FINAL REPORT

Hierarchical Petascale Simulation Framework for Stress Corrosion Cracking

Scientific Discovery Through Advanced Computing
Grant Number: DE-FC02-06ER25788
SciDAC Science Application: Materials Science & Chemistry

Priya Vashishta, Principal Investigator
Collaboratory for Advanced Computing and Simulations
Departments of Chemical Engineering & Materials Science,
Computer Science, and Physics & Astronomy
University of Southern California
3651 Watt Way, Vivian Hall of Engineering (VHE 608)
Los Angeles, CA 90089-0242
Phone: 213-821-2662
Fax: 213-821-2664
E-mail: priyav@usc.edu

DOE Office of Science Program Office:
Office of Advanced Scientific Computing Research

DOE Technical Program of Contact:

NNSA POC: Dr. Njema Frazier
Phone: 202-586-5789, Email: Njema.Frazier@nnsa.doe.gov

Office of Advanced Scientific Computing Research, SciDAC Division POC: Dr. Lali Chatterjee
Phone: 301-903-7379, Email: Lali.Chatterjee@ascr.doe.gov

Collaborating Institutions and Investigators:

University of Southern California	**Rajiv K. Kalia, Aiichiro Nakano**
Harvard University	**Efthimios Kaxiras**
Purdue University	**Ananth Grama**
California State University-Northridge	**Gang Lu**
Los Alamos Nat'l Laboratory	**Stephan Eidenbenz, Arthur F. Voter**
Lawrence Livermore Nat'l Laboratory	**Randy Q. Hood, John A. Moriarty, Lin H. Yang**

Stress Corrosion Cracking:
Molecular Dynamics Simulations Reveals Mechanisms of Nickel Fracture

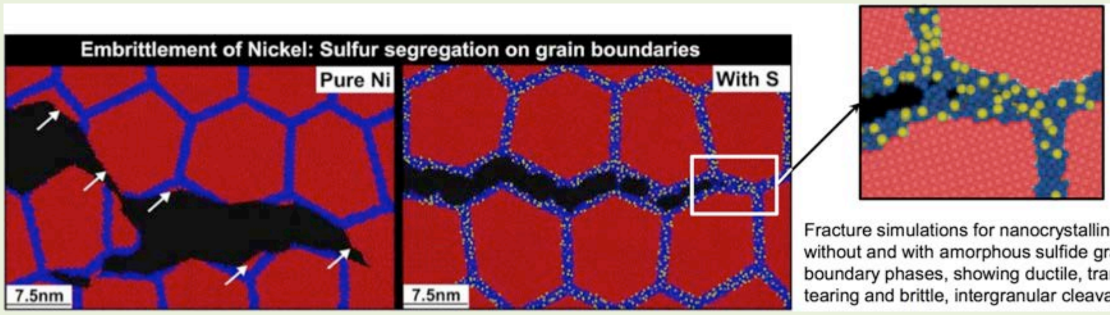

Fracture simulations for nanocrystalline nickel without and with amorphous sulfide grain-boundary phases, showing ductile, transgranular tearing and brittle, intergranular cleavage.

Highlighted in the FY 2012 Budget Request to Congress for DOE's Office of Science by William F. Brinkman, Director of Science, DOE (February 14, 2011) [H. Chen, R. K. Kalia, E. Kaxiras, G. Lu, A. Nakano, K. Nomura, A. C. T. van Duin, P. Vashishta, and Z. Yuan, "Embrittlement of metal by solute segregation-induced amorphization," *Physical Review Letters* **104**, 155502 (2010)].

§1 PROJECT OVERVIEW AND OBJECTIVES

The performance and lifetime of materials widely used in nuclear technologies and in advanced power generation technologies such as turbines, combustors, and fuel cells are often severely limited by corrosion under stress loads. Most critical here is premature and catastrophic failure of materials resulting from chemically influenced corrosion. The basic requirements for the operation of structural systems exposed to corroding conditions under stress loads are safety and reliability. Safe and reliable operation is endangered by the uncertainties in stress corrosion cracking (SCC). To prevent SCC and to predict the lifetime beyond which SCC may cause failure requires that we understand the atomistic mechanisms underlying SCC.

The goals of this SciDAC project are presented in the figure on the previous page. We are designing a hierarchical petascale simulation framework to address the challenging problem of corrosion-induced failure in materials. Our multidisciplinary team consisting of computational material scientists, applied mathematicians, and computer scientists from four universities and two DOE labs is developing a scalable parallel and distributed computational framework consisting of modeling techniques, algorithms, analytical underpinnings, and release quality software for 1) multi tera-to-petascale simulations with quantum-level accuracy; 2) multimillion-to-multibillion atom molecular dynamics (MD) simulations based on density functional theory (DFT) and temperature dependent model generalized pseudopotential theory; 3) quasicontinuum (QC) method embedded with quantum simulations based on DFT; and 4) accelerated molecular dynamics (AMD) coupled with hierarchical atomistic/QC simulations to reach macroscopic length and time scales relevant to SCC. This computational framework entails linear-scaling algorithms and performance-optimization techniques to achieve scalability beyond 10^5 processors, and automated model transitioning to embed higher fidelity simulations concurrently inside coarser simulations on demand with well-defined error bounds and controlled error propagation across the scales and models.

We are using this hierarchical petascale simulation framework to perform large-scale DFT and multimillion-atom, microsecond MD simulations of SCC in metals, alloys, glasses, and high-temperature structural ceramics in gaseous and aqueous environments. These petascale simulations provide the basic mechanisms of SCC, and also essential input for atomistically informed continuum models to make contact with engineering SCC data. We are also designing an extensible simulation integration framework on parallel/distributed platforms to enable 1) collaborative construction and execution of complex, multi-component, computationally demanding numerical simulations based on first-principles models of SCC; and 2) optimization of the performance of structures and systems in chemically aggressive environments, motivated by the challenge and opportunities for the development of innovative SCC inhibitors and SCC-resistant materials.

The integrated SCC framework, as well as each simulation methodology individually packaged, will be made available to the public domain under a GNU GPL, and released through a web portal linked to the NERSC and other DOE supercomputing centers. This will serve a variety of applications of critical interest to the DOE mission. The highlight of our education program is a dual-degree offering that provides graduate students the opportunity to obtain a PhD in the physical sciences or engineering and an MS in computer science with specialization in high performance computing and simulations.

Our SciDAC project is tightly coupled with the SciDAC Performance Engineering Research Institute (PERI) to optimize our petascalable MD and QM applications, as well as a wide range of other scientific applications, on emerging multicore parallel supercomputers combining our metascalable message passing + multithreading + data parallelism parallelization framework and PERI's auto-tuning and code-generation framework. **We are requesting support for one postdoc at USC for outreach to SciDAC Center collaboration over to the next level by demonstrating the auto-tuning for a broader range of high-end computing applications including our MD, ReaxFF-MD and EDC-DFT (see §5 Interactions with SciDAC Centers and §6 Budget Overview).**

§2 RESEARCH HIGHLIGHTS

In the two-and-a-half years since this project started, we have developed a scalable, parallel and distributed computational framework consisting of methods, algorithms, and integrated software tools for multi tera-to-petascale simulations. We have designed scalable parallel algorithms for quantum molecular dynamics (QMD) and large-scale MD based on reactive force fields (F-ReaxFF). Scalability has been achieved on 10^5 processors through linear-scaling algorithms and performance-optimization techniques. We have demonstrated 218 billion-atom

MD, 1.72 billion-atom reactive MD, and 1.68 trillion degrees-of-freedom (19.2 million-atom) QMD in the framework of divide-and-conquer density functional theory on adaptive multigrids with the parallel efficiency exceeding 0.95 on 212,992 Blue Gene/L processors; see Fig. 1.

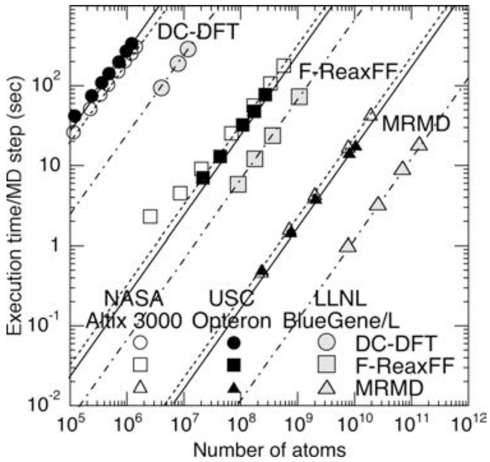

Fig. 1: Execution time per time step vs. number of atoms for our ReaxFF and effective force-field (EFF) MD and DFT algorithms on 212,992-CPU BG/L at LLNL, 1,920 CPUs of SGI Altix 3000 at NASA-Ames, and 2,048 AMD Opteron CPUs at USC. Lines show linear scaling.

We have performed ab-initio MD simulation of S embrittlement of Ni with the "Car-Parrinello-like approach to Born-Oppenheimer MD", to accelerate density functional theory (DFT)-based MD simulations. This has been achieved by propagating the electronic degrees of freedom in terms of the contra-covariant density matrix in a nearly time-reversible manner. We are developing a new hybrid QMD-MD/model generalized pseudopotential theory (MGPT) simulation algorithm for materials properties to be used in SCC applications. This algorithm directly links QMD with large-scale MD through quantum-based model generalized pseudopotential theory interatomic potentials.

We have performed multimillion-atom MD simulations to investigate 1) damage initiation and growth mechanisms in mode-1 fracture over microsecond time scales (see Fig. 2); 2) deformation, plasticity, and break-up of voids in shearing silica glass; and 3) defect nucleation and migration in nanoindentation of a-SiO_2. We have also developed a hybrid MD and mesoscale lattice Boltzmann (LB) simulation algorithm to study fluid flow through a network of cracks in silica (Fig. 3). **The microsecond fracture simulations are being carried out on the IBM BG/L at LLNL**.

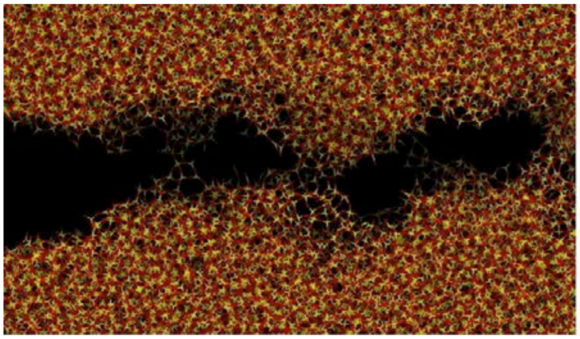

Fig. 2: Snapshot of an MD configuration showing a nanocavity ahead of the pre-crack. Only a small part of the system is shown here. Red and yellow are O and Si atoms, respectively.

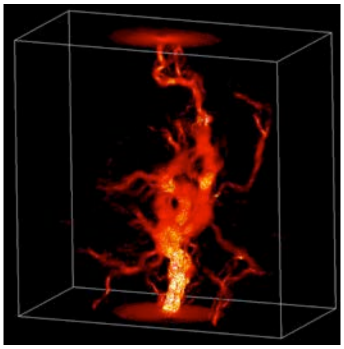

Fig. 3: Lattice Boltzmann flow simulation through nanopores and cracks in silica glass. The silica simulations are performed with MD for systems ranging between 10^6 and 10^9 atoms.

We have performed 48 million-atom reactive MD simulations of sulfur segregation-induced embrittlement of nickel on BlueGene/P at ANL (Fig. 4). The simulation results have revealed a link between sulfur-induced intergranular amorphization and embrittlement. We find that an order-of-magnitude reduction of grain-boundary shear strength due to amorphization, combined with tensile-strength reduction, allows the crack tip to always find an easy propagation path. This mechanism explains an experimentally observed crossover from transgranular to intergranular fracture and suppressed plastic activities due to sulfur segregation.

We have generalized two accelerated MD (AMD) methods, namely parallel replica dynamics (ParRep) and hyperdynamics, using the key concept of a super-state to extend the applicability of AMD to solid-liquid interfaces. Using this idea, we have adapted the statistical mechanics concepts that underpin the AMD methods so that they can be applied at a coarse-grained level.

We have developed a concurrent multiscale method that makes it possible to simulate multi-million-to-billions atoms effectively based on the DFT (Fig. 5). The method termed QCDFT is formulated within the framework of quasicontinuum method, and with DFT as its sole energetic input, *i.e.*, there is no empiricism in the formulation.

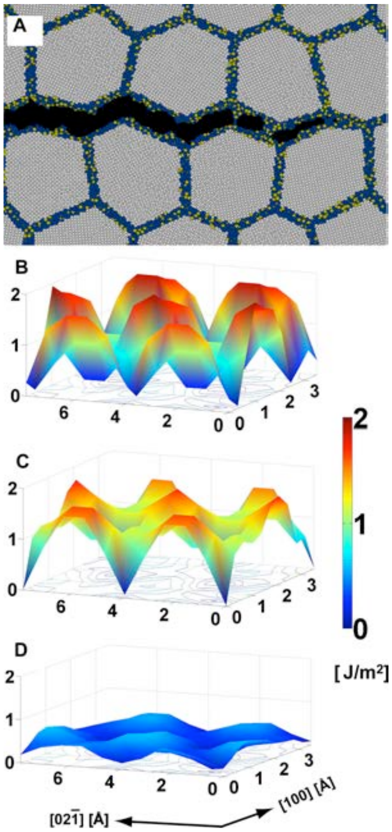

Fig. 4: (**A**) Close-up of fracture simulation in nanocrystalline Ni with amorphous sulfide GB layers (gray, blue and yellow colors represent Ni atoms inside grains, Ni atoms within 5 Å from GBs, and S atoms). Generalized stacking fault energy (GSFE) of Ni ∑5(012) GB without S (**B**), with a monolayer of segregated S (**C**), and with amorphous sulfide GB phases (**D**).

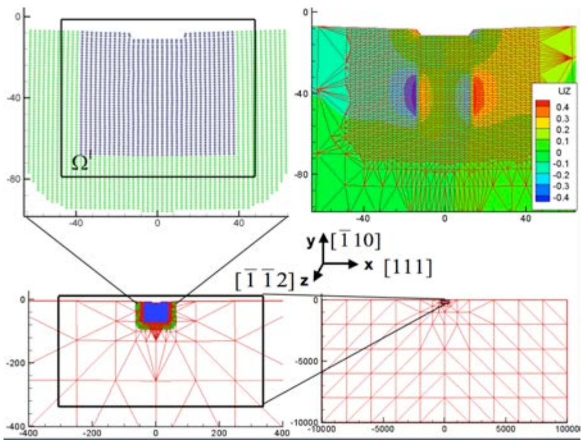

Fig. 5: The overview of the entire system and domain partition in QCDFT with nanoindentation as an example. DFT atoms are shown in blue and buffer atoms are in green. The buffer atoms provide environment for which the quantum problem of DFT atoms is solved. The energy of the finite-elements and the interaction energy between the buffer atoms and DFT atoms are calculated by DFT as well.

§3 DESCRIPTION OF RESEARCH ACCOMPLISHMENTS

Many challenging materials applications require a hierarchy of simulation methods from first-principles predictions of the fundamental processes to simulations of the performance of engineering structures. Stress corrosion spans length scales from interatomic distances to the macroscopic scale. Chemically reactive domains in SCC dynamically evolve and changes in the reaction zone occur on time scales ranging from femtoseconds to seconds. Simulations over such vast length and time scales require a continuous re-evaluation of the domains over which the computational methods are applied and straightforward communication across the quantum mechanical (QM), classical molecular dynamics (MD), and continuum boundaries is vital.

The problem at the QM/MD interface lies in the fact that the charge flow and bond breaking events at the QM electronic scale are difficult to translate to the rigid bonding/fixed point charges regime usually employed in the classical MD approach. To facilitate this translation process, it is advantageous to insert another computational method that has the ability to dissociate and create chemical bonds and can handle charge flow by employing a chemical-environment dependent charge model.

First-principles generalized pseudopotential theory (GPT) provides a fundamental basis for interatomic potentials in metals and alloys through rigorous expansions of the electron density and total energy within DFT. First-principles GPT and simplified model-GPT (MGPT) potentials, including explicit 3- and 4-body multi-ion interactions, have been successfully developed for prototype systems and applied in MD simulations with as many as 100 million atoms on the IBM BG/L supercomputer at LLNL. We are expanding these capabilities in both scope and accuracy through the seamless coupling of MD/MGPT simulations with first-principles quantum molecular dynamics (FPMD) simulations. Our approach is to develop and optimize *temperature-dependent* MGPT interatomic potentials via hybrid MD/MGPT–FPMD simulations on small systems in which the potentials are improved on-the-fly in an iterative fashion. **This will bring for the first time quantum-level accuracy to multi-million-atom, nanosecond simulations of SCC**.

The other development required for application to SCC is embedding of the accelerated molecular dynamics (AMD) simulation into finite-temperature quasi-continuum method to reach length and time scales relevant to SCC. This development will impact many other technologically relevant processes, including vapor-deposited film growth, dislocation motion during plastic flow, grain-boundary motion during grain growth and recrystallization, and annealing after radiation damage events.

Lastly, in order to make contact with engineering data and applications, we are performing quasicontinuum simulations of SCC incorporating DFT and classical MD approaches. The quasicontinuum model enables the elucidation of propagation mechanisms occurring at the mesoscale, *e.g*., intermittent crack growth and the validation of the approach based on engineering SCC data.

§3.1 PETASCALABLE HIERARCHICAL SIMULATION ALGORITHMS AND TOOLS

We have designed a computational framework consisting of methods, algorithms, and integrated software tools for 1) petascale simulations with quantum-level accuracy; 2) multimillion-atom MD simulations based on reactive force fields; and 3) semi-empirical, non-reactive force fields to reach time scales of up to a microsecond. We have developed linear-scaling algorithms and performance-optimization techniques to achieve scalability beyond 10^5 processors. Automated model transitioning has been employed to embed higher fidelity DFT simulations concurrently inside quasicontinuum simulations only when and where they are required.

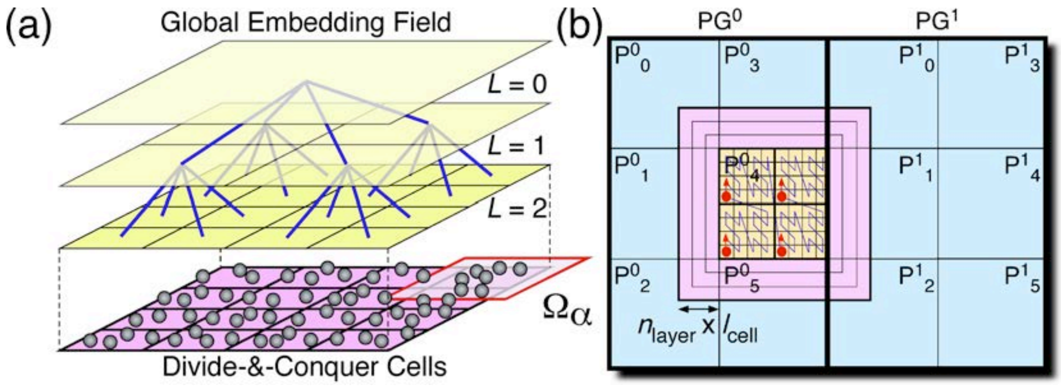

Fig. 6: (a) Embedded divide-and-conquer (EDC) algorithm. The physical space is subdivided into spatially localized cells, which are embedded in a global field solved with a tree-based algorithm. (b) In tunable hierarchical cellular decomposition (HCD), the physical volume is recursively subdivided into process groups, PG^γ, processes, P^γ_π, and computational cells.

A. Metascalable Dwarf

We have developed a metascalable (or "design once, scale on new architectures") parallel computing framework for large-scale atomistic simulations of materials based on spatiotemporal data locality principles, which is expected to scale on future multipetaflops architectures. The metascalable framework (Fig. 6) consists of 1) an embedded divide-and-conquer (EDC) framework based on spatial locality to design linear-scaling algorithms for high complexity problems; 2) a space-time-ensemble parallel (STEP) approach based on temporal locality to predict long-time dynamics, while introducing multiple parallelization axes; and 3) a tunable hierarchical cellular decomposition (HCD) framework to map these $O(N)$ algorithms onto a multicore cluster combining message passing, critical section-free multithreading, data-level parallelism, and optimal use of memory hierarchy. The "seven dwarfs" (a dwarf is an algorithmic method that captures a pattern of computation and communication), which were first identified by Phillip Colella of the Lawrence Berkeley National Laboratory, have been used widely to develop scalable parallel programming models and architectures. The EDC-STEP-HCD framework encompasses all seven dwarfs and is expected to serve as a "metascalable dwarf" to represent broad, large-scale scientific and engineering applications (*e.g.*, equation solvers, constrained optimization, search, visualization, and graphs involving massive data). The following summarizes specific simulation algorithms developed for this project within the metascalable framework.

First-principles QM simulation: Our embedded divide-and-conquer density functional theory (EDC-DFT) algorithm represents the physical system as a union of overlapping spatial domains, $\Omega = \cup_\alpha \Omega_\alpha$, and physical properties are computed as linear combinations of domain properties. For example, the electronic density is expressed as $\rho(\mathbf{r}) = \Sigma_\alpha p^\alpha(\mathbf{r}) \Sigma_n f_n^\alpha |\psi_n^\alpha(\mathbf{r})|^2$, where $p^\alpha(\mathbf{r})$ is a support function that vanishes outside the α-th domain Ω_α, and f_n^α and $\psi_n^\alpha(\mathbf{r})$ are the occupation number and the wave function of the n-th electronic state (*i.e.*, Kohn-Sham orbital) in Ω_α. The domains are embedded in a global Kohn-Sham potential, which is a functional of $\rho(\mathbf{r})$ and is determined self-consistently with $\{f_n^\alpha, \psi_n^\alpha(\mathbf{r})\}$. We use the multigrid method to compute the global potential. The DFT calculation in each domain is performed using a real-space approach, in which electronic wave functions are numerically represented on grid points. The real-space grid is augmented with coarser multigrids to accelerate the convergence of iterative solutions. Furthermore, a finer grid is adaptively generated near every atom in order to accurately operate norm-conserving pseudopotentials that describe electron-ion interactions. The exchange-correlation energy is calculated in a generalized gradient approximation. The convergence of the $O(N)$ DFT algorithm, as well as its energy conservation in MD simulations, have been fully verified.

Reactive molecular dynamics: We have developed a fast reactive force-field (F-ReaxFF) algorithm to perform 10^9-atom ReaxFF MD. The variable N-charge problem in ReaxFF amounts to solving a dense linear system of equations to determine atomic charges $\{q_i \mid i = 1,\ldots,N\}$ at every MD step. F-ReaxFF reduces the $O(N^3)$ complexity to $O(N)$ by combining the fast multipole method (FMM) based on spatial locality and iterative minimization to utilize the temporal locality of the solution. To accelerate the convergence, we use a multilevel preconditioned conjugate-gradient (MPCG) method that splits the Coulomb-interaction matrix into short- and long-range parts and uses the sparse short-range matrix as a preconditioner. The extensive use of the sparse preconditioner enhances data locality and improves parallel efficiency.

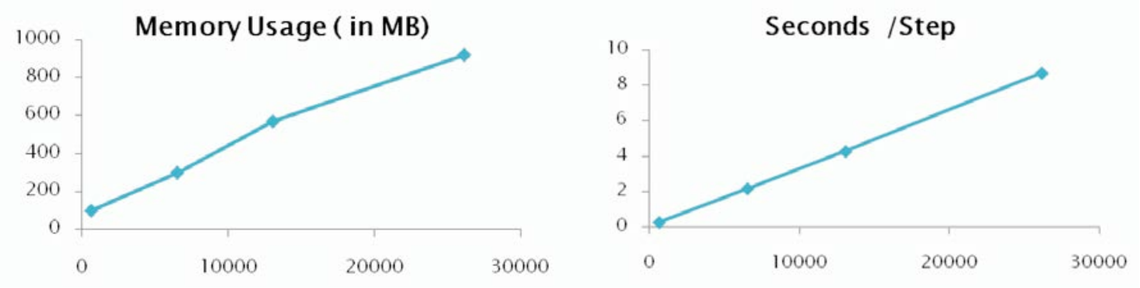

Fig. 7: Memory footprint and node computation time for bulk water ensemble at temperature 300 K. (a) Memory usage as a function of the number of atoms, and (b) computation time as a function of the number of atoms.

To support highly sustained computation rates, scalable parallel performance, and flexible data structures and design, we have further developed a number of computational techniques for reactive MD. We begin with a comprehensive characterization of both the runtime and memory footprint of our code. Figure 7(a) shows the node memory scaling for our code. Specifically, even for systems with up to 26,000 atoms, the code requires under 1 GB of memory (bulk water). The node execution time for the same systems is shown in Fig. 7(b). Two important observations are: 1) Even for complex systems like water, our code is extremely efficient, capable of running 10,000 atoms/node at under 3 seconds/step; and 2) for the 26,000-atom system, the simulation time is approximately 9 seconds/step. Considering that this fits in the memory of most conventional scalable parallel systems, this allows us to amortize significant communication overhead, thus yielding excellent parallel scalability.

