FROM THE EDITOR-IN-CHIEF

United Semiconductors of America

Hsien-Hsin S. Lee, *Intel Corporation, Boxborough, MA, 01719, USA*

Over the past several months, investments, partnerships, and strategic intervention in the semiconductor industry—particularly within the United States—have dominated headlines across various mainstream media. Driven by the lightning speed of artificial intelligence (AI) development, hardware companies keep throwing in staggering capital expenditures to help build and scale AI infrastructure at an unprecedented pace. Cloud service providers (CSPs), along with a rising class of new players referred to as *neoclouds*, are procuring hundreds of thousands, if not millions, of GPUs and AI accelerators, to meet the forecast surge in demand for AI computing power anticipated over the next five to 10 years. Meanwhile, software powerhouses such as OpenAI, Microsoft, Google, Anthropic, Meta, and so on are already locked in fierce competition on the AI battleground. Against this backdrop, OpenAI, Nvidia, and Oracle forged a strategic alliance. Their shared goal is to enable OpenAI with access to 10 GW of AI data center capacity built on Nvidia's systems. Through the Stargate Initiative, OpenAI committed to contracting $300 billion worth of Oracle's computing power over five years, while Nvidia has pledged to invest up to $100 billion in OpenAI. Complementing this, Oracle will spend $40 billion to purchase GB200 GPUs from Nvidia to power OpenAI's new data centers. The partnership is not limited to Nvidia. Most recently, OpenAI struck a deal with AMD, taking up to 160 million shares of AMD common stock, roughly a 10% stake in the company, in exchange for 6 GW of GPUs. These dramatic activities will surely propel AI-enabled economies.

Another notable trend has also emerged: both CSPs and software companies are increasingly moving into hardware and system design to become more independent and self-sustained. Unsatiated by the lower-than-expected utilization of peak floating-point operations per second (flops), they are recklessly investing in efficiency optimization for AI workloads. By maximizing model flops utilization, these firms not only lower the operational expenditures but also mitigate their strong reliance on Nvidia's solutions.

This landscape shift is underscored by major developments. *The Financial Times* recently reported on a $10 billion hardware supply deal between OpenAI and Broadcom, indicating OpenAI's ambition to build their own AI data center fleet besides their recent alliance with Nvidia, Oracle, and AMD. More recently, Meta was reported to acquire Rivos, a start-up commencing in server-class RISC-V cores that has since pivoted their efforts to AI accelerator products. Previously, Meta has already demonstrated their in-house Meta Training and Inference Accelerator (MTIA), exhibiting their ambition to establish their own hardware and system capabilities for their service applications.

There have been even more headline-grabbing developments on the Intel front. In August, SoftBank Group announced a $2 billion investment in Intel's common stock, followed shortly by a historic agreement between Intel and the Trump administration. Under this new deal (different from what the CHIPS Act promised under the previous administration), the U.S. government invested $8.9 billion in Intel's common stock, giving the U.S. government almost 10% ownership in Intel's market capital. The investment spree did not end there. Just two weeks later, Nvidia and Intel announced a program to jointly develop new products that more tightly integrate x86 CPU cores with Nvidia GPU—targeting both AI data center GPUs as well as client PC products. As part of the collaboration, Nvidia also agreed to invest $5 billion in Intel's common stock. This joint development of data center AI products stands to benefit both companies and is closely in line with what Nvidia's NVLink Fusion program announced in May and what was discussed in my previous article.[1] Adding further momentum, unconfirmed reports have circulated about AMD's potential adoption of Intel Foundry services, and a possible investment from Apple in Intel. Coupled with the Trump administration's proposal to Taiwan Semiconductor Manufacturing Company for a 50–50 production split between Taiwan and the United States, there has never been a more "united" push in the United States to bolster the

Digital Object Identifier 10.1109/MM.2025.3620859
Date of current version 31 October 2025.

