FROM THE EDITOR-IN-CHIEF

Taiwan Semiconductor Manufacturing Company's $165 Billion Bet

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

The semiconductor industry is undergoing a profound transformation. On 3 March, history was made in the White House Roosevelt Room with an unprecedented foreign investment announcement. Introduced as a legend by President Donald Trump, Dr. C.C. Wei, the CEO of Taiwan Semiconductor Manufacturing Company (TSMC), stood alongside Commerce Secretary Howard Lutnick and crypto czar David Sacks to unveil a staggering $100 billion investment from TSMC in Phoenix, AZ, over the next four years. This commitment includes the construction of three cutting-edge logic fabrication plants (fabs) and two advanced packaging facilities. More significantly, TSMC pledged to establish a 1,000-person R&D center in Phoenix to continuously optimize operational efficiencies and enhance production quality in these fabs. With this new pledge, TSMC's total investment in Arizona, including its previously announced advanced logic fabs in the same location, now amounts to an unprecedented $165 billion—highlighting the company's effort to bring its world-leading semiconductor manufacturing capabilities to American soil. To further drive down production costs, TSMC is expected to push its ecosystem partners—material suppliers, downstream testing firms, and packaging vendors—to relocate to Arizona, creating a robust and self-sustaining semiconductor hub in Arizona.

Many view this move as TSMC's strategic response to the looming threat of a 100% tariff on Taiwanese semiconductor imports, a proposal floated by the Trump administration. Meanwhile, in a somewhat unexpected yet potentially game-changing development, TSMC was reported in discussions with leading chip design firms, including AMD, Broadcom, and Nvidia, to form a joint venture investing in Intel foundry. According to an exclusive Reuters report, this joint venture aims to get Intel's foundry business off the ground to manufacture chips for external customers beyond products designed by Intel. At first glance, such an alliance may seem implausible. However, it presents a compelling win-win scenario, in my opinion, for all participating parties. If such speculation turns out to be true, TSMC can offload excess demand to its partner (Intel) while profiting from this investment: Intel foundry will build up its credibility as a legitimate foundry by attracting major design companies through TSMC's network; chip design companies will have more options in choosing manufacturing partners. Meanwhile, the U.S. government will achieve its ambitious objective of securing domestic production of the world's most advanced semiconductors, thereby minimizing reliance on Taiwan.

Critics argue that manufacturing in the United States will drive up costs significantly—some have even claimed by a factor of four to five—compared to manufacturing in Taiwan. However, I always take such assertions with a grain of salt. The capital expenditure (CapEx) for a state-of-the-art 3nm fab today already reaches $20 billion, with a substantial allocation to the procurement of equipment and tools, which cost roughly the same worldwide regardless of the geographical locations of the fabs. Even if TSMC were to pay each of its U.S.-based engineers a hefty, additional $500,000 over its Taiwanese counterparts, the total labor premium for 1,000 engineers would amount to just an extra $500 million—a mere fraction of the overall $20 billion CapEx. A recent cost analysis on TSMC's 4nm Fab 21 in Phoenix published by *TechInsights*, a leading semiconductor information platform, supports my skepticism by showing that the wafer production costs in Arizona are only 10% higher than in Taiwan, challenging the narrative that U.S.-based semiconductor operations are prohibitively expensive. If the conclusion from this analysis is closer to reality, relocating semiconductor manufacturing operations to the United States will be more economically feasible than previously thought.

Digital Object Identifier 10.1109/MM.2025.3557486
Date of current version 29 April 2025.

APPENDIX: RELATED ARTICLES

A1. W. Schonbein and J. Schuchart, "Special issue on Hot Interconnects 31 [Guest Editors' Introduction]," *IEEE Micro*, vol. 45, no. 2, pp. 6–7, Mar./Apr. 2025, doi: 10.1109/MM.2025.3555461.