We next examine the computational cost of the algorithm, in addition to constants at specific problem sizes. There are four major components of the computation—neighbors, bonded atoms, non-bonded atoms, and charge equilibration. While the first three correspond to computation of interatomic potential functions, charge equilibration corresponds to computation of partial charges to minimize electrostatic energy under charge-neutrality constraint. This computation can be reduced to a single linear system solver, with the matrix kernel corresponding to shielded electrostatics (1/*r*). The linear system is solved using generalized minimal residual (GMRES) or conjugate gradient (CG) method. Since the total computation time is essentially determined by the solver as well as the stopping criteria, we first examine the required stopping criteria. In Fig. 8, we plot the total energy of an NVE ensemble with two different stopping criteria (residuals of 10^{-4} and 10^{-6}). The plots are for bulk water at 300 K over 30,000 time steps (of 0.2 fs each). The figure shows a drift in energy at higher tolerance (residual of 10^{-4}). This effect is observed in a variety of other systems as well (we have experimented with silica and silica-water), implying that a residual of 10^{-6} or lower is necessary for high-fidelity simulations.

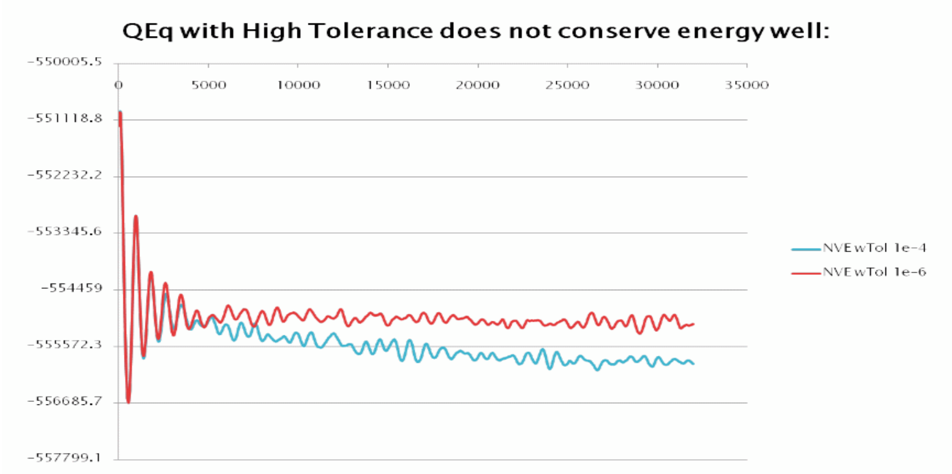

Fig. 8: Energy drift in bulk water NVE ensemble with high tolerance (high residual norm). In general, a residual norm of 10^{-6} or lower is necessary for minimizing this drift.

Based on the aforementioned requirements on charge equilibration, we present a detailed breakdown of our node execution time in Fig. 9. The figure presents execution profile for an NVE bulk water system at 300 K. We observe that while at higher residuals ~10^{-4}, charge equilibration takes just over a third of computation time, at more realistic residuals (10^{-6} and lower), this time increases to between half and three-quarters of total time. This clearly motivates development of efficient solvers (preconditioners) capable of reducing computational cost of charge equilibration. It is also possible to show that the iteration count for charge equilibration is, in fact, a function of the number of atoms in the system. This implies that the per-time step simulation cost is in fact super-linear in the number of atoms. This must be addressed through effective preconditioners that can maintain constant iteration counts with increasing numbers of atoms.

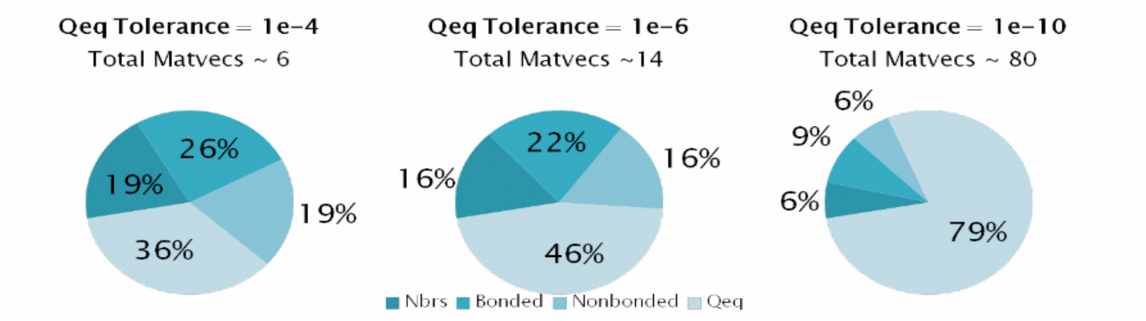

Fig. 9: Detailed breakdown of simulation time for various tolerances associated with charge-equilibration linear solve.

A significant fraction of our algorithm development has focused on efficient solvers for charge equilibration. Here, we leverage a number of key observations:

- The partial charge on atoms does not change significantly across time steps. Therefore, it is possible to extrapolate charges from previous time steps to derive a good initial solution.
- It is possible to reuse Krylov subspaces from previous iterations, thus alleviating overheads of orthogonalization at each time step.
- Diagonal and block-Jacobi-type accelerators can provide some reduction in iteration counts. Block-Jacobi-type preconditioners simply compute inverses of (decoupled) short-range interactions and use these as preconditioners. Note that this is identical to block-Jacobi, if a consistent numbering of atoms is possible. We precompute the inverses at each time step and use it over a small number of subsequent steps as well. We are currently investigating a more aggressive preconditioner based on our Spike toolkit, which explicitly handles the off-block-diagonal ‘spikes’.
- The charge-equilibration coupling decays with distance—for this reason, we can use a preconditioner similar to a multipole operator, which aggregates far-field coupling. We show that, while such preconditioners are somewhat expensive to compute, they are capable of delivering constant iteration counts. For this reason, while we do not use this preconditioner for small systems, this is available as an option for large systems at low residuals.

Based on these approaches, we have developed and implemented a highly efficient parallel charge equilibration solver, along with the neighbor, bonded, and non-bonded potential computations.

We have also developed a highly efficient and scalable MPI-based implementation of the ReaxFF code, with key optimizations including 1) minimization of message counts by trading off redundant computations; 2) efficient use of shadow domains and the midpoint method to minimize redundant computations; 3) optimizations with respect to orthogonalizations in charge equilibration solve; and 4) optimizations for selected platforms. A limited release parallel version is currently in validation phase. An AIX optimized version of our code has also been developed by Prof. Jorn Amundsen's group at the Norwegian University of Science and Technology.

B. Massively Parallel Accelerated Molecular Dynamics

A challenging problem is to predict long-time dynamics because of the sequential bottleneck of time. Due to temporal locality, the system stays near local minimum-energy configurations most of the time, except for rare transitions between them. In such cases, the transition state theory (TST) allows the reformulation of the *sequential long-time dynamics* as computationally more efficient *parallel search* for low activation-barrier transition events. We also introduce a discrete abstraction based on graph data structures, so that combinatorial techniques can be used for the search.

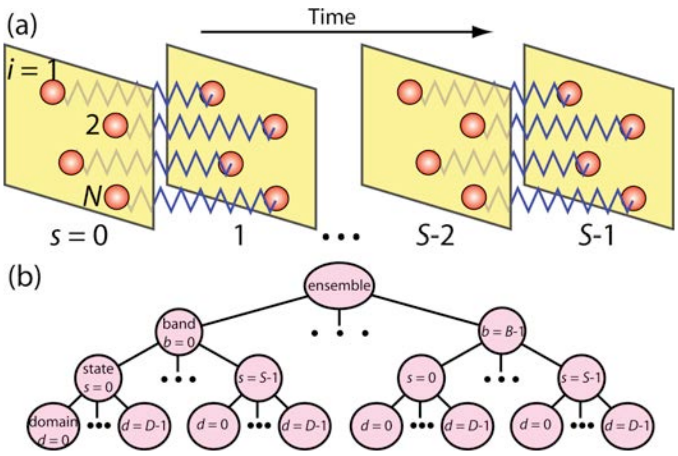

Fig. 10: Space-time-ensemble parallel (STEP) approach. (a) A nudged elastic band consists of a sequence of S states (yellow parallelograms), R_s ($s = 0, ..., S-1$), where each state consists of N atoms (red spheres). Corresponding atoms in consecutive states interact via harmonic forces represented by blue wavy lines. (b) Tree-structured processor organization in the STEP approach. An ensemble consists of B bands, each consisting of S states; each state in turn contains D spatial domains.

We have developed a directionally heated nudged elastic band (DH-NEB) method, in which a NEB consisting of a sequence of S states, $\mathbf{R}_s \in \Re^{3N}$ ($s = 0, ..., S-1$, $\Re$ is the set of real numbers, and N is the number of atoms), at different temperatures searches for transition events (Fig. 10(a)). We use an ensemble consisting of B bands to perform long-time simulation in the framework of kinetic Monte Carlo simulation. Here, our space-time-ensemble parallel (STEP) approach combines a hierarchy of concurrency, *i.e.*, the number of processors is $P = BSD$: 1) Spatial decomposition within each state (D is the number of spatial subsystems); 2) temporal parallelism across S states within each band; and 3) ensemble parallelism over B bands (Fig. 10(b)).

C. Petascalability of MD and QM Algorithms

To map our $O(N)$ QM, MD and AMD algorithms onto parallel computers, we have developed a hierarchical cellular decomposition (HCD) framework, including wavelet-based computational-space decomposition for adaptive load balancing. Scalability of the F-ReaxFF MD and EDC-DFT simulation algorithms, as well as that of our non-reactive space-time multiresolution MD (MRMD) algorithm, has been tested on a number of parallel architectures including the IBM BlueGene/L at the Lawrence Livermore National Laboratory, IBM

BlueGene/P at the Argonne National Laboratory, NASA Columbia, and AMD Opteron cluster at USC. Benchmark tests carried out on 212,992 BlueGene/L processors have achieved unprecedented scales of quantum-mechanically accurate and validated, chemically reactive atomistic simulations—1.72 billion-atom F-ReaxFF MD and 1.68 trillion grid points (19.2 million-atom) EDC-DFT MD—in addition to 218 billion-atom nonreactive MRMD. The metascalable EDC-STEP-HCD framework exposes maximal concurrency and data locality, and consequently the parallel efficiency is well over 0.95 on 212,992 BlueGene/L processors at Lawrence Livermore and 131,072 IBM BlueGene/P processors at Argonne; see Fig. 11. The spatiotemporal scale covered by direct MD simulation on a sustained petaflops computer per day (*i.e.* petaflops•day of computing) is estimated to be *NT* = 2.14 (*e.g.*, *N* = 2.14 million atoms for *T* = 1 microseconds).

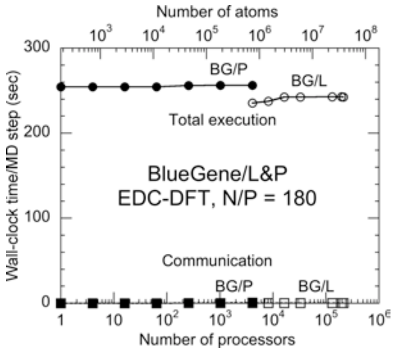

Fig. 11: Total execution (circles) and communication (squares) times per MD time step as a function of the number of processors *P* of BlueGene/L at the Lawrence Livermore National Laboratory (open symbols) and BlueGene/P at the Argonne National Laboratory (solid symbols) for EDC-DFT MD for 180*P* atom alumina systems.

Scalability of the STEP-MRMD algorithm (note that the STEP approach can be combined with any of the MRMD, F-ReaxFF and EDC-DFT algorithms or their combinations to compute interatomic forces) is tested on a cluster of dual-core, dual-processor AMD Opteron nodes with Myrinet interconnect. We first test the scalability of temporal decomposition, where we fix the number of bands *B* = 1 and the number of domains per state *D* = 1. Here, the simulated system is amorphous SiO_2 consisting of *N* = 192 atoms. The measured speedup on 1,024 processors is 980.2, and thus the parallel efficiency is 0.957; see Fig. 12(a). Next, we test the scalability of ensemble decomposition, where we fix the number of states per band *S* = 4 and the number of spatial domains per state *D* = 1. On 1,024 processors, the measured speedup is 989.2, and thus the parallel efficiency of ensemble decomposition is 0.966, which is slightly higher than that of temporal decomposition on the same number of processors; see Fig. 12(b).

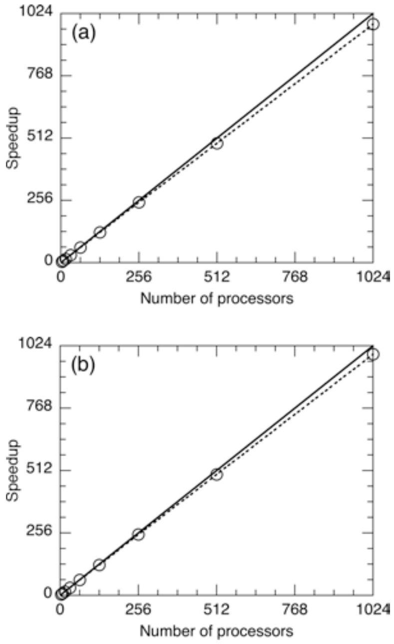

Fig. 12: (a) Speedup of temporal decomposition in the STEP-MRMD algorithm as a function of the number of processors *P*. The circles are measured speedups, whereas the solid line denotes the perfect speedup. (b) Speedup of ensemble decomposition in the STEP-MRMD algorithm as a function of *P*.

D. Multicore Optimizations

We are extending the metascalable framework to capture salient features of emerging multicore parallel architectures. The extended hierarchical parallelization framework combines 1) internode level parallelism by spatial decomposition using message passing; 2) intra-node (inter-core) level parallelism through a master-worker paradigm and cellular decomposition using critical section-free multithreading; and 3) intra-core level parallelism via single-instruction multiple-data (SIMD) techniques.

In EDC algorithms, computations at the finest level consist of computational cells. On a multicore compute node, these computational cells are distributed among multiple cores in an interleaved fashion (see Fig. 13). One problem in the interleaved thread parallelization is that multiple threads may update a common lattice site. To avoid such a critical section on a multicore architecture, we have designed a double-layered data structure. On the Cell Broadband Engine (CBE) architecture, for example, this safely separates computation and date access via direct memory access (DMA). Furthermore, we have developed a multithreading scheme that takes into account cache coherency to maximize performance.

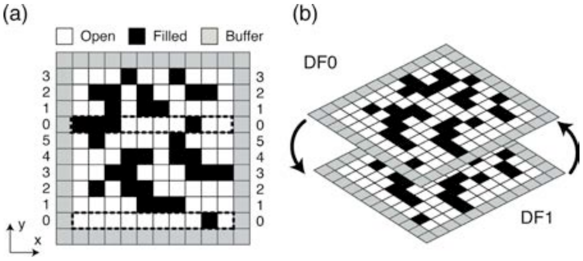

Fig. 13: (a) Two-dimensional illustration for each domain in spatial decomposition. Small squares are computational cells. The simulation system is divided into N_y computational chunks, each of which consists of N_yN_z cells, and the chunks are interleavingly assigned to cores. The numerals show thread ID responsible for each chunk. (b) Schematic of a double-layered calculation comprising of two arrays *DF0* and *DF1*.

We have tested the multithreading scalability of MRMD on a dual Intel Xeon quadcore platform. Figure 14 shows the speedup of the multithreaded code over the single-thread counterpart as a function of the number of threads. The code scales quite well up to 8 threads on the 8-core platform. The multithreading efficiency of two-body force calculation is 0.927, while that for three-body force calculation is 0.436, for 8 threads. The low efficiency of the three-body force calculation may be due to the redundant computations to eliminate critical sections. Nevertheless, the efficiency of the total program is rather high (0.811), since the fraction of the three-body calculation is about one third of the two-body force calculation. In a test calculation for a 12,228-atom silica system, the running time is 13.6 milliseconds per MD time step.

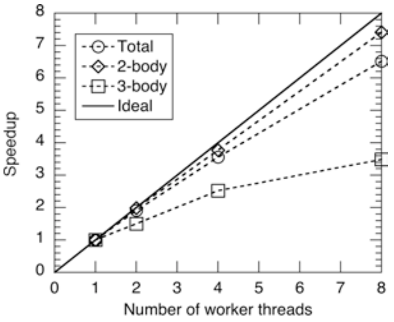

Fig. 14: Speedup of the multithreaded MRMD algorithm over a single-threaded counterpart for the total program (circles), the two-body force calculation (diamonds), and three-body force calculation (squares). The solid line shows the ideal speedup.

For data-level parallelism *via* SIMD, we have developed various code transformation techniques of general use, *e.g.*, translocated statement fusion, vector composition *via* shuffle, vectorized data layout reordering (*e.g.*, matrix transpose), zero padding for solving the alignment issue for complex data type (*e.g.*, array), and simple data-type reformatting for solving the alignment issue for data with irregular memory accessing; see Fig. 15.

Figure 16 shows the performance of SIMDized lattice Boltzmann method (LBM) for fluid-flow simulation on CBE. We test the SIMD effect with various problem sizes both for the collision part (which is fully SIMDized) and the whole pLBM application on a Cell, where the number of threads is set to one. Figure 16(a) shows the SIMD speedup of the collision function (*i.e.*, the ratio of the running time without SIMD and that with SIMD). It is an ascending function of the problem size and reaches 3.72 for 2.61×10^5 grid points, which is close to the ideal speedup of 4 (SIMD efficiency is 0.930). Figure 16(b) shows the SIMD effect for the whole pLBM application. The SIMD speedup of pLBM is again an ascending function of the problem size and reaches 3.12 for 2.614×10^5 grid points (SIMD efficiency 0.780). The SIMD optimization is thus highly effective for LBM.

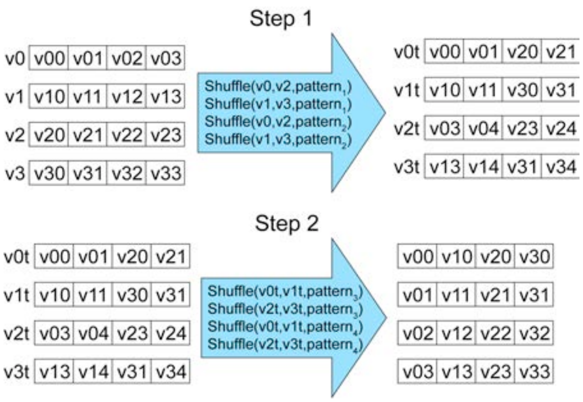

Fig. 15: An example of SIMD transformation, in which a 4×4 array is transposed by a sequence of shuffle operations. Function *shuffle*(*vec1*, *vec2*, *pattern*) selects bytes according to the *pattern* (a control vector) from source vectors *vec1* and *vec2* and places them in a target vector.

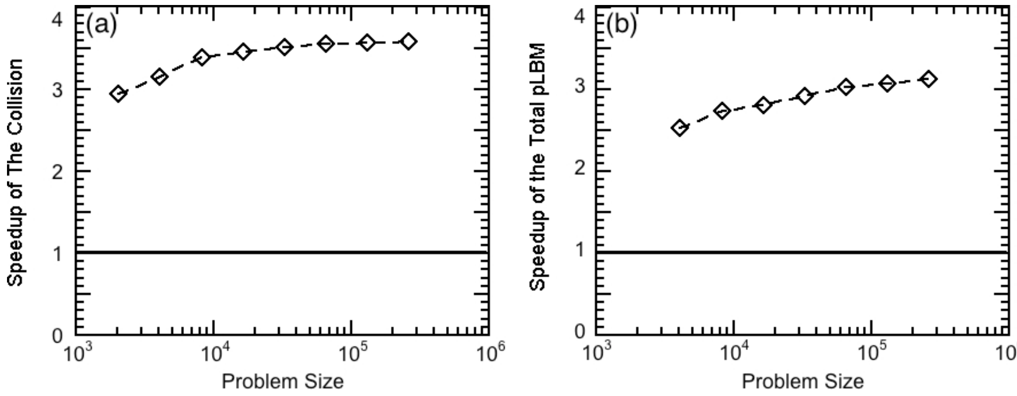

Fig. 16: Effect of SIMD with problem size on a Cell. The speedup of collision function due to SIMD (a) and the speedup of LBM due to SIMD (b) are plotted as a function of the problem size.

For LBM, we have found that a low-cost PlayStation3 cluster based on Cell outperforms a Xeon cluster when the problem size is large enough in compute and memory-access intensive applications, and the performance enhancement is an increasing function of the problem size. This indicates that the DMA efficiency increases with the data size. For the largest problem size, the performance enhancement of PlayStation3 over Xeon for LBM is 11.0.

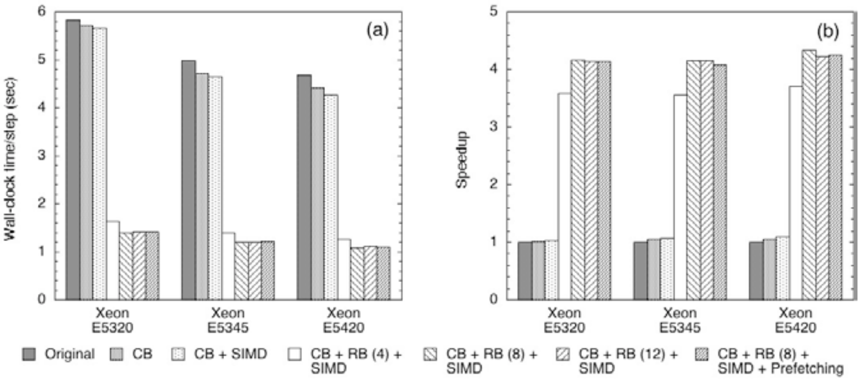

Fig. 17: Wall-clock time (a) and corresponding speedup over the original code (b) due to our optimization techniques: cache blocking (CB); register blocking (RB); single-instruction multiple-data (SIMD); and prefetching.

We have also applied a number of in-core optimization techniques to high-order stencil computations in real-space DFT calculation, including 1) cache blocking for efficient L2 cache use; 2) register blocking and data-level parallelism via SIMD techniques to increase L1 cache efficiency; and 3) software prefetching. Our generic approach is tested with a 6th-order stencil on a suite of Intel Xeon architectures. Cache blocking and prefetching are found to achieve modest performance improvement, whereas register blocking and SIMD implementation reduce L1 cache line miss dramatically accompanied by moderate decrease in L2 cache miss rate. Optimal register blocking sizes are determined through analysis of cache performance of the stencil kernel for different sizes of register blocks, thereby achieving 4.3-fold speedup on Intel Harpertown (Fig. 17).

We have also implemented LBM on a graphic processing unit (GPU) based on the compute unified device architecture (CUDA). We have compared the comparative performance of the GPU (NVIDIA 8800 GTS) using CUDA and a conventional CPU (Sempron 3500+) for pLBM. Here, the running time of each function as well as the speed-up of GPU over CPU are measured for various problem sizes from 8^3 to 64^3 lattice points.

Figure 18(a) shows the running time of the entire program on GPU and CPU, and Fig. 18(b) shows the corresponding speedup of GPU over CPU. The performance of GPU is much better than that of CPU for all problem sizes we have tested, despite the large preparation time of GPU. The best speed-up we obtain is 8.76.

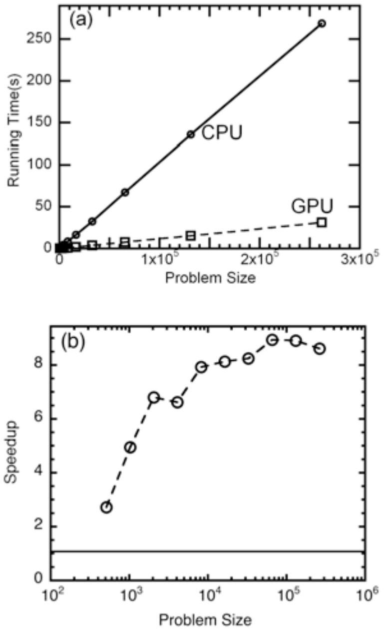

Fig. 18. (a) Total running time of pLBM flow simulation on CPU and GPU with different problem sizes. (b) Speed-up of the total pLBM flow simulation of GPU over CPU with different problem sizes.

