Generalized Irving-Kirkwood formula for the calculation of continuum quantities in molecular dynamics models

Jerry Zhijian Yang[1], Xiaojie Wu[2], Xiantao Li[2]

[1]Wuhan University, Wuhan, 430072, China
[2]Penn State University, University Park, PA 16802, USA

Continuum mechanics quantities can be computed from molecular dynamics (MD) models based on the classical Irving-Kirkwood (IK) formalism. Practical implementations of IK formulas involve a spatial averaging using a smooth kernel function. The obtained results usually need to be further processed to reduce the fluctuation, e.g., by ensemble or time averaging. In this talk the IK formalism is extended to systematically incorporate both spatial and temporal averaging into the expression of continuum quantities. We will present the results both in Lagrangian and Eulerian coordinates.

[1] J. Z. Yang, X. Wu and X. Li, J. Chem. Phys. **137**, 134104 (2012).

A concurrent parallel multiscale algorithm for large 3d continuum/atomic simulations at finite temperature using lammps

Fabio Pavia, William A. Curtin

EPFL, Lausanne, Switzerland

Deformation and fracture processes in engineering materials combine complex phenomena and often require simultaneous descriptions over a range of length and time scales, with each scale using a different computational technique. Here we present a high-performance parallel 3D computing framework for large multiscale studies that couples an atomic domain, modeled using molecular dynamics, and a continuum domain, modeled through explicit finite element. The method ensures mechanical equilibrium at the interface between the different domains of the model, with the finite elements providing boundary conditions for the atomic simulations. Most importantly, the entire method is implemented using the large-scale parallel molecular dynamics code LAMMPS, thus taking advantage of all the tools associated with this popular open-source code. We demonstrate operation of the code with a study of dislocation glide and bowing around obstacles, for dislocation lengths of ~50nm using fewer than 1,000,000 atoms but with no spurious image forces or errors due to coupling.

Analysis of an energy localization method used in kinetic Monte Carlo simulations of heteroepitaxial growth

Kyle Golenbiewski[1], Tim P. Schulze[1], Peter Smereka[2]

[1]Department of Mathematics, University of Tennessee, Knoxville, TN 37909-1300, USA
[2]Department of Mathematics and Michigan Center for Theoretical Physics, University of Michigan, Ann Arbor, MI 48109-1109, USA

Heteroepitaxial growth consists of slowly depositing one material onto a crystalline substrate formed from a second material. An important feature of this process is that the natural lattice spacing of the deposited material may differ from that of the substrate's, resulting in elastic strain. Simulation of such growth using kinetic Monte Carlo is often based on rates determined by differences in elastic energy between two configurations. This, however, is computationally challenging due to the long range nature of elastic interactions. Adopting an atomistic approach, we consider a method in which the elastic field is updated using highly accurate local approximations, though the energies themselves are far less accurate. Namely, we approximate the elastic energy barrier by constraining the displacement field with an atom removed to agree with the displacement field when the atom is present outside some local region. In order to gain insight into this energy localization method, we appeal to a continuum analogue of the discrete mechanical system. For the scenario of an isolated island sitting on an unbounded and otherwise flat film on a flat substrate, we extend earlier results done on a 2D system to 3D.

Frozen Gaussian approximation for wave propagation in periodic media

Xu Yang

UC Santa Barbara

Unlike the Gaussian beam method (GBM), frozen Gaussian approximation (FGA) is a semiclassical approximation to high frequency waves with Gaussians of fixed width. It was called Herman-Kluk approximation in quantum chemistry, and used for solving the Schrodinger equation. In this talk, we generalize the idea, derive and analyze the accuracy of FGA for waves propagating in periodic media, modeled by the Schrodinger equation with a lattice potential. Assuming the existence of band gap, we are able to obtain the effective dynamic equations of electrons within each energy band, and prove the first order accuracy of FGA.

Semiclassical Models for Quantum Systems and Band Crossings

Lihui Chai, Shi Jin, Qin Li

UCSB

We derive the coupled semi-classical Liouville system for the Schrödinger equations with band crossings to describe the inter-band transitions semi-classically. The main idea is using the Wigner transform and the basis representation, retaining the off-diagonal terms in the Wigner matrix which cannot be ignored near the point of band crossing. The coupled Liouville systems are able to describe the inter-band transition phenomena very well and consistent to the adiabatic approximation if the adiabatic assumption holds. A domain decomposition method that couples these non-adiabatic models with the adiabatic Liouville equations is also presented for a multiscale computation. Solutions of these models are numerically compared with those of the Schrödinger equations to justify the validity of these new models for band-crossings.

