CSPi

OPINION

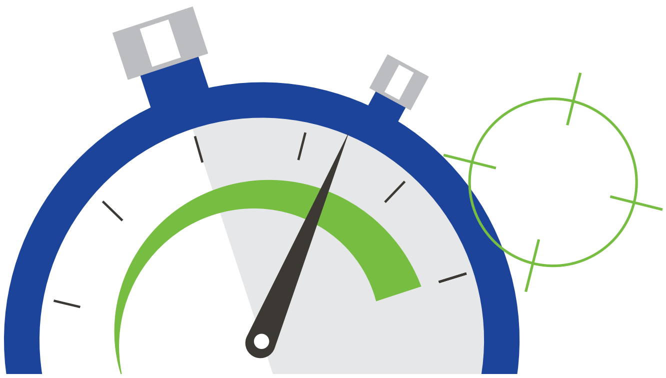

Get Accurate Timestamps for MiFID II with a Full Function Network Adapter

Accurate Timestamping Is No Longer a Luxury

Craig Lund
Vice President &
General Manager
High Performance
Products

For Automated Financial Trading, precise and accurate timestamping is no longer a luxury. There are new regulations like MiFID II* which require market participants to "prove" they can synchronize timestamps to UTC to within "1µ or less granularity" (this is the most extreme requirement in a regulation with a broad range of requirements).

The demand for accuracy is very real but meaningful data is often elusive. Industry datasheets are often deceptive regarding accuracy, quoting some narrow parameter under vendor control and completely ignoring the system-wide issues. Rather than comparing ill-defined data points, a better place to start is by understanding the factors affecting timestamp accuracy.

Sources of uncertainty

What are the sources of uncertainty that lead to lower accuracy?
Let us examine the flow of a packet through the system.

1. To meet MiFID II record-keeping requirements, you will need to apply a timestamp when a packet enters your adapter. If a vendor says that the uncertainty associated with applying their timestamp is +- 3 ns, they are saying that their ASIC or FPGA is using a 322.5 MHz clock (what 10 Gbit Ethernet uses) and examining the bits flowing off the Ethernet wire at the Unit Interval level. A simple statement like this overlooks the highly complex nature of dealing with the "elastic buffers" that compensate for Ethernet's permitted clock

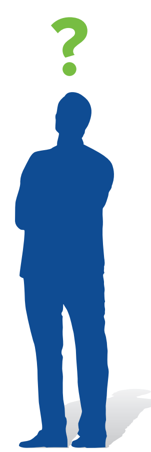

*(Marketing in Financial Instruments II) is a new set of detailed regulations for securities markets.

drift. Myricom could claim the 3 ns number but we prefer to use +- 5 ns. This is a very difficult number to actually measure and, in the system-wide scheme of things, is too small to be very important.

2. The timestamp cannot be any more accurate than the oscillator on the card that advances the ASIC or FPGA's idea of time (its "timebase"). You can generally ignore the oscillator's specifications if you are using a single oscillator to measure the latency of something quick, like a switch delay or an automated trade. However, if you are trying to compare timestamps taken hours apart in time, the oscillator on the adapter card can become a meaningful source of error. To avoid this discrepancy, all Myricom ARC adapters contain a high-quality, temperature-compensated oscillator. Some adapter vendors seek higher long-term accuracy by supplying an oscillator with an oven. At Myricom we take a different approach, offering a 10 MHz input so you can replace our oscillator with something better, including an atomic clock, if you need the best possible, long term accuracy.

3. If you want to compare timestamps taken by separate cards, you will need to synchronize the FPGA or ASIC clock with some source of external time. There are many choices. You can use an external Precision Time Protocol (PTP) or Network Time Protocol (NTP) master over Ethernet. You can also enhance your accuracy further by using a 1 Pulse Per Second (PPS) signal applied to an adapter card's timing coax input. The 1 PPS marks the beginning of a second, but does not tell you which second, which is why you need time from the PTP or NTP master. Once every second the PPS signal brings the adapter's clock accuracy to the PPS accuracy; this accuracy varies, based on the PPS source (see item 5 below).

4. To achieve consistent accuracy in Myricom ARC adapters, the FPGA will "trim its clock" to match the external source to within +- 10 ns. This synchronization integrates observations taken over a period 100's of seconds long to make small, incremental adjustments. Some adapters don't trim their clock and instead occasionally make their internal timebase suddenly jump forward or backward in time to stay in synchronization.

5. However, don't get too focused on that +- 10 ns accuracy number. It is added to whatever error might exist in the external signal you are synchronizing with. The best possible source is a 1 PPS signal emitted by a high-end GPS that has its own, high-quality, local oscillator. Over long periods of time a GPS signal is as accurate as the cesium atomic clock on the GPS satellite. However, over short periods of time the atmosphere gets in the way. Look to the GPS unit's specifications for short-term accuracy. The best you are likely to find is about +- 30 ns.