E. Performance Modeling, Analysis, Optimization, and Hybrid MPI/OpenMP Parallelization of Molecular Dynamics

We have developed a performance prediction model for molecular dynamics (MD) simulations, thereby predicting the optimal cell dimension in a linked-list cell method. The model expresses computation time in terms of the number and unit computation time of key operations. The model accurately estimates the number of operations during the simulations with the maximum standard error of 10.6% compared with actual measurements. Then, the unit computation times of the operations are obtained by bisquare regression. Analysis of this model reveals that the optimal cell dimension to minimize the computation time is determined by a trade-off between decreasing search space and increasing linked-list cell access for smaller cells. The model predicts the optimal cell dimension correctly for 80% of all tested cases, resulting in an average speedup of 10% and 52% for the cutoff radius of interaction, 6.6 and 10.0 Å, respectively.

We have parallelized the calculation of the long-range Coulombic potential in MD simulation based on a hybrid message-passing (MPI)/multithreading (OpenMP) scheme. The explicit pair kernel of the particle-particle/particle-mesh algorithm is multi-threaded using OpenMP, while communication between multicore nodes is handled by MPI. We have designed a load balancing spanning forest (LBSF) partitioning algorithm, which combines:

1. Fine-grain dynamic load balancing.
2. Minimal memory-footprint data privatization via nucleation-growth allocation.

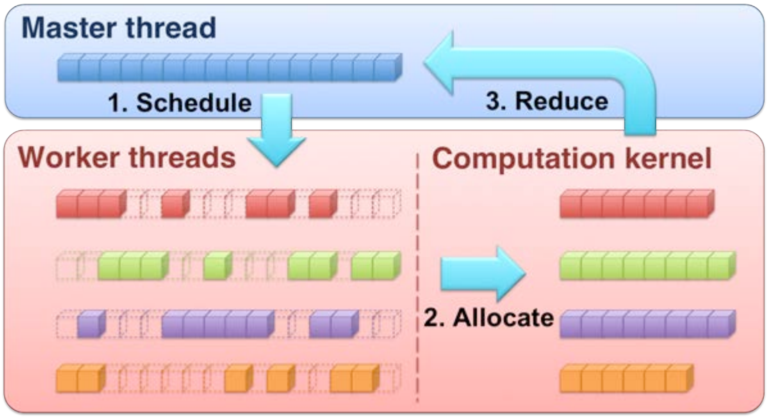

Fig. 19: Memory layout and three-step algorithm for scheduling algorithm.

The three-step LBSF algorithm (Fig. 19) consists of the following steps: (1) the scheduler assigns computation units to threads and then determines which subset of the global data each thread requires; (2) each thread allocates its private memory as determined by the scheduler; and (3) private force arrays from all threads are reduced into the global force array. The algorithm reduces the memory requirement for thread-private data from $O(np)$ to $O(n + p^{1/3}n^{2/3})$ — amounting to 75% memory saving for p = 16 threads working on n = 8,192 atoms, while maintaining the average thread-level load-imbalance less than 5%. Strong-scaling speedup for the kernel is 14.4 with 16-way threading on a four quad-core AMD Opteron node. In addition, our MPI/OpenMP code shows 2.58× and 2.16× speedups over the MPI-only implementation, respectively, for 0.84 and 1.68 million particles systems on 32,768 cores of Blue Gene/P (Fig. 20).

F. Intelligent Auto-tuning

The hierarchy of computational cells provides an efficient mechanism for performance optimization of both computation and communication. The layout and size of the cells are treated as tunable parameters and are optimized on each computing platform. Our EDC-STEP algorithms are implemented as hybrid message-passing + multithreading programs in the tunable HCD framework, in which the numbers of message passing interface (MPI) processes and POSIX threads are also tunable parameters. The HCD framework thus maximally exposes data locality and concurrency. **We are collaborating with compiler and artificial intelligence (AI) research groups in the SciDAC-Performance Engineering Research Institute (PERI) to use 1) knowledge-representation techniques for expressing the exposed concurrency; and 2) machine-learning techniques for optimally mapping the expressed concurrency to hardware**.

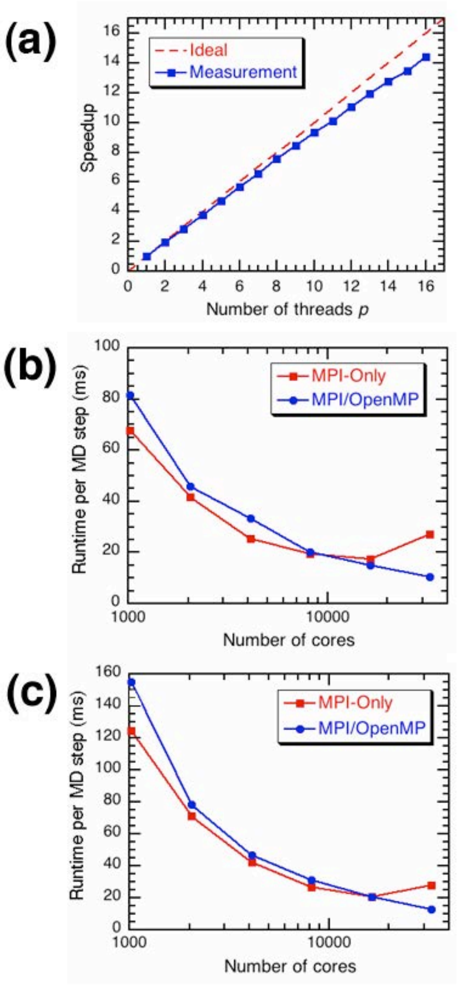

Fig. 20: (a) Thread-level strong scalability of the parallel section on a four quad-core AMD Opteron 2.3 GHz with fixed problem size at n = 8,192 particles. (b) Total running time per MD steps of 1,024 – 32,768 Power PC 450 850 MHz cores on Blue Gene/P for a fixed problem size at N = 0.84-million particles and (c) 1.68-million particles.

Specifically, the SciDAC-PERI is developing automatically tuned computational kernels, using model-guided empirical optimization (MGEO), with our applications as an example. Model-guided empirical optimization is a compiler-based strategy for automatic performance tuning that combines the complementary strengths of model-based optimization and empirical search. Models and compiler heuristics can limit the search to a small number of the most promising candidate implementations. Empirical measurements provide the most accurate information to the compiler to select among candidates and tune optimization parameter values. The MGEO approach will include intelligent search techniques using machine learning to systematically prune the space of alternative implementations to reduce search time.

G. Data-Mining Approach for Fast DFT-Quality Molecular Dynamics

We are developing new approaches for calculating DFT-informed interatomic forces at a speed similar to a fast empirical potential. Our method is based on data-mining and database search algorithms. The method stores results of high-quality DFT force evaluations in a large, self-adaptive and self-updating database. Each entry, which we call a chunk, consists of an atom (called the central atom), the DFT force on that atom, and all the atom's neighbors out to some cutoff radius R (a few neighbor shells). The chunks have been carved out of a larger system on which DFT was performed. During an MD simulation, the force on each atom i in the simulated system S will be computed as follows: 1) A chunk with atom i as a central atom is carved out of the system using the same radius R as the chunks in the database. We call this chunk a target chunk. 2) The database is scanned to collect a set of chunks C that are geometrically "similar" to this target chunk. 3) The force on atom i in S is then determined using a physics-based interpolation over the DFT forces on the central atoms using the chunks in C. 4) If the size of C is too small and/or the error estimation from step 3 is too large, the force on atom i is computed directly from DFT, either using the entire system, or a suitable subset with passivated boundaries. The target chunk with its associated force computed in Step 3 is added to the database.

A key part of this procedure is the high-dimensional, physics-based interpolation, which consists of parametrizing a local, empirical interatomic potential form to obtain the best fit to the DFT forces over the set C of chunks. This potential is then differentiated to obtain the force on the central atom of the target cluster. Here, the key computational challenge is the rapid extraction of geometrically similar chunks from the database. In order to formalize that, represent each chunk of k atoms as a set of k-1 vectors, where the i-1st vector, $i > 1$, is obtained by subtracting the coordinates of the first atom, which is the central atom, from the coordinates of the i-th atom. Let c_1 and c_2 be two chunks and for simplicity assume that they have the same number of atoms. Let c_i, for i=1,2, consists of the vectors, v_{i1}, ..., v_{ik}. The distance between c_1 and c_2 is defined as

$$d(c_1,c_2) = \max_{1 \le j \le k} \left\| v_{1j} - v_{2j} \right\|$$

where $\| \ \|$ is the Euclidean norm in 3D. The distance $d(c_1, c_2)$ does not take into consideration rotations and permutations on the set of the vectors, which do not change the shape of the chunk. Hence, we define similarity between $c_1$and c_2 as

$$\sigma(c_1,c_2) = \min_{p \in \Pi, r \in R} d\big(p(r(c_1)),c_2\big) \qquad (1)$$

where Π and R are the sets of all permutations of order k-1 and all rotations in 3D, respectively, and where permutation are applied on the order of the vectors in the chunk. The shape-matching problem we consider is, given a pair of chinks c_1 and c_2, compute $\sigma(c_1,c_2)$. However, computing similarity between chunks directly using formula (1) is computationally hard since there is an exponential number of permutations and the space of all possible rotations is also quite large. Therefore, we are developing an algorithm not directly using formula (1) that is able to compute efficiently a good approximation of the similarity and will then generalize it so that it can be also applied to pairs of chunks of different sizes.

Our shape matching algorithm computes for each chunk a signature, which is a fixed-length vector uniquely characterizing the chunk. Each element of the vector is the measurement of a statistical characteristic of the chunk. In our particular implementation, described below, the characteristic is an eigenvalue of a matrix constructed using the vectors representing the chunk. Instead of using formula (1), similarity between two chunks will be defined as the Euclidean distance between their signatures, and could be computed in time linear to the signature size. In order to describe our approach to matching, let $c = v_i, \ldots, v_{ik}$ be any chunk and let (x_i, y_i, z_i) be the coordinates of v_i. Define the vectors

$$X = (x_1,\ldots,x_k), \quad Y = (y_1,\ldots,y_k), \quad Z = (z_1,\ldots,z_k)$$

and the matrix

$$M = \begin{bmatrix} \langle X,X \rangle & \langle X,Y \rangle & \langle X,Z \rangle \\ \langle Y,X \rangle & \langle Y,Y \rangle & \langle Y,Z \rangle \\ \langle Z,X \rangle & \langle Z,Y \rangle & \langle Z,Z \rangle \end{bmatrix}$$

which is called a covariance matrix for X, Y, and Z. The covariance matrix is a measure of the statistical dispersion of the x, y and z coordinates and hence its elements could be considered as a possible choice for a signature for c. Since M is not invariant with respect to rotations of the chunk, we instead use the set of all eigenvalues of M. The symmetric M has three real eigenvalues, which we denote by λ_1, λ_2, λ_3. Sort λ_1, λ_2, λ_3. in increasing order and let $\lambda_1', \lambda_2', \lambda_3'$,be the resulting sequence. Then we can use the vector $(\lambda_1', \lambda_2', \lambda_3')$ as a signature for c. Our experiments have shown that, in many cases, such choice for a signature for the chunk is appropriate. Its advantage is its very short size and, hence, very small times for computation and for measuring chunk similarity. However, the small signature size has a drawback of not guaranteeing unique identification. In order to increase the accuracy, we extend the signature by generating s different copies of c using appropriate transformations on c, then computing a signature of length 3 as described above for each of the s copies, and finally concatenating them into a single signature for c. Specifically, for generating the i-th copy, we multiply each vector v of c by $k_i(\| v \|)$, where k_i is a weight function. In the current implementation, we set $s = 100$.

Our approach allows efficient database querying. By appropriately sorting the database elements with respect to the elements of the signature, we are able to extract closest elements to a query chunk in $O(s\log N)$ time, where s is the size of the signature and N is the size of the database (which will be of order millions). This bound is much better than the bounds of most of the existing alternative methods, which are linear in N. Moreover, the above bound is a worst-case bound and we believe that in practice the algorithm will work even faster.

H. Petascale Data Visualization and Analysis

We incorporate real-time visualization capabilities in petascale simulations. Our visualization lab has a 14'×8' tiled display and an immersive and interactive *ImmersaDesk*, which are used to render large datasets at a near interactive speed. To achieve real-time petascale visualization capability, we have developed a system based on 1) a parallel and distributed approach to efficiently select a data subset within the field-of-view using the octree data structure; 2) a probabilistic approach to efficiently remove hidden atoms; and 3) a machine-learning approach to predict the user's next movement and prefetch data from a cluster to the graphics server. This visualization system renders a billion-atom dataset in an immersive virtual reality environment at nearly real time.

The visualization software incorporates data analysis tools based on graph algorithms that abstract materials structure as a chemical-bond network or graph $G = (V, E)$, in which atoms constitute the set of vertices V, and the edge set E consists of chemical bonds. We use our scalable, parallel, shortest-path circuit analysis algorithm with small memory footprint, based on dual-tree expansion and spatial hash-function tagging (SHAFT). SHAFT utilizes the vertex-position information to design a compact collision-free hash table, thereby avoiding the degradation of cache utilization for large graphs. The graph-based data mining approach allows automated identification of key chemical reactions and topological changes of chemical bond networks.

§3.2 PETASCALE SIMULATIONS WITH QUANTUM-LEVEL ACCURACY

A. Ab-Initio Molecular Dynamics Simulation of Sulfur Embrittlement of Nickel

Finite temperature effects are tackled by probing the system under tensile strain using the "Car-Parrinello-like approach to Born-Oppenheimer Molecular Dynamics", which was recently developed by one of us in order to accelerate DFT-based MD simulations. In the spirit of the Car-Parrinello approach, the method circumvents the need to perform fully self-consistent electronic structure calculations. However, in contrast to the original scheme, large integration time steps can be used. To achieve this, we propagate the electronic degrees of freedom in terms of the contra-covariant density matrix in a nearly time-reversible manner by applying the always stable predictor-corrector method. The corrector consists of a single preconditioned minimization step using the idempotency conserving orbital transformation method. In this way sizable minimization steps can be taken, which guarantees that the propagation is performed within a small skin very close to its instantaneous ground state. Deviations from the Born-Oppenheimer surface are small and, more importantly constant, which implies that energy differences are portrayed with great accuracy. In spite of these excellent properties, the incomplete energy minimization and the use of a non-symplectic integrator for the electronic degrees of freedom induces a small dissipative drift. We show how this can be rigorously corrected using a properly modified Langevin equation, which leads to correct sampling of the Boltzmann distribution. Furthermore the friction term is low enough, so that the dynamics is also reproduced with excellent approximation. The method works well irrespective of system type and band gap, so it is very efficient even in simulating large metallic systems. We implemented these ideas in the mixed Gaussian Plane Wave (GPW) code *QUICKSTEP* which is part of the publicly available suite of programs CP2K. Using this code we have begun to explore proof-of-concept studies at ambient conditions.

In order to investigate the ductile-to-brittle transition in Ni, we have performed a study of the mechanical behavior of fcc-Ni as a function of added sulfur impurities in the context of Rice's macroscopic theory. Specifically we have investigated bulk fcc-Ni and fcc-Ni with a $\Sigma5(012)$ grain boundary (GB) defect. We have calculated the decohesion energy γ_s and the unstable stacking fault energy γ_{us} for both systems and as a function of S impurity concentration. This allowed us to calculate the ductility D, for each case and hence enabled us to quantify if a change from ductile-to-brittle behavior is expected to occur. Our DFT calculations for the bulk structure were performed for the (111) surface of bulk Ni. This contains the unstable stacking energy and lowest decohesion energy for fcc-Ni. We have extended the Rice formalism to the case of a grain boundary and assume here that when a GB is involved in determining the mechanical behavior of a solid then it serves as both the cleavage and slip reference plane. This is reasonable since it is the weak plane of the system on which plastic deformation or decohesion occurs. Our results show that the decohesion energy for bulk Ni decreases for impurity addition. However, the unstable stacking energy was found to decrease as well, giving an overall increase in the ductility D by 24%. In principle, this would indicate that this system becomes more ductile. In contrast to this, the ductility for fcc-Ni with a GB defect changes dramatically—it is reduced by 86% when the GB surface has a S concentration of 100% of the available interstitial sites. The decohesion energy for the GB decreases and the unstable stacking energy increases as shown in Figs. 21(a) and (b), respectively. In light of the experimental evidence for sulfur-induced embrittlement of Ni, this strongly suggest that S impurities in the bulk crystal do not play a dominant role in embrittlement, but instead this effect is primarily caused by S impurities that accumulate at GB defect sites and then weaken the overall structure mechanically.

Table I: Summary of decohesion and unstable stacking fault energy for bulk Ni and Ni with a $\Sigma5(012)$ GB. The percentage given is the percent of available impurity sites occupied at the relevant surfaces.

	bulk Ni	bulk - 25% S	GB - pure	GB - 25% S	GB - 100% S
γ_s [J/m^2]	1.91	1.40	1.78	1.53	1.12
γ_{us} [J/m^2]	0.29	0.17	0.30	-	1.18
D	6.59	8.24	5.93	-	0.95

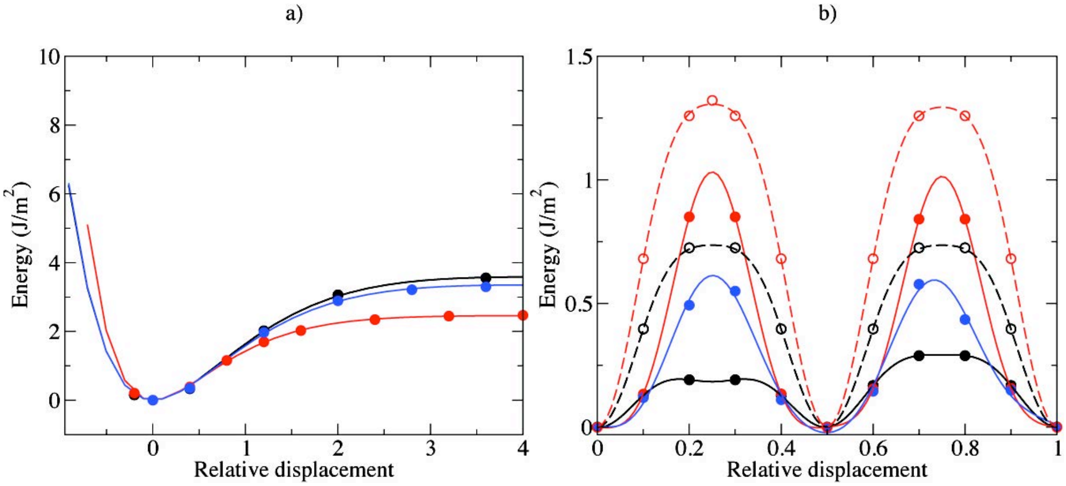

Fig. 21: (a) Decohesion energy for fcc-Ni with a GB. (b) Generalized stacking fault energy for fcc-Ni with a GB. Black, blue and red points are for 0, 25% and 100% S surface concentration in the available GB sites. Dashed and solid lines are for calculations without or with surface relaxation, respectively.

The dependence on finite temperature is studied by means of ab-initio MD (AIMD) simulations, where tensile stress is applied while measuring the stress along the *z*-axis. The initial sample was obtained by carefully relaxing the previously obtained structure at 300 K within a NPT simulation, where we allowed isotropic cell deformations. In this first proof-of-concept study a rather high strain velocity of 50 m/s was applied. The measured stress is shown in Fig. 22 and exhibits sharp drops when atoms change positions to accommodate the applied strain. These simulations can provide insight into the process of crack formation due to the presence of impurities.

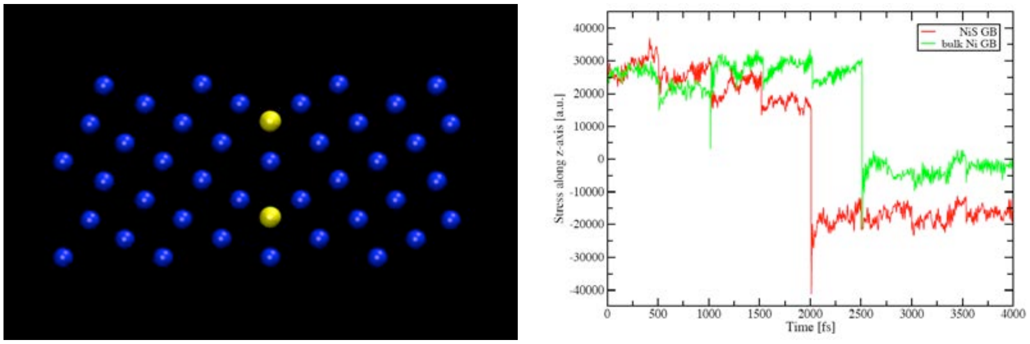

Fig. 22: (Left) Ni Σ5(012) grain boundary, where blue and yellow spheres are Ni and S atoms, respectively. (Right) Time evolution of stress during AIMD simulation of Ni GB under tension with and without a layer of segregated S.

B. Quantum-based Model Generalized Pseudopotential Theory (MGPT) Interatomic Potentials

In this project, we are also developing a new hybrid QMD-MD/MGPT simulation algorithm for materials properties to be used in SCC applications. This algorithm directly links first-principles DFT quantum simulations, namely, quantum molecular dynamics (QMD), with large-scale atomistic simulations, namely, molecular dynamics (MD) with quantum-based model generalized pseudopotential theory (MGPT) interatomic potentials. By developing this capability for small systems in the high-temperature solid and in the liquid, one can dramatically extend the effective QMD time scale and harvest optimized *temperature-dependent* MGPT potentials in the long time limit, as shown schematically below in Fig. 23. In this process, there is a real-time information feedback loop between the alternate QMD and MD/MGPT segments of the simulation, with both snapshot and accumulated QMD information being used to achieve optimization of the potentials. The optimized potentials can then be re-used in ultra large *linear-scaling* MD/MGPT atomistic simulations of materials properties. The initial goal is to simulate one billion atoms at high temperature in prototype systems (*e.g*., Ni and Ni_3Al) via MD/MGPT with the accuracy to a direct QMD simulation on a few hundred to thousand atoms.

In the first two years of the SCC project, we have fully adapted our existing quantum simulation code to perform efficient hybrid QMD-MD/MGPT simulations on state-of-the-art parallel platforms, including Atlas and BG/L at LLNL. The resulting parallel QMD-MD/MGPT code is now called P^3MD and is the primary workhorse in developing and applying the QMD-MD/MGPT algorithm, as well as in performing complementary DFT electronic structure calculations needed for SCC applications. In QMD mode, the P^3MD code uses the well established plane-wave pseudopotential method to solve the finite-temperature DFT quantum-mechanical equations for the valence and outer core electrons in the system at each time step in the simulation, thereby providing accurate instantaneous forces on the ions to perform MD simulations. Efficient parallelization is obtained by distributing the plane-wave basis functions, electronic energy bands and reciprocal-space k points across a group of nodes in an optimum way. A custom parallel 3D Fast Fourier Transform (FFT) routine has been developed and used to transfer information between reciprocal space, where the DFT equations are solved, and real space, where the forces are needed. In MD/MGPT mode, P^3MD evaluates the real-space, finite-temperature, many-body MGPT potentials at each time step, including two-, three- and four-ion interactions, to provide corresponding quantum-based energies and forces. This is done using the fast matrix version of MGPT, which also provides the most general functional form of the potentials currently available based on non-canonical *d* bands. In the QMD and MD/MGPT modes, the same molecular dynamics equations are solved and only the input forces are different. In the hybrid QMD-MD/MGPT simulation, P^3MD performs constant-volume, constant-temperature molecular dynamics in a chosen fixed-shape computational cell. Ion positions and velocities are passed forward from one segment of the simulation to the next to provide initial conditions. QMD snapshot and accumulated information is passed forward to test and refine the MGPT potentials at the beginning of each MD/MGPT simulation segment.

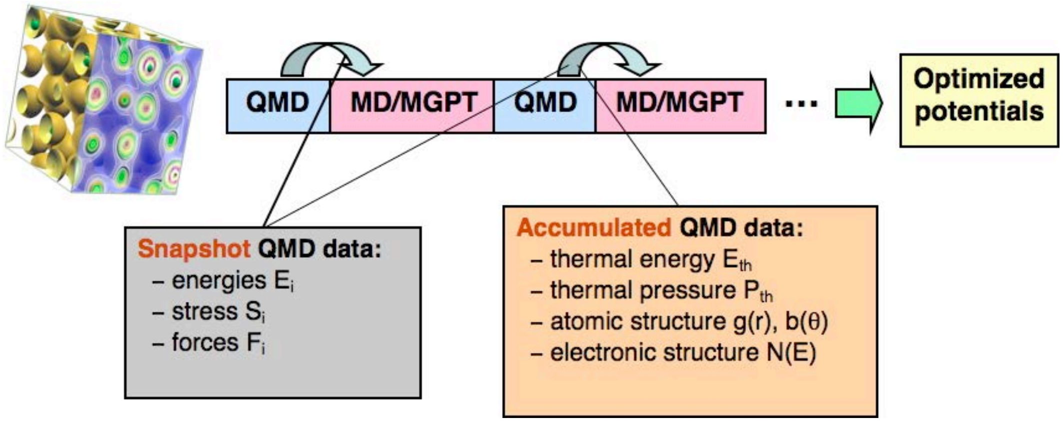

Fig. 23: Schematic representation of the present hybrid QMD-MD/MGPT simulation algorithm. A real-time information feedback loop between the alternate QMD and MD/MGPT segments of the simulation will allow optimized and re-useable MGPT potentials to be obtained in the long time limit.

