

SCALING UP QUANTUM FIDELITY

ARGONNE SYSTEMS USED TO SMOOTH OUT QUANTUM MOLECULAR DYNAMICS SIMULATIONS

About the Customer

Argonne National Laboratory is a U.S. Department of Energy (DOE) multidisciplinary science and engineering research center, where talented researchers work together to confront the biggest scientific questions facing humanity. Argonne's computing resources, including the powerful Polaris system and Aurora exascale computer, located at the Argonne Leadership Computing Facility (ALCF), leverage several innovations to support cutting-edge machine learning and data science workloads alongside traditional modeling and simulation capabilities. This drives progress in a variety of projects, from cosmology to drug discovery and beyond.



Improved fidelity scaling is central to ensure that meaningful scientific knowledge is extracted from large-scale simulations on leadership-scale parallel computers.

Hikaru Ibayashi et al.,
University of Southern
California, Sony Group
Corporation, and Argonne
National Laboratory



Their Challenge

Materials science and engineering is getting a boost from quantum methods, as neural-network quantum molecular dynamics (NNQMD) simulations, based on machine learning, are revolutionizing atomistic materials simulations. A state-of-the-art (SOTA) model, called “Allegro” (a musical term meaning “fast and lively”), demonstrates increased accuracy and speed — but it **struggles with fidelity scaling on massively parallel supercomputers**, a big problem in the era of exascale.

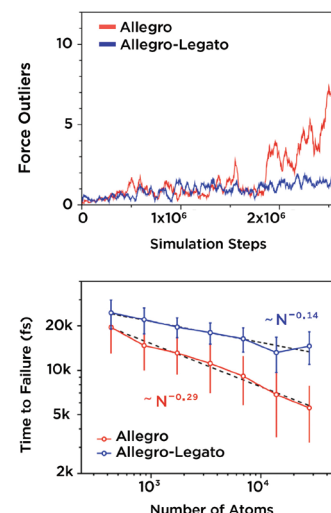
A research group of experts from the University of Southern California, Sony Group Corporation, and Argonne were determined to find a way to scale successfully.

Our Solution

To address the fidelity scaling challenge, the researchers used ALCF’s powerful computing resources, including the 34-petaflop Polaris supercomputer, to solve the problem by combining the Allegro model with sharpness-aware minimization (SAM) to increase its smoothness and robustness. Polaris is GPU-equipped and its workloads are orchestrated by Altair® PBS Professional®, a fast, powerful workload manager designed to improve productivity, optimize utilization and efficiency, and simplify resource administration for even the biggest high-performance computing workloads, including demanding materials modeling. The resulting “Allegro-Legato” model — Legato means “smooth” — increases time to failure without compromising speed or accuracy. **Allegro-Legato is much less dependent on problem size than the SOTA Allegro model, allowing larger and longer NNQMD simulations.** In addition, Allegro-Legato retains Allegro’s SOTA accuracy and computational speed while enabling large spatio-temporal NNQMD simulations on leadership-scale computers like Aurora. An example is the study of the vibrational properties of ammonia, which has a higher energy density than liquid hydrogen but can be stored at a far less energy-intensive temperature (-33 C vs. -253 C). To develop new technologies based on ammonia, we need to understand and model its complex physical and chemical interactions.

Results

The delayed time to failure of the Allegro-Legato model has dramatically improved fidelity scaling, making it easier to extract insightful scientific knowledge from large-scale simulations on leadership-scale computers. **Accurate models are vital in predicting materials’ thermodynamic behavior and how they can be used in energy, biological, and pharmaceutical systems.** The researchers’ scalable parallel implementation of Allegro-Legato with excellent computational scaling and GPU acceleration on Polaris combines accuracy, speed, robustness, and scalability, allowing practical large spatio-temporal-scale NNQMD simulations for challenging applications on exascale computing platforms like Aurora.



TOP: While the Allegro model becomes unstable and eventually fails, Allegro-Legato maintains a nearly constant number of outliers, remaining stable. **BOTTOM:** Time to failure lengthens by a factor of almost two with the smoother Allegro-Legato model.