

# The Daily Transcript

www.sdtranscript.com

THURSDAY, AUGUST 2, 2018

## Startup aims to revolutionize computer speed, processing *Local entrepreneur, 2 physicists get results with unconventional approach*

By Mark Armao  
*Daily Transcript*

Imagine you're a salesperson. You plan on traveling to multiple cities, peddling your wares in each of them before returning to your original location.

Determining the shortest possible route that allows you to hit every city may not be difficult with three or four cities. But, as the number of destinations grows, the complexity of the problem skyrockets.

Even multi-million dollar supercomputers are hard-pressed to find solutions to such problems when the number of cities balloons to thousands or millions.

The problem is that, with each city you add to the list, the number of possible roundtrips grows exponentially. Even with 10 cities, the number of possible routes explodes to more than 3 million. And as the list grows, the number of possible routes to compute begins to outstrip the capabilities of the fastest computers. The traveling salesman problem is an example of what computer scientists refer to as a combinatorial optimization problem, and it is a well-known benchmark for various optimization methods developed over the years.

Although computer scientists have come up with shortcuts to cut down on the vast amount of time and resources needed to solve—or come close to solving—such problems, processing certain optimization problems can still take supercomputers days, weeks, even years to compute.

This means that, without a significant jump in capabilities, modern computers will hit a technological wall when faced with important problems that require heavy-duty computational power, from predicting global weather patterns to combing the human genome for insights that could spur the development of precision treatments for certain diseases.

Some think that next-generation quantum computers may hold the key to solving certain complex problems, while others contend that processing the most difficult problems in a reasonable amount of time is simply impossible. But, by approaching the problem in



Mark Armao/The Daily Transcript

Physicists Fabio Traversa (left) and Massimiliano Di Ventra (center) teamed up with serial entrepreneur John Beane two years ago to tackle some of the world's most complex computational problems using a novel approach.

a completely different way, a pair of Italian physicists have developed a technology that can solve certain optimization problems in a fraction of the time it takes the fastest computers.

Two years ago, they joined forces with a serial entrepreneur in San Diego to form MemComputing, Inc., a startup that has since tackled some of the toughest benchmarks in optimization and found solutions with unparalleled speed and efficiency.

"They looked at these problems and realized they could convert these computational problems into a physics problem, and it turns out that solving that physics problem requires a fraction of the resources and time," said MemComputing CEO John Beane. "Their research is groundbreaking to the point that there are people who think they're mad scientists, because what they're claiming is something that people don't believe can happen."

Beane is no stranger to turning long-shots into successful ventures. The local entrepreneur was an original member of Pyxis Corp., which later became CareFusion, and ecoATM,

which operates automated device-recycling kiosks. All told, his companies have returned more than \$1.5 billion to investors, he said.

In early 2016, Beane was serving as an entrepreneur in residence at the University of California, San Diego when he met two PhD physicists—Massimiliano Di Ventra and Fabio Traversa—who were working on an unconventional computing technology.

The Italian-born scientists began collaborating after Traversa came to UCSD to work on an unrelated project in the lab of Di Ventra, a professor of physics at the university. By the time they met Beane, they had produced multiple papers exploring the mathematical foundations for a theoretical device they called a universal memcomputing machine.

Hearing that the scientists needed "a little bit of cash to build a prototype," Beane thought to himself: "If I don't invest in it, then they're going to get grant money for something else and this will go on a shelf—and this could be groundbreaking technology," he said.

Beane took a chance on the fledgling en-

terprise. Just weeks after meeting, the trio founded the company and licensed the technology from UCSD.

### Busting benchmarks

While the team hopes to one day build their reimagined computer architecture in hardware, they can simulate its properties in software, said Traversa, who serves as the company's chief technology officer.

"Memcomputing, ultimately, is a circuit design," Traversa said. "But, we use a conventional computer to simulate [the circuit]. So the computational burden is not in the computer, but in the thing that you are simulating."

To vet the technology, Beane went to UCSD's San Diego Supercomputer Center, where researchers pitted the technology against a slew of well-known benchmarks in computer science. In one test, the software emulation was given a problem that takes the best solvers about 30 minutes to work out.

The MemComputing prototype solved it in under a second, Beane said.

Recognizing that the technology was no computational pushover, the Supercomputer Center upped the ante.