Our applications research accomplishments during the first two years of the project have been primarily in developing first-principles DFT databases of fundamental bulk and defect properties for SCC materials, including Ni, Ni_3S_4 and the hydrated amorphous silica surface, and the development and application of $T = 0$ MGPT potentials for Ni using the DFT data base.

First-principles Bulk and Defect Property Databases for SCC Materials

In addition to dynamic high-temperature QMD data, it is very useful to have static $T = 0$ DFT databases of fundamental properties for interatomic-potential development on SCC relevant materials. In this regard, we have calculated extensive databases for pure non-magnetic Ni over an extended pressure range and for polymidymite Ni_3S_4 and the hydrated amorphous silica surface at ambient pressure. These calculations have been done with the same DFT plane-wave pseudopotential method used in the QMD simulations.

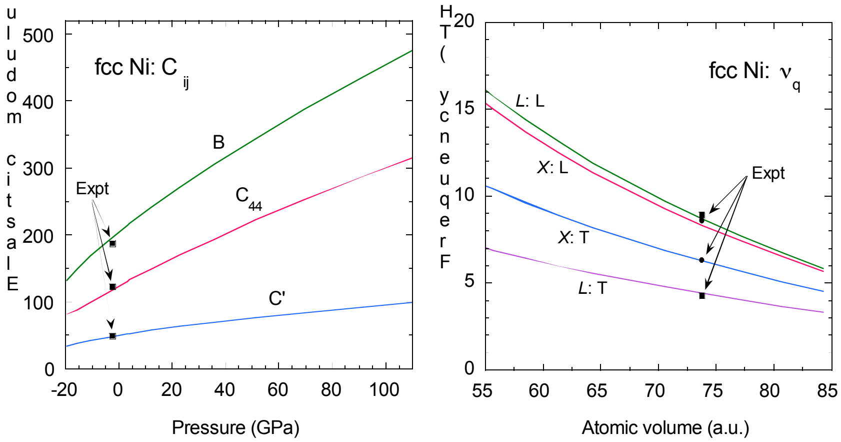

Fig. 24: Bulk and shear elastic moduli (left panel) and *L*- and *X*-point zone-boundary phonons (right panel) for non-magnetic fcc Ni, as calculated from the present first-principles DFT electronic-structure calculations and compared with experiment.

The pressure-dependent structural and mechanical properties of non-magnetic Ni have been calculated over a substantial pressure range extending to above 100 GPa. The calculations use modern generalized gradient approximation corrections to the exchange-correlation potential and energy within density-functional theory together with suitably developed pseudopotentials valid over this pressure range. The calculated zero-temperature structural and mechanical properties include the cold equation of state; fcc elastic moduli and zone-boundary phonons; fcc, bcc and hcp structural energies; and defect properties relevant to SCC applications. The observed ambient-pressure fcc phase is found to be thermodynamically stable throughout the pressure range studied. The calculated bulk and shear moduli and the L - and X-point zone-boundary phonons are displayed in Fig. 24. The elastic moduli and phonons are in very good agreement with experiment at ambient pressure and remain real and positive at high pressure, demonstrating that the fcc phase is mechanically stable in this regime. We also calculated the fully relaxed vacancy formation energy for fcc Ni at zero pressure, as well as {111}<112> generalized stacking faults and the Σ5 (012) grain boundary. These latter results will be discussed below in comparison with corresponding MGPT calculations.

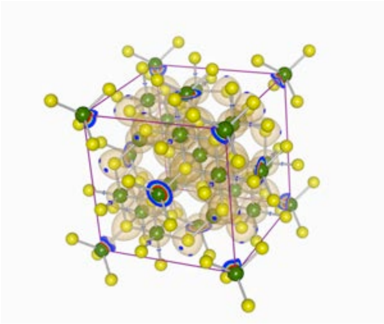

Fig. 25: Crystal structure and electronic charge density of polymidymite Ni_3S_4. There are two types of Ni atoms in this system: the Ni^I atom occupies 1/8 of the tetrahedral site and the two Ni^{II} atoms occupy half of the octahedral sites. There are a total of 56 atoms in the supercell, including 24 Ni atoms and 32 S atoms.

Polymidymite Ni_3S_4 is a sulfur rich compound and a good prototype system for the fundamental understanding of the electronic and structural properties associated with the SCC problem of sulfer concentrated at or near Ni grain boundaries. Polydymite Ni_3S_4 has two types of Ni atoms and one type of S atoms, $Ni^INi_2^{II}S_4$, in a cubic structure. The Ni^I atom occupies 1/8 of the tetrahedral site, the two Ni^{II} atoms occupy half of the octahedral sites, which produces a trigonal distortion along the [111] direction, and the four S atoms are on an approximate close-packed fcc lattice (Fig. 25). We have performed first-principles DFT electronic-structure calculations on the Ni_3S_4 solid at ambient pressure. The bulk properties of this system are important material parameters that will be used for interatomic-potential development on S-rich Ni-S systems. The calculated lattice constant for the ideal structure is 9.489 Å and three cubic elastic constants are C_{11} = 201 GPa, C_{12} = 81 GPa and C_{44} = 115 GPa. In addition, we have calculated the electronic properties of this system. In particular, the density of the states at the Fermi level $N(E_F)$ = 31.1 states/eV-cell, with 56 atoms per cell. This leads to a band-structure electronic-heat-capacity constant of g = 6.542 mJ/mol-K^2, where 1 mol refers to 1 mol of Ni_3S_4.

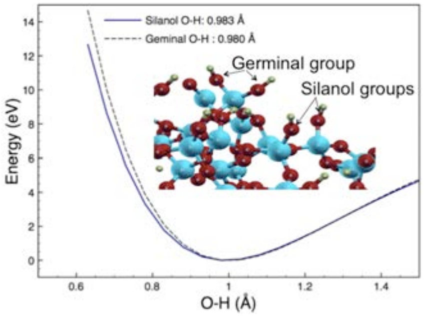

Fig. 26: Stretching energy profiles of O-H bonds for the silanol and germinal groups as calculated from the first-principles DFT calculations.

In close collaboration with USC, we have also established a first-principles database for passivated silica surface that can be used for SCC interatomic potential development. The interatomic potentials developed using this database are being integrated with the existing USC MD codes for large-scale simulations of SCC in silica glass in the presence of water. To establish the first-principles database, we have constructed a H-terminated a-SiO_2 surface model with a total surface silanol density of 6.41 nm^{-2}, which is in line with that quoted in a previous study. The stretching energy profiles of O-H bonds were then obtained with DFT calculations by stretching one of the O-H group in the system. We have calculated a total of six energy profiles, which include 3 silanol and 3 germinal groups, respectively. The average energy profiles for the silanol and germinal groups are shown in Fig. 26. Because of the differing bonding environments for these two groups, when the O-H separation is less than the mean equilibrium bond length of around 0.98 Å, the energy profiles differ by about 15%. When the separation is greater than 0.98 Å, on the other hand, the energy difference between these two groups is within 1%, indicating that a single set of database interatomic potential parameters should be sufficient to describe the O-H interactions.

Quantum-Based Interatomic Potentials for Ni

Using the above DFT database for Ni, we have developed $T = 0$ quantum-based MGPT potentials for this metal over the −20 to 100 GPa pressure range. In Ni, the *sp-d* hybridization is of primary importance and the three- and four-ion interactions are of secondary importance. We have, therefore, now developed an advanced form of MGPT which restores the *sp-d* hybridization from the full first-principles generalized pseudopotential theory in the two-ion pair potential, in a suitably screened form, and we have applied the new approach to Ni. In both the standard and advanced forms on MGPT, the potentials are volume dependent in the bulk, although the inclusion of the hybridization necessarily imparts a longer range to the two-ion potentials. Advanced MGPT potentials calculated for Ni at five representative atomic volumes, including equilibrium, are displayed in Fig. 27.

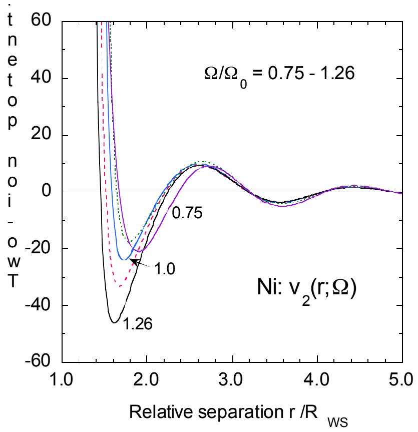

Fig. 27: MGPT pair potential, $v_2(r;\Omega)$, for Ni at five atomic volumes, where $\Omega_0 = 73.82$ a.u. is the observed equilibrium volume in the fcc material. The long-range oscillations in the potentials result primarily from *sp-d* hybridization.

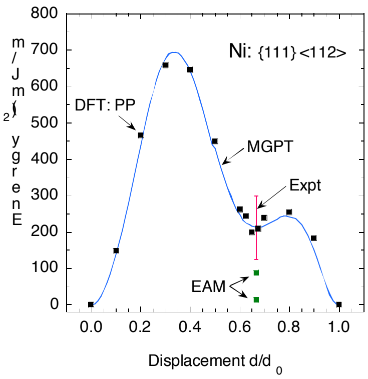

Fig. 28: The {111}<112> generalized stacking fault energy in Ni calculated at the observed equilibrium volume of $\Omega_0 = 73.82$ a.u., as obtained from the present MGPT interatomic potentials and from first-principles DFT calculations. For comparison, experimental measurements and previous EAM calculations of the observed intrinsic stacking fault energy γ_{isf} at $d/d_0 = 2/3$ are shown.

We have applied the $T = 0$ MGPT potentials for Ni to SCC-relevant point and extended defects at ambient pressure and to the calculation of the melt curve below 100 GPa. Calculated {111}<112> generalized stacking fault energies from both the MGPT potentials and the present DFT electronic-structure calculations are displayed in Fig. 28 and seen to be in excellent agreement. The calculated intrinsic stacking fault energy γ_{isf} at a displacement of $d/d_0 = 2/3$ is within the experimentally measured range and also agrees well with previous DFT calculations, as shown in Table I. In contrast, calculations of $\gamma_{\rm isf}$ with simple embedded-atom method (EAM) interatomic potentials tend to underestimate this quantity, as shown in Fig. 28. This is, at least in part, because of the inherent short-range nature of EAM potentials. Also given in Table II are MGPT and DFT calculations of the relaxed vacancy formation energy $E_{\rm vac}$ at ambient pressure and the Σ5(012) grain boundary energy at an atomic volume of 76.55 au. Again there is a good level of consistency between the MGPT and DFT results.

The melting curve of Ni has been calculated with the $T = 0$ MGPT potentials in the standard weak coupling limit. In the weak coupling limit the free energy is partitioned into cold, ion-thermal and electron-thermal contributions. The ion-thermal contribution is calculated in terms of the $T = 0$ electronic structure and forces, in this case from the MGPT potentials using MD simulations in the fcc solid and liquid. The electron-thermal contribution is calculated from the standard low-T expansion in terms of the density of electronic states at the Fermi level, which we obtain from the present DFT calculations. The calculated melt curve both with and without the electron-thermal contribution is plotted in left panel of Fig. 29 and compared with DAC experimental data. The MGPT pair correlation function in the liquid at 4,000 K is compared with that obtained from a QMD simulation, as shown in the right panel of Fig. 29; the pair correlation function based on the $T = 0$ potentials is too sharply peaked.

Table II. Calculated MGPT and DFT defect properties of Ni at ambient pressure.

Quantity	MGPT	Present DFT	Previous DFT	Experiment
$E_{\rm vac}$ (eV)	1.79	1.95	1.67-1.77	1.79
$\gamma_{\rm isf}$ (mJ/m^2)	213	195	180-183	125-300
Σ5(012) (J/m^2)	1.39	1.09	1.19	

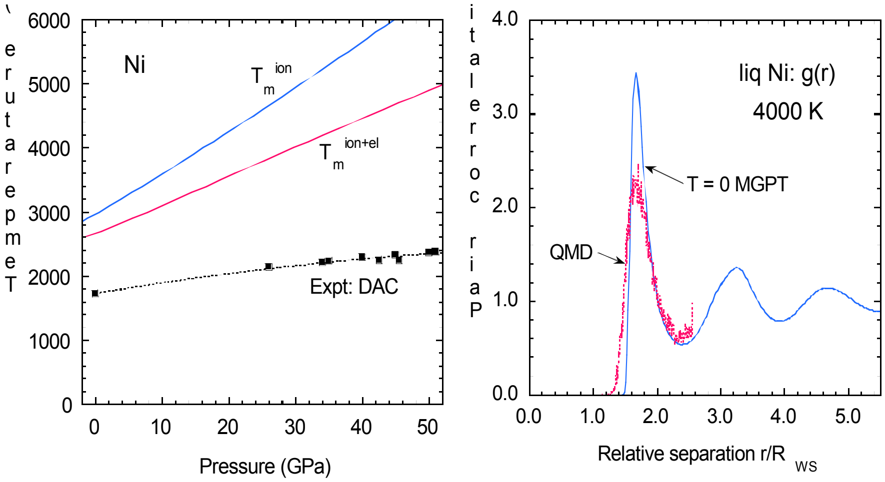

Fig. 29: Left panel: calculated melt curve for Ni via $T = 0$ MGPT potentials in the weak coupling limit, with (red line, $T_m^{\rm ion+el}$) and without (blue line, $T_m^{\rm ion}$) an electron-thermal contribution. Right panel: pair correlation function in the liquid at 4,000 K, as obtained from the MGPT potentials and from a QMD simulation.

§3.3 DEFORMATION, NANOINDENTATION AND FRACTURE IN AMORPHOUS SILICA

We have performed multimillion atom MD simulations to investigate 1) damage initiation and growth mechanisms in mode-1 fracture over nanosecond to microsecond time scales; 2) deformation, plasticity, and break-up of voids in shearing silica glass; and 3) defect nucleation and migration in nanoindentation of a-SiO_2. The simulations are based on an interatomic potential, which incorporates ionic and covalent effects through a combination of two-body and three-body terms. The interatomic potential was validated by comparing the MD simulation results for structural and mechanical properties of a-SiO_2 with experimental measurements and quantum mechanical calculations based on density functional theory, which were performed by Lin Yang at LLNL.

A. Microsecond MD Simulation of Fracture in Silica Glass

About five years ago, two experimental groups probed SCC in amorphous silica with an atomic force microscope (AFM) and arrived at entirely different conclusions. Bouchaud and co-workers claimed "Glass Breaks like Metal, but at the Nanometer Scale". They observed damage in the form of nanocavities ahead of the crack, growth of nanocavities with an increase in mode-1 loading, and fracture resulting from coalescence between cavities and the advancing crack front. Wiederhorn's group, on the other hand, did find any evidence of damage nanocavities in amorphous silica.

On the IBM BlueGene/L at LLNL, we are performing five-million atom MD simulations approaching microsecond time scales to examine the following issues: 1) Is there a process zone in silica glass and, if there is, what is the nature of damage in the process zone; 2) how does the strain rate affect the damage zone; and 3) what are the atomistic mechanisms of crack extension and failure in silica glass? Our simulations span three decades of strain rates—from 6×10^7 to 6×10^4 sec^{-1}. The smaller strain rates in our MD simulations are comparable to those employed in dynamic fracture experiments.

Figure 2 shows a pre-crack (black region on the left) and a nanocavity ahead of the pre-crack just before they merge to extend the crack. In the figure, red and yellow represent O and Si atoms, respectively. Some of these atoms re-bond the crack surfaces even after the crack front advances.

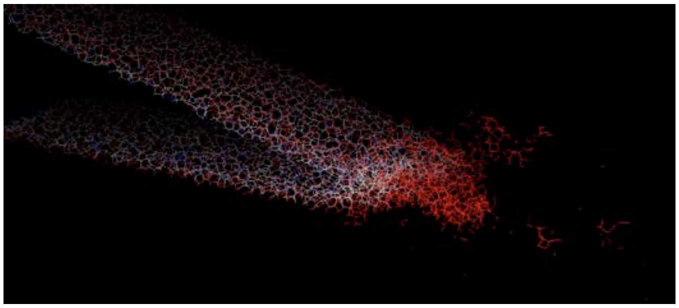

Fig. 30: Low-density region (red) just before the nucleation of a cavity at the crack tip.

The cavitation process finds its origin in the amorphous structure of silica, which contains SiO_4 terahedra connected at the corners to form rings of various sizes—from 3-membered to 9-membered –O-Si-O-Si-O-Si– rings. In the vicinity of the crack tip, the stresses create density fluctuation with low density regions surrounded by regions of normal density (i.e. bulk a-SiO_2 density of 2.2

g/cc). Low-density regions ahead of the crack tip are weaker or "less tough", and cavitation occurs preferentially in those regions. Figure 30 shows such a region (red) in front of the pre-crack tip, where the mass density is less than 1.9 g/cc.

The MD simulations performed at smaller strain rates reveal intermittent crack growth. At a strain of ε = 3.5% the crack front oscillates back and forth, while nanocavities nucleate in the damage zone (Fig. 31). At ε = 4.0% we observe 1) the crack front has advanced by merging with damage cavities, 2) many more cavities ahead of the crack front, 3) crack branching, and 4) unbroken material regions inside the crack (white regions in the dark blue crack). With further increase in strain (ε = 4.5%), the crack front advances by coalescing with damage cavities while new cavities continually nucleate ahead of the crack front.

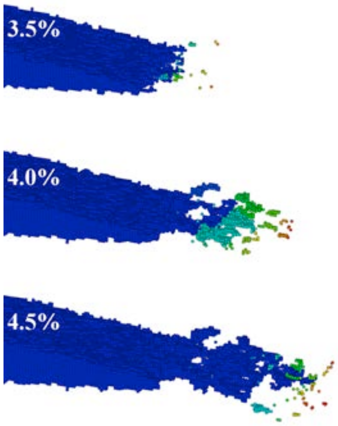

Fig. 31: Crack (dark blue) and damage cavities (red, green and light blue) at applied strains of 3.5%, 4.0% and 4.5%. Here the strain rate is 6×10^5 sec^{-1}.

We have examined the size distribution and morphologies of damage cavities. The average size of cavities is estimated from

$$R^2 = \frac{1}{N_s} \sum_{i=1}^{N_s} (\bar{r}_i' - \bar{r}_0)^2 , \qquad (2)$$

where $\bar{r}_0$ is the center of the cavity and $\{\bar{r}_i'\}$ denote the positions of the centers of voxels on the cavity surface. The surface morphology of cavities is obtained by calculating,

$$W = \left\{ N_s^{-1} \sum_{i=1}^{N_s} \left(\left| \bar{r}_i' - \bar{r}_0 \right| - R \right)^2 \right\}^{1/2} . \qquad (3)$$

The MD data show that cavity radii and interface widths scale as $V^{1/2}$, where V denotes the cavity volume.

Figure 32 illustrates the effect of strain rate on cracks and damage cavities. The snapshots of cracks and damage zones in Figs. 32(a) and 32(b) are taken from systems that are subjected to strain rates of 6×10^7 sec^{-1} and 6×10^5 sec^{-1}, respectively. The applied strain is the same in both systems (ε = 5%). The simulations reveal that the strain rate has a significant effect on the process zone but much less on crack growth. The slower strain rate generates a larger process zone with bigger cavities and much rougher crack front. We have also examined "temperature" distribution around cracks and damage cavities by calculating the kinetic energy of atoms in voxels on the crack front and cavity surfaces.

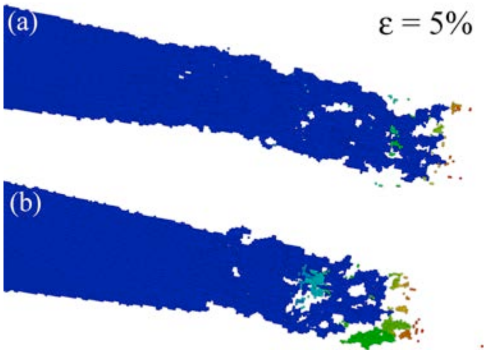

Fig. 32: Crack (dark blue) and damage cavities (light blue, green, red) at a strain of ε = 5% in MD simulations with strain rates of (a) 6×10^7 and (b) 6×10^5 sec^{-1}.

The "temperatures" around the crack and cavities are significantly higher at higher strain rates.

B. Damage and Cracking in Shearing Silica Glass

We have investigated shear-induced deformation and breakup of nanovoids with and without water in silica glass. MD simulations reveal that shear-induced void deformation, damage and flow in silica glass have the same underlying mechanism involving Si-O bond breaking and the migration of three-fold coordinated silicon and non-bridging oxygen defects on –Si-O-Si-O– rings. The density of these defects remains small and nearly constant in the elastic deformation regime. With the onset of plasticity at a shear strain of 15%, the defect density increases and reaches a plateau soon after a void breaks up. Despite the enormous differences in spatio-temporal scales, the observed shape changes and fragmentation of voids in silica glass are remarkably similar to the deformation and breakup of macroscopic inviscid drops at high shear rates.

Shear simulations in a-SiO_2 involve up to 633 million atoms. The initial diameter of the void ranges between 3 and 50nm, which covers nearly the entire range of void sizes observed in the damage zone in dynamic fracture simulations and quasi-static stress corrosion cracking experiments on silica glass.

Figure 33 shows the degree of deformation for a void in silica glass at room temperature. The initial diameter of the void is 10 nm. As time evolves, the void shape changes to ellipsoidal and damage in the form of nanocracks nucleates on the void surface at the two ends of the short axis of the ellipsoid; see the snapshot in Fig. 33(a). The same kind of deformation and damage are also observed for voids of initial diameters 3 nm and 50 nm. Stress calculations reveal that nanocracks form along the direction of maximum tensile stress.

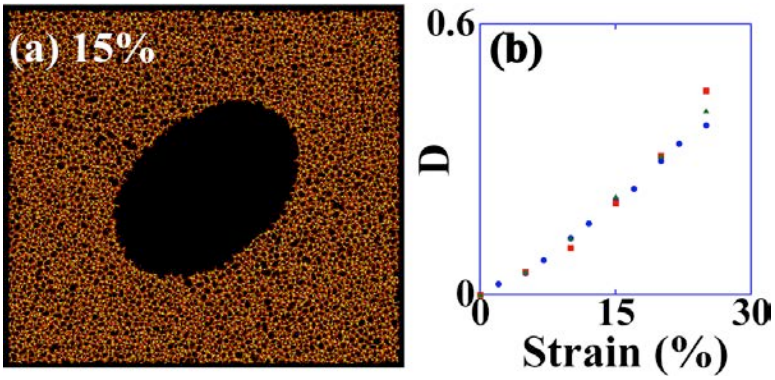

Fig. 33: (a) Shear deformation of a spherical void of diameter 10nm; (b) Strain variation of the deformation parameter, *D*, for voids of initial diameters 3nm (blue circles), 10nm (red squares), and 50nm (green triangles).

Figure 33(b) shows the time variation of the deformation parameter, *D*, for initially spherical voids of diameters 3 nm (blue), 10 nm (red) and 50 nm (green). The deformation parameter is defined as $D = (L-B)/(L+B)$, where L and B are the major and minor axes of the ellipsoid, respectively. In all three cases D increases linearly with strain up to 15%, which is an indication that the void deformation is elastic up to a shear strain of 15%. To distinguish between elastic and plastic void deformations, we switch off the shear strain rate at a certain value of the strain and let the system relax without shear. We find that if the shear strain rate is dropped to zero before the strain reaches 15% voids recover their initial spherical shapes and sizes. This confirms that the void deformation is elastic so long as the strain is below 15%.

Plastic deformation in shearing silica glass appears for ε > 15% as the deformation parameter begins to increase non-linearly with time. Figures 34(a) and (b) are snapshots showing long, thin voids with nanocracks on the surfaces near the middle of the long axis. Initially, these were spherical voids of diameters 3 nm and 50 nm. In both cases, the voids become threadlike before breakup, which occurs above a strain of 40%. Figure 34(c) shows plastic deformation of a void, which was initially a sphere of diameter 10 nm. The snapshot shows a long, thin void with ends deformed roughly like an *S* shape. There are nanocracks on the surface and near the ends of the void. As the strain increases, nanocracks grow and the void becomes more elongated, and at 40% the void fragments; see Fig. 34(d).

