Engineers build a lower-energy chip that can prevent hackers from extracting hidden information from a smart device

February 18 2022, by Adam Zewe

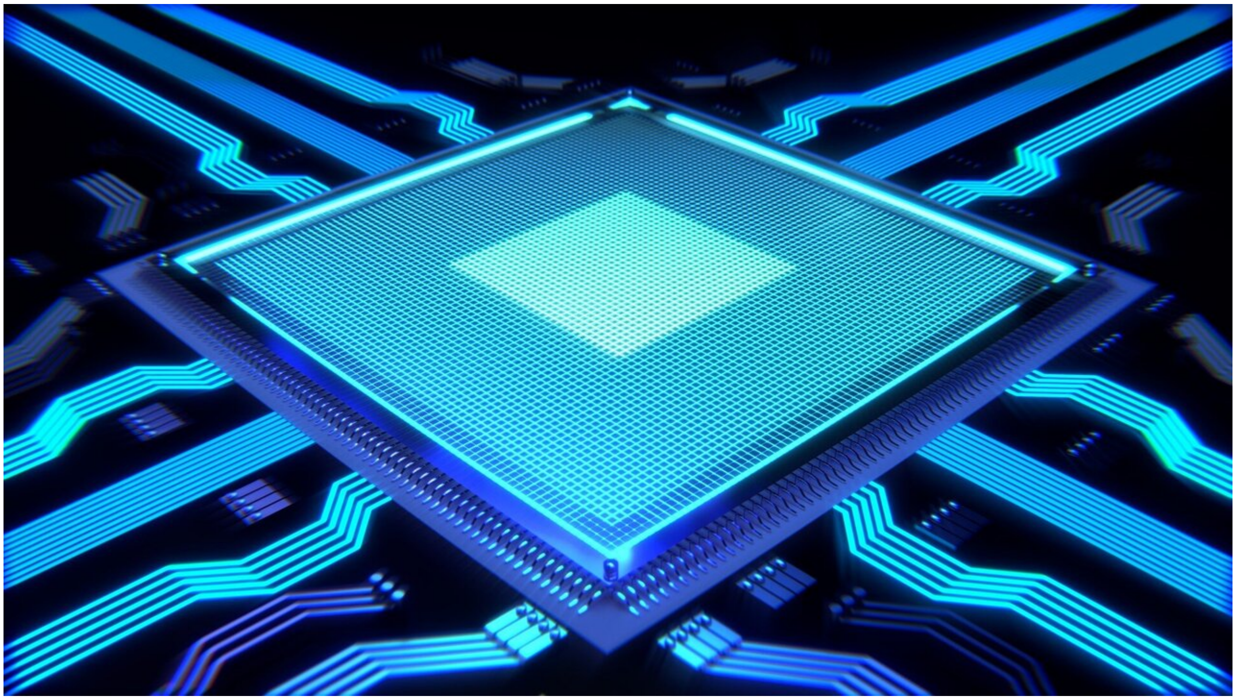

Credit: CC0 Public Domain

A heart attack patient, recently discharged from the hospital, is using a smartwatch to help monitor his electrocardiogram signals. The smartwatch may seem secure, but the neural network processing that health information is using private data that could still be stolen by a malicious agent through a side-channel attack.

A side-channel attack seeks to gather secret information by indirectly exploiting a system or its hardware. In one type of side-channel attack, a savvy hacker could monitor fluctuations in the device's power consumption while the neural network is operating to extract protected information that "leaks" out of the device.

"In the movies, when people want to open locked safes, they listen to the clicks of the lock as they turn it. That reveals that probably turning the lock in this direction will help them proceed further. That is what a side-channel attack is. It is just exploiting unintended information and using it to predict what is going on inside the device," says Saurav Maji, a graduate student in MIT's Department of Electrical Engineering and Computer Science (EECS) and lead author of a paper that tackles this issue.