"They created this problem [for which] the estimate for current technology to solve it is beyond the age of the universe," Beane said. "We solved it in an hour and a half."

Perhaps more impressive is that the prototype was running on a \$17,000 server. The most advanced supercomputers, on the other hand, can cost upward of \$100 million, not to mention the multi-million-dollar annual energy costs.

Based on those results, the startup set its sights on refining its software-based approach, and has since developed multiple iterations of its MemCPU coprocessor technology.

On Wednesday, the company announced that it solved a problem known as the f2000 problem, a computational benchmark that has remained unsolved for nearly



Mark Armao/The Daily Transcript

MemComputing CEO John Beane sits with adviser Massimiliano Di Ventra and CTO Fabio Traversa in the offices of EvoNexus, a local incubator in which the startup is a resident company.

a decade. The problem, which is part of a well-known set of unsolved problems called the Mixed Integer Programming Library, was previously considered intractable, meaning that finding a solution would use too many resources to be useful.

---

**"Their research is groundbreaking to the point that there are people who think they're mad scientists, because what they're claiming is something that people don't believe can happen."**

*—John Beane, MemComputing CEO*

---

Running on an average Linux-based server, the MemComputing engine cranked out a feasible answer within 60 seconds.

A third party involved in the project also evaluated the technology against a set of problems that represent real-world challenges in various industries. According to the company, the technology found solutions at least 10 to 100 times faster than the leading commercial solver software.

The technology demonstrated unrivaled scalability when tasked with an efficiency planning problem directed at open-pit mining. In a challenge dealing with bioinformatics, the technology solved the problem 100 times faster than a leading solver, and produced a better solution to boot.

"We performed significantly better in finding high-quality results as compared to the leading commercial solver within the same time threshold," Traversa said in a prepared statement. "Moreover, depending on the complexity of the instance, we found them in seconds or minutes, versus hours or days."

To date, the technology has not come up against a problem with too many variables for the software to handle, Beane said.

But, while the executive is pleased with the lightning speed and unsurpassed power the technology has demonstrated in each test, he said its capabilities can prove problematic when trying to attract investors.

"If you create a technology that's a thousand times faster than what's out there today, that's a business," he said. "But if I tell you that I can solve something that can't be solved in the age of the universe, then you're going to say, 'Ok, you're smoking something.' So we have this credibility barrier to overcome because it seems too good to be true."

### A new architecture

Despite the mind-boggling speed and processing power of modern computers, they have some serious limitations when faced with a class of problems categorized as NP-hard or NP-complete.

Computer scientists have traditionally attacked these problems using algorithms, the procedural instructions that drive modern computers.

"Memcomputing is radically different," Di Ventra said.

As opposed to procedural algorithms, the MemComputing approach works by transforming the problem into a physical system made of circuits that are unlike those used in classical computing.

Standard computer circuits operate by constantly shuffling information—collections of 0s and 1s—between a central processing unit and a memory area. This division of labor stems from the nature of conventional components; memory areas can't do processing, and processors can't store memory.

It takes just a tiny bit of electricity and a fraction of a second to bounce the data between the two locations, but when the process is multiplied millions of times over, it becomes a massive undertaking.

The circuit elements in a memcomputer, though, can simultaneously process information and store memory in the same place—hence the name memcomputing. By cutting out the time and energy required to send information back and forth, the dual ability of each "memprocessor" leads to an enormous leap in efficiency, Di Ventra said.

The ability to store and process information in the same units is analogous to the way the neurons in a human brain work, Di Ventra said, adding that the capability gives rise to other unique properties.

In traditional computers, operations are performed by "logic gates" that have input terminals and output terminals. Signals pass through collections of logic gates sequentially in one direction, from input to output.

Memcomputers rely on a different type of logic gate, in which all terminals can function as both input and output, Beane said.

"The gates are wired up in such a way that they can talk to each

other in forward and reverse at the same time,” Beane said in an email.

This allows memcomputers to employ the collective behavior of all, or large chunks, of its units simultaneously, Di Ventra said.

As opposed to today’s computers, which solve optimization problems by systematically checking every possible permutation in sequential cycles, the components in a memcomputer communicate with each other to compute the solution in unison.

Di Ventra explained the process using a real-world computational problem: finding a way out of a maze embedded into a circuit. While classical computers would use an algorithm to explore the maze in small consecutive steps, a memcomputing machine could compute and store the solution in one shot.

