



**DEPARTMENT OF COMPUTER
SCIENCE AND AUTOMATION**
Indian Institute of Science Bangalore



Accelerator-level Parallelism

Mark D. Hill

Gene M. Amdahl and John P. Morgridge Professor Emeritus
Computer Sciences Department
University of Wisconsin-Madison

Computer system performance has improved due to creatively using more transistors (Moore's Law) in parallel via bit-, instruction-, thread-, and data-level parallelism. With the slowing of technology scaling, a way to further improve computer system performance under energy constraints is to employ hardware accelerators. Each accelerator is a hardware component that executes a targeted computation class faster and usually with (much) less energy. Already today, many chips in mobile, edge and cloud computing concurrently employ multiple accelerators in what we call accelerator-level parallelism (ALP).

This talk develops our hypothesis that ALP will spread to computer systems more broadly. ALP is a promising way to dramatically improve power-performance to enable broad, future use of deep AI, virtual reality, self-driving cars, etc. To this end, we review past parallelism levels and the ALP already present in mobile systems on a chip (SoCs). We then aid understanding of ALP with the Gables model and charge computer science researchers to develop better ALP "best practices" for: targeting accelerators, managing accelerator concurrency, choreographing inter-accelerator communication, and productively programming accelerators. This joint work with Vijay Janapa Reddi of Harvard. See also: <https://www.sigarch.org/accelerator-level-parallelism/>



Mark D. Hill (<http://www.cs.wisc.edu/~markhill>) is Gene M. Amdahl and John P. Morgridge Professor Emeritus in Computer Sciences at the University of Wisconsin-Madison. He has long also held a courtesy appointment in Electrical and Computer Engineering. His work targets computers with complex memory systems, multiple processing cores, and systems which do not yet exist so they have to be simulated. Over three decades, he has collaborated with over 160 co-authors, has over 40 patents, and has held several visiting positions in the computer industry, most recently as a 2018 Google Intern. He serves as Chair Emeritus of a national computing think tank--the Computing Community Consortium--and he was Wisconsin Computer Sciences Department Chair 2014-2017. Mark won the highest award in computer hardware--Eckert-Mauchly--in 2019. Mark is a fellow of ACM and IEEE, and has a PhD from the University of California, Berkeley.

About CSA

The Department of Computer Science and Automation was created in 1969. The vision of the department is to advance the frontiers of research in computer science and offer world-class research and pedagogical experience to its students. CSA faculty and students contribute to cutting edge research in topical areas of computer science, and are actively engaged in high-impact collaborative projects with academia and industry, both from within India and abroad. CSA has been ranked highest among all Indian institutions in terms of its publication profile and citations. The goal of CSA is to be among the top 25 computer science departments in the world.

CSA alumni have done exceedingly well in both academia and industry with as many as 8 CSA alumni receiving the Distinguished Alumnus Award of IISc. More than 50 new (Startup) companies have been founded by CSA and EECS alumni during the course of the past five years.

The CSA department is celebrating its Golden Jubilee year. As part of the Golden Jubilee celebrations, the Department is planning several events and initiatives. One of the initiatives is to organize the "CSA Golden Jubilee Frontier Lecture Series" to introduce the audience to a wide gamut of research themes in computer science. More details on the Golden Jubilee initiatives are available at CSA's webpage

<https://www.csa.iisc.ac.in/>

October 15, 2020 | 6:30 PM

Join by MS Teams link: <https://tinyurl.com/y2vzddcf>