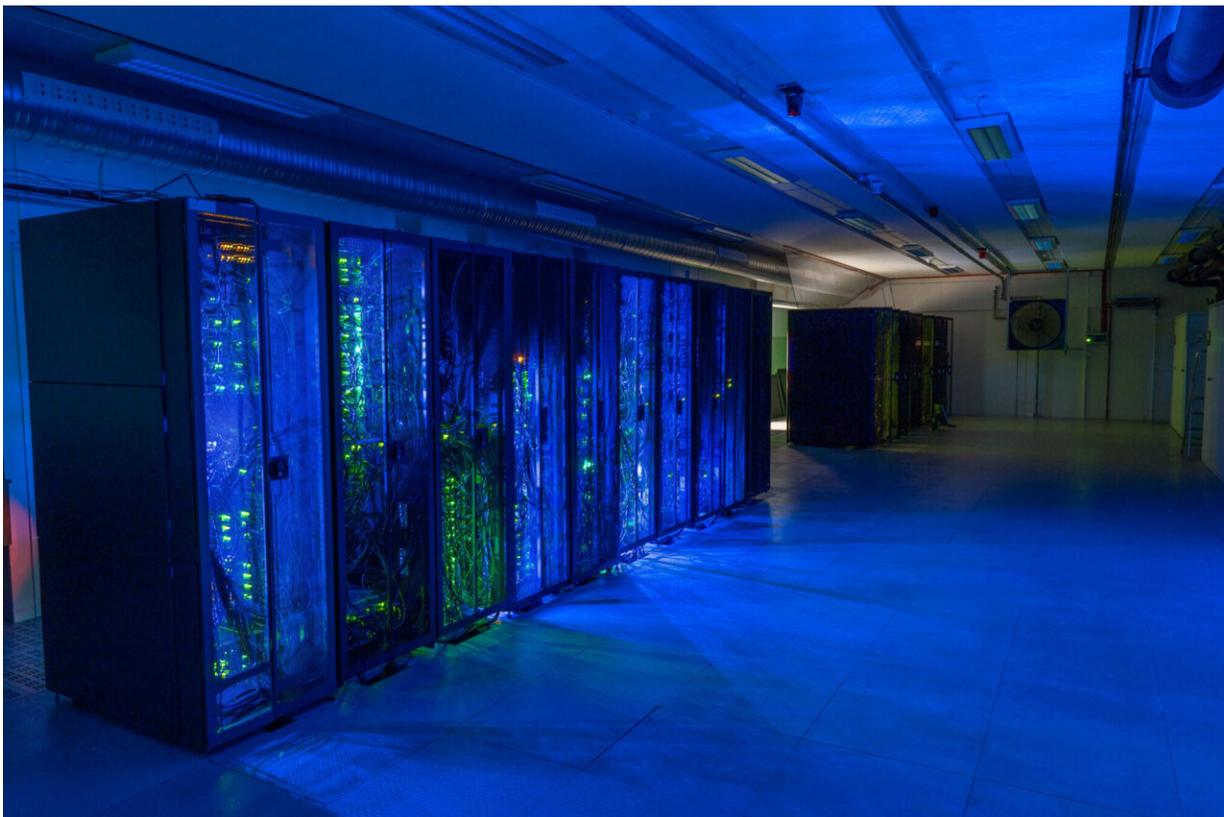


New, super-smart hardware could provide drastic cuts in power consumption

December 19 2023



The high-performance computing facility at the Norwegian University of Science and Technology (NTNU) has helped researchers to run realistic simulations using enormously large data sets. They used the simulations to develop hardware to cut computer program runtimes. Credit: NTNU

Greener data processing requires systems that work smarter, faster, and

are more energy efficient. Researchers from the Norwegian University of Science and Technology (NTNU) have developed a tiny piece of super-smart hardware that enables all of the above.

Anyone who knows anything about running [computer programs](#) knows that they take time, and everyone wants the program they are running to work as quickly as possible.

A new, tiny piece of hardware designed by [computer scientists](#) can be integrated into processors, which are the computer's brains, making it easier for developers to write programs that run faster and more efficiently.

Better performance using less energy

The hardware "tells a developer what the program they have written spends time on and why," says Professor Magnus Jahre from NTNU's Department of Computer Science. "When developers have this information, they can change their programs to do the same job more efficiently."

In conjunction with doctoral research fellow Björn Gottschall and Professor Lieven Eeckhout from Ghent University, Jahre has developed two variants of the hardware, [TIP](#) and [TEA](#). NTNU's Technology Transfer AS has made patent applications for both.

More sustainable data processing

Time is a critical factor in enabling more sustainable data processing. One way to reduce [power consumption](#) is to use computers for a shorter period of time. This requires programs that use the hardware as efficiently as possible.

"And that is exactly what TEA and TIP contribute to," Jahre says.

Once the time thieves have been identified and removed from the software, other major problems can also be solved. For example, more efficient computer equipment can help reduce the need for more electricity. The constant need to create new servers is then reduced. As a result, society will require far fewer power-hungry data centers than we might think we do.

18 data centers and 1 million servers

Currently, there are 18 data centers in Norway. Together, these house approximately 1 million computers, or servers. Magnus Jahre has calculated that TIP and TEA make it possible to cut CO₂ emissions by approximately 63,000 tons simply by reducing program execution time. The calculation assumes that all the processors operate using the new hardware and that the programs are as efficient as those used by NTNU researchers.