When a small pulse of electricity is applied across the maze’s entrance and exit points, “all units react collectively so that the electrical current can only flow through the paths that represent the solution [or solutions] of the maze,” Di Ventra said in an email.

By altering the properties of those gates that belong to the correct solution, the current leaves a “memory trace” along the correct path or paths, Di Ventra said.

“Whether we fabricate these electrical circuits as hardware or simulate them in software, we let the physical system find the solution for us without the need to explore a vast space of possibilities as in traditional algorithms,” Di Ventra said. “This allows us to greatly accelerate the computational time and find more accurate solutions faster than using traditional solving methods.”

### Cracking the code

Aside from their academic accomplishments, the team recently collaborated with a San Diego-based startup on a project they say

further validates the technology.

Canvass Labs helps software developers scan the source code used in their products to identify what parts of the code are “open source.”

Open source software operates using code that anyone can modify or integrate into their own software. The system is intended to streamline and optimize software development, but it carries various legal and cybersecurity implications.

In the wake of multiple high-profile cases in which the use of open source code led to proprietary disputes or security breaches, large software companies are spending hundreds of thousands of dollars to analyze their code to ensure it is legally compliant and secure from cyber-attacks.

Canvass Labs CEO Peter Shin founded the company last year with the aim of automating the time-consuming process of code scanning, which often produced numerous false positives that had to be corrected manually.

Even using Canvass Labs’ unique approach, it was still taking an exorbitant amount of computing power to analyze the code.

Scanning a set of about 1,000 open source packages, it was taking the startup more than five minutes to get meaningful results on the source code.

The results were commendable, but because the goal is to eventually match a company’s code with a library of hundreds of thousands of open source packages, the exponential increase in compute power would make scaling the technology impractical.

Shin heard about MemComputing from a mutual acquaintance, and enlisted the team’s help in the hopes of accelerating the process. By pairing the MemComputing software

with the Canvass Labs solution to scan the same 1,000 open source packages—which likely contain more than 50 million lines of code—the system cut the processing time from several minutes to less than a second.

Not only were the results more accurate than the previous results, but they allowed Canvass Labs to explore the data in new ways.

“Their technology is allowing us to look at a problem in a way that is not possible otherwise,” said Shin, who has held high-level positions at the Supercomputer Center and Qualcomm. “It’s allowing us to do something that traditionally would require an extremely large amount of resources.”

### Speed bumps

While the digital engines have demonstrated remarkable capabilities, the question of whether an analog version of the technology could scale up to tackle even more complex problems remains uncertain. Some in the computer science world have argued that, for very technical reasons, the proposed memcomputing machine suffers from technological limitations that stunt its scalability.

Those limitations are what forced the team to pursue a digital approach in the first place, but, owing to the apparent success of their software solution, the startup is content to stick with software for now.

“When I first met these guys,” Beane said, “I thought we were going to be a hardware company. But, by testing the prototype, we saw that the software simulation is orders of magnitude faster than everything else.

That was well beyond what my expectations were.” And, because the speed of the software simulation is partly dependent on the hardware running it, Beane said pairing the technology with more advanced hardware would lead to

increases in speed and computing power.

“One of the advantages of the technology is—if you’re solving a problem in 20 minutes and you want it to take 20 seconds—there’s the possibility of throwing additional hardware at the problem and bringing it down to 20 seconds,” Beane said. Despite the limitations of an analog approach, Beane said, the company has had “conversations” with semiconductor companies about the possibility of creating chips with memcomputing components.

Di Ventra noted that a hardware version of the technology would be required for certain tasks, such as the routing of autonomous vehicles fleets in real time. “For certain applications where you need real-time computing, you really need the hardware,” he said. “If you want immediate response in an autonomous vehicle, you cannot simulate the machine [because it may take seconds to compute].”

### Moving forward

For now, the company will forge ahead with its software solution.

Beane said the company will continue to improve on its coprocessor technology while formalizing its go-to-market strategy, before pursuing Series A funding sometime next year.

And, while some in academia may think the startup’s claims are far-fetched, Beane is confident that early adopters will come forward, knowing the technology could help their companies shave hundreds of thousands of dollars’ worth of computing time from their bottom line.

“Industry may think we’re crazy too. However, industry knows they have problems, and they haven’t been able to solve those problems,” he said. “What I love about industry is industry has the guts to give it a try.”