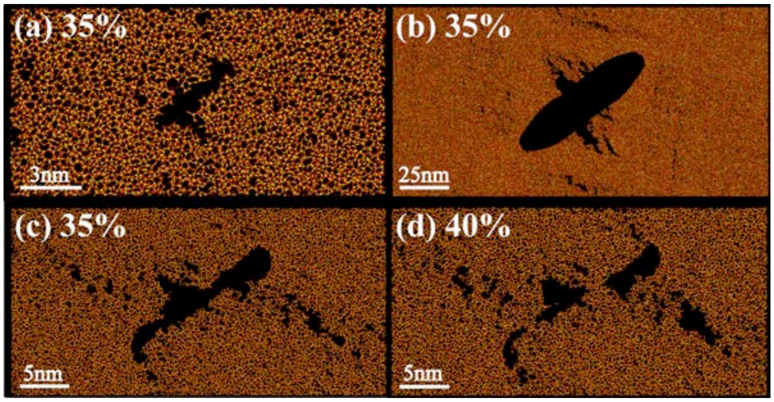

Fig. 34: Panels (a) and (b) show plastic deformations of voids of initial diameters 3 nm and 50 nm, respectively. Panel (c) shows a plastically deformed void which was initially a sphere of diameter 10nm; and panel (d) shows its breakup at $\varepsilon = 40\%$.

Detailed analyses of elastic-to-plastic void deformation and crack initiation and growth reveal a novel mechanism involving strain-enhanced defect transport; see Figs. 35(a) - (c). In the unstrained a-SiO_2, each Si atom (yellow) is connected to four O atoms (red) in the form of a SiO_4 tetrahedron and these tetrahedra are linked into nanometer size –Si-O-Si-O– rings through corner-sharing O atoms. In Fig. 35(a), green and blue regions represent 6- and 9-membered rings, respectively, at a strain of 5%. The magenta atom is a bridging O and the blue atom is a 3-fold coordinated Si atom in the 9-membered ring. Figure 35(b) shows that at a strain of 8% the magenta O atom becomes non-bridging while the blue Si atom remains under-coordinated. The grey ring with these Si and O defects is a 13-membered ring. At a strain of 10%, the green Si and magenta O atoms become fully coordinated by bonding with each other. The blue Si is still under-coordinated, but now belongs to a 14-membered ring (yellow). Non-bridging oxygen and three-fold coordinated silicon atoms play pivotal roles in the nucleation of cracks through the enlargement of –Si-O-Si-O– rings, and in void deformation and shear flow in silica glass.

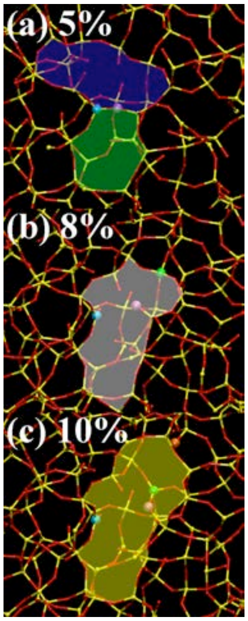

Fig. 35: Defects responsible for shear deformation and flow in silica glass.

We have also investigated changes in the void shape and elastic-to-plastic transition in silica glass at 1,200 K. In contrast to the case at room temperature, the voids are partially filled with atoms released from the void surfaces and the density of these atoms inside the voids increases with time. The voids still undergo spherical-to-ellipsoidal shape transformation, but the time variation of the deformation parameter is slightly less than that at room temperature. Nanocracks appear on void surfaces with the onset of plasticity, and voids deform into long ellipsoidal shapes with pointed ends and their breakup is again preceded by threadlike structures just as at room temperature.

C. Defect Migration and Recombination in Nanoindentation of Silica Glass

We have performed multimillion-atom MD simulations to investigate plasticity and flow mechanisms in nanoindentation of amorphous silica. The simulations reveal migration of non-bridging oxygens and their recombination with undercoordinated silicons in the densified plastic region under and the material pileup region around the indenter. The migration involves chains of defects switching bonds in string-like processes. These defects were postulated by Mott to explain shear flow in silica glass.

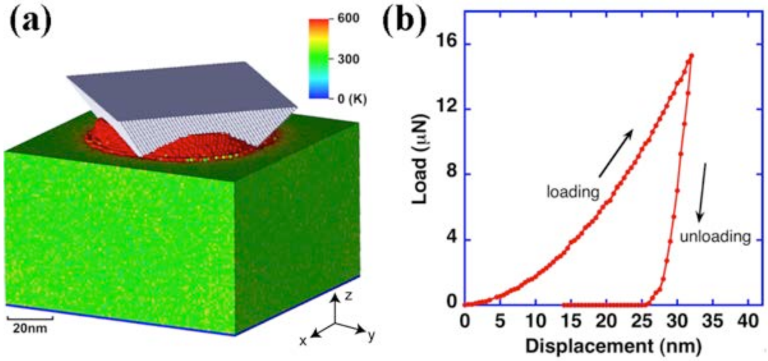

Fig. 36: (a) Pileup around the indenter at the maximum indentation depth. Indentation causes localized damage creating high-temperature area underneath the indenter, which facilitates pile-up through plastic flow. The colors represent temperature. (b) Load versus displacement during loading and unloading phases of indentation.

Figure 36(a) shows a snapshot of the system at the maximum indentation depth. The amorphous film contains 40 million atoms and has dimensions 100×100×60 nm^3. The indenter, a square-based rigid pyramid with an apical angle of 90°, interacts with the substrate through a repulsive potential. We have used two kinds of indenters to study the effect of the indenter curvature: an atomistically sharp indenter and an indenter whose end is spherical with a diameter of 10 nm. Figure 36(b) displays the applied load in the MD simulation as a function of the penetration depth of the indenter. From the load and residual indentation impression, we obtain the hardness, H, of a-SiO_2. The hardness values are found to be 10.6 GPa and 12.0 GPa for the atomistically sharp and blunt indenters, respectively. These results compare favorably with the AFM measurement of hardness (10 GPa) in fused silica.

Figure 37 shows snapshots of the indented silica film at indentations depths h = 11 nm (panel (a)) and 32 nm (panel (b)). For clarity, only a slice of the system without the indenter is shown here. At low indentation depths (*e.g.*, h = 11 nm), the material close to the bottom of the indenter densifies but there is no material pile-up around the indenter. At the maximum indentation depth of 32nm, we observe significant pileup and a 20% increase in the mass density of a-SiO_2 beneath the indenter relative to the mass density of bulk silica glass under ambient conditions (2.2 g/cc). In the pileup region shown in panel (b), the density is 18%

lower (1.8 g/cc) than in the bulk system. During the unloading phase of the simulation, the silica system exhibits slight elastic recovery and a decrease in the density under the indenter. Panel (c) is a snapshot of a slice of the system after the indenter is pulled out of the substrate. Here the density just beneath the indenter goes back to the normal bulk a-SiO_2 density (2.2 g/cc), but there is still a considerable pileup around the indenter.

Detailed analyses of atomistic configurations reveal how plastic deformation and pileup occur in the nanoindentation of a-SiO_2. Under-coordinated Si and non-bridging O point defects in –Si-O-Si-O– rings play a central role in the plastic deformation and pileup observed in both nanoindentation simulations. The defects are involved in bond-switching events, where either silicon or oxygen severs its bond with one of its neighbors and shortly thereafter (~ 10ps in simulation 1) forms a new bond with a different neighbor atom. Figures 38 (a)-(c) show how a bond-switching event occurs in the case of a non-bridging oxygen atom. The three atoms involved in this event are labeled as O1 (initial oxygen defect), Si2 (bond switching silicon) and O3 (the final oxygen defect). Initially, O1 is a non-bridging oxygen whereas Si2 and O3 have normal coordinations (4 O neighbors for Si2 and 2 Si neighbors for O3). At the transition state, Si2-O3 bond breaks and these two atoms become undercoordinated. Subsequently, Si2 forms a new bond with O1 and they both become fully coordinated whereas the oxygen labeled O3 remains a defect.

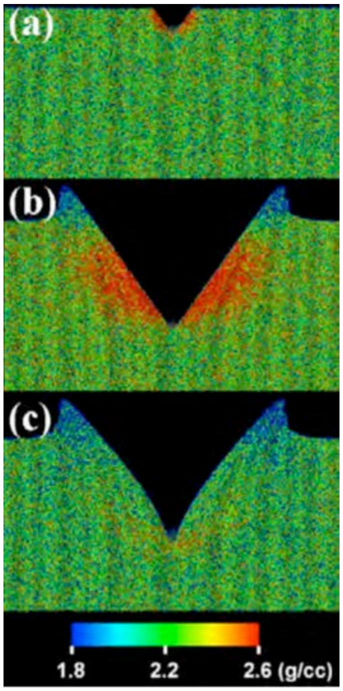

Fig. 37: Density distribution before pileup (**a**), at the maximum indentation-depth (**b**), and after unloading (**c**).

Another bond-switching event we have observed involves the annihilation of a non-bonded pair of undercoordinated silicon and oxygen atoms. Figures 38 (d)-(f) show this pair-annihilation event. The four atoms involved in this event are labeled as Si1 (an undercoordinated silicon), O2 (bond-switching oxygen), Si3 (bond-switching silicon) and O4 (a non-bridging oxygen). Initially, O2 is bonded with Si3 and they both have normal coordinations of 2 and 4, respectively. In the transition state, O2-Si3 bond breaks and they both become undercoordinated. In the final state, Si1 forms a bond with O2 and Si3

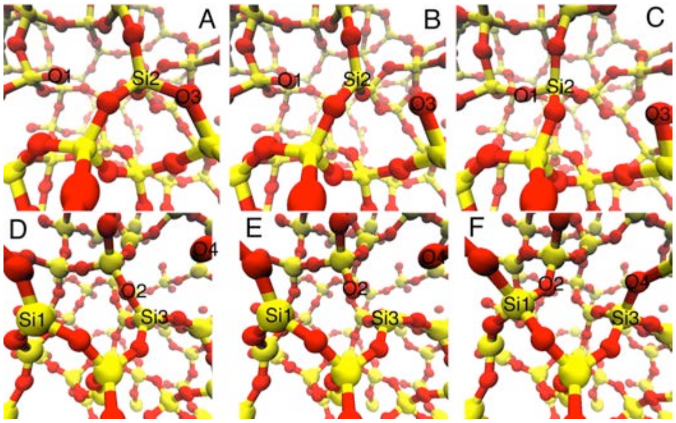

Fig. 38: (a)-(c): Atomic configurations showing transport of oxygen defects; (a) the initial state, (b) the transition state, and (c) the final state. The oxygen defect migrates by switching covalent bond from Si2-O3 to O1-Si2. (d)-(f): Atomic configurations showing an annihilation event for a pair of defect atoms; (d) the initial state, (e) the transition state, and (f) the final state.

with O4 and now these four atoms have normal coordinations. Our simulations also reveal a defect transport mechanism in which defects migrate a considerable distance via a chain of bond-switching events.

D. Hybrid Atomistic/Mesoscale Simulations of Fluid Flow through Cracks

We have designed an algorithm to interface MD with a mesoscale approach to fluid flow simulation, i.e., the lattice Boltzmann (LB) method which solves a minimal form of the Boltzmann kinetic equation ignoring all details of molecular motion except those needed to recover hydrodynamic behavior at the macroscopic scale. For hierarchical fluid flow simulations, we have developed an extensible LB computational framework for tera-to-petascale parallel/distributed platforms. The LB framework includes scalable algorithms and tools as well as capabilities for handling, analyzing, and visualizing petabyte simulation datasets. It allows collaborative construction and execution of complex, multi-component, computationally demanding hybrid numerical simulations of fluid flow. We have achieved scalability beyond 10^5 processors through linear-scaling algorithms and performance-optimization techniques, and performed fluid-flow simulations through a network of cracks in amorphous silica on a cluster of PlayStation3 consoles (Fig. 3).

Currently, we are developing automated model transitioning to embed higher fidelity MD simulations concurrently inside LB simulations only when and where they are required (Fig. 39). We are using a domain decomposition algorithm based on Schwartz alternating method to implement hybrid MD-LB simulations. In this scheme, the system is decomposed into 1) a purely atomistic region described by MD, 2) a continuum region described solely by LB, and 3) an overlapping MD-LB region. Since the MD region is open and particles move out of the MD into the LB region, it is necessary to design an efficient algorithm to insert particles back into the MD region to keep the average particle density constant. This is a challenging task, especially in the case of a dense fluid.

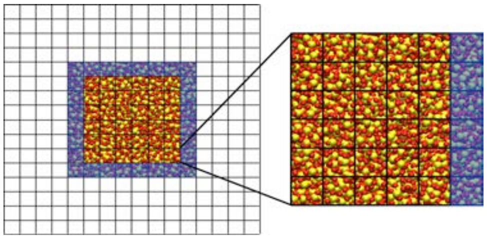

Fig. 39: Domain decomposition for the hybrid MD/LB simulation, where squares are LB cells, yellow and red spheres are atoms in MD, and the shaded region represents the Schwartz alternating MD/LB coupling.

We are using a highly efficient algorithm called USHER to introduce particles effectively even in a dense fluid. The attractive features of USHER are: 1) If the initial random position of the inserted particle is too close to an existing particle, USHER solves this problem in one update; 2) the displacement in each step is bounded by an upper limit; and 3) the final energy of the inserted particle is close to the average energy per particle in the system.

§3.4 AMORPHIZATION-INDUCED INTERGRANULAR SULFUR EMBRITTLEMENT IN NICKEL

Chemistry of a small amount of impurities segregated to grain boundaries (GBs) essentially controls the mechanical properties of materials. A prime example of such GB mechanochemistry is sulfur segregation-induced embrittlement of nickel. Fundamental understanding and control of embrittlement mechanisms in the Ni-S system are important not only scientifically but also technologically, *e.g.*, for the development of next-generation nuclear reactors. Despite decades of intense experimental and theoretical efforts, however, mechanisms of S-induced fracture phenomena are not well understood. Experiments by Heuer *et al.* have provided key insights into the problem. They found the existence of a critical S concentration at the Ni GB that causes a transition from ductile, mixed transgranular and intergranular fracture to brittle, intergranular fracture. Furthermore, the critical concentration for GB embrittlement was found to coincide with that for the amorphization of Ni single crystal under ion implantation. A theoretical framework by Rice and Wang explains GB embrittlement from the energetics of solute segregation. For example, Yamaguchi *et al.* performed first-principles quantum-mechanical (QM) calculations of the decohesion of Ni GB with 1-2 monolayers of segregated S atoms, and proposed S-induced reduction of GB tensile strength as an embrittlement mechanism. However, this mechanism does not explain the experimentally found relation between embrittlement and amorphization.

Understanding amorphization-induced intergranular S embrittlement in Ni poses a so called multiscale simulation challenge, *i.e.*, coupling QM accuracy to describe solute chemistry with large length scales to incorporate long-range stress fields and microstructures such as grains and amorphous GB phases. Recent developments in chemically reactive atomistic simulation methods and parallel computing technologies that are scalable over 10^5 processors have set the stage to address this challenge. Here, we employ a multiscale simulation approach that involves 48 million-atom molecular dynamics (MD) simulations based on reactive force fields (ReaxFF), which are trained and validated by first-principles QM calculations based on the density functional theory (DFT), to study the effect of S segregation on the fracture of nanocrystalline Ni. Our simulations reveal a direct link between S-induced GB amorphization and embrittlement: We have found an order-of-magnitude reduction of GB shear strength due to amorphization, which, combined with tensile-strength reduction, always provides an easy cleavage path at the crack tip. This amorphization-induced mechanism also explains an experimentally observed crossover from transgranular to intergranular fracture as well as suppression of plastic activities. The computation has taken 10 days on 65,536 IBM Blue Gene/P processors at the Argonne National Laboratory.

A. Validation through Tensile Tests

To validate the ReaxFF for Ni-S system, we performed a tensile test of a $\sum 5(012)$ symmetric tilt grain boundary (STGB) of face-centered cube (fcc) Ni crystal. Sulfur atoms were inserted in the vacancy sites on the GB plane as illustrated in Fig. 40(A). All available empty sites on GB were filled with S atoms. The system consists of 52,800 atoms of dimensions 173.16 Å × 157.72 Å × 21.12 Å in the x (parallel to the $[02\bar{1}]$ axis of the fcc crystal), y ([012]) and z ([100]) directions, respectively. After fully relaxing the GB structure, the tensile test was performed as follows: The upper grain was displaced upward instantaneously without relaxation, and we calculated the system energy at each displacement d. Numerical derivative of the energy-displacement relation with respect to *d* provided tensile stress as a function of grain separation. Figure 40(B) compares the tensile test results for Ni GB with and without monolayer of segregated S atoms. The figure exhibits the reduction of GB cohesive strength by 3.5 GPa due to S segregation. To further validate the ReaxFF results, we also performed a tensile test for a smaller system consisting of 40 atoms. The DFT result, 4.6 GPa, for the S-induced reduction of GB cohesive strength agrees well with the ReaxFF result.

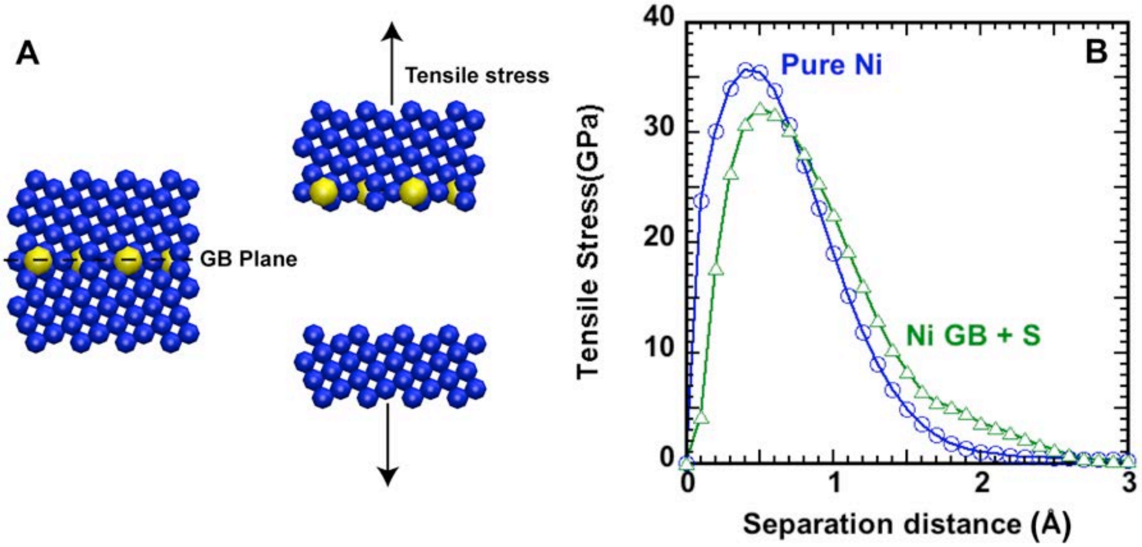

Fig. 40: **(A)** Schematics of the tensile test. **(B)** Stress-displacement curve.

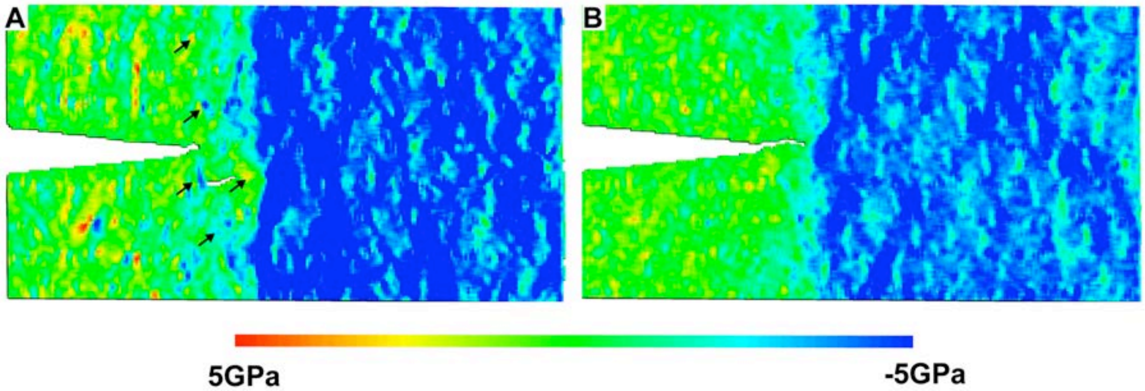

Fig. 41: Stress distribution during fracture simulations of nanocrystalline Ni (A) without and (B) with (B) amorphous sulfide GB phases.

Ductility of nanocrystalline Ni (nc-Ni) is also manifested as inhomogeneous, localized stress concentrations away from the crack tip. Figure 41 shows stress distribution in the direction of tension (negative value is tensile, and positive is compressive). Arrows in Fig. 41(A) (nc-Ni) points localized stress concentrations, whereas stress distribution with sulfide GB phases (nc-Ni+20%S) in Fig. 41(B) is more homogeneous.

B. Amorphization and Intergranular Sulfur Embrittlement

To understand amorphization-induced intergranular S embrittlement, we first investigate the effect of S doping on the structure of Ni crystal. Here, we perform ReaxFF MD simulation, validated by DFT calculations, of face-centered cubic (fcc) Ni crystal consisting of 4,000 atoms in a cube of side 35.2 Å with periodic boundary conditions, where S atoms are randomly substituted for a fraction of Ni atoms in the range of 0-24%. For each S concentration, the system is relaxed by a steepest-descent procedure and is then relaxed for 5 ps at 10 K, before it is gradually heated up to 300 K. Figure 42(A) compares the radial distribution functions, $g(r)$, for 0% and 24% S substitutions. We observe broadening of the peaks, representing disordering of the fcc structure due to S doping. The full-width at half-maximum of the first peak of $g(r)$ as a function of S concentration in Fig. 42(B) exhibits a sudden increase at 16%, which is close to the experimental threshold (14.2±3.3%) for amorphization.

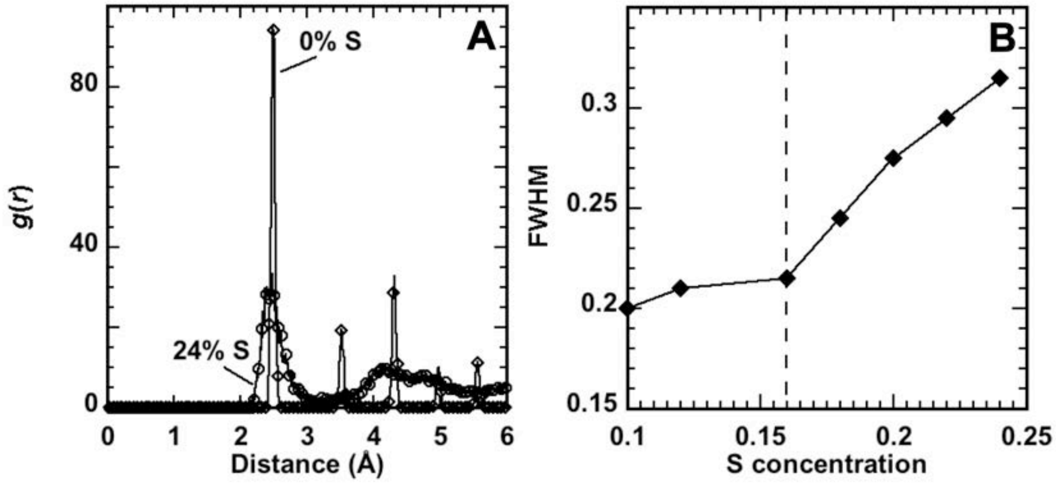

Fig. 42: **(A)** Radial distribution function $g(r)$ of Ni fcc crystal with 0 and 24% S doping. **(B)** Full-width at half-maximum (FWHM) of the first peak of $g(r)$ as a function of S concentration.

To study the effect of S-induced GB amorphization on the fracture behavior of Ni, we have performed ReaxFF MD simulations for two systems: Pure Ni nanocrystal (nc-Ni) and Ni nanocrystal with 20% S-doped GB layers of thickness 5 Å (nc-Ni+20%S). The two nanocrystalline configurations consisting of randomly oriented grains of average size 7.5 nm (in total of 2,048 grains) are generated using the Voronoi tessellation method with the same grain geometry. Samples with two different grain geometries have been made to provide statistics. Each sample consists of 48,393,282 atoms, with dimensions 468.0 nm × 234.1 nm × 4.9 nm (nc-Ni) or 470.4 nm × 235.2 nm × 4.9 nm (nc-Ni+20%S) in the x, y and z directions.

We first relax the atomic configurations and the simulation box size with periodic boundary conditions over 50 ps until the pressure reaches zero, *i.e.*, 2.5% (nc-Ni) or 2.0% (nc-Ni+20%S) compression, followed by further relaxation for 50 ps at 10 K, and subsequent heating up to 300 K. As expected, with 20% S doping, GB sulfide layers amorphize during the sample preparation. A notch of length 150 nm is inserted in the xz plane of each sample, which is then fully relaxed again by removing the periodic boundary condition in the *x* direction (Fig. 43(A)). Fracture simulations are performed by applying uniaxial tension along the *y* direction at a strain rate of 4×10^{8} sec^{-1}. The tensile strain is applied by changing the simulation box length in the *y* direction, while periodic boundary conditions are maintained in both *y* and *z* directions, until the system fragments into two separate pieces (at strain of 12.3% for nc-Ni and 6.4% for nc-Ni+20%S, respectively).

