RISC-V, the Linux of the chip world, is starting to produce technological breakthroughs

RISC-V, the open standard for chip instructions, is leading to some impressive technical innovation, one of its creators says.

By Tiernan Ray | November 30, 2020 -- 18:47 GMT (10:47 PST) |

A decade ago, an idea was born in a laboratory at the University of California at Berkeley to create a lingua franca for computer chips, a set of instructions that would be used by all chipmakers and owned by none.

It wasn't supposed to be an impressive new technology, it was merely supposed to get the industry on the same page, to simplify chip-making in order to move things forward.

But a funny thing has happened on the way to a global chip standard: RISC-V, as the Berkeley effort is known, has begun to produce some technical breakthroughs in chip design.

As just one example, a recent microprocessor design using RISC-V has a clock speed of 5 gigahertz, well above a recent, top-of-the-line Intel Xeon server chip, E7, running at 3.2 gigahertz. Yet the novel RISC-V chip burns just 1 watt of power at 1.1 volts, less than one percent of the power burned by the Intel Xeon.

"The potential was always there" for innovation in chip design, said Dr. David Patterson, co-creator of the open standard chip instructions known as RISC-V. "Maybe because of all the competition, we are starting to see some really interesting points in the design space being realized."

Tiernan Ray for ZDNet

The speed and power efficiency of the RISC-V part also top the specs for Exynos 4, a top-of-the line part made by Samsung Electronics for its smartphones, based on the computing core provided by ARM Holdings Plc, Intel's chief rival.

"It's kind of amazing," said David Patterson, a professor at the University of California at Berkeley who helped create RISC-V, in an interview with ZDNet, describing his impression of a demo he was given of the chip recently. "I think IBM mainframes have a 5-gigahertz product that's liquid-cooled, and takes 100 watts" to run.

"I've also heard some impressive numbers around people doing FPGAs, around 600 megahertz," said Patterson, referring to re-programmable chips. "For a soft core, that seems pretty fast."

Patterson is surprised, he said, that technological innovation is cropping up. "The potential was

always there" for innovation, he said, but that wasn't the main expectation when he and a fellow Berkeley professor, Krste Asanović, first wrote their manifesto for RISC-V, back in 2011.

"One thing I thought would happen is, because it's open, we would see all this competition," he reflected.

"Maybe because of all the competition, we are starting to see some really interesting points in the design space being realized," said Patterson, who also serves as a distinguished engineer at Google.

The new 5-gigahertz processor, which is merely a prototype, is not the creation of a garage startup. It was made by Micro Magic Inc., a Silicon Valley intellectual property designer for chips that has been consulting to all the big Valley firms for twenty-five years. The ability of a small but seasoned crew of chip designers to accomplish such a task suggests a design renaissance that could be on the horizon.

Not only is the chip faster at lower power, it scores higher than the Intel and the Samsung chips on a benchmark score, called CoreMark, of raw CPU performance, as recorded by the Embedded Benchmark Microprocessor Consortium. The RISC-V chip has a score of 13,000, more than double the per-core performance score of the ARM-based Exynos. While the Intel Xeon is nominally higher per core, at 26,009, the Xeon part takes many more threads of execution, 120, to reach that performance.

Dr. Andy Huang, a longtime chip industry executive, works as the business liaison for Micro Magic. He explained to ZDNet in an interview by phone that the breakthrough lies in the way the CPU and memory interact. The two founders of Micro Magic, Mark Santoro and Lee Tavrow, had patented in the early Nineties an SRAM computer-memory chip that was the fastest such memory ever invented.

The RISC-V prototype eliminates a bottleneck that can exist with fast memory and slower chips.

"If memory is running at five gigahertz and logic is running at one gigahertz, who's the bottleneck?" Huang teased, without disclosing details.

The point, says Huang, is that because RISC-V is open, unlike CISC, the complex instruction-set architecture of Intel's chips, or even the version of RISC that's in ARM chips, things can be done in chip design to resolve that bottleneck in ways not possible if the chip's instructions were locked down.

The analogy he used is Android versus iOS.

"I ask my son why he prefers Samsung [smartphones] to Apple, and he says it's because if he wants something changed, he can just ask one of his programming friends to do it for him, because Android is open, unlike iOS," said Huang.

"This is why we attribute all of our success to Dr. Patterson," said Huang. "He created the most efficient, the most elegant RISC architecture, by far."

"We should call him Saint Patterson," said Huang.

It's not merely the ability to tinker with the instruction set that makes possible the kind of part that Micro Magic has shown. There is an economic element that comes into play.

Unlike CISC or ARM's instruction set, which each have over 1,000 instructions, RISC-V has fewer than one hundred instructions, Dr. Huang emphasized.

Because of the simplicity of the instruction set of RISC-V, Micro Magic was able to have its chip produced using a standard silicon wafer with no special tweaking. That makes it possible to use what is called a shuttle run, where the chip is grouped together in the manufacturing process with other people's chips, on the same wafer. That can be vastly cheaper because the wafer's cost is shared among so many parties.

"People talk about $100 million to do a custom ASIC," observed Patterson, meaning, a chip that is tuned for a particular application. "Well, they did not spend a hundred million dollars to do that," said Patterson of Micro Magic's effort.

