ERDC MSRC RESOURCE

Major Shared Resource Center

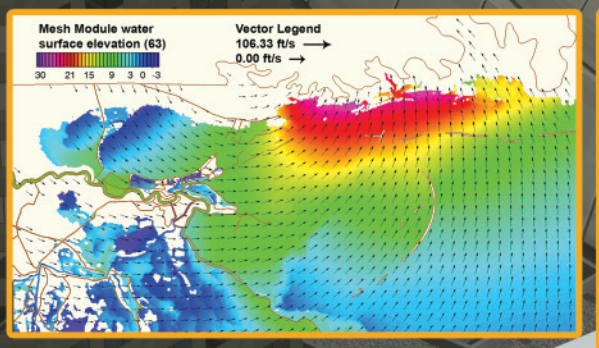

HIGH PERFORMANCE COMPUTING FOR THE WARFIGHTER. SPRING 2006

Contents

Nonlinear Response of Materials to Large Electric Fields: A Capabilities Application Project on the Cray XT3

By Dr. James Freericks, Department of Physics, Georgetown University

Nonlinear effects in nature are often the most important and most interesting effects in science. Many electronic devices either have their operation based on their nonlinear characteristics (like the transistor) or have the regime where they stop operating properly determined by where nonlinear effects become too strong (like intermodulation distortion in passive microwave filters). There is much interest, from both a basic science and an applications perspective, in understanding nonlinear effects in strongly correlated materials. This is because strongly correlated materials have many properties that can be tuned by simply applying pressure, changing the temperature, or by chemical doping. Examples of these materials include the high-temperature superconductors, rare-earth magnets, Kondo-effect and heavy Fermion materials, Mott insulators, and so on. Strongly correlated materials are characterized by having a strong enough electron-electron interaction that the independent particle picture (i.e., band theory) does not apply, and their behavior is determined by complex quantum-mechanical dynamics. The most interesting materials are those that lie close to a phase boundary between different types of materials, like the boundary between a metal and an insulator, or between a magnetic and a nonmagnetic system.

As the military becomes more and more "high-tech," it relies on electronics in both the analog and digital realms for all forms of operation. These electronics devices are susceptible to attack or damage from large electric fields (either from natural sources or from high-energy, pulsed-beam weapons). In addition, the military is moving toward "smart" electronics, which have elements that can be tuned "on the fly" to respond to different needs of the user. Strongly correlated materials, because of their inherent high tunability, are solid candidates for use in such devices. Unfortunately, little is known about their response to these kinds of electric fields, and modeling can provide useful answers to help engineer and design the next generation of smart electronics devices.

Guest writer Dr. James Freericks received his bachelors degree (*summa cum laude*) in Physics from Princeton University in 1985 and his Ph.D. in Physics from the University of California, Berkeley, in 1991. After spending time as a postdoctoral fellow at UC's Santa Barbara and Davis campuses, he went to Georgetown University, where he has been a Professor of Physics since 1994. He has been funded by the Office of Naval Research, starting with a Young Investigator Program Award, since 1996, and has been involved in high performance computing on DoD machines since 2001. This was his first CAP.

This Capabilities Application Project (CAP) provides the first numerically exact solution to this long-standing problem.

Strongly correlated materials require the solution of the so-called many-body problem to understand their behavior. For years, the many-body problem was viewed to be essentially intractable, but significant progress has been made over the past two decades as computational algorithms became more sophisticated, and as new ideas emerged for successfully solving the quantum-mechanical problem. Dynamical mean-field theory [1] was introduced in 1989 as a new way to approach this problem. In the ensuing 17 years, it has been employed to solve most of the "classic'" models in many-body physics (Hubbard, periodic Anderson, Falicov-Kimball, Holstein, etc.) [2] in equilibrium. These solutions have provided tremendous insight into the Mott metal-insulator transition and the role of strong electron correlations in real materials.

Our project extended this highly successful approach from the equilibrium case to the nonequilibrium case (in the presence of an external electric field), which allows the nonlinear response of the material to be calculated. The formalism for solving the nonequilibrium many-body problem (in an arbitrary strength external field) was worked out in the early 1960s [3,4]. We solve the simplest many-body model that illustrates this behavior. The model has two kinds of particles on a crystal lattice: mobile electrons, which can move from one lattice site to the nearest neighbor site, and localized electrons, which do not move [5]. When a mobile electron sits on the same site as a conduction electron, there is an energy cost (of magnitude U) because of

the mutual electrical repulsion between the particles. If U is large enough, and the number of mobile plus localized electrons equals the number of lattice sites in the crystal, then *all* electrons become frozen, and the material cannot conduct electricity – the system undergoes a metal-to-insulator transition as a function of U.

In nonequilibrium calculations, we need to work with objects called Green's functions, which describe the distribution in energy of available quantum states in the system, and the statistical occupation of those states as they are driven by the external electric field. Our formulation works in real time, where the transient effects are included, and the steady-state response only builds up after a long period. The Green's functions are continuous matrix operators defined on the so-called Kadanoff-Baym-Keldysh contour, which involves evolving the system forward in time from some starting point to some ending point, and then de-evolving the system backward in time from the ending point to the starting point. One of the numerical challenges of this approach is that we need to discretize the continuous matrix operator so we can perform matrix operations on it via computer-based linear-algebra techniques. We are limited, by computational time (and to some degree by memory), in how large a matrix we can use in the calculations, which ultimately limits how far out in time the calculations can go.