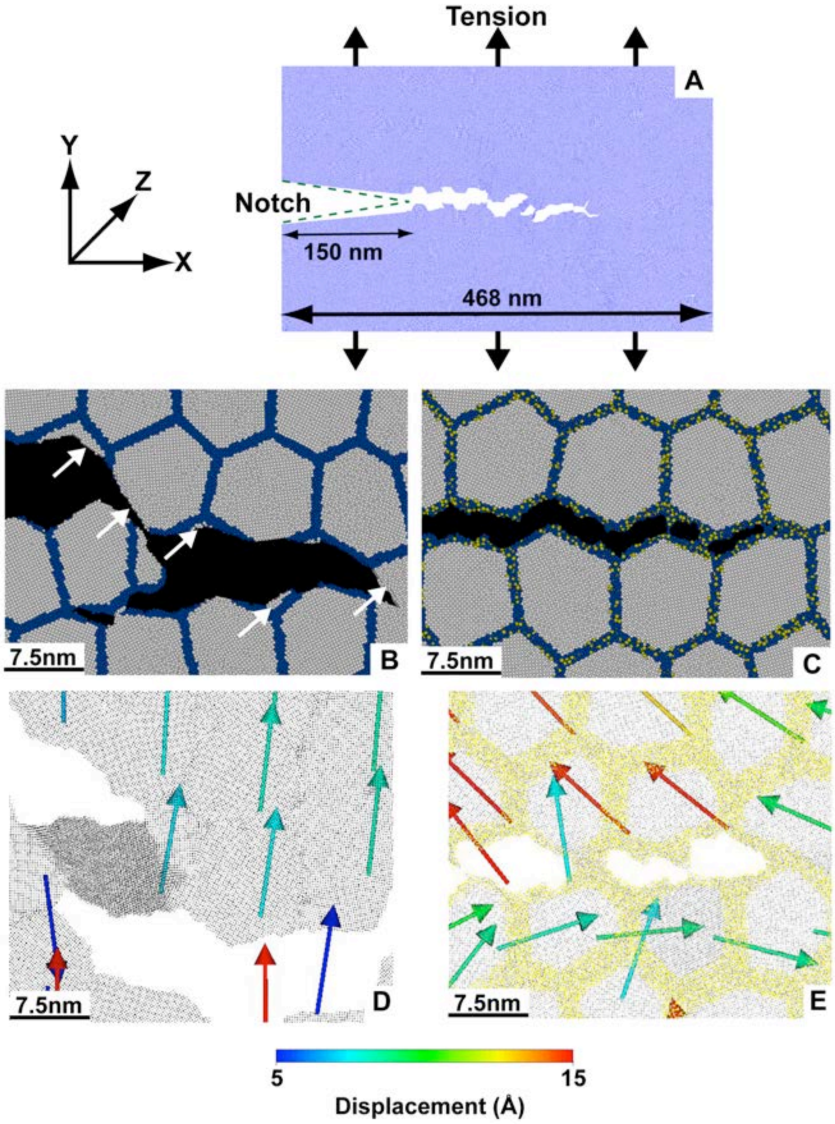

Fig. 43: **(A)** A snapshot of simulation, where pre-notched nanocrystalline Ni is fractured by applying uniaxial tension. **(B)** and **(C)** are close-ups of fracture simulations for nc-Ni and nc-Ni+20%S, respectively, where gray, blue and yellow colors represent Ni atoms inside grains (>5 Å from GBs), Ni atoms within 5 Å from GBs, and S atoms, respectively. White arrows in **(B)** point to transgranular fracture surfaces. **(D)** and **(E)** are close-ups showing displacement vectors of grains in nc-Ni and nc-Ni+20%S, respectively.

Simulation results show a crossover from mixed transgranular and intergranular fracture in nc-Ni (Fig. 43(B)) to purely intergraunlar fracture in nc-Ni+20%S (Fig. 43(C)). In nc-Ni, the crack propagates through both grains (transgranular fracture indicated by white arrows in Fig. 43(B)) and GBs (intergranular fracture), resulting in 56% intergranular fracture. Common ductile fracture behaviors such as crack-tip blunting, void formation in front of the crack tip and its coalescence with the main crack, and necking are observed in nc-Ni. In contrast, the crack in nc-Ni+20%S propagates in a cleavage manner only through GBs, resulting in 100% intergranular fracture. The transgranular-to-intergranular fracture transition with S segregation is in agreement with experiments.

Our simulations also reveal another type of crossover of fracture behavior. Namely, fracture occurs predominantly in opening mode (or mode I) in nc-Ni, whereas there is mixture of sliding mode (or mode II) in nc-Ni+20%S. To quantify this transition, Figs. 43(D) and (E) use an arrow to represent a displacement vector of each grain from its position at 0.2% smaller strain. Grain displacements in nc-Ni are perpendicular to the crack surfaces, indicating pure opening mode of fracture, whereas displacements in nc-Ni+20%S have tangential components to the surface, indicating the mixture of sliding mode.

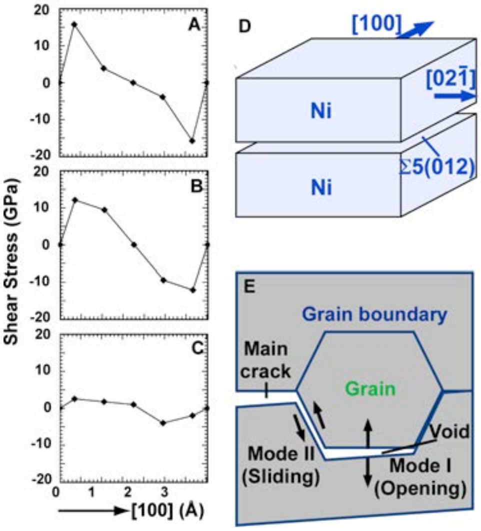

Fig. 44: Shear stress calculated from the derivative of the GSFE with respect to displacement along the [100] direction without S **(A)**, with a monolayer of segregated S **(B)**, and with amorphous sulfide GB phases **(C)**. **(D)** Schematic showing sliding one grain with respect to the other to calculate GSFE. **(E)** Schematic of opening (mode-I) and sliding (mode-II) fractures in front of the main crack.

A clue for additional embrittlement mechanism lies in the crossover of fracture modes. The mode II (sliding-mode) cracks in nc-Ni+20%S (Fig. 44(E)) imply the existence of shear stress acting on them. To study the response of GB to shear, we calculate a generalized stacking fault energy (GSFE), i.e., the energy required to slide Ni Σ5(012) GB, without S segregation (Fig. 44(B)), with a monolayer of segregated S atoms (Fig. 44(C)), and with 5 Å-thick amorphous sulfide ($Ni_{0.8}S_{0.2}$) GB layer (Fig. 44(D)). For each case, the GSFE is calculated by displacing the upper grain with respect to the lower grain, while atoms are allowed to relax only in the vertical direction (Fig. 44(D)). In Figs. 44 (A)-(C), shear stress is calculated from the derivative of the three GSFEs (Figs. 44(B)-(D)) with respect to displacement along the [100] direction. The GB shear strength (*i.e.*, the maximum shear stress) drastically decreases from 17.3 GPa (without S, in Fig. 44(A)) to 2.5 GPa (with amorphous sulfide GB phase, in Fig. 44(C)), due to GB amorphization. In contrast, S monolayer does not decrease the GB shear strength as much (Fig. 44(B)). This amorphization-induced GB shear-strength reduction provides a hitherto missing link between GB amorphization and embrittlement. As Fig. 44(E) illustrates, the S amorphization-induced shear-strength reduction provides mode-II cleavage paths in front of the crack tip, while the S-induced tensile-strength reduction as in the Rice-Wang model provides mode-I paths. Together, the two mechanisms always provide weak cleavage paths right in front of the crack tip, obviating void opening, which is a signature of ductile fracture in pure Ni (Figs. 44(B) and (D)), and making the fracture with S segregation purely intergranular and brittle.

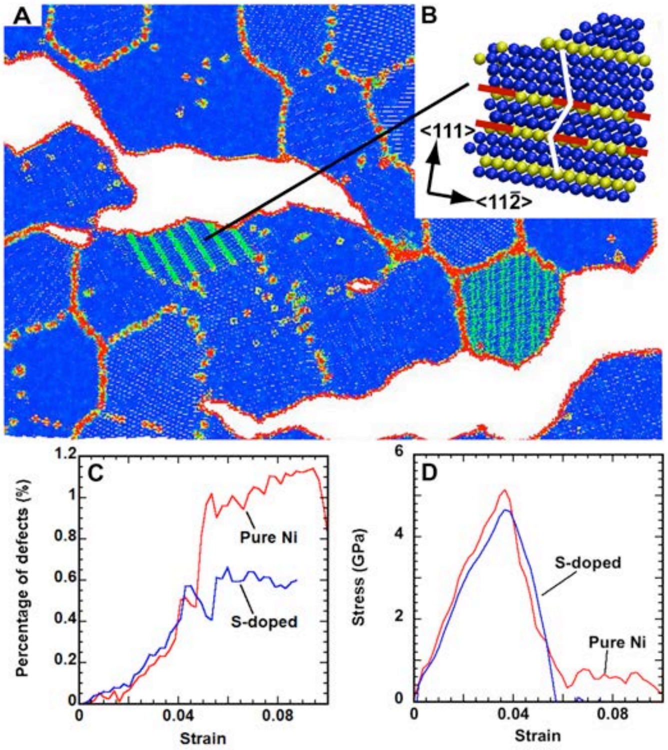

Fig. 45: **(A)** Close-up of fractured nc-Ni, where the atoms are color-coded according to the common neighbor parameter. Blue color represents atoms with perfect fcc structure, colors other than blue represent atoms with imperfect fcc structure. Green stripes represents twinning boundaries. **(B)** The inset magnifies a deformation twinning in one of the grains, where yellow color denotes the twinning boundary, blue and yellow colors denote perfect atoms and defect atoms, respectively. **(C)** Number of intragrain defect atoms for nc-Ni (red) and nc-Ni+20%S (blue). **(D)** Stress-strain relation ahead of the notch for nc-Ni (red) and nc-Ni+20%S (blue).

We have also found that amorphization reduces the GB shear strength comparable to the critical stress to move dislocations (*i.e.*, Peierls stress) in Ni crystal. Consequently, plastic activities within grains are suppressed. We use the common neighbor parameter (CNP) to identify topological defects such as dislocations and twins during the fracture simulations in nc-Ni and nc-Ni-20%S. In nc-Ni, dislocations formed at a GB (dislocations source) are found to be absorbed by another GB (dislocations sink). Figure 45(A) is a snapshot (color-coded by CNP) of ductile fracture in nc-Ni, where the green stripes show grains that have deformation twinning. The atomic configuration of the twinning structure shown in Fig. 45(B) is similar to those observed experimentally. Such plastic activities inside grains are largely suppressed in nc-Ni-20%S, where plastic deformations are confined in soft amorphous GB phases. The number of intragrain defect atoms (as identified by the CNP) is plotted as a function of time in

Fig. 45(C) to provide an evidence of the amorphous layer confined plasticity. The large number of defects in nc-Ni represents massive intragrain slip activities, where the ups and downs of the curve correspond to the development of dislocations and their absorption at GBs, respectively. In contrast, limited plasticity inside the grains for nc-Ni-20%S is reflected in the small number of defects in Fig. 45(C).

Another signature of embrittlement is the reduction of toughness, which is a consequence of reduced plastic activities. Figure 45(D) compares the stress-strain relations calculated ahead of the notch for nc-Ni and nc-Ni+20%S. The toughness (the area under the stress-strain curve) is 0.101 GJ/m^3 and 0.0898 GJ/m^3 for nc-Ni and nc-Ni+20%S, respectively. Furthermore, the stress-strain curve for nc-Ni+20%S exhibits a characteristic of brittle fracture, *i.e.*, a sudden drop of the stress to zero at a critical strain. In contrast, the stress-strain curve of nc-Ni retains nonzero stress up to much larger strains—a characteristic ductile behavior. In addition, stress distribution exhibits a marked difference between nc-Ni and nc-Ni+20%S. In experiments, much larger time scales compared to our simulations lead to more extensive plastic activities, resulting in macroscopic elongation in pure Ni and more pronounced drop of toughness value due to S segregation than those reported here.

The above simulation results reveal a fundamental role of GB chemistry, *i.e.*, the formation of amorphous GB phases above a critical concentration (16%) of segregated S in determining fracture behaviors of Ni. Deformation and failure are confined in the soft amorphous sulfide GB phases with an order-of-magnitude smaller shear strength, thereby leaving the interior of grains intact. This mechanism explains all experimental observations. Similar roles of soft amorphous intergranular phases for unique mechanical properties are found in other systems such as superhard nanocrystalline SiC.

§3.5 NANODUCTILITY IN HIGH STRAIN-RATE DEFORMATION AND FRACTURE OF CERAMICS

Advances in experimental methods, such as *in situ* micro compression and atomic force microscopy, have provided new insights into the role of nanoplasticity in known brittle materials. In contrast with bulk materials, experiments have shown surprising existence of ductility at the nanoscale, *e.g.* metallic glass nanopillars may behave similar to ductile metals,[2] and nanovoid coalescence was shown to play an important role in the ductile fracture mechanism in silica. Similarly, nanoductility has been shown to play a crucial role in the damage initiation in high-strength ceramics under high-velocity impact. Nanoplasticity could be a key ingredient to understand and improve high performance ceramics such as AlN, Al_2O_3, SiC and their composites that are known for their outstanding mechanical properties and make these materials ideal for structural applications involving high temperature and pressure. It is essential to clarify the role of nanoplasticity and to evaluate intrinsic mechanisms of damage leading to fracture at extreme conditions of temperature and pressure. Shock induced damage by projectile impact is an essential tool in this regard, since it provides an ideal environment of extreme conditions of high gradients of pressure, temperature and density.

We performed a 302 million-atom MD simulation of projectile impact on SiC. To study material properties and processes encompassing diverse spatiotemporal scales, it is essential to develop accurate interatomic potentials reflecting first-principles information. We developed a many-body interatomic potential for SiC. In the simulations a 3C-SiC target is impacted at 15 km/s and the atomic trajectories are calculated as a function of time allowing an accurate description of the spontaneous shock-induced damage dynamics. The dimensions of the zinc blende single crystal 3C-SiC target are 160 × 158 × 124 nm^3 in the *x*, *y* and *z* directions, respectively. The *z*-axis, parallel to the [111] orientation of the zinc blende crystal, is chosen as the impact direction. Periodic boundary conditions are imposed in the *x* and *y* directions, parallel to the $[\bar{2}11]$ and $[0\bar{1}1]$ directions. In the *z* direction the target has two free surfaces; see Fig. 1(a). The projectile is a hexagonal cylinder consisting of 563,000 atoms, which is 15 nm wide and 30 nm long. The projectile is five times stiffer and its atoms are five times heavier than those of the target. The interparticle potential between the target and projectile atoms is purely repulsive. A Langevin damping of the shock waves is applied to the *x* and *y* boundary conditions to avoid reflection of the waves. Defects in the system are dynamically data mined during the whole simulation using shortest path rings selection rules in a highly efficient implementation. Local stresses are calculated using the atomic Virial averaged spherically around each atom position.

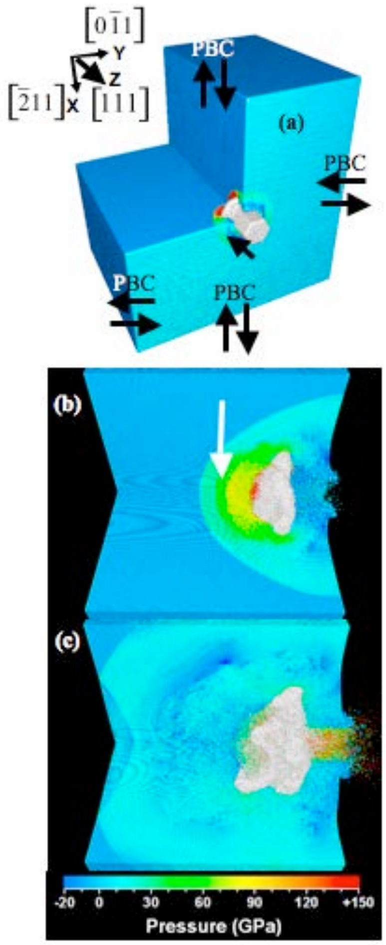

Fig. 46: Shock wave evolution during the impact of a 0.56 million atom projectile on a 302 million-atom silicon carbide ceramic target (160 x 158 x 124 nm^3) at the speed of 15 km/s. **(a)** Shock wave generation at 1.05 ps. Projectile atoms are shown in grey and target atoms are color-coded according to the local pressure. Only a 3/4 cut of the target is shown to better visualize the internal pressure distribution. **(b)** Splitting of the initial single shock wave (indicated by the white arrow) into an elastic wave (in light green) and a structural phase transformation wave (in yellow/dark green) at 5.17 ps. **(c)** SiC target at 10.35 ps showing the reflection of the elastic compressive wave from the back surface generating a tensile release wave into the target.

During the first picosecond of penetration the projectile releases a large amount of kinetic energy into the target (Fig. 46(a)), which results in a high local energy density. The large gradients of pressure and temperature causes localized melting and vaporization, and creates a strong shock wave with peak pressure of 150 GPa and speed of up to 24 km/s. During the shock propagation into the target, wherever the pressure of the shock front reaches a critical value of 89 GPa, the material undergoes a structural phase transformation (SPT) from the 4-fold-coordinated zinc blende structure to the 6-fold-coordinated rock salt structure. This SPT accompanies the shock wave as long as the pressure is above the critical value. In about 7.5 ps, the transformation wave slows down significantly and the SPT stops when the local pressure drops below the threshold value of 89 GPa. The elastic component of the shock wave continues to travel through the system at an average speed of 12.5 km/s, which is close to the longitudinal sound speeds of 3C-SiC. The shock wave split is evident in Fig. 46(b) where it is indicated by the white arrow. The elastic wave continues to propagate until it reaches the back surface and generates a release wave back into the system; see Fig. 46(c).

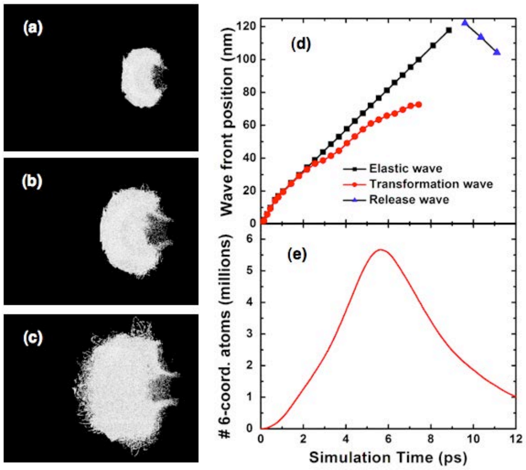

Fig. 47: Shock-induced structural phase transformation in the silicon carbide target. (a) Cluster of 6 coordinated atoms at 3 ps. To better visualize the transformation dynamics, only atoms with coordination of 6 in the target are shown in panels (a)-(c). (b) Six-fold coordinated cluster at 5.5 ps and (c) at 9 ps. (d) Evolution of the shock wave front position with simulation time. The curves show the split of the initial single shock wave into one elastic component traveling with constant speed and a structural transformation wave (STW) travelling with decreasing speed. The STW transforms the SiC 4 coordinated zinc blende structure to the more compact 6-coordinated rocksalt structure. As the shock compression is attenuated, the structural transformation wave stalls and the pressure-induced transformation is reverted. (e) Evolution of the number of 6 coordinated atoms with time. At 5.5 ps the forward structural transformation reaches a maximum and is reversed, the number of 6 coordinated atoms then drops continuously with the decreasing shock pressure. Following the curve in (e) the configurations illustrated in (a), (b), and (c) show the 6 coordinated cluster during the forward transformation and increasing number of atoms (a), with the maximum number of atoms (b), and during the reverse transformation and decreasing number of atoms (c).

Figures 47(a)-(c) show that the SPT wave has an almost spherical shape with a rough surface formed by the nucleation of rock salt clusters with a heterogeneous grain structure of different sizes and orientations. At about 9 ps the decreasing stress at the shock front decrease the rate of transformation and the SPT shock front starts to become rough and undefined as shown in Fig. 47(c). Figure 47(d) shows that from 5.5 ps the split in the shock fronts increases until the SPT wave virtually stops. The evolution curve of the number of 6 coordinated atoms shown in Fig. 47(e) indicates a maximum at 5.5 ps. After that the reverse transformation occurs at a fast rate as the local pressure drops with the dissipation of the shock wave energy. With the attenuation of the shock pressure, the SPT stalls and the generation of plastic deformation becomes the dominant mechanism of damage. In the central region surrounding the impact path high levels of shear stress are generated and they are released by the generation of dislocations. The results indicate that initially the plastic zone is rooted at the zinc blende-rock salt interface, which is a favorable spot for stress concentration and defect nucleation. When the hydrostatic pressure is not high enough to drive the SPT, the interface becomes rough and the high local shear stress causes spontaneous emission of dislocations into the zinc blende phase. Figure 48(a) illustrates the emission of dislocations from this region at 11.1 ps.

The zinc blende cubic structure of SiC has the {111} family of planes as the preferred, lowest energy, set of planes for dislocation gliding. The distribution of high levels of shear stress around the shock front generates dislocation lines on all {111} planes. An estimation of the shear stress levels in the central region along the impact *z* direction by $\tau = 0.5[\sigma_{zz} - 0.5(\sigma_{xx} + \sigma_{yy})]$, gives values up to 45 GPa. As no plane in the {111} set is aligned with the *z* direction the generated dislocations glide away from the central region to adjacent regions of the system while propagating to the back surface; see Figs. 48(a)-(c). The high shear stress and the rough surface of the interface of rock salt and zinc blende structures favor the emission of a large number of dislocations, forming a complex configuration. As dislocations are emitted they propagate until they reach the back surface, leaving surface steps; see Figs. 48(c) and (d).

Dislocations show rich dynamics while propagating to the back surface of the target. The flow of atoms in the central region of the target creates shear stress that favors the glide of dislocations in the *z* direction to the back surface of the system. As this shear stress is not aligned with {111} planes, several dislocation dynamic processes have a propitious environment, such as dislocations climb, cross slip, and multiplication of dislocations. All these processes are active and enhanced by the overlap of compressive and tensile shock waves. Dissociation, multiplication and cross slip of dislocations are illustrated in Fig. 48(b). In the process of cross slip the results show the formation of dislocation locks and structural debris; see Fig 48(d).

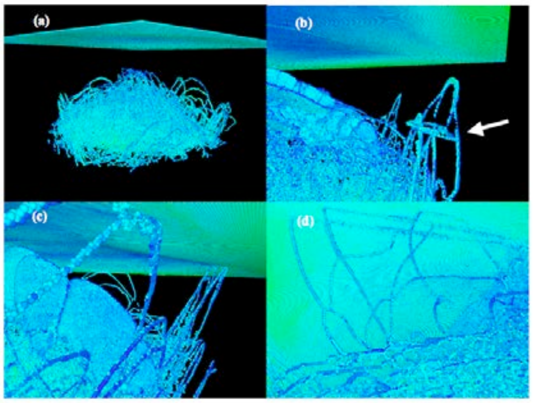

Fig. 48: Shock-induced plastic structures in SiC. For better visualization only atoms in the left half system of Fig. 1(b)-1(c) are shown here, with the back target surface on top. Also, only those atoms that are part of defect structures, i.e., dislocations, system surfaces and crack surfaces, selected on the basis of the number of shortest path rings are shown. The selection criterion is to show all atoms with the number of shortest path rings different than 12. **(a)** Snapshot at 11.1 ps shows extensive dislocation loops emitted from the interface of the transformed and untransformed phases growing into the zinc blende crystal. **(b)** Detail of a dislocation dissociation and multiplication, indicated by the white arrow, at 14.4 ps. **(c)** Dislocation lines propagating in {111} slip planes reaching the back surface at 17.7 ps **(d)** Snapshot of dislocation structures connected with the back surface at 31.2 ps. Dislocations here release the shear stress which is produced by the shock wave during the impact. (**a**) Shows dislocations structures generated around the whole damaged region surrounding the structural transformation cluster shown in Fig. 2**a**-2**c**. (**b**)-(**c**) however show that dislocations develop and reach the back surface only in the adjacent regions of the system where the shear stress is maximum. In the central regions of the system brittle cracks nucleate and grow releasing the tensile stress of the release wave.