Although not emphasized by Patterson nor by Huang, there is a second economic element at play: It is a lot easier to use a shuttle run when you don't have the overhead of paying for an ARM license, which then has to be amortized across many parts.

In a sense, then, RISC-V can potentially promote the micro-batch approach seen in a lot of modern product lines, from breweries to cheese to clothing.

The question of economics is a provocative one considering that Nvidia, one of the biggest chip makers in the world, is in the process of buying ARM for $40 billion. That sale would put Nvidia in a position to reap the royalty stream from ARM's intellectual property, and to dictate the roadmap of development for the world's most widely used chip instructions.

Nvidia CEO Jensen Huang has described ambitious plans for what he would do with ARM, and he has assured Wall Street that ARM's licensees, his competitors, won't mind his buying their biggest supplier.

But the move obviously presents a new opportunity for alternatives. Patterson, asked about the deal, is circumspect. Nvidia is a member of the RISC-V ecosystem and has endorsed the technology whole-heartedly.

"I think people have tended to think of RISC-V often as just an academic idea," Patterson told ZDNet. "And then, when it's demonstrated that proprietary instruction sets can be bought and sold, that becomes another argument for an open architecture."

Micro Magic's Huang says he has had inbound approaches from tech giants since Micro Magic posted its brief chip announcement.

"I have received an email from two of the four trillion-dollar, publicly listed companies already," said Huang, without disclosing names.

Huang offered hypothetical scenarios in which the chip could be used by Apple or Google to make breakthroughs in energy consumption.

"Google already owns the mobile open-source software Android, imagine the benefit to all mobile customers if they would also own the most power-efficient, highest-performance, open-source RISC core," Huang told ZDNet.

“Imagine the newest Apple Watch not having to be recharged overnight,” is another possibility Huang offered.

With or without such mega-deals, Micro Magic, said Huang, hopes to get its RISC-V intellectual property into more and more designs in order to make a substantial dent in the the world’s electricity consumption.

“Our intention with this IP is to help the world, help the PC, the laptop world, the tablet world, the mobile phone world, the wearables, the gaming, the electric car, and the IoT — you name it, everything, our goal is to contribute to reducing the carbon output of the world by half.”

One prototype CPU does not a revolution make. Comparisons to actual shipping product by Intel and others leaves out the fact that a lot more parts are required to make a finished chip design.

That is where the ecosystem of companies around RISC-V becomes important. The number of announced companies saying they’re going with RISC-V is small but growing.

“All the products you can think of, all the way up to data center, there are people thinking very seriously about RISC-V now,” said Patterson.

“There is a sense we have turned the corner,” he said, “from, Why would I ever use RISC-V” a few years ago, “to, Why wouldn’t I use RISC-V?”

Prominent among the ecosystem parties are SiFive, a Silicon Valley startup that has for several years been developing chip intellectual property exclusively based on RISC-V. In August, the company started a business unit dedicated to producing custom chips for a variety of applications including AI and edge computing, known as OpenFive. Patterson’s collaborator, Professor Asanović, is chief architect for SiFive.

Another is Andes Technology of Taiwan, a maker of embedded processors, which has over the years collectively sold billions of CPU designs to makers of electronics products.

Both SiFive and Andes Technology last month presented new chip designs for AI using RISC-V at a prominent chip technology conference, the Linley Fall Processor Conference.

SiFive told ZDNet it now has over 200 design wins with more than 80 companies, including six of the top ten semiconductor makers. “With design wins in FADU, Huami, Qualcomm, Samsung, and Synaptics, SiFive has tens of millions of cores shipping currently,” SiFive told ZDNet.

Andes told investors in its quarterly report this month that roughly a third of its revenue this year has been from RISC-V based parts.

Both Seagate Technology and Western Digital, large makers of disk drives, are sponsors of next month’s RISC-V Summit, the third annual technology gathering of the ecosystem. That event is sponsored by the RISC-V International Association, a non-profit corporation that now represents over 750 parties working to advance the standard, including Chinese smartphone vendor Huawei, chip makers Xilinx and Qualcomm, and IBM. Asanović is chair of the group, and Patterson is vice-chair.

No matter how successful RISC-V proves, however, the world may never know the full extent of its usage. That is because while ARM and other commercial technology providers make their licensees

sign documents, no one using RISC-V has to disclose usage.

RISC-V International asks vendors to voluntarily disclose usage, but does not compel such disclosure.

For that reason, “It’ll be hard to see concrete evidence” of the extent of RISC-V usage, said Patterson.

Nevertheless, evidence of technological progress in a part like Micro Magic’s suggests to Patterson and others that the impact of RISC-V could ultimately be large.

Recently, Patterson conducted a series of one-on-one interviews with his collaborators, via video, to produce a virtual celebration for the tenth anniversary of RISC-V’s creation.

One collaborator offered a striking perspective, Patterson told ZDNet.

“In five or ten years, RISC-V could be the most important instruction set” in the world, Patterson recalled the individual saying.

“On one hand, it sounds crazy,” he said, “but it’s not impossible.”