

The New York Times

"All the News That's Fit to Print"

SATURDAY, JUNE 30, 2018

Reprinted With Permission

National

Whittling a New York Minute To 100 Billionths of a Second

By JOHN MARKOFF

SAN FRANCISCO — Computer scientists at Stanford University and Google have created technology that can track time down to 100 billionths of a second. It could be just what Wall Street is looking for.

System engineers at Nasdaq, the New York-based stock exchange, recently began testing an algorithm and software that they hope can synchronize a giant network of computers with that nanosecond precision. They say they have built a prototype, and are in the process of deploying a bigger version.

For an exchange like Nasdaq, such refinement is essential to accurately order the millions of stock trades that are placed on their computer systems every second.

Ultimately, this is about money. With stock trading now dominated by computers that make buying and selling decisions and execute them with blazing speed, keeping that order also means protecting profits. So-called high frequency trading firms place trades in a fraction of a second, sometimes in a bet that they can move faster than bigger competitors.

The pressure to manage these high-speed trades grows when the stock market becomes more volatile, as it has been in recent months, in part to prevent the fastest traders from taking unfair advantage of slower firms. High frequency traders typically account for more than half of daily stock trading volume in the United States, according to data from the Tabb Group.

"The financial industry has easily become the most obsessed with time," said Balaji Prabhakar, a Stanford University electrical engineer who is one of the designers of the new synchronization system.

Because the orders are placed from locations around the world, they frequently arrive at the exchange's computers out of sequence. The new system allows each computer to time stamp an order when it takes place.

As a result, the trades can be sorted and executed in correct sequence. In a networked marketplace, this precision is necessary not only to prevent illicit trading on advance information known as "front-running," but also to ensure the fair placement of orders.

The importance of technical advances in measuring time was underscored by European regulations that went into effect in January and that require financial institutions to synchronize time-stamped trades with microsecond accuracy.

Being able to trade at the nanosecond level is vital to Nasdaq. Two years

ago, it debuted the Nasdaq Financial Framework, a software system that it has envisioned eventually trading everything from stocks and bonds to fish and car-sharing rides.

The new synchronization system will make it possible for Nasdaq to offer "pop-up electronic markets on short notice anywhere in the world, Mr. Prabhakar said. He cited the World Cup as a hypothetical example of a short-term electronic marketplace.

"There are tickets needed, housing, people will need transportation," he said. "Think of an electronic market almost like a massive flea market hosted by Nasdaq software.

To go from trading equities to managing all sorts of financial transactions will require more than an order of magnitude speedup in the company's networks of computers. It will be possible only if all of the exchange's computers agree on time with nanosecond accuracy.

A generation ago, computing usually took place in a single mainframe or personal computer. Now it is routinely spread across thousands of independent processors in machines that can be separated by a few feet or entire continents.

Chip designers have long struggled to maintain the precise timing needed to order mathematical operations inside individual computing chips. And synchronizing these vast ensembles of them has become the limiting factor in the speed and processing power of what Google describes as "planetary-scale" computers.

"It's kind of mind-boggling," said Peter Hochschild, a Google software engineer who specializes in the challenges associated with spreading software and data across networked computers. "Inside a processor, an enormous amount of stuff happens in a billionth of a second."

A billionth of a second is roughly the time it takes light to travel one foot. It has long been viewed as a crucial measure in computing. In the 1960s, the computing pioneer Grace Murray Hopper would hand out 11.8-inch lengths of wire to illustrate how designing smaller electronic parts would create faster computers.

Distance has become even more significant as software has begun to escape the boundaries of individual computers and make its way into the cloud — the web of giant computer data centers that have come to blanket the planet.

They are near dams to take advantage of cheap hydroelectric power and in cold climates to save on cooling costs. Microsoft has even begun submerging them in the ocean to take advantage of power generated by tidal surges.

Because software and data are no longer in the same place, correctly calculating the order of the events that may be separated by feet or miles has become the dominant factor in the speed with which data can be processed.

"So much of our expectation about computing being correct depends essentially on knowing this order," said Krishna Palem, a theoretical computer scientist at Rice University.

In the world of cloud computing, entire databases are scattered among different computers and data centers.

That has created tremendous challenges for the designers of electronic commerce systems. The new software synchronization standard under which Nasdaq's system would work, known as Huygens, is intended to replace a 33-year-old Network Time Protocol, or N.T.P., as well as more expensive approaches that have relied on atomic clocks and global positioning satellites.

Huygens, named for the Dutch physicist Christiaan Huygens, who invented the pendulum clock in 1656, uses so-called machine-learning techniques to synchronize a network of computers to within 100 billionths of a second. In contrast, the N.T.P. standard can synchronize computers no more accurately than a millisecond, or one thousandth of a second.

To ensure that buyers and sellers are treated fairly, Nasdaq has for decades looked for ways to ensure that trades are processed in the order they are placed.

While building a network for Nasdaq in the 1990s, Brian Reid, a computer scientist at Digital Equipment Corporation, experimented by coiling large rolls of cables of different lengths in a Massachusetts warehouse in order to insert tiny delays in the time it took data to travel in a network to make sure that messages were delivered fairly. He then employed timing information from satellites to synchronize clocks at different locations.

Google would later use this method to synchronize computers based on GPS data and atomic clocks to make sure that their database system could correctly order transactions. But since the system requires super-accurate clocks and satellite receivers, it is more costly than the software-based Huygens approach.

Mr. Reid built his original system in an era when the Securities and Exchange Commission required that all stock sales be entered by humans.

"Five millisecond accuracy in clock synchronization pleased everyone," he said. "It took much longer than five milliseconds to press the 'Enter' key on the big green terminals that people used.

(#S098875) Copyright © 2018 by The New York Times Company. Reprinted with permission.

For subscriptions to *The New York Times*, please call 1-800-NYTIMES. Visit us online at www.nytimes.com.

For more information about reprints and licensing visit www.parsintl.com.

**CLOCK
WORK**

Corporate HQ
3000 El Camino Real Suite 2-100
Palo Alto, CA 94306
United States

Contact Us
1-415-275-0431
hello@clockwork.io
www.clockwork.io