RIKEN

arm

How supercomputer performance and power efficiency can coexist

Building the world's fastest supercomputer doesn't mean you need to break the power budget, as engineers at Fujitsu and RIKEN proved.

Case Study

Company snapshot

Name: RIKEN Center for Computational Science

Description: RIKEN, Japan's largest comprehensive research institution, completed Fugaku, a supercomputer powered by an Arm-based processor from Fujitsu, A64FX. Fugaku is currently the fastest supercomputer on the planet.

Website: www.r-ccs.riken.jp/en/postk/project

Goal

Build a high-performance supercomputer with world class energy efficiency to deliver powerful research tools into the hands of scientists.

Solution

Researchers and engineers at RIKEN had, as early as the early 2000s, been researching power efficiency in computing. Their initial effort, the K supercomputer – named for the Japanese word/numeral "kei" (京), meaning 10 quadrillion – was manufactured by Fujitsu and installed at the RIKEN Advanced Institute for Computational Science campus in Kobe, Japan.

K contained 84,000 nodes, ran at 11.28 petaflops peak performance and used 1.27 petabytes of memory, leveraging their proprietary Tofu interconnect. The machine, based on Fujitsu's SPARC64 V chip, won the ACM Gordon Bell prize two years running in 2011–12 and was retired in August 2019 as No. 1 on the Graph 500 list.

"It's amazing that people have such passion for a box full of silicon and wires," says Professor Satoshi Matsuoka, of the Tokyo Institute of Technology (Tokyo Tech), RIKEN.

RIKEN engineers were at work as early as 2010 on the successor to the K computer, and, in laying out their objectives, wanted to include features to monitor and control power usage, which K lacked. Working with Fujitsu with the official launch of the project in 2014, they devised a fine-grained approach to the processor, which would include technologies such as power gating (turning on only the parts of the device needed to accomplish a given task at a given moment).

"The power technologies were based partly on what we did with embedded research, and we knew we could put them into production because they were based on mobile processors," Matsuoka says. Fujitsu settled on the Arm architecture as the foundation for what would become the Fujitsu A64FX chip.

The nearly 9-billion-transistor A64FX is the first processor to use the Arm SVE (scalable vector extension) 512-bit instruction set.

Objectives

At a high level, the objective of most HPC systems is to deliver massive amounts of compute resources to enable researchers to model the world around us – the notion of building "digital twins" of real-world objects and systems – so that we might better solve problems that face us today and avoid those we anticipate tomorrow. And one of the world's key objectives in recent years has been to cross the exascale compute barrier – a quintillion (10^{18}) calculations each second – which is the level at which it is believed the human brain can be simulated.

But compute solutions can't be limitless. They are constrained not only by development cost but increasingly by the cost of power. Whether it's a university, a research center or a government lab, each has to pay a hefty utility bill to run their supercomputer. And taking this a step further, most energy delivered into the world comes from non-eco-friendly sources, such as coal-fired power plants, so HPC systems that run more efficiently are doing their part in the global warming fight.

This is where the concept of compute efficiency comes into play and how the story of the Fugaku supercomputer becomes so compelling.

Fugaku (an alternative name for the iconic Mt. Fuji) began to be shipped to RIKEN on December 2, 2019 and installation was completed in the spring of 2020. The machine runs Red Hat Enterprise Linux unmodified. On May 13, RIKEN announced its detailed specifications: 158,976 A64FX processors (each containing 48 Arm cores for a total of more than 7.5 million Arm cores) running at 2.2GHz spread over 432 racks.

When the Top 500 list was announced, Fugaku was ranked the world's fastest supercomputer, with a peak of 2.1 exaflops of FP16 and 4.3 exaflops of INT8 performance for AI workloads. Performance is 537 petaflops on 64-bit FP, 1.07 exaflops on 32-bit FP as well. Its high-bandwidth memory (HBM) is 32 GB per node at approximately 1 TB/sec of bandwidth per socket, for a total of 4.85 PB of memory capacity and 163 PB/second aggregate memory bandwidth.

Just as impressive is the power efficiency. The prototype version in November showed it to be the most power efficient in the world according to the Green500 supercomputers list at 16.876 GFlops/watt.

Matsuoka describes the machine as being a factor of 3 times more energy efficient than supercomputers based on conventional CPUs, and even supersedes those based on special accelerators such as GPUs. This is all the more astonishing when considering that most modern HPC machines are massively parallel, and power efficiency limits the scale of the machine.

"We can run same compute at one-third the power requirement," he says. "Had we not designed our own chip and met expectations and had to purchase existing chips from outside, the machine would have been three times bigger."

Translated into money, that means that RIKEN's energy bill would have been nearly $200 million a year, instead of the $60 million the institute spends with the energy-efficient Fugaku.

In the spring of 2020, as the world grappled with the spread of the novel coronavirus, the research center has opened up the machine to the world's leading researchers to help model possible drug solutions.

"We're running the COVID-19 program in our resources that are already set up. That's about one-sixth of the machine so far, but even so this is several times faster than any other machine in Japan," Matsuoka said. "We run those side by side other development work. Getting those massive resources to scientists is important."

Future Impact

Matsuoka hopes that the lessons of the Fugaku design be extrapolated elsewhere, into areas like data center servers.

"We hope that the chip we developed for this machine or what we design for future generations of machines makes it into mainstream market and sells into the millions," he said. "That's how we assess part of our impact; in other words, we advanced the state of the art in IT. That's why Arm Neoverse is important. You need powerful, efficient machines in the data center."

Additionally, energy efficiency is important for bigger end markets, such as the Internet of Things (IoT). The goal of achieving a trillion connected devices in the coming years "is only achievable if we get massive power savings. You can't sustain a trillion processors worldwide without it," he says. "Achieving a higher degree of savings is paramount to this ubiquitous computing vision."

Learn more about Fugaku and about Arm Neoverse computing solutions at www.neoverse.com