


AI models microprocessor performance in real-time

December 10 2021, by Ken Kingery



Credit: Duke University

Computer engineers at Duke University have developed a new AI method for accurately predicting the power consumption of any type of computer processor more than a trillion times per second while barely using any computational power itself. Dubbed APOLLO, the technique has been validated on real-world, high-performance microprocessors and could help improve the efficiency and inform the development of new

microprocessors.

The approach is detailed in a paper published at *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, one of the top-tier conferences in computer architecture, where it was selected the conference's best publication.

"This is an intensively studied problem that has traditionally relied on extra circuitry to address," said Zhiyao Xie, first author of the paper and a Ph.D. candidate in the laboratory of Yiran Chen, professor of electrical and computer engineering at Duke. "But our approach runs directly on the microprocessor in the background, which opens many new opportunities. I think that's why people are excited about it."

In modern computer processors, cycles of computations are made on the order of 3 trillion times per second. Keeping track of the power consumed by such intensely fast transitions is important to maintain the entire chip's performance and efficiency. If a processor draws too much power, it can overheat and cause damage. Sudden swings in power demand can cause internal electromagnetic complications that can slow the entire processor down.

By implementing software that can predict and stop these undesirable extremes from happening, [computer engineers](#) can protect their hardware and increase its performance. But such schemes come at a cost. Keeping pace with modern microprocessors typically requires precious extra hardware and [computational power](#).

"APOLLO approaches an ideal power estimation algorithm that is both accurate and fast and can easily be built into a processing core at a low power cost," Xie said. "And because it can be used in any type of processing unit, it could become a common component in future chip design."

The secret to APOLLO's power comes from artificial intelligence. The algorithm developed by Xie and Chen uses AI to identify and select just 100 of a processor's millions of signals that correlate most closely with its power consumption. It then builds a power consumption model off of those 100 signals and monitors them to predict the entire chip's performance in real-time.

Because this learning process is autonomous and data driven, it can be implemented on most any computer processor architecture—even those that have yet to be invented. And while it doesn't require any human designer expertise to do its job, the algorithm could help human designers do theirs.

"After the AI selects its 100 signals, you can look at the algorithm and see what they are," Xie said. "A lot of the selections make intuitive sense, but even if they don't, they can provide feedback to designers by informing them which processes are most strongly correlated with power consumption and performance."

The work is part of a collaboration with Arm Research, a computer engineering research organization that aims to analyze the disruptions impacting industry and create advanced solutions, many years ahead of deployment. With the help of Arm Research, APOLLO has already been validated on some of today's highest performing processors. But according to the researchers, the algorithm still needs testing and comprehensive evaluations on many more platforms before it would be adopted by commercial [computer](#) manufacturers.

"Arm Research works with and receives funding from some of the biggest names in the industry, like Intel and IBM, and predicting power consumption is one of their major priorities," Chen added. "Projects like this offer our students an opportunity to work with these industry leaders, and these are the types of results that make them want to

continue working with and hiring Duke graduates."

More information: Zhiyao Xie et al, APOLLO: An Automated Power Modeling Framework for Runtime Power Introspection in High-Volume Commercial Microprocessors, *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture* (2021). [DOI: 10.1145/3466752.3480064](https://doi.org/10.1145/3466752.3480064)

Provided by Duke University

Citation: AI models microprocessor performance in real-time (2021, December 10) retrieved 26 April 2024 from <https://techxplore.com/news/2021-12-ai-microprocessor-real-time.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.