6. Synchronizing against that same grandmaster hardware using NTP instead of PTP is generally believed to synchronize to roughly a millisecond. This is clearly not good enough to meet MiFID II requirements for systems in a colocation center.

7. Synchronizing against a PTP signal is an even more complex discussion. For example, an ARC adapter could be deployed in the datacenter's PTP grandmaster connected to a very high-end GPS unit. Doing so results in maybe ±40 ns of accuracy at the grandmaster. You are not going to duplicate that accuracy at the slave (ordinary) PTP clock in your host. It is generally accepted that a very good PTP implementation can synchronize adapter clocks like the ARC implementation across a datacenter to one microsecond or less. Achieving this requires expert help.

You need to consider every item on the list above to both achieve and be able to verify extremely accurate timestamping, as demanded by MiFID II.

Today's practices fall short

Today most HFT firms meet their regulatory and contractual reporting requirements by deploying a stand-alone packet capture device in a data collection system. It watches the Ethernet connection to all the trading system servers, it sees the market feed enter them and sees the trades exit them. The data collection runs without effecting trade latency and reports can be run as needed. However, the new MiFID II reporting requirements cannot be easily met by simply snooping packets, as the reports require data not found within a packet.

Some trading firms also use packet capture devices to measure the latency of their internal systems by tagging trade packets with some indication of which UDP feed packet triggered an identified trade. (It's worth noting that some markets don't like this extraneous information appended to trade packets.)

Our opinion is that the easiest place to generate MiFID II audit trails is the trading system itself but to do that traders will need both a very accurate, low latency timestamp capability and the capability to implement not just Receive but also Send timestamps. MiFID II compliance will also involve consultation with timekeeping experts who can set up and document processes that prove a complete, system-level timestamp implementation actually meets MiFID II accuracy requirements.

Some packet capture devices do provide accurate timestamps and could be configured to receive MiFID data not in their timestamps by some side channel to the trading system. But using these devices for that purpose is not cost effective and unnecessarily complex. A full featured network adapter can not only deliver a market feed to the trading system but also implement highly accurate, low latency Receive and Send timestamp. It is even possible to put a real-time, latency “speedometer” onto the trading system status screen.

More than just accuracy – other considerations

Most network adapters on the market timestamp only IEEE 1588 PTP packets but do so on both ingress and egress. The egress side is required for the most accurate IEEE 1588 work. There are also some specialized packet capture cards that timestamp everything on the ingress side. As far as we know, the ARC network adapters are the only cards that timestamp either all or just specifically selected packets on egress too and feeds those stamps back to the application.

Why is an egress timestamp valuable for more than IEEE 1588 packets? Because it is the only way to accurately measure latency of a device under test. The financial Tick-To-Trade benchmark is one example – see the Application Note on that test.

You can also measure the latency of a device like an Arista switch. Just send a packet into the device under test, and receive it back, using a single card. Doing that with a single card means that you are using the same oscillator and can ignore all the errors that creep when clocks are synchronized. Those errors often exceed in magnitude the latency of something like an Ethernet switch.

Implementing Timestamping – Time APIs

Myricom ARC adapters implement the full suite of functionality called for in the Linux PTP Project. Thus your applications can use the SO_TIMESTAMPING socket option through both the kernel stacks and in the kernel bypass stacks provided with our DBL offering. We have an implementation of /dev/ptp which allows ptp4l to synchronize the FPGA’s timebase with a PTP signal. That signal can flow into any of the host’s Ethernet ports. In fact, for most of our customers the PTP synchronization is done through their motherboard’s 1 Gbit Ethernet port, not through our adapter. Also ph2sys synchronizes the CPU’s timebase to the FPGA timebase. All of this becomes even more accurate if you run a 1 PPS signal from the grandmaster to our adapter.

The story on Windows is more complex. Microsoft’s sockets have no equivalent to the SO_TIMESTAMPING socket option. Nevertheless, if you use Myricom DBL on Windows, the DBL kernel bypass stacks do provide the SO_TIMESTAMPING option. Microsoft offers no kernel APIs to ease the implementation of a slave (ordinary) PTP clock.

Summary

Applying accurate timestamping to your operations, whether for regulatory or other purposes, involves a careful assessment of both your needs and your system configuration. The Myricom ARC network adapters are designed to deliver the highest possible accuracy in timestamping for all types of packets, IEEE 1588 or not, on both the ingress and the egress sides of packet transfers. We would welcome the opportunity to discuss your specific timestamp requirements.

About CSPi

CSPi (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

Network Products