APPENDIX: RELATED ARTICLES

A1. C. Wilkerson, "Special Issue on Contemporary Industry Products 2025," *IEEE Micro*, vol. 45, no. 5, pp. 6–8, Sep./Oct. 2025, doi: 10.1109/MM.2025.3617427.

A2. J. Yi, "A review of *Wisconsin Alumni Research Foundation v. Apple*—Part VI," *IEEE Micro*, vol. 45, no. 5, pp. 134–137, Sep./Oct. 2025, doi: 10.1109/MM.2025.3619665.

A3. S. Greenstein, "Wireless commercial breakthrough," *IEEE Micro*, vol. 45, no. 5, pp. 138–140, Sep./Oct. 2025, doi: 10.1109/MM.2025.3603656.

A4. G. Singer, "Exploring trends to ride artificial intelligence's explosive growth trajectory," *IEEE Micro*, vol. 45, no. 5, pp. 127–132, Sep./Oct. 2025, doi: 10.1109/MM.2025.3617424.

semiconductor industry to accelerate infrastructure construction and service deployment.

This special issue presents the latest work on contemporary industry products. Launched in 2020, this special issue is published annually in collaboration with the Industry Track of the International Symposium on Computer Architecture (ISCA). The goal is to highlight works that evaluate real industry products or products under development for our audience. I would like to extend my heartfelt gratitude to the guest editor, Chris Wilkerson, of Intel—who also served as the technical program chair of the ISCA 2025 Industry Track—and his outstanding program committee members for their service, dedication, and hard work. The selection process followed the standard conference review model: the program committee members reviewed submissions, authors provided their rebuttals, online discussions were conducted prior to the program committee meeting, and final decisions were jointly reached at the program committee meeting. This year's outcomes include several contemporary AI/data center accelerator papers from emerging startups, along with contributions from leading companies such as Google, Huawei, Intel, Meta, and Samsung. For more details, please read the Guest Editor's Introduction,[A1] which provides an overview of the review and selection process, as well as brief introductions of these works.

In Part VI of the *Wisconsin Alumni Research Foundation (WARF) v. Apple* series in the Micro Law column,[A2] Dr. Joshua Yi delves into the intricacies of the motion to compel Apple to produce an executable version of its A6, A7, and A8 processor simulators described in Part V. He discusses the implications of the judge's decision, the concerns raised by Apple's counsel, and WARF's counterarguments. In the Micro Economics column,[A3] Prof. Shane Greenstein explores how Wi-Fi became a commercial breakthrough through a mix of technological innovation, strategic corporate actions, and favorable public policy. He traces its evolution from narrow early applications to today's pervasive role in consumer devices, business products, and online services such as social network, video sharing, and streaming. This issue also inaugurates a new Micro AI column,[A4] led by our new associate editor, Gadi Singer, to provide perspectives of the past, present, and future of AI technology and their impact on computing system designs. Drawing on more than 40 years of leadership experience—from the Intel Pentium processor to directing research at Intel's Emerging AI Research Labs—Singer offers invaluable insights to our readers. In his inaugural article, Singer reflects on AI's explosive growth trajectory and how current trends are shaping the future.

I hope that you enjoy this special issue and the articles.

LAUNCHED IN 2020, THIS SPECIAL ISSUE IS PUBLISHED ANNUALLY IN COLLABORATION WITH THE INDUSTRY TRACK OF THE INTERNATIONAL SYMPOSIUM ON COMPUTER ARCHITECTURE.

REFERENCE

1. H.-H. S. Lee, "Toward disaggregated and heterogenous AI systems," *IEEE Micro*, vol. 45, no. 3, pp. 4–5, May/Jun. 2025, doi: 10.1109/MM.2025.3575180.

HSIEN-HSIN S. LEE is an Intel Fellow at Intel Corporation, Boxborough, MA, 01719, USA. Contact him at lee.sean@gmail.com.