This pales, of course, in comparison to data centers in the United States. The US Department of Energy says that "Data centers are one of the most energy-intensive building types, consuming 10 to 50 times the energy per floor space of a typical commercial office building. Collectively, these spaces account for approximately 2% of the total U.S. electricity use, and as our country's use of information technology grows, data center and server energy use is expected to grow too."

85,000 fewer computers

Since Norwegian data centers use renewable energy, the CO₂ emissions in Jahre's calculations will mostly come from the production of the servers used in the data center. If everything works as it should, faster

runtimes will allow data centers to do the same job with about 85,000 fewer servers. This is equivalent to approximately half of [the emissions produced by Stjørdal Municipality in 2021](#) (in Norwegian).

This improvement will also free up around 400-gigawatt hours (GWh) of energy. Based on [Statistics Norway's statistics](#) (in Norwegian), that's the same used by approximately 23,000 households over the course of a year.

Good for weather forecasts and medication

Better performance means more can be done at the same time. Weather forecasts can become more precise because they can be run at higher resolution, and medications can become much more effective when computers use less time to make many more analyses of how they work.

The researchers say that Chat GPT and other language models can provide us with faster, better answers, even if they might have to conduct much larger searches than they do today.

Extreme difference in power consumption

Currently, the performance of a program, meaning how much time and power it consumes, can vary a lot. One design may have a performance that is tens of thousands of times better than another version that does exactly the same job.

The fact that runtime can vary from a few seconds to many minutes shows the need for smarter, more sustainable solutions.

Moore's Law soon out of date

It is almost 60 years since Gordon E. Moore, [who was also one of Intel's founders](#), came up with the idea that the circuit performance of one computer chip would double every 24 months.

The discussion is whether this prediction, also known as Moore's Law, is now obsolete. In any case, Jahre believes that the trend is starting to run out of steam.

"Since the processors do not automatically become faster to the same extent as before, it is even more important that software uses hardware more efficiently," he said.

A good idea in practice—but extremely challenging to achieve.



Time and power consumption in computer programs can vary extremely. One design can have performance several tens of thousands of times higher than another that does exactly the same job. Two new pieces of hardware can cut the time it takes to run programs and thus cut power consumption. Credit: NTNU

Chips are everywhere

All major chip manufacturers such as Intel, IBM, ARM and AMD make powerful, high-performance processors. The market is formidable.

Chips are used in cloud infrastructure, in all data centers, as well as in laptops, desktop computers, and smartphones. There are several of them in your TV, and your car is full of them.

Simultaneously executing 500 different orders

The most powerful processors work with up to 500 instructions at once, doing a tiny bit of work on each one at the same time. Software, by contrast, is based on the instructions being executed one at a time, much like following a recipe step by step.

"The key problem is how to look at all of this at the same time and find out exactly which of the 500 instructions the program is spending time on at any given point. This is exactly what TIP is able to do," says the computer professor.

Wasting time and giving the wrong answer

Manufacturers have equipped their chips so that they can analyze what the programs spend time on. Everyone has their own solutions. The NTNU researchers have tested four of them. They have shown that all of them make systematic errors, which means their answers are wrong.

"In fact, they are so inaccurate that it becomes difficult for the developer to understand why the program performs the way it does. It is very challenging to tease apart what it spends time on. As a result, they are unable to understand what they can do about it. By replacing these units with TIP and TEA, the developers' job becomes much easier," says Magnus Jahre.

High-performance computing and a super-simulator

The IDUN high-performance computing facility at NTNU has made it possible for researchers to develop TIP and TEA. They have also used the FireSim simulator developed at Berkeley.

Researchers who are testing new hardware solutions often use software that simulates how the hardware behaves. They can then look at a few hundred million instructions at a time. The setup at NTNU makes it possible to look at hundreds of billions of instructions. There aren't many places in the world that can do that.

Customers benefit the most

NTNU owns the rights to TIP and TEA, and NTNU Technology Transfer (TTO) has applied for a patent and is looking for commercial partners.

Strictly speaking, the invention does not solve any problems for chip manufacturers. Magnus Jahre says that it is their customers who will primarily benefit from the invention.

As long as Intel, Arm, and other manufacturers manage to keep up with the competition, they will probably think that the processors they have are efficient enough and will be unwilling to spend money, time, and resources on something completely new.

The burden of proof lies with the researchers

Even if a manufacturer decides to use the NTNU invention today, it will still take several years before TEA and TIP are found in anything you buy in the store.

"A high-performance processor is an extremely complicated product. Thousands upon thousands of person-years have been spent developing them. When we scientists come creeping out of the undergrowth and tell the manufacturers that they have to do things completely differently, the burden of proof is on us," he said.

We need to explain to them that it works and that they should use it. It is quite a challenging exercise," says Jahre.

Amazon, Google, and Apple

Business developer Lodve Berre from TTO says that they are targeting organizations that make their own hardware as well as software, such as Amazon, Google, and Apple.

"These organizations will be able to exploit the value of TIP and TEA directly in their dedicated cloud solutions and server parks. It will give them an advantage over other actors that operate using traditional hardware," he says.

Better performance and longer lifespan

The NTNU technology will mean that companies such as Apple can identify errors and bottlenecks more easily—in operating systems and other software.

"This will enable them to increase both the performance and battery life of their devices such as MacBooks, mobile phones, and tablets," says Lodve Berre.

And where will TIP and TEA be in five years?

"I think they will be in use and be on their way to a store near you," says Jahre.

Provided by Norwegian University of Science and Technology

Citation: New, super-smart hardware could provide drastic cuts in power consumption (2023, December 19) retrieved 9 May 2024 from <https://techxplore.com/news/2023-12-super-smart-hardware-drastic-power-consumption.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.