When the compressive shock wave reaches the back surface, it generates a tensile release wave back into the target, which reduces the local hydrostatic pressure to negative values, locally reaching up to -20 GPa. That causes the spontaneous emission of cracks from the plastic deformed region. Our analyses indicate that when the release wave reaches the dislocation structures, as illustrated in Fig. 48(a), it causes the immediate emission of micro cracks from dislocation line cores; see Fig. 49(a). The process occurs in all {110} planes, which are the lowest surface energy, preferred planes for crack cleavage. ***This process of micro crack emission from single dislocation cores is one of the most striking features of this work, since it is the first time that it has been numerically demonstrated as a possible mode of micro crack nucleation***.

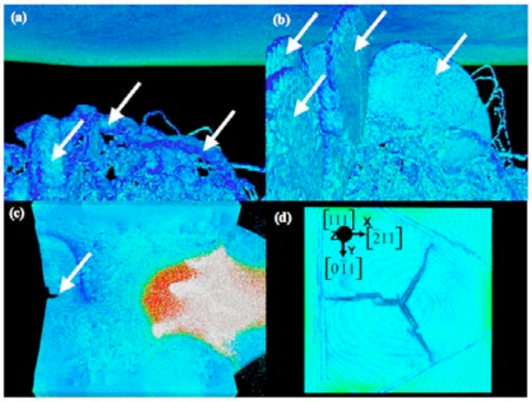

Fig. 49: Nucleation and growth of shock-induced brittle cracks in SiC. Atoms are color-coded as in Figs. 1 and 3. **(a)** System and back surface at 13.05 ps. White arrows indicate brittle cracks nucleating directly from dislocation lines when the tensile wave reflected from the back surface drops the local pressure to -20 GPa. **(b)** At 14.4 ps several cracks indicated by the white arrows develop, cleaving {110} planes of the zinc blende crystal in the direction of the back surface. **(c)** ¾ system view of the cracked configuration at 27 ps. **(d)** Snapshot of the cracked back surface at 31.2 ps showing extensive fractures in {110} planes. Surface steps created by dislocations reaching the surface can be seen in lines perpendicular to cracks.

Nucleated micro cracks grow cleaving initially all the six planes of the {110} family. As the tensile stress is released, micro cracks on planes $\{0\bar{1}1\}$, $\{10\bar{1}\}$, and $\{\bar{1}10\}$, growing in the [111] direction to the back surface, merge with other micro cracks forming larger cracks cleaving extensive planes of the zinc blende crystal until they reach the back surface; see Fig. 49(b). The growth of the cracks is not shielded by the emission of dislocations and it results in clean cleavage surfaces. Because of confinement the crack partially heals as the tensile stress is relaxed. Figure 49(c) shows the system in a ¾ configuration at 27 ps showing the whole fractured system. Figure 49(d) shows the cracked back surface of the target at 31.2 ps with three major cracks on $\{0\bar{1}1\}$, $\{10\bar{1}\}$, and $\{\bar{1}10\}$ planes.

The role of plastic structures in the process of micro crack nucleation in ceramics is overlooked in the literature. In brittle high-strength ceramics, plasticity is restricted and failure generally occurs with little plastic deformation, except at high temperatures or high stresses, where plasticity is favored. Because of this restricted plasticity, the role of plastic structures in the damage response of ceramics at extreme conditions of stress is rarely discussed in the literature. Here we demonstrate that plastic structures may have a complex dynamics in ceramics and have a central role in the failure by direct micro crack nucleation. In the literature only Fu, Evans, and Kriven, to the best of our knowledge, have suggested that single dislocation line cores can in principle be a source for stress concentration and micro crack nucleation. In their work, however, the micro crack initiation was analysed only in twinned

ZrO_2 ceramics. ***Our results demonstrate for the first time that micro crack initiation from dislocation lines is a plausible mechanism, in monolithic SiC ceramics, with just the presence of dislocation lines***. This behavior is very surprising because the common mechanisms for micro crack nucleation, *i.e.* from micro voids or dislocation coalescence, are both active in AlN ceramics. Several details related to the results presented here would require deeper analysis or new studies. In particular the stress and temperature conditions necessary to trigger dislocation dynamics and its role in the generation of micro cracks in ceramics deserve further studies. Finally the atomistic details from this study may be incorporated into continuum simulations and enable predictive modeling of damage processes in brittle ceramics.

§3.6 ACCELERATED MOLECULAR DYNAMICS METHODS

A. Accelerated Molecular Dynamics at the Solid-Liquid Interface

In many situations where SCC occurs, the crack is in contact with a corrosive solution. The capability to simulate chemical reactions at the solid-liquid interface, as well as other effects on the surface kinetics caused by the liquid, is thus critical to understanding SCC. While the presence of a liquid phase does not pose a problem to traditional MD simulations, it makes the application of AMD methods considerably more subtle. Indeed, these methods rely on the fact that the conformation of the solid changes slowly compared to its vibrational motion, but addition of a liquid introduces new relevant timescales into the problem. While the kinetics of a bulk liquid is generally fast enough to decouple from the kinetics of solid surfaces, wetting layers at a solid interface are known to evolve much more slowly, so that the two kinetics might couple in a such a way that acceleration would be difficult to achieve.

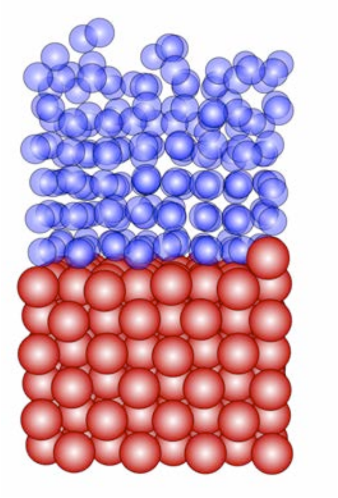

Fig. 50: Snapshot of the EAM-LJ model system at 400K; Silver atoms are red and liquid atoms are blue.

To extend the applicability of AMD to solid-liquid interfaces, we have generalized two accelerations methods, namely parallel replica dynamics (ParRep) and hyperdynamics, using the key concept of a super-state. Super-states are defined by a partition of configuration space such that internal sampling occurs more rapidly than transitions to neighboring super-states. Using this idea, we have adapted the statistical mechanics concepts that underpin the AMD methods so that they can be applied at a coarse-grained level. The new methods were tested on a model system composed of a silver adatom on a silver (100) substrate wetted by a Lennard-Jones fluid parametrized to interact more strongly with the solid than water would (Fig. 50). By defining super-states solely using the configuration of the solid atoms, Fig. 51 demonstrates that ParRep is able to accurately reproduce the transition rates of the adatom obtained by molecular dynamics between 600 and 700 K. This shows that the kinetics in the wetting layer is fast enough to decouple from the solid's. Using the parallel speedup provided by parallel replica dynamics, we extended the simulations to lower temperatures and showed that the rate follows an Arrhenius behavior, but with a significantly larger barrier and prefactor than for the dry surface. As the temperature decreases, the scalability of the method increases so that hundreds or thousands of processors can be efficiently used at 450 K. The super-states hyperdynamics also shows reasonable agreement with parallel replica results, while providing 10× and 30× speedup over MD at 500 K and 450 K, respectively. The two methods were also combined to increase the reach of MD by 400 and 1,200 times at the corresponding temperatures, using only 40 additional processors. The slight underestimation of the rates observed using hyper+ParRep might however indicate a bias due to the selective acceleration of the solid's degrees of freedom, which we are now investigating. ParRep, properly applied, is immune to this effect.

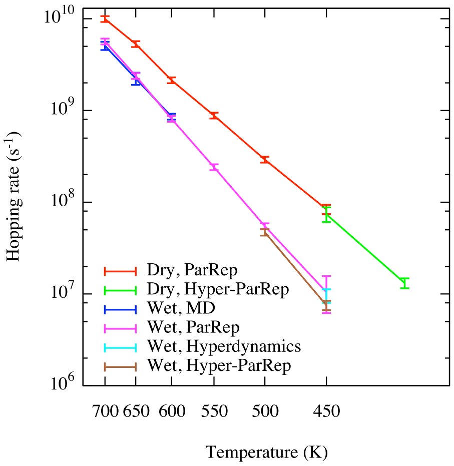

Fig. 51: Temperature dependence of the transition rate of the adatom using different simulation techniques. Error bars show the 90% confidence interval.

These results demonstrate that with suitable generalization, the AMD methods can substantially accelerate the dynamics of solid-liquid interfaces; these methods will play a vital role in the SCC multiscale framework.

B. Accelerated Molecular Dynamics near Grain Boundary and Crack Tip

In this subproject we are testing the use of AMD, temperature accelerated dynamics (TAD) in particular, on the important cases of activated processes near a crack tip, near a Σ5 grain boundary, and in the bulk, with and without applied stress. In TAD we use a constrained high-temperature (T_{high}) trajectory to accurately follow the state-to-state dynamics at the lower, desired temperature, T. Currently our focus is on vacancy migration in EAM Ni. When a high-quality Ni-S potential is ready (under development at USC), we will also study the kinetic characteristics of sulfur. Ab-initio calculations performed by Kaxiras indicate that sulfur atoms prefer substitutional sites. Thus, sulfur probably migrates by a vacancy-assisted mechanism, so learning about the vacancy kinetics also starts to tell us something about sulfur behavior.

Figure 46 shows results from one of our TAD crack-tip simulations. We placed a vacancy 8.8 Å (2.5 times the Ni lattice parameter) in front of a 3-layer wide notch tip and applied a 3% tensile strain to the whole system. The existence of the notch creates local stress concentration near the notch tip (quantified in Fig. 46(b)). Running TAD at T = 300K (T_{high} = 1,200K) on a 4,678-atom subsystem with frozen boundaries, we obtained a computational boost factor of roughly 10^7 compared to direct MD, which is very encouraging for the use of TAD in our SCC multiscale modeling. Over a time of 0.11 seconds, the vacancy diffused to the tip of the notch and created a small kink (Fig. 52(a)). In the unstretched perfect crystal, the vacancy diffuses isotropically by exchanges with first neighbors. When applying a uniaxial stress in the <001> direction, we observe that diffusion is confined to the {001} plane perpendicular to the load. In the case of the model notch, it creates a strong confinement of the vacancy in the crack plane, together with an enhanced diffusion towards the tip. The barriers decrease from 0.97eV ($\sigma_{zz} = 0$ bulk) to 0.67eV (σ_{zz} = 10 GPa) (Fig. 52(b)). The final jump to the tip lowers the energy considerably, so this is essentially an irreversible process. Such a confinement effect of point defects, mostly efficient in the damage zone ahead of a blunted crack, such as cracks obtained by stress assisted anodic dissolution, might be an important element in understanding the micro-crystallography of SCC fracture surfaces. Indeed, transgranular fracture surfaces of Cu in $NaNO_2$, a system where dissolution dominates, are not always composed of well-identified facets and a {100} component might be helpful.

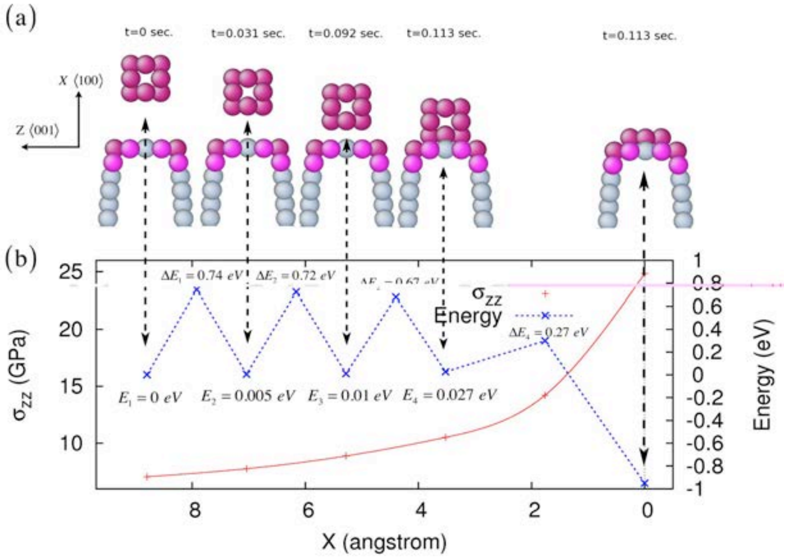

Fig. 52: (a) The reaction pathway for the vacancy migrating to the Ni crack tip. Only under-coordinated atoms are shown to highlight the location of the vacancy and the tip. The vacancy migration is confined to one (002) atomic plane. (b) The system energy and σ_{zz} as a function of X, the distance to the tip. ΔE_i denotes the barrier for the vacancy to hop from state i to i+1; E_i denotes the energy of state i with respect to initial state 1.

To understand the impact of the stress in the (310)[001] grains around the Σ5 GB, we performed TAD simulations of vacancy diffusion in Ni single crystals uniformly stretched in the <310> direction. The vacancy migration barriers as a function of the applied σ_{zz} stress are shown in Fig. 53(b). The out-of-plane (tilted (002)) migration is suppressed while in-plane diffusion is largely enhanced, in agreement with the <001> orientation discussed above. To obtain a significant effect, the stress should be large: between 1 and 10 GPa, in comparison to the elastic limit of the Ni single crystals (several tens of MPa). This stress range is relevant in the vicinity of the crack tip. Figure 53(a) shows the stress profile ahead of a crack along the Σ5 (310)[001] GB. Due to the stress singularity, stress levels are between 12 GPa and 6 GPa in a region of length 10 nm ahead of the tip. The confinement effect that we report, largely crystallographic, has an unexpected effect on the behavior of the vacancies in this zone which is relevant for intergranular damage in SCC: Elasticity would predict that the vacancies be attracted to the tip, as in the case of the notch in the <001> orientation, while in the case of the inclined (002) enhanced diffusion planes, the vacancy sees higher and higher barriers as it has to perform expensive out-of-plane diffusion jumps to move directly toward the tip. In the meantime, diffusion is enhanced along the inclined (002) plane where the vacancy belongs, which increases the probability to meet the GB, where we find it stays trapped. In the GB it presumably diffuses readily, although we are still studying this. The TAD simulations on a vacancy in this particular GB suffer from extreme loss of boost due to very low barrier processes. This complex interplay between the crystallography, the crack stress field and the trapping is worth studying further, most likely by suitably parameterized kinetic Monte Carlo simulations.

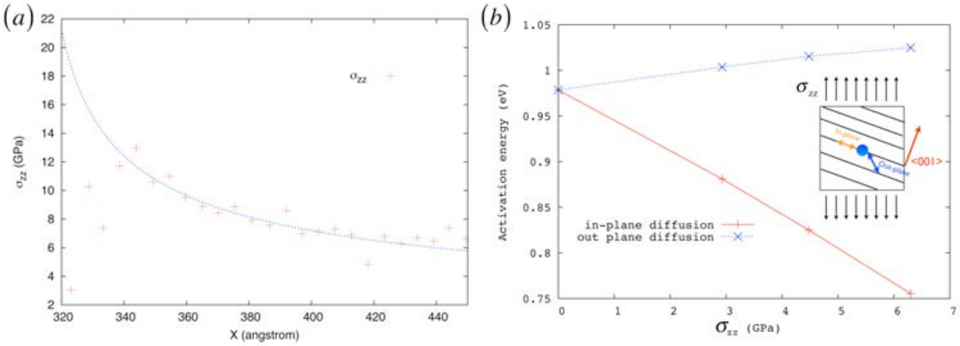

Fig. 53: Enhanced vacancy diffusion in the stress field of a crack along a Σ5(310)[001] symmetrical tilt GB. (a) Traction profile ahead of a crack loaded in mode I. The values are obtained by an average of the virial stress in the GB core over elementary volumes of size 5.3×8.8 A^2, centered on the GB plane. The maximum stress, compatible with linear continuum, is 12 GPa. The dotted line is a fit to $k/\mathrm{sqrt}(x-c)$, where k is the applied stress intensity factor and c the crack centre. The load is just below k_{Ie}, the critical load for emitting a partial dislocation. This stress field is therefore the steepest profile the crack can experience in Ni. (b) The activation barriers of vacancy migration as the function of applied σ_{zz}. The inset is the schematic diagram of the simulation system. In-plane diffusion denotes vacancy hopping confined in the tilted (002) plane (4 possible hop directions), whereas out-of-plane diffusion denotes vacancy hopping out to the neighboring (002) plane (8 possible hop directions).

§3.7 QUASICONTINUUM DENSITY FUNCTIONAL SIMULATIONS

We have developed a concurrent multiscale method that makes it possible to simulate multi-million to billions atoms effectively based on the density functional theory (DFT). The method termed QCDFT is formulated within the framework of Quasicontinuum method, and with DFT as its sole energetic input, *i.e.*, there is no empiricism in the formulation. The only underlying energy functional in QCDFT is density functional theory, either Kohn-Sham DFT or Orbital-free DFT. The unimportant degrees of freedom are systematically coarse-grained out by finite-element interpolation scheme. The development of QCDFT represents a major progress in quantum simulation of materials properties and opens door for many applications that are beyond the reach of existing quantum simulation methods. The code has been parallelized with MPI. The QCDFT method has been applied for a nanoindentation of an Al thin film. The results suggest that QCDFT is an excellent method for quantum simulation of materials properties at length-scales that are relevant to experiments. Fig. 54 shows the schematic representation of the sample and the indentor. Fig. 5 presents an overview of the entire system and domain partition in the nanoindentation simulation.

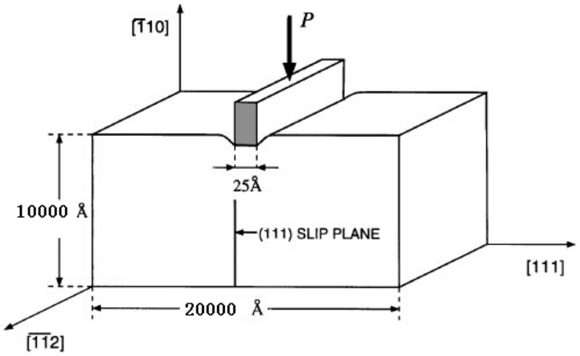

Fig. 54: Schematic representation of nanoindentation of Al thin film: geometry and orientation.

A. Validation of EAM Potentials Using First-Principles Calculations

We demonstrated that the sparse grid method could be used to greatly reduce the necessary first-principle DFT calculations in order to validate the accuracy of empirical potentials. We have examined three different EAM potentials of Al by computing energy and stress as a function of uniform hydrostatic deformations. It is found that the EAM results deviate significantly from the corresponding DFT values for large compressions, while the results agree reasonably well to the DFT values for expansions (Fig. 55). The failure of the EAM potential is attributed to the uniform electron density approximation adopted in the EAM theory. In compressed configurations, the local electron density gradient increases considerably, and the uniform electron density approximation is no long valid, hence EAM fails. We identified a critical value of the electron density gradient beyond which EAM energetics deviate significantly from the DFT values. We then test this critical value for the case of the generalized stacking fault energy surface calculations, and found that for the run-on stacking fault where the maximum local density gradient exceeds the critical value, EAM indeed fails in giving the correct energy (Fig. 56). For all other stacking faults where the critical

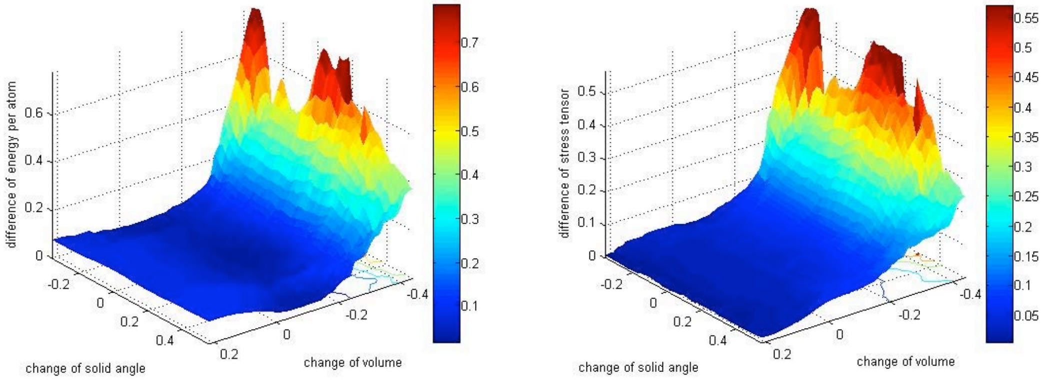

Fig. 55: The difference between Ercolessi-Adams embedded atom method (EAM) potential and first-principles (VASP) calculations for the energy ($dE = E_{EAM}-E_{VASP}$) vs. change of solid angle of the deformed cell and the change of cell volume in the left panel. The difference of stress tensor $dS = ||S_{EAM}-E_{VASP}||$ is shown in the right panel. The results are represented in a sparse grid with only 8,433 points to represent 1,291,467,969 points in a uniform grid with the same precision. The sparse grid is being incorporated in QCDFT method to perform DFT calculations for local region.

value is not exceeded, EAM gives reasonable results. We proposed a solution to improve the transferability of EAM potentials by considering the gradient of electron density in the analytical expression of EAM energy—the embedding function should not only depend on electron density, but also on its gradient.

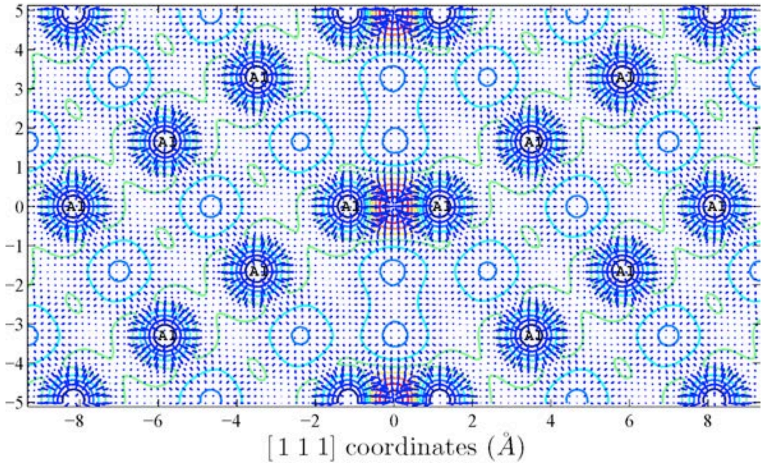

Fig. 56: The valence charge density and its gradient for a run-on stacking fault of Al. The length and direction of the arrows represent the magnitude and direction of the charge density gradient. Two layers of Al atoms stack right on top of each other in [111] direction (in the middle of the figure) constitute a run-on stacking fault.

§4 GRADUATE EDUCATION AND TRAINING

We envision a new generation of graduate students who have the expertise developed through interdisciplinary education, research, and training to attack complex system-level problems. Our educational program is fundamentally different from other programs, in which students are trained from the "bottom up" in increasing detail in a narrow field. The new training is from the "top down", in which students combine theory, modeling and petascale simulations to study complex systems and then decide which details are important and how they can be incorporated.

The SciDAC program offers a unique opportunity to train graduate students in this new cross-disciplinary paradigm. Our students participate in a **dual-degree program** in which they fulfill Ph.D. requirements within their own discipline (*e.g.*, physics, materials science, chemical engineering, computational biology and bioinformatics) and participate in a program of study towards an M.S. in computer science with specialization in high performance computing and simulations (MSCS–HPCS).

To prepare the dual-degree students for interdisciplinary research projects, we are developing a well-defined sequence of courses that are cross-listed between the participating departments so as not to inflate students' loads. Four courses covering theoretical and simulation aspects provide atomistic-to-continuum level perspectives.

- A new course on hybrid atomistic-mesoscale simulations are being developed to prepare students for research. The new course will be based on the following three simulation courses at USC: 1) MASC 575—Basics of Atomistic Simulation of Materials includes building a parallel computer from components, molecular dynamics method, and computation of structural, thermodynamics and transport properties; simulation projects; 2) MASC 576—Molecular Dynamics Simulations of Materials and Processes covers atomistic simulations of materials and processes using parallel computing, correlation functions for structural and dynamical properties; and 3) a course on quantum-mechanical simulations is being developed.

The MSCS-HPCS courses provide the necessary computer science (CS) background for research projects. The following courses address CS research issues.

- CSCI 503—Parallel Programming explores parallel programming languages (conventional, Global Address Space, SIMD, and streaming models) for the growing diversity of modern parallel architectures, with hands-on programming assignments on supercomputers and commodity multi-core systems.
- CSCI 595—Compiling for High Performance involves optimizations in high-performance parallel architectures, including VLIW microprocessors, systems-on-a-chip, novel memory systems, FPGA-based systems, multimedia extension architectures and chip multiprocessors.
- CSCI 596—Scientific Computing and Visualization uses MD and QM simulations to teach basic elements of high performance parallel scientific computing and visualization.
- CSCI 653—High Performance Computing and Simulations covers advanced techniques—hierarchical particle/continuum and deterministic/stochastic simulation algorithms on parallel and distributed computers, and massive data visualization and mining in virtual environments.

