“Ultrascale System Interconnects at the end of Moore’s Law”

John Shalf
Department Head for Computer Science, Lawrence Berkeley National Lab

The tapering of lithography advances that have been associated with Moore’s Law will substantially change requirements for future interconnect architectures for large-scale datacenters and HPC systems. Architectural specialization is creating new datacenter requirements such as emerging accelerator technologies for machine learning workloads and rack disaggregation strategies will push the limits of current interconnect technologies. Whereas photonic technologies are often sold on the basis of higher bandwidth and energy efficiency (e.g. lower picojoules per bit), these emerging workloads and technology trends will shift the emphasis to other metrics such as bandwidth density (as opposed to bandwidth alone) and reduced latency, and performance consistency. Such metrics cannot be accomplished with device improvements alone, but require a systems view of photonics in datacenters.

Taking this systems-view, the path towards realizing next-generation exascale HPC and mega datacenters is increasingly dependent on building scalable interconnects that capture the communication requirements of ultrascale applications. The talk will walk through one of the broadest studies to date of high-end application communication requirements, whose computational methods include: finite-difference, lattice-Boltzmann, particle-in-cell, sparse linear algebra, particle mesh ewald, FFT-based solvers, and numerous commercial workloads from FaceBook and Google. Using derived communication characteristics, has enabled us to build a fit-tree approach for designing network infrastructure that is can be dynamically tailored to application requirements at runtime.

Bio: John Shalf is CTO for the National Energy Research Supercomputing Center and also Department Head for Computer Science and Data Sciences at Lawrence Berkeley National Laboratory (LBNL). Shalf is a coauthor of over 60 publications in the field of parallel computing software and HPC technology, including three best papers and the widely cited report “The Landscape of Parallel Computing Research: A View from Berkeley” (with David Patterson and others). He also coauthored “ExaScale Software Study: Software Challenges in Extreme Scale Systems,” which sets the Defense Advanced Research Project Agency’s (DARPA’s) information technology research investment strategy for the next decade. He was a member of the Berkeley Lab/NERSC team that won a 2002 R&D 100 Award for the RAGE robot. (Taken from the ISC 2018 Program)