An efficient rescaling algorithm for simulating the evolution of precipitates in an elastic media

Shuwang Li[1], Amlan Barua[1], Xiaofan Li[1], John Lowengrub[2]

[1]LDepartment of Applied Mathematics, Illinois Institute of Technology, Chicago 60616, USA
[2]Department of Mathematics, University of California at Irvine, Irvine CA 92697, USA

In this talk, we present a space-time rescaling scheme for computing the long time evolution of multiple precipitates in an elastically stressed medium. The algorithm is second order accurate in time, spectrally accurate in space and enables one to simulate the evolution of precipitates in a fraction of the time normally used by fixed-frame algorithms. In particular, we extend the algorithm recently developed for single particle by Li [1] to the multiple particle case, which involves key differences in the method. Our results show that without elasticity there are successive tip splitting phenomena accompanied by the formation of narrow channels between the precipitates. In presence of applied elastic field, the precipitates form dendrite-like structures with the primary arms aligned in the principal directions of the elastic field. We demonstrate that when the far-field flux decreases with the effective radius of the system, tip-splitting and dendrite formation can be suppressed, as in the one particle case. Depending on the initial position of the precipitates, we further observe that some precipitates grow while others may shrink, even when a positive far field flux is applied. We will also explore the possibility that like single precipitates, multiple precipitates may tend to attractive self-similar configurations under appropriate growth conditions

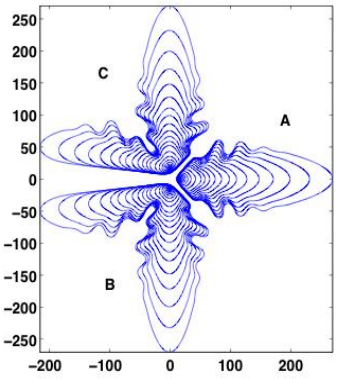

Figure 1: Three precipitates develop dendritic structures under applied shear strain.

The authors thank support from NSF.

[1] S. Li, J. Lowengrub and P. Leo, J. Comput. Phys. **225**, 554 (2007)

Scalable parallel kinetic Monte-Carlo without approximation

Tomas Oppelstrup[1], Vasily V. Bulatov[1], David Jefferson[1], Luis Zepeda-Ruiz[1] and George H. Gilmer[2]
[1]Lawrence Livermore National Laboratory, Livermore, California 94550, USA
[2]Colorado School of Mines, Colorado 80401, USA

Summary: We employ the Time Warp optimistic parallelization scheme to an atomistic kinetic Monte-Carlo simulation. This method of parallelization is free of approximation and delivers the same trajectory as a serial execution, while mitigating communication latencies and temporal load imbalance through speculative execution. On even moderately sized lattice gas models we obtain 1000-fold speedup compared to serial execution.

Kinetic Monte-Carlo (KMC) is a widely used simulation method in materials science, statistical mechanics, biology, etc. A KMC simulation propagates by discrete events, which move or otherwise change the state of one or more of the simulated objects (e.g. atoms). The events happen at rates given by distribution functions prescribed by each particular KMC model. There is no global time stepping, and in most cases each event require updating only a small subset of the simulated system. This structure makes KMC adapt automatically to the time scales present in the system. This particular feature has made it possible to simulate for example coarsening and radiation damage over time scales spanning many orders of magnitude, and makes KMC an ideal tool for simulation of systems with a frequently changing and broad spectrum of time scales.

The event driven nature and often inherent asynchronicity of KMC simulation makes it difficult to parallelize, and large scale parallelization has so far been approximate in nature. This has largely limited KMC simulation to serial computation, which severely limits accessible length scales.

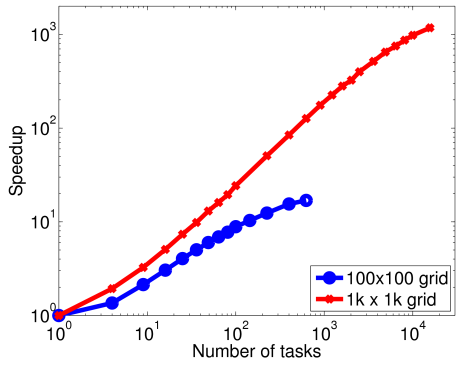

Figure 1: Speedup of lattice gas model compared to serial execution.

In this work, we employ the optimistic parallelization scheme Time Warp[1,2] where the synchronization problem inherent in any parallel computation is handled through speculative execution, and the promise to undo incorrect computations when causality errors are detected. Initial results show that a 1000-fold speedup can be attained in even moderately sized lattice gas problems (see Figure 1). The availability of an exact method for atomistic KMC, free of potential artifacts resulting from approximate schemes, will allow multi-micron length scales to be reached, while retaining the inherent multi-scale time adaptivity of KMC.

[1] D. Jefferson, "Vrtual Time", ACM Transactions on Programming Languages and Systems **7**, 404-425 (1985)

[2] P. Barnes, C. D. Carothers, D. Jefferson, "Warp Speed: Executing Time Warp on 1,966,080 Cores", ACM SIGSIM Principles of Advanced Discrete Simulation (PADS), Montreal (2013)

OptiDis: a MPI/OpenMP Dislocation Dynamics Code for Large Scale Simulations

Arnaud Etcheverry[1], Pierre Blanchard[1], Olivier Coulaud[1], Laurent Dupuy[2]

[1]HiePACS team, Inria Bordeaux Sud-Ouest
[2]SRMA, CEA-Saclay, 91191, Gif-sur-Yvette, France

Dislocation dynamics simulation is a powerful and widely used tool to model the collective behavior of large dislocations ensemble in crystalline materials. Massively parallel codes, based on message passing paradigm (MPI) have been proposed [2], [3] , and are able to simulate millions of segments over thousand of cores. Nowadays, such approach are not well suited to take advantages of new multicore architectures. The goal of OptiDis project is to develop a new parallel code for such architectures. This code is based on the physics and algorithm of the Numodis code [1].