Distance Education Network (DEN): Many components of the dual-degree program are accessible to students anywhere and anytime, exploiting the infrastructure of USC's top-rated distance-education learning program called DEN (US News and World Report, 2002).

The students supported by this project are involved in our outreach activities for undergraduate students from underrepresented groups in science and engineering. They help us organize annual computational science workshops for underrepresented groups (CSWUG), which are held annually at USC. Twenty-four undergraduate students and twelve faculty mentors from Historically Black Colleges and Universities and Minority Serving Institutions will be invited to participate in a special one-week workshop, where we provide them with intense hands-on experience in multi-core parallel computing and immersive and interactive visualization. At the workshop, our Ph.D. students help the participants build a parallel computer from components and then lead them through parallel computing exercises. Experts are invited to make presentations on career opportunities in science and engineering. At the end of the workshop, each student/faculty participant is given a node to bring back to his/her institution. Follow-on activities include bringing students back for summer research, and students and faculty will be involved in mentoring those undergraduate students.

§5 INTERACTIONS WITH SCIDAC CENTERS

As we described in §2, we are collaborating with compiler and artificial intelligence (AI) research groups in the SciDAC-Performance Engineering Research Institute (PERI) to optimize our metascalable parallel MD and QM applications using 1) knowledge-representation techniques for expressing the exposed concurrency; and 2) machine-learning techniques for optimally mapping the expressed concurrency to hardware.

Currently, we are testing the use of PERI's auto-tuning and code-generation tools for the optimization of our hierarchical message passing + multithrading + single-instruction multiple-data (SIMD) parallel MD and DFT codes in collaboration with Mary Hall at the University of Utah. For example, we have initiated the use of a code transformation tool, CHiLL, to generate optimized code variants of a high-order stencil computation kernel of our real-space EDC-DFT code (a joint publication with the PERI team is under review). CHiLL is a transformation framework that allows the user to specify code transformations using a high-level script interface. CHiLL supports code transformations such as loop tiling, interchanging, unrolling, fusion and distribution, data copying and data prefetching. CHiLL takes as input the original code and a script specifying code transformations, and automatically generates a transformed code. Optimization parameters such as tile sizes can be specified as integer values or as unbound parameters to be determined later. **We are requesting support for one postdoc to carry this SciDAC Center collaboration over to the next level by demonstrating the auto-tuning for a broader range of high-end computing applications including our MD, ReaxFF-MD and EDC-DFT.**

The significant contribution of this work over the first three-year duration of this project is the development of novel data structures, algorithms, and numerical techniques for reactive atomistic simulations. This work has resulted in a state-of-the-art software infrastructure, now in limited public-domain release. Table below shows our partners at the National Labs, as well as academia, worldwide, using this code for a broad class of simulations. The software release is limited only by our ability to maintain and support a large and diverse user base.

Lab/Research Group	Institution	Application/Development
Aidan Thompson Steve Plimpton	Sandia National Labs	LAMMPS integration, force-field optimization
Markus Buehler	MIT	Silica-water interface
Alejandro Strachan	Purdue	SiGe nanorods
Narayan Aluru	UIUC	Damped Si cantilevers
Adri van Duin	Penn State	Force-field development
William Goddard	CalTech	Method development
Jorn Amundsen	Norwegian University of Science and Technology	Parallel code optimization

§6 PUBLICATIONS

1. Y. Chen, Z. Lu, K. Nomura, W. Wang, R. K. Kalia, A. Nakano, and P. Vashishta, "Interaction of voids and nanoductility in silica glass," *Phys. Rev. Lett.* **99**, 155506 (2007).

2. P. Vashishta, R. K. Kalia, A. Nakano, E. Kaxiras, A. Grama, G. Lu, S. Eidenbenz, A. F. Voter, R. Q. Hood, J. A. Moriarty and L. H. Yang, "Hierarchical petascale simulation framework for stress corrosion cracking," *J. Phys.: Conf. Series* **78**, 012036 (2007).

3. C. Zhang, S. Callaghan, T. Jordan, R. K. Kalia, A. Nakano, and P. Vashishta, "ParaViz: a spatially decomposed parallel visualization algorithm using hierarchical visibility ordering," *Int'l J. Comput. Sci.* **1**, 407 (2007).

4. B. Bansal, U. Catalyurek, J. Chame, C. Chen, E. Deelman, Y. Gil, M. Hall, V. Kumar, T. Kurc, K. Lerman, A. Nakano, Y. L. Nelson, J. Saltz, A. Sharma, and P. Vashishta, "Intelligent optimization of parallel and distributed applications," in *Proc. of Next Generation Software Workshop, International Parallel and Distributed Processing Symposium* (IEEE, Long Beach, CA, 2007).

5. X. Zhang and G. Lu, "Quantum mechanics/molecular mechanics methodology for metals based on orbital-free density functional theory," *Phys. Rev. B* **76**, 245111 (2007).

6. S. Banerjee, N. Ghoniem, G. Lu and N. Kioussis, "Non-singular descriptions of dislocation cores: a hybrid *ab initio*-continuum approach," *Phil. Mag*. **87**, 4131 (2007).

7. L. H. Yang, R. Q. Hood, J. E. Pask and J. E. Klepeis, "Large scale quantum mechanical simulations of high-Z metals," *J. Computer-Aided Mater. Design.* **14**, 337 (2007).

8. P. Vashishta, R. K. Kalia, A. Nakano, E. Kaxiras, A. Grama, G. Lu, S. Eidenbenz, A. F. Voter, R. Q. Hood, J. A. Moriarty, and L. H. Yang, "Hierarchical petascale simulation framework for stress corrosion cracking," *J. Phys.: Conf. Series* **125** 012060 (2008).

9. A. Nakano, R. K. Kalia, K. Nomura, A. Sharma, P. Vashishta, F. Shimojo, A. C. T. van Duin, W. A. Goddard, III, R. Biswas, D. Srivastava, and L. H. Yang, "De novo ultrascale atomistic simulations on high-end parallel supercomputers," *Int'l J. High Performance Comput. Appl.* **22**, 113 (2008).

10. F. Shimojo, R. K. Kalia, A. Nakano, and P. Vashishta, "Divide-and-conquer density functional theory on hierarchical real-space grids: Parallel implementation and applications," *Phys. Rev. B* **77**, 085103 (2008).

11. F. Shimojo, R. K. Kalia, A. Nakano, K. Nomura, and P. Vashishta, "Metascalable molecular dynamics simulation of nano-mechano-chemistry," *J. Phys.: Condens. Matter* **20**, 294204 (2008).

12. K. Nomura, R. K. Kalia, A. Nakano, and P. Vashishta, “A scalable parallel algorithm for large-scale reactive force-field molecular dynamics simulations,” *Comput. Phys. Commun.* **178**, 73 (2008).
13. L. Peng, K. Nomura, T. Oyakawa, R. K. Kalia, A. Nakano, and P. Vashishta, “Parallel lattice Boltzmann flow simulation on emerging multi-core platforms,” *Lecture Notes in Computer Science* **5168**, 763 (2008); in *Proc. of 14th International European Conference on Parallel and Distributed Computing* (Las Palmas de Gran Canaria, Spain, August 26-29, 2008).
14. K. Nomura, S. W. de Leeuw, R. K. Kalia, A. Nakano, L. Peng, R. Seymour, L. H. Yang, and P. Vashishta, “Parallel lattice Boltzmann flow simulation on a low-cost PlayStation3 cluster,” *Int’l J. Comput. Sci.* **2**, 437 (2008).
15. X. Zhang and G. Lu, “Electronic origin of void formation in f.c.c. metals,” *Phys. Rev. B* **77**, 174102 (2008).
16. Q. Peng, X. Zhang, L. Hung, E.A. Carter, and G. Lu, “Quantum simulation of materials at micron scales and beyond,” *Phys. Rev. B* **78**, 054118 (2008).
17. Z. Chen, G. Lu, N. Kioussis, and N. Ghoniem, “Effect of local environment on the mobility of dislocations in refractory bcc metals: Concurrent multiscale approach,” *Phys. Rev. B* **78**, 134102 (2008).
18. X. Zhang, C.Y. Wang and G. Lu, “Electronic structure analysis of self-consistent embedding theory for quantum/molecular mechanics simulations,” *Phys. Rev. B* **78**, 235119 (2008).
19. R. Q. Hood, L. H. Yang and J. A. Moriarty, “Quantum molecular dynamics of uranium at high pressure and temperature,” *Phys. Rev. B* **78**, 24116 (2008).
20. J. A. Moriarty, L. H. Yang, R. Q. Hood, P. Vashishta, E. Kaxiras and A. F. Voter, “Toward petascale atomistic simulations with quantum-level accuracy,” *3rd Institutional Computing Grand Challenge Proposal* (LLNL, February, 2008).
21. J. A. Moriarty, J. N. Glosli, R. Q. Hood, J. E. Klepeis, D. A. Orlikowski, P. Söderlind and L. H. Yang, “Quantum-based atomistic simulation of metals at extreme conditions,” in *TMS 2008 Annual Meeting Supplemental Proceedings Volume I: Materials Processing and Properties* (TMS, Warrendale, PA, 2008) p. 313.
22. K. Nomura, H. Dursun, R. Seymour, W. Wang, R. K. Kalia, A. Nakano, P. Vashishta, F. Shimojo, and L. H. Yang, “A metascalable computing framework for large spatiotemporal-scale atomistic simulations,” in *Proc. of International Parallel and Distributed Processing Symposium (IPDPS 09)* (IEEE, Rome, Italy, 2009).
23. L. Peng, R. Seymour, K. Nomura, R. K. Kalia, A. Nakano, P. Vashishta, A. Loddoch, M. Netzband, W. R. Volz, and C. C. Wong, “High-order stencil computations on

multicore clusters," in *Proc. of International Parallel and Distributed Processing Symposium (IPDPS 09)* (IEEE, Rome, Italy, 2009).

24. K. Nomura, W. Wang, R. K. Kalia, A. Nakano, P. Vashishta, and F. Shimojo, "Large spatiotemporal-scale material simulations on petaflops computers, in *Multiscale Simulation Methods in Molecular Sciences*, edited by J. Grotendorst, N. Attig, S. Blügel, and D. Marx (John von Neumann Institut für Computing, Jülich, Germany, 2009) p. 321.

25. Y. Song, Y. Tanaka, H. Takemiya, H. Nakada, S. Sekiguchi, A. Nakano, and S. Ogata, "The development and evaluation of an integrated framework supporting sustainable execution for large-scale computations on Grids," *Int'l J. Comput. Sci.* **3**, 18 (2009).

26. G. Wu, G. Lu, C.J. Garcia-Cervera, and W. E, "Density-gradient-corrected embedded atom method," *Phys. Rev. B* **79**, 035124 (2009).

27. D. Perez, Y. Shim, J. G. Amar, B. P. Uberuaga, and A. F. Voter, "Accelerated molecular dynamics methods: Introduction and recent developments," *Annual Reports in Computational Chemistry* **5**, 79 (2009).

28. L. Peng, M. Kunaseth, H. Dursun, K. Nomura, W. Wang, R. K. Kalia, A. Nakano, and P. Vashishta, "A scalable hierarchical parallelization framework for molecular dynamics simulation on multicore clusters," in *Proc. of International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 09)* (Las Vegas, NV, 2009).

29. H. Dursun, K. Nomura, W. Wang, M. Kunaseth, L. Peng, R. Seymour, R. K. Kalia, A. Nakano, and P. Vashishta, "In-core optimization of high-order stencil computations," in *Proc. of International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 09)* (Las Vegas, NV, 2009).

30. K. Nomura, Y. Chen, W. Wang, R. K Kalia, A. Nakano, P. Vashishta, and L. H Yang, "Interaction and coalescence of nanovoids and dynamic fracture in silica glass: multimillion-to-billion atom molecular dynamics simulations," *J. Phys. D* **42**, 214011 (2009).

31. H. Chen, R. K. Kalia, E. Kaxiras, G. Lu, A. Nakano, K. Nomura, A. C. T. van Duin, P. Vashishta, and Z. Yuan, "Embrittlement of metal by solute segregation-induced amorphization," *Phys. Rev. Lett.* **104**, 155502 (2010).

32. M. Vedadi, A. Choubey, K. Nomura, R. K. Kalia, A. Nakano, P. Vashishta, and A. C. T. van Duin, "Structure and dynamics of shock-induced nanobubble collapse in water," *Phys. Rev. Lett.* **105**, 014503 (2010).

33. P. S. Branicio, R. K. Kalia, A. Nakano, and P. Vashishta, "Nanoductility induced brittle fracture in shocked high performance ceramics," *Appl. Phys. Lett.* **97**, 111903 (2010).

34. M. Kunaseth, R. K. Kalia, A. Nakano, and P. Vashishta, "Performance modeling, analysis, and optimization of cell-list based molecular dynamics," in *Proc. International Conference on Scientific Computing, CSC'10* (Las Vegas, NV, 2010).

35. L. H. Yang, M. Tang and J. A. Moriarty, "Dislocations and plasticity in bcc transition metals at high pressure," in *Dislocations in Solids*, Vol. 16, eds. J. P. Hirth and L. Kubin (2010) pp. 1-46.

36. M. Manguoglu, M. Koyuturk, A. Sameh, and A. Grama, "Weighted matrix ordering and parallel banded preconditioners for iterative linear system solvers," *SIAM J. Sci. Comput*. **32**, 1201 (2010).

37. K. Nomura, Y.-C. Chen, R. K. Kalia, A. Nakano, and P. Vashishta, "Defect migration and recombination in nanoindentation of silica glass," *Appl. Phys. Lett.* **99**, 111906 (2011).

38. Z. Yuan, H.-P. Chen, W. Wang, K. Nomura, R. K. Kalia, A. Nakano, and P. Vashishta, "Sulfur-impurity induced amorphization of nickel," *J. Appl. Phys.* **110**, 063501 (2011).

39. P. Vashishta, R. K. Kalia, A. Nakano, and J. P. Rino, "Interaction potential for aluminum nitride: a molecular dynamics study of mechanical and thermal properties of crystalline and amorphous aluminum nitride," *J. Appl. Phys.* **109**, 033514 (2011).

40. L. Peng, M. Kunaseth, H. Dursun, K. Nomura, W. Wang, R. K. Kalia, A. Nakano, and P. Vashishta, "Exploiting hierarchical parallelisms for molecular dynamics simulation on multicore clusters," *J. Supercomput.* **57**, 20 (2011).

41. M. Kunaseth, D. F. Richards, J. N. Glosli, R. K. Kalia, A. Nakano, and P. Vashishta, "Scalable data-privatization threading for hybrid MPI/OpenMP parallelization of molecular dynamics," in *Proc. Int'l Conf. Parallel and Distributed Processing Tech. Appl., PDPTA'11* (Las Vegas, NV, 2011).

42. K. Nomura, R. K. Kalia, A. Nakano, and P. Vashishta, "Mechanochemistry of shock-induced nanobubble collapse near silica in water," *Appl. Phys. Lett.* **101**, 073108 (2012).

43. H. Dursun, M. Kunaseth, K. Nomura, J. Chame, R. F. Lucas, C. Chen, M. Hall, R. K. Kalia, A. Nakano, and P. Vashishta, "Hierarchical parallelization and optimization of high-order stencil computations on multicore clusters," *J. Supercomput.* **62**, 946 (2012).

44. M. Kunaseth, K. Nomura, H. Dursun, R. K. Kalia, A. Nakano, and P. Vashishta, "Memory-access optimization of parallel molecular dynamics simulation via dynamic data reordering," *Lecture Notes in Computer Science* **7484**, 781 (2012); in *Proc. Int'l European Conf. Parallel and Distributed Comput., Euro-Par 2012* (Rhodes Island, Greece, 2012).

45. M. Aktulga, S. Pandit, and A. Grama, "Scalable parallel formulations of reactive atomistic simulations," *Parallel Comput*. **38**, 245 (2012).

§7 INVITED PRESENTATIONS

1. “Multiscale modeling for dislocations”, Mini-symposium on Modeling and Simulation of Dislocations at the International Congress on Industrial and Applied Mathematics (ICIAM07), Zurich, Switzerland, July 2007.
2. “Multiscale simulations of complex materials for engineering and biological applications,” International Discussion Meeting on Molecular and Structural Basis of Functional Systems, Johannes Guttenberg University, Mainz, Germany, September 2007.
3. “Extending the size scale in accelerated molecular dynamics methods,” Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, MN, July 2007.
4. “Recent advances and ongoing challenges in accelerated molecular dynamics methods,” National Meeting of the American Chemical Society, Boston, MA, August 2007.
5. “Petascalable reactive molecular dynamics simulation of mechanochemical processes,” CECAM (Centre Européen de Calcul Atomique et Moléculaire) Workshop on “Linear-Scaling Ab Initio Calculations: Applications and Future Directions,” Lyon, France, September 2007.
6. “Billion-electron quantum simulations for metals,” Workshop on “Multiscale Modeling of Materials,” Princeton University, Princeton, NJ, October 2007.
7. “Metascalable molecular dynamics simulation of mechanochemical processes,” Third International Workshop on High-Performance Computing for Nano-science and Technology, Reno, NV, November 2007.
8. “Extending the size scale in accelerated molecular dynamics methods,” MRS Fall Meeting, Boston, MA, November, 2007.
9. “Toward petascale quantum molecular dynamics simulations for high-Z metals,” International Conference on the Enhancement & Promotion of Computational Methods in Engineering and Science, Kyoto, Japan, December 2007.
10. “Metascalable reactive atomistic simulation of nano-mechano-chemistry,” International Workshop on Large-scale Matrix Computation and Applications in Physics and Engineering Science, Tokyo, Japan, December 2007.
11. “A metascalable atomistic simulation framework: Spatiotemporal data localities,” Grid Technology Research Center, National Institute of Advanced Industrial Science and Technology, Tsukuba, Japan, December 2007.
12. “Extending the size scale in accelerated molecular dynamics methods,” (plenary talk) Workshop on Mathematical and Computational Nanoscience, Victoria University of Wellington, Wellington, New Zealand, December 2007.

13. “Multimillion-to-billion atom molecular dynamics simulations of deformation, flow, fracture & nanoindentation in silica glass,” International Symposium on AMO and HPC: A Seamless Frontier, Kolkatta, India, January 2008.

14. “QCDFT: A route to quantum simulations at macroscale,” NSF Workshop on “Dislocation Dynamics”, UC San Diego, La Jolla, CA, January 2008.

15. “Quantum-based atomistic simulation of metals at extreme conditions,” TMS Annual Meeting, New Orleans, Louisiana, March 2008.

16. Workshop on Atomistic Models of Materials: Mathematical Challenges, Mathematisches Forschungsinstitut Oberwolfach, Germany, April 2008.

17. One-Week Course: “Accelerated molecular dynamics methods,” University of Oxford, Oxford, UK, April 2008,

18. “Multimilion-to-bilion atom molecular dynamics simulations of deformation, fracture and nanoductility in silica glass,” IUTAM Symposium on Modeling Nanomaterials and Nanosystems, Aalborg, Denmark, May 2008.

19. “Multiscale simulations of complex materials for engineering and biological applications,” SIAM Conference on Mathematical Aspects of Materials Science, Symposium on Multiscale Models and Techniques for Defects in Solids, Philadelphia, PA, May 2008.

20. “Real-time dynamics of electrons and ions in nanostructures and biomolecules,” SIAM Conference on Mathematical Aspects of Materials Science, Symposium on Electronic Structure, Philadelphia, PA, May 2008.

21. “Large spatiotemporal-scale atomistic simulations of materials on petaflops computers,” International Conference of the Grand Challenge to Next-Generation Integrated Nanoscience, Tokyo, Japan, June 2008.

22. “Parallel banded preconditioners for non-symmetric linear system solvers,” SIAM Parallel Matrix Algorithms and Applications, Neuchatel, Switzerland, June 2008.

23. IPAM Workshop on Bridging Time and Length Scales in Materials Science and Bio-Physics, Lake Arrowhead, CA, June 2008.

24. Symposium on “Orbital-free Density Functional Theory”, the 6th Congress of the International Society for Theoretical Chemical Physics (ISTCP-VI), University of Vancouver, Canada, July, 2008.

25. Applied Math Colloquium, University of Arizona, AZ, September 2008.

26. Special Symposium on “Defects in Materials”, the 4th International Conference on Multiscale Modeling of Materials, Florida State University, October 2008.

27. Workshop on Multiscale Modeling of Defects in Solids, 45th Annual Technical Meeting of Society of Engineering Science, University of Illinois at Urbana-Champaign, October 2008.

28. "A metascalable computing framework for large spatiotemporal-scale atomistic simulations," Workshop on Computational Science Frontiers, Nagoya, Japan, November 2008.

29. "Accelerated molecular dynamics methods," Materials Design Annual Users Group Meeting, Taos, NM, November 2008.

30. "Multibillion-atom simulations of nano-mechano-chemistry on petaflops computers," First International Symposium on Global Center of Excellence for Mechanical Systems Innovation, Tokyo, Japan, February 2009.

31. "Large spatiotemporal-scale material simulations on petaflops computers," (opening keynote), Winter School on Multiscale Simulation Methods in Molecular Sciences, Jülich, Germany, March 2009.

32. "QCDFT: a route to quantum simulations at micron scale", Symposium on Science and Engineering of Nanoscale Systems: Scientific Discovery through Advanced Computation and Experimental Validation at the International Conference on Computational & Experimental Engineering and Sciences (ICCES'09), Phuket, Thailand, April 2009.

33. "Metascalable atomistic simulations of nano-mechano-chemistry on petaflops computers," European Materials Research Society 2009 Spring Meeting, Strasbourg, France, June 2009.

34. "Large reactive molecular dynamics simulations of nano-mechano-chemistry on petaflops computers," Workshop on "Molecular Dynamics with Reactive Potentials", Memphis, TN, August 2009.

35. "Reactive molecular dynamics simulations of nano-mechano-chemistry on petaflops computers," Fourth Meeting on "Reactive Potential Development", Madison, WI, September 2009.

36. "Petascalable molecular dynamics simulations of nano-mechano-chemistry," Workshop on "Towards Microscale Molecular Simulations with High-Performance Computing, Dublin, Ireland, October 2009.

37. "Metascalable molecular dynamics simulations for materials and energy applications," 13th International Conference on "Theoretical Aspects of Catalysis", Matsushima, Japan, June 2010.

38. "Multimillion-atom photo-mechano-chemistry simulations of metal-oxide nanotechnology on high-end parallel computers," The 6th International Workshop on "Zinc Oxide and Related Materials", Changchun, China, August 2010.

39. “Multimillion-atom photo-mechano-chemistry simulations for energy nanoscience on high-end parallel computers,” DOE SciDAC Performance Engineering Research Institute , Salt Lake City, UT, September 2010.
40. “Metascalable molecular dynamics simulations of photo-mechano-chemistry for energy applications,” International Workshop on “Photocatalysis and Environmental Remediation Materials”, Tsukuba, Japan, January 2011.
41. “A metascalable approach to large reactive molecular-dynamics simulations for energy and nanoscience applications,” Symposium on “Computational Aspects of Large Scale Molecular Dynamics Simulations” at SIAM Conference on Computational Science and Engineering, Reno, NV, March 2011.
42. “A metascalable approach to reactive molecular dynamics simulations for energy applications,” Conference on “Materials for Energy Applications—Experiment, Modeling and Simulations”, Rancho Palos Verdes, CA, March 2011.
43. “Petascale molecular dynamics simulations of photo-mechano-chemistry,” Symposium on “Atomistic-to-continuum-scale Interaction between Flow and Material” at International Conference on Computational and Experimental Engineering and Sciences, Nanjing, China, April 2011.
44. “A metascalable algorithmic framework for large spatiotemporal-scale molecular-dynamics simulations of photo-mechano-chemistry,” Symposium on “Quantum Modeling in Molecular Simulation” at International Congress on Industrial & Applied Mathematics, Vancouver, Canada, July 2011.

§8 PERSONNEL

Postdoctoral Researchers

Maria Fyta (Harvard)

Thomas Kuehne (Harvard)

Sagar Pandit (Purdue)

Weiqiang Wang (Purdue)

Gang Wu (CSUN)

Danny Perez (LANL)

Chun-Wei Pao (LANL)

Graduate Students

Yi-Chun Chen (USC, graduated May 2008)

Chih-Ying Lin (USC)

Mohammad Vedadi (USC)
Georg Schusteritsch (Harvard)
Ee Hou Yong (Harvard)
Hassan Metin Aktulga (Purdue)
Xu Zhang (CSUN)

Visiting Scientists

Döme Tanguy (CNRS, Ecole des Mines de Saint-Etienne, France)
Fuyuki Shimojo (Kumamoto University, Kumamoto, Japan)