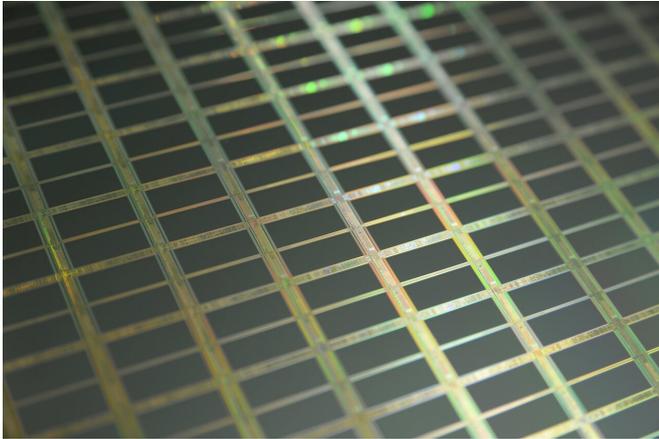


Trillion-transistor chip breaks speed record

26 November 2020, by Peter Grad



Credit: Unsplash/CC0 Public Domain

The biggest computer chip in the world is so fast and powerful it can predict future actions "faster than the laws of physics produce the same result."

That's according to a post by Cerebras Systems, a [startup company](#) that made the claim at the online SC20 supercomputing conference this week.

Working with the U.S. Department of Energy's National Energy Technology Laboratory, Cerebras designed what it calls "the world's most powerful AI compute system." It created a massive chip 8.5 inch-square chip, the Cerebras CS-1, housed in a refrigerator-sized computer in an effort to improve on deep-learning training models.

Cerebras team leader Michael James and the Department of Energy's Dirk Van Essendelft said in their paper that the CS-1, powered by 1.2 trillion transistors, performed at 200 times the speed of a Joule supercomputer in a simulation of powerplant combustion processes. They said the chip's performance cannot be matched by current supercomputers regardless of the number of CPUs and GPUs they house.

The Joule uses Intel Xeon chips with 20 cores

each, a total of 16,000 cores. It took six milliseconds to calculate the combustion processes. Cerebras accomplished the task in 28 microseconds.

The [chip](#) is composed of 84 virtual chips along a single silicon wafer and 4,539 computing cores, which means there are effectively 381,276 computing cores to tackle mathematical processes in parallel. Packed with 18 GB RAM, the cores are connected with a communications fabric called Swarm that runs at 100 petabits per second.

"This work opens the door for major breakthroughs in scientific computing performance," Cerebras researchers wrote in a blog post. "The CS-1 is the first ever system to demonstrate sufficient performance to simulate over a million fluid cells faster than real-time. This means that when the CS-1 is used to simulate a power plant based on data about its present operating conditions, it can tell you what is going to happen in the future faster than the laws of physics produce that same result."

"We can solve this problem in an amount of time that no number of GPUs or CPUs can achieve," said Cerebras's CEO, Andrew Feldman. "This means the CS-1 for this work is the fastest machine ever built, and it's faster than any combination of clustering of other processors."

In a paper distributed at the conference Tuesday, titled "Fast Stencil-Code Computation on a Wafer-Scale Processor," Cerebras and Department of Energy researchers explained that shared memory is the critical difference between the CS-1 and Joule machines.

"It is interesting to try to understand why this striking difference arises," the reports states. The Intel Xeon caches used in the Joules computer "seem to be less effective at deriving performance from the available SRAM," even though Joules has significantly more memory. One reason is that the Xeon cores run at only approximately 40 percent of the peak rate of the Cerebras cores.

The Cerebras cores, in contrast, have no need to compete for shared RAM memory, maximizing efficient usage of the 48KB housed in each [core](#).

More information: Fast Stencil-Code Computation on a Wafer-Scale Processor, arXiv:2010.03660 [cs.DC]
arxiv.org/abs/2010.03660

© 2020 Science X Network

APA citation: Trillion-transistor chip breaks speed record (2020, November 26) retrieved 25 September 2021 from <https://techxplore.com/news/2020-11-trillion-transistor-chip.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.