Current methods that can prevent some side-channel attacks are notoriously power-intensive, so they often aren't feasible for internet-of-things (IoT) devices like smartwatches, which rely on lower-power computation.

Now, Maji and his collaborators have built an integrated circuit chip that can defend against power side-channel attacks while using much less energy than a common security technique. The chip, smaller than a thumbnail, could be incorporated into a smartwatch, smartphone, or tablet to perform secure machine learning computations on sensor values.

"The goal of this project is to build an integrated circuit that does machine learning on the edge, so that it is still low-power but can protect against these side channel attacks so we don't lose the privacy of these models," says Anantha Chandrakasan, the dean of the MIT School of

Engineering, Vannevar Bush Professor of Electrical Engineering and Computer Science, and senior author of the paper. "People have not paid much attention to security of these machine-learning algorithms, and this proposed hardware is effectively addressing this space."

Co-authors include Utsav Banerjee, a former EECS graduate student who is now an assistant professor in the Department of Electronic Systems Engineering at the Indian Institute of Science, and Samuel Fuller, an MIT visiting scientist and distinguished research scientist at Analog Devices. The research is being presented at the International Solid-States Circuit Conference.

Computing at random

The chip the team developed is based on a special type of computation known as threshold computing. Rather than having a neural network operate on actual data, the data are first split into unique, random components. The network operates on those random components individually, in a random order, before accumulating the final result.

Using this method, the information leakage from the device is random every time, so it does not reveal any actual side-channel information, Maji says. But this approach is more computationally expensive since the neural network now must run more operations, and it also requires more memory to store the jumbled information.

So, the researchers optimized the process by using a function that reduces the amount of multiplication the neural network needs to process data, which slashes the required computing power. They also protect the neutral network itself by encrypting the model's parameters. By grouping the parameters in chunks before encrypting them, they provide more security while reducing the amount of memory needed on the chip.

"By using this special function, we can perform this operation while skipping some steps with lesser impacts, which allows us to reduce the overhead. We can reduce the cost, but it comes with other costs in terms of neural network accuracy. So, we have to make a judicious choice of the algorithm and architectures that we choose," Maji says.

Existing secure computation methods like homomorphic encryption offer strong security guarantees, but they incur huge overheads in area and power, which limits their use in many applications. The researchers' proposed method, which aims to provide the same type of security, was able to achieve three orders of magnitude lower energy use. By streamlining the chip architecture, the researchers were also able to use less space on a silicon chip than similar security hardware, an important factor when implementing a chip on personal-sized devices.

"Security matters"

While providing significant security against power side-channel attacks, the researchers' chip requires 5.5 times more power and 1.6 times more silicon area than a baseline insecure implementation.

"We're at the point where security matters. We have to be willing to trade off some amount of energy consumption to make a more secure computation. This is not a free lunch. Future research could focus on how to reduce the amount of overhead in order to make this computation more secure," Chandrakasan says.

They compared their chip to a default implementation which had no security hardware. In the default implementation, they were able to recover hidden information after collecting about 1,000 power waveforms (representations of power usage over time) from the device. With the new hardware, even after collecting 2 million waveforms, they still could not recover the data.

They also tested their chip with biomedical signal data to ensure it would work in a real-world implementation. The chip is flexible and can be programmed to any signal a user wants to analyze, Maji explains.

In the future, the researchers hope to apply their approach to electromagnetic side-channel attacks. These attacks are harder to defend, since a hacker does not need the physical device to collect hidden information.

More information: "A Threshold Implementation-based Neural-Network Accelerator Securing Model Parameters and Inputs Against Power Side-Channel Attacks" International Solid-States Circuit Conference (2022).

Provided by Massachusetts Institute of Technology

Citation: Engineers build a lower-energy chip that can prevent hackers from extracting hidden information from a smart device (2022, February 18) retrieved 25 April 2024 from https://techxplore.com/news/2022-02-lower-energy-chip-hackers-hidden-smart.html