.28 microseconds. TCP is 1.83.

This UDP number is less than the competition. The TCP number is slightly higher. However, our focus is on optimizing the latency for data paths and functions actually used by our customers. Interactive discussions with our customers are key to directing our optimization efforts and we welcome ongoing input on how we can help improve overall trading application performance.

Numbers you should look for – getting close to zero

As shown in Table 1, our most recent Myricom ARC network adapter is 453 ns or 18.4% faster than the competition's 99% result. The ARC Series FPGA-based architecture enables continual enhancements and our models show that new generations of silicon and firmware will drive this number down even further.

Summary

The Tick-to-Trade latency measurements, impressive as they are, should just be the 1st step in your selection of a network adapter. They are really just 'table stakes', what it takes to get in the game. You also need an adapter that will help you compress that big middle layer by accelerating your application. To learn more, please visit our website or contact your CSPi representative for further information on our Myricom network adapters and the integrated DBL firmware + software solution.

1.675 µsec.

About CSPi

CSP, Inc. (NASDAQ: CSPi) is a global technology innovator driven by a long history of business ingenuity and technical expertise. A market leader since 1968, we are committed to helping our customers meet the demanding performance, availability, and security requirements of their complex network, applications and services that drive success.

CSPi Corporate Headquarters
175 Cabot Street - Suite 210
Lowell, MA 01854
800.325.3110 (US & Canada)

CSPi High Performance Products
800.325.3110 (US & Canada)
us-hpp-sales@cspi.com

CSPi Technology Solutions
800.940.1111
us-ts-sales@cspi.com
us-hpp-sales@cspi.com

www.linkedin.com/company/csp-inc
@ThisIsCSPi
us-hpp-sales@cspi.com
www.cspi.com

MYRICOM Network Products