

MOLECULAR DYNAMICS

EXAALT: Molecular Dynamics at the Exascale

Molecular dynamics (MD) is a cornerstone of computational science. However, MD is prevented from achieving complete scientific success by an inability to simultaneously reach the necessary length and timescales while maintaining sufficient accuracy. While the raw computing power available at the exascale should allow for a dramatic extension of the range of applicability of MD, conventional massively parallel codes suffer from poor strong scalability. In other words, a simple scale-up of current practices would enable only the simulation of much larger systems (i.e., containing billions or trillions of atoms) and would do little to improve current timescales (ns) and accuracy (empirical potentials). As most challenging problems require accessing different regions in the accuracy (A), length (L), and time (T) simulation space (ALT), one of the scientific community's key tools, MD, is in danger of missing out on the exascale revolution. EXAALT strives to fill this gap.

The EXAALT project combines three state-of-the-art codes—LAMMPS, LATTE, and ParSplice—into a unified tool that will leverage exascale platforms efficiently across all three dimensions of the ALT space. The new integrated capability is composed of three software layers. First, a task management layer enables the creation of MD tasks, their management through task queues, and the storage of results in distributed databases. It is used to implement various replica-based accelerated MD techniques, as well as to enable other complex MD workflows. The second layer is a powerful MD engine based on the LAMMPS code. It offers a uniform interface through which various physical models can be accessed. The third layer provides a wide range of physical models from which to derive accurate inter-atomic/molecular forces. In addition to the large number of empirical potentials implemented in LAMMPS, it provides high-performance implementations of quantum MD at the Density Functional Tight Binding level, as well as to Spectral Neighbor Analysis Potentials (SNAP), a set of high-accuracy machine-learned potentials.

The first challenge problem is related to nuclear fission. Nuclear energy based on fission provides about 16% of the world's electricity. However, only 4–6% of the uranium atoms in the primary fuel, UO_2 , are burned, leaving behind a vast energy resource and creating a greater-than-necessary nuclear waste problem. One of the primary reasons

is material integrity: as the fuel burns, radiation damage and fission gases accumulate, causing swelling of the fuel, pellet–clad interactions, and increased pressure on the clad. Because current burnup levels are predicated on our understanding of how the fuel evolves, improved models of fission gas evolution offer the potential for extracting more energy from the fuels.

Solving this grand challenge will require a significant advance in the scientific community's ability to carry out high-accuracy, electronic structure–driven MD simulations on the timescales needed to observe diffusion of defects. Given the size of these defects, relatively small systems (~100 atoms) are sufficient. However, given the high barriers for U-defect evolution, very long timescales will be required for the defects to move at the temperatures of interest. On petascale platforms, simulation rates are estimated at only 10 ns per day; thus, solving this problem requires development of a new simulation capability for the exascale.

The second challenge problem relates to nuclear fusion. Realizing the promise of fusion as a commercially attractive energy source for the 21st century requires advanced structural materials capable of sustained operation in an extreme environment with high temperatures and high fluxes of helium, hydrogen isotopes, and neutrons. The performance demands on plasma-

facing components of future fusion power plants are beyond the capability of current materials. Tungsten will be the divertor material in ITER and is the leading candidate material for future fusion reactors. However, experiments indicate the possibility of substantial surface modification in tungsten exposed to low-energy plasma containing helium. Experiments show that nanostructured fuzz, a nanoporous phase with tendrils on the order of tens of nm in diameter, forms on the surface when the surface temperature is between 1000 and 2000 K and the incident ion energies are between 20 and about 100 eV. Such surface features will impact heat transfer and fuel retention, increase the rates of erosion through both sputtering and dust formation, and embrittle the divertor. These modifications to the microstructure can lead to premature materials failure or may quench the fusion reaction by cooling and destabilizing the plasma. However, at this time, the fundamental mechanisms that lead to fuzz formation are unknown, which makes the development of mitigation strategies very difficult. Solving this problem requires a dramatic extension of the reach of large-size extended-time MD simulations. The project simulates the evolution of a tungsten first wall in conditions typical of fusion reactor operation. The primary target is to simulate a 10^5 atom system with a quantum-trained SNAP potential.

Progress to date

- Released baseline physical models for both challenge problems (i.e., tungsten for fusion, UO_2 for fission).
- Fission science-at-scale demonstration on homogeneous nodes (~270,000 cores on ALCF/Theta).
- Fusion science-at-scale demonstration on homogeneous nodes (~270,000 cores on ALCF/Theta).

EXAALT is creating an exascale MD application that will answer fundamental material problems for energy-relevant materials, especially for fission and fusion energy production. Knowledge gained will be used to design accident-tolerant fuels and will aid in the design of the first wall of fusion reactors.

PI: Danny Perez, Los Alamos National Laboratory

Collaborators: Los Alamos National Laboratory, Sandia National Laboratories, University of Tennessee