On the numerical side, we consider the non-singular force field formulation introduced by Cai et al. [4] and we use a fast multipole method based on Lagrange polynomials to compute the force with a linear complexity. This kernel was introduced inside our fast multipole library ScalFMM [5].

On the computer science side, we developed new data structures and algorithms to be more efficient with hierarchical architectures, both in term of memory and computing units. First of all, we introduced an adaptive cache-conscious data structure to manage large set of segments and nodes. It combines good computing performance and data management. Second of all, thanks to the octree used by the FMM, we split the domain on each processor, we use a message passing paradigm to communicate between sub-domains. Each sub-domain is mapped on one multicore machine. Moreover, inside a subdomain (i.e. one multicore architecture) we consider a task parallelism, based on the OpenMP standard to speed up computation inside a node. Thanks to the hybrid parallelism (MPI/OpenMP) we are able to improve computation efficiency, and to decrease the memory used for the parallelism.

On the physical side, particular attention was given to the junction formation algorithm that allows to explicitly treat the individual interactions between dislocations and radiation-induced loops. This approach was validated by carefully comparing individual interactions obtained with dislocation dynamics [6] and available molecular dynamics simulation results [7].

Finally, two kinds of results will be shown. The first one will validate our strategies in term of parallelism. Then, we will introduce a large scale simulation involving more than one million segments in order to study the clear band formation mechanism in irradiated zirconium.

[1] Numodis: http://www.numodis.fr/
[2] Z. Wang, N. M. Ghoniem, S. Swaminarayan, and R. LeSar, “A parallel algorithm for 3D dislocation dynamics,” J. Comput. Phys., vol. 219, no. 2, pp. 608–621, 2006. [3] A. Arsenlis et al, Modelling Simul. Mater. Sci. Eng. 15, 553-595 (2007).
[4] W. Cai et al, J. Mech. Phys. Solids 54, 561-87 (2006).
[5] ScalFMM: Parallel Fast Multipole Library for Large Scale Simulations http://scalfmm-public.gforge.inria.fr/
[6] J. Drouet et al, J. Nuclear Materials (2013) in press
[7] A. Serra, D.J. Bacon, Modelling Simul. Mater. Sci. Eng. 21, 045007 (2013)

Fast Multipole Method with Application in Dislocation Dynamics

Chao Chen[1], Sylvie Aubry[2], Clif Dudley[1]
Tom Arsenlis[2], Eric Darve[1]

[1]Stanford University, Stanford, California 94305, USA
[2]Lawrence Livermore National Laboratory, Livermore, California 94550, USA

In the dislocation dynamics (DD) code ParaDis [1], as in other DD codes, the exact computation of forces on the dislocation nodes using the anisotropic elasticity formalism [2] has a complexity of $O(N^2)$, where N is the number of segments. This calculation can be sped up to O(N) using the fast multipole method (FMM). The black box FMM [3] using the Chebyshev interpolation requires a small pre-computation time. We combined it with the singular value decomposition to guarantee a minimal number of coefficients under the L^2 norm. In our tests, we considered both the isotropic as well as the recently implemented anisotropic version of the ParaDis. For DD, we found that a fifth order Chebyshev interpolation is sufficient and that it achieves an accuracy of three digits even for very anisotropic cases.

One special difficulty with DD is that, because of the tensorial nature of the interaction kernel in the FMM, the memory requirements of the method are very large. For that reason, we also implemented uniform grids with Lagrange interpolation which cuts the storage from $O(n^6)$ to $O(n^3)$, where n is the number of grid points in every dimension. In addition, the new scheme enables using the fast Fourier transform in the multipole-to-local operations, which is the most expensive step in the FMM. According to our experiments, the new scheme has a much lower increasing rate.

We have also developed a reference FMM implementation based on the Taylor expansion of the kernel. The Chebyshev FMM is a general purpose method that leads to O(N) FMMs for any analytic kernel, while the Taylor expansion method was specially developed for DD. Having both methods fully implemented, we will present comparisons of both methods in terms of memory footprint, accuracy and computational cost.

[1] Arsenlis, A and Cai, W and Tang, M and Rhee, M and Oppelstrup, T and Hommes, G and Pierce, T G and Bulatov, V V. “Enabling strain hardening simulations with dislocation dynamics,” Modelling and Simulation in Materials Science and Engineering (2007)
[2] Aubry, S and Arsenlis, A. “Use of spherical harmonics for dislocation dynamics in anisotropic elastic media,” Modelling and Simulation in Materials Science and Engineering (2013)
[3] Fong, W and Darve, E. “The black-box fast multipole method,” Journal of Computational Physics (2009)