A2. R. Egawa and Y. Wada, "Special issue on COOL Chips [Guest Editors' Introduction]," *IEEE Micro*, vol. 45, no. 2, pp. 65–66, Mar./Apr. 2025, doi: 10.1109/MM.2025.3555464.

A3. J. J. Yi, "A review of *Wisconsin Alumni Research Foundation v. Apple*—Part III [Micro Law]," *IEEE Micro*, vol. 45, no. 2, pp. 114–117, Mar./Apr. 2025, doi: 10.1109/MM.2025.3557487.

A4. S. Greenstein "Artificial intelligence and the Jevons paradox [Micro Economics]," *IEEE Micro*, vol. 45, no. 2, pp. 118–120, Mar./Apr. 2025, doi: 10.1109/MM.2025.3548921.

A5. D. Burger et al., "Derek Chiou [In Memoriam]," *IEEE Micro*, vol. 45, no. 2, pp. 122–124, Mar./Apr. 2025, doi: 10.1109/MM.2025.3544666.

We are witnessing one of the most dynamic and transformative periods in semiconductor history. This seismic shift is not merely about technological advancements or talent acquisition, it is deeply rooted in national security and geopolitical strategy. With artificial intelligence (AI) computing at the forefront of innovation and semiconductor policy dominating discussions in business, government, and state-to-state rivalry, chips have become more than just a piece of silicon—they are now a matter of economic and technological sovereignty and their discussion can be seen everywhere.

This issue of *IEEE Micro* highlights two special topics: *Hot Interconnects* and *COOL Chips*, both critical subjects for designing high-efficiency computing and communication systems, especially in the era of data-intensive, accelerator-centric computing for AI and machine learning. For the *Hot Interconnects* theme, Dr. Whit Schonbein from Sandia National Laboratories and Prof. Joseph Schuchart from Stony Brook University did an exceptional job selecting six outstanding works from the *IEEE Hot Interconnect Symposium 2024*. Five of these contributions, from both industry and academia, explore various aspects and their implications of designing interconnect and collective communication for transformer-based large language models and high-performance computing. Additionally, one article discusses a congestion control mechanism aimed at improving system performance. Please read the Guest Co-Editors' Introduction[A1] for a preview of these insightful works. For the *COOL Chips* theme, I would like to extend my gratitude to Guest Co-Editors Prof. Ryusuke Egawa from Tokyo Denki University and Prof. Yasutaka Wada from Meisei University for selecting three top works from the *IEEE Symposium on Low-Power and High-Speed Chips and Systems 2024*, also known as *COOL Chips*. These articles cover a range of innovative topics, including an AI-based solution for processing lidar data, a reconfigurable cryptographic accelerator, and an approach to approximate computing using GPUs. Please read the Guest Co-Editors' Introduction for a short description of these articles.[A2]

Beyond these two themes, this issue also features a technical article that presents a Compute Express Link (CXL) 2.0 memory expansion solution from SK Hynix, along with its empirical study on Meta's Caching Services (CacheLib) to optimize key-value cache performance. In the Micro Law column,[A3] Dr. Joshua Yi brings readers Part III of the *Wisconsin Alumni Research Foundation (WARF) vs. Apple* case, which centers on a computer architecture patent dispute. In his article, he provides an overview of the plaintiff (WARF, the patent's inventors) allegations of infringement, and the legal rulings from both the district court and appellate panel. Meanwhile, in the Micro Economics column,[A4] Prof. Shane Greenstein responds to Microsoft CEO Satya Nadella's recent social media post by drawing a connection between the rising demand for AI and Jevons' paradox, a classic economic observation that increased efficiency often leads to higher resource consumption. As we conclude this issue, we honor the memory of Prof. Derek Chiou, a beloved and influential figure in the computer architecture community. In a heartfelt tribute,[A5] six esteemed computer architects share their recounts, celebrating Derek's remarkable contributions and mourning his sudden passing last December.

We hope that you enjoy the articles we have curated for you in this issue.

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