

## **Berkeley Lab's Leadership in Deep Learning, Exascale Computing and Beyond on Tap at SC17**

Lawrence Berkeley National Laboratory staff will be sharing their expertise reflecting the lab's work in pushing the envelope to exascale computing and beyond, as well as Berkeley Lab's leadership in deep learning.

In the area of deep learning, lab staff co-authored papers on "[Deep Learning at 15PF: Supervised and Semi-Supervised Classification for Scientific Data](#)" and "[Scaling Deep Learning on GPU and Knights Landing Clusters](#)." Prabhat will also give a plenary talk on Sunday, Nov. 12 at the Intel Developer Conference held in conjunction with SC17.

In the exascale arena, Kathy Yelick will be part of a presentation on recent accomplishments of DOE's Exascale Computing Project at 2:30 p.m. Tuesday, Nov. 14 in the DOE booth (613). Two simulations running in the DOE booth were created using adaptive mesh refinement software developed by AMReX, the ECP Block-Structured Adaptive Mesh Refinement Co-Design Center. Lab staff also co-organized a Nov. 14 Birds-of-a-Feather session on [Reconfigurable Computing in Exascale](#)."

Looking beyond exascale, in a DOE booth presentation at 3:15 p.m. Tuesday, Nov. 14, Jonathan Carter will discuss Berkeley Lab's Advanced Quantum-Enabled Simulation (AQuES) Testbed, highlighting recent and future developments in our sc-qubit platform, control electronics and coupling to classical computing. Examples of the project's eight-qubit rings will be displayed in the booth. George Michelogiannakis will join a Nov. 16 panel discussion on "[Post Moore Supercomputing](#),"

Lab staff will also share their expertise with the global HPC committee through other tutorials, technical papers and focused workshops in the SC17 conference technical program, as well as numerous demos in the DOE booth (613). Here are several examples of leading-edge presentations by Berkeley Lab experts:

**The Roofline Model for Performance Tuning:** Roofline Model developer Sam Williams will co-lead a Nov. 12 tutorial on "[Performance Tuning of Scientific Codes with the Roofline Model](#)" and host a demo on "Using the Roofline Model and Intel Advisor" at 2 p.m. Tuesday, Nov. 14 in the DOE booth (no. 613).

**The Cori Burst Buffer:** NERSC pioneered the deployment of an SSD "burst buffer" with the installation of Cori, using this added storage layer to accelerate I/O. NERSC's Deborah Bard will co-present the "[Getting Started with the Burst Buffer: Using DataWarp Technology](#)" tutorial on Monday, Nov. 13, then host an hour-long "Hands-on with Cori's Burst Buffer," demo at 4 p.m. Wednesday, Nov. 15, in the DOE booth (no. 613) This will be followed by a roundtable discussion, also in the DOE booth.

[View a day-by-day guide to sessions featuring Berkeley Lab staff.](#)