

CSCI 653: HIGH PERFORMANCE COMPUTING AND SIMULATIONS

Fall 2019 (section: 30398D—lecture; 30399R—discussion)

- Instructor:** Aiichiro Nakano; office: VHE 610; phone: (213) 821-2657; email: anakano@usc.edu
Lecture: 3:30-4:50pm M W, LVL 13—lecture; 3:30-4:20 pm F, VPD 106—discussion
Office Hours: 4:30-5:20pm F, VHE 610
Course Page: <http://cacs.usc.edu/education/cs653.html>
Textbooks: A. Grama, A. Gupta, G. Karypis, and V. Kumar, *Introduction to Parallel Computing, 2nd Ed.* (Addison-Wesley, 2003)—recommended
D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications, 2nd Ed.* (Academic Press, 2001)—recommended
Prerequisites: (1) CS596 (Scientific Computing and Visualization) or (2) basic knowledge of numerical methods, parallel computing, and 3D graphics.

Course Description

Provide students with advanced techniques that are common to high performance computer simulations in science and engineering. Both deterministic and stochastic simulation algorithms for particles and continuum will be implemented on massively parallel and distributed computing platforms, and the simulation datasets will be visualized and analyzed in immersive and interactive virtual environment.

Syllabus

1. Deterministic particle simulation algorithms: Survey of molecular dynamics (MD) simulation—spatiotemporal data locality in MD; fast computation of electrostatic interaction— $O(N)$ fast multipole method; multiple time stepping—fuzzy cluster dynamics
2. Parallel computing frameworks: Parallel algorithm design—divide-conquer—“recombine” parallelization, spatial vs. particle vs. force vs. tuple decomposition, data-driven parallelization; load balancing—wavelet-based computational space decomposition, recursive spectral bisection, spacefilling-curve decomposition, load diffusion; scalability analysis; optimization of parallel scientific applications; new architectures—multicore and GPU acceleration
3. Deterministic continuum simulation algorithms: Survey of quantum dynamics (QD) simulation; fast solutions of partial differential equations (PDE)— $O(N \log N)$ fast Fourier transform, $O(N)$ wavelet transform, $O(N)$ multigrid method; $O(N)$ Lanczos and Davidson algorithms for the eigenvalue problem; Newton Krylov-subspace solvers for nonlinear equations
4. Multiscale particle-continuum simulation: Hybridization techniques—minimizing model-boundary artifacts, modular algorithm design, adaptive hybridization; $O(N)$ multiscale optimization; space-time multiscaling
5. Stochastic simulation algorithms: Survey of Monte Carlo (MC) simulation—estimator, importance sampling, Markov chain, Metropolis algorithm; simulated annealing; kinetic MC—master equation, Poisson process
6. Distributed scientific computing: Grid/cloud programming—Grid-enabled message passing interface (MPI-G2), Grid remote procedure call (Ninf-G), MapReduce; Grid/cloud-enabled applications—virtualization-aware scientific algorithms based on data-locality principles; Distributed MC applications—parallel replica and replica exchange MC
7. Scientific data visualization and analytics: Interactive visualization of large datasets in immersive virtual environment—hierarchical/probabilistic culling algorithms; topology analysis—shortest-path circuits, parallel graph algorithms; scientific data mining; data compression; singular value decomposition for low-rank approximations; integration of simulation, data visualization and analytics workflows on Grid/cloud
8. Advanced scientific computing methods: Local and global optimization in MD—physics-based preconditioning of iterative solvers, basin-hopping algorithms, disconnectivity-graph analysis of the energy landscape; accelerated long-time dynamics—path-integral sampling, ensemble mean-field method, hyper dynamics, activation-relaxation metadynamics; explorative search—pathfinders

Grading Scheme

Programming assignments (4-5 assignments), 60%; paper presentations, 20%; final “GitHub repository”, 20%
A (100-90%); A- (90-85%); B+ (85-80%); B (80-75%); B- (75-70%); C (70-60%); D (60-50%)