13TH HEIDELBERG LAUREATE FORUM
SEPTEMBER 13–18, 2026

Master Classes

Jack J. Dongarra
An Overview of High-Performance Computing

In this talk we examine how high-performance computing has changed over the last 10 years and look toward the future in terms of trends. These changes have had and will continue to have a major impact on our numerical scientific software. A new generation of software libraries and algorithms are needed for the effective and reliable use of (wide area) dynamic, distributed, and parallel environments.

Torsten Hoefler
Accelerated Ethernet Networking for AI Systems

AI supercomputers increasingly combine cloud-scale reliability requirements with the extreme communication demands of HPC, making scalable and resilient Ethernet-based networking essential. This talk introduces Ultra Ethernet, an open standard intended to unify AI and HPC networking without proprietary fabrics. We explain why standard RoCE falls short and discuss key innovations including per-packet multipathing, flexible congestion control, zero-RTT reliability, scalable addressing, workload-optimized communication semantics, and integrated security. We also cover compatibility with existing Ethernet infrastructure and libfabric, along with emerging capabilities such as programmable congestion control and in-network computation. Finally, we present Multipath Reliable Connection (MRC) as a precursor toward the Ultra Ethernet vision. MRC demonstrates resilient multipath networking during large-scale training runs of 50,000–75,000 GPUs, sustaining frequent link flaps, NIC–top-of-rack link outages, a complete 512-port switch failure, and eight simultaneous link failures. These results illustrate a practical path toward efficient, interoperable, and failure-resilient Ethernet networks for next-generation AI and HPC systems.

Robert Melancton Metcalfe
Ethernet

Ethernet was invented on May 22, 1973. Dave Boggs RIP and I built the first Ethernet. I received the 2022 ACM A.M. Turing Award for leading the invention, standardization, and commercialization of Ethernet. One of Ethernet's lessons is that invention is a flower, but innovation is a weed. Which is another way of saying that when your invention starts to matter, the status quo moves in to kill it. Ethernet was received by the status quo as an unnecessary increase in the connectivity among networks of dumb terminals and mainframes. Wrong. Then came personal computers on the Internet, Ethernet's sweet spot. To launch Ethernet, we built some on an internet inside the Xerox Palo Alto Research Center (PARC). Then we gave talks and wrote papers. Then we decided to make Ethernet an open standard for Internet plumbing. We (Xerox, DEC, Intel, (and 3Com)) submitted Ethernet specifications to the IEEE. The status quo followed us in. After a few years, GM and IBM did not kill Ethernet. But they did succeed in getting the IEEE to standardize all three local-area networks (LANs): Ethernet, General Motor's Token Bus, and IBM's Token Ring. Try telling people that your competitors are GM and IBM. It took a decade or two for Ethernet to win the LAN Wars. Next came commercialization. My billion-dollar Silicon Valley startup, 3Com Corporation, today part of HP Enterprise, got to engage early customers of Ethernets among IBM PCs. We uncovered Ethernet's first round of killer apps. They were PFMTS -- shared Printer, shared Files, Mail, Terminal emulation, and Stubs (APIs). Then came multi-user accounting systems. Spreadsheets. And then … the unanticipated World Wide Web, the killer app of all time. The Web disrupted publishing, then search, then commerce, social, then fake everything, artificial intelligence, and reduction of poverty. Sometimes the status quo loses.

Brian Schmidt
The Creation of the Elements in the Universe

More than 13 billion years ago, the Big Bang initiated the evolution of our Universe from a simple, hot plasma into today's complex structure of stars, planets, and galaxies. A primary tracer of this cosmic evolution is the creation of chemical elements over time. This talk explores the mechanisms of cosmic nucleosynthesis across epochs, tracing matter from the first few minutes of the Universe through to nuclear fusion in stars. It will show how elemental abundances can serve as a key tracer of cosmic history.

Dennis Parnell Sullivan
Understanding 3D Fluid Motion and Anticipated Applications

3D fluid motion is puzzling in theory and practice. Here we describe a few ideas that combine in a strategy for a better understanding of the theory and perhaps in some cases for more useful predictions of the phenomena the Navier–Stokes theory models. The ideas include:
a) transport of structures embodied in the model
b) a tautologous Riemann metric on a 5D symmetric space related to a)
c) the moduli space of shapes of the standard three torus containing the fluid
d) probability theory uniting the distortion of b) by the fluid motion on the torus with the geodesic flow on the symmetric space
e) finally based on d) the logarithm law for geodesics mentioned in c) and the above indicated strategy.

Robert Endre Tarjan
Double Recursion

Double recursion is an algorithmic design paradigm that uses two recursive calls. The first one solves the problem on an appropriate subset of the input. This solution allows one to reduce the size of the original problem. The second recursion solves the reduced version of the original problem. This technique can result in a linear-time algorithm. We shall explore the use of this idea in algorithms for median-finding, computing minimum spanning trees, and (if time permits) analysis of an algorithm for disjoint set union.

Leslie G. Valiant
Educability for Humans and AIs
We seek to define the capability that has enabled humans to develop the civilization we have, and that distinguishes us from other species. For this it is not enough to identify a distinguishing characteristic – we want a capability that is also explanatory of humanity's achievements. “Intelligence” does not work here because we have no agreed definition of what intelligence is or how an intelligent entity behaves. We need a concept that is behaviorally better defined. The definition will need to be computational in the sense that the expected outcomes of exercising the capability need to be both specifiable and computationally feasible. This formulation is related to the goals of artificial intelligence research but is not synonymous with it, leaving out, for example, the many capabilities we share with other species.

We make a proposal for this essential human capability, which we call “educability.” It synthesizes abilities to learn from experience, to learn from others, and to chain together what we have learned in either mode and apply that to particular situations. It starts with the now standard notion of learning from examples as captured by the Probably Approximately Correct model and used in machine learning. The demonstrated ability of large language models learning to generate smoothly flowing prose is a clue that pursuing computationally well-defined tasks, as is done in machine learning, constitutes a promising approach to simulating human capabilities. The basic question then is how to extend this approach to encompass broader human capabilities beyond learning from examples. This is what the educability notion aims to answer.

What we ask computers to do, in the main, reflects human capabilities. Hence a better understanding of human capabilities can be expected to provide goals for our future technology.