The formalism is straightforward, but complex. The noninteracting solution in arbitrarily large electric field is known [6], so we need to add in the effects of the electron-electron interactions. We use the dynamical mean-field theory approach to do this, which requires

The formalism is straightforward, but complex.

us to solve a set of nonlinear matrix-valued equations self-consistently. Describing the quantum-mechanical basis for this algorithm is beyond the scope of this article. Instead we focus here on the numerical issues and the use of the XT3 for large-scale parallel programs. The most challenging numerical part of the algorithm involves evaluating a two-dimensional Gaussian-weighted matrix-valued integral to determine the Green's function. The integrand requires a matrix inversion and a matrix multiplication to determine its value (the matrix is a general complex matrix). We use Gaussian integration and average the results of a discretization of 100 points per dimension and 101 points per dimension for a total of 20,201 quadrature points, and hence 20,201 matrix inversions and multiplications per iteration step. Since the matrices in the integrand are independent of one another, this part of the algorithm is easily parallelized in the master/slave format, by sending the matrix inversions to different slave nodes. After the integration is completed, the master node needs to perform an additional four matrix inversions to complete one step of the iteration. These calculations must proceed in a serial fashion, because results from the previous calculation are required for the next. Because of this structure, we expect the computational time for the code to be expressed as a sum of these two pieces $T = T_0+T_1/N$, where T_0 is the time for the serial part of the code, T_1 the total time for the parallel part, and N the number of slave nodes during the run. Of course, during the time that the serial

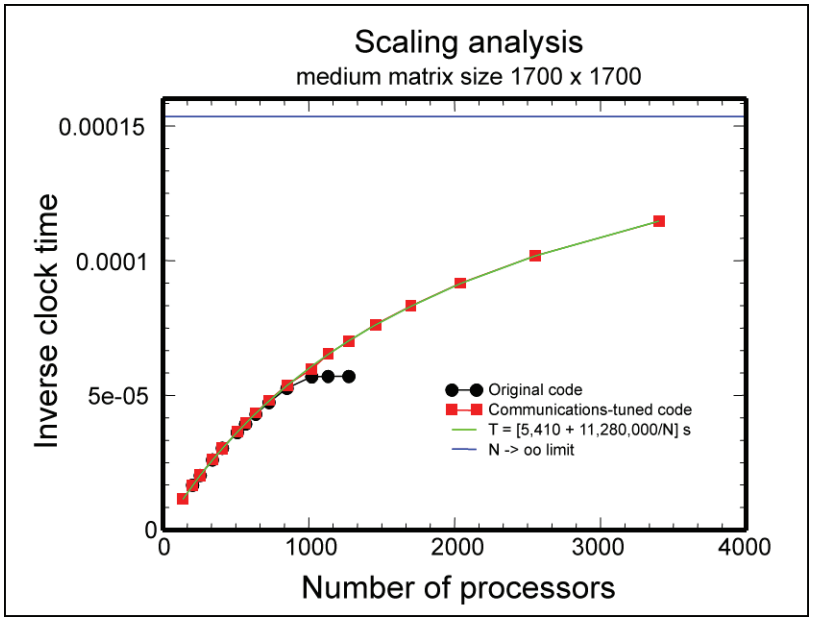

Figure 1. Run time scaling analysis for a medium-sized code on the Cray XT3. The black symbols are the original code; the red are the communications-tuned code. The green is the fit to the expected form because of the serial part of the code. The blue is the limiting result for N = 10,201 processors or more

part of the code is running, the slave nodes are idle, and the relative fraction of time for this part of the code increases as the number of nodes increases. Hence the code cannot display complete linear scale-up with the number of nodes, but we will see that we are able to perform a linear scale-up of the parallel part of the code up to the maximal number of nodes that it makes sense to run on. Since the serial part of the code is at most only about half of the total computational time (and often is much less), we run with an efficiency of 50 percent or more.

The code is quite robust, having been run on Beowulf-type machines, the T3E, and the X1, prior to being ported to the XT3, so it was believed that there would be little required to get the code to scale well to large numbers of slave nodes. This turned out not to be true, as can be seen from our scaling attempts on a medium-sized problem of 1700X1700 matrices, with the black symbols in Figure 1. Once we reached about 1,000 processors, the code stopped scaling with the number of slave nodes. We quickly realized the problem here was a communications bottleneck of using a many-to-one communications method to send the matrix results from each slave node to the master node for accumulation of the quadrature. This problem was easily fixed by using a recursive-binary-gather operation to accumulate results on the nodes, in turn, and then send the final results to the master node. One can see the scaling behavior returns by looking at the red symbols for the communications-tuned code in Figure 1. The recursive gather scheme involves taking all the active slave nodes, dividing them in half, and sending the results from one half of the nodes to the other half. Once a slave node has sent its results, it becomes inactive. The process is repeated until only one active slave node remains. At this point, it sends the accumulated results to the master node, and the serial part of the code starts. The serial code takes the local Green's function matrix (which resulted from the two-dimensional integration) and performs four more matrix inverses to construct the next approximation to the self-energy matrix. The system is checked for convergence, and if not converged, the iterations continue. Typically, we need between 10 and 50 iterations to reach convergence. Once the calculations have converged for the Green's functions, we can directly extract the electron distribution functions. These data tell us how the electrons directly evolve in response to the electric field. In addition, we also investigate the current as a function of time, to see the transient effects of oscillations.

The results of the scaling analysis show that the code runs well on a large number of processors. The rollover as N becomes large occurs because of the serial part of the code. If we follow the initial slope, for the linearly scaling part of the curve, we can see that in virtually all results, the efficiency is better than 50 percent; so choosing an optimal number of processors requires us to balance the efficiency factor versus the actual clock time for the code. On the XT3, we find a medium-sized problem, with 1700X1700 matrices taking about 1,100 hours per iteration, with between 15 and 30 iterations needed for convergence. Increasing the size to 2200X2200 resulted in an increase of the computational time to about 2,000 hours per iteration (because the matrix inversion is the limiting part of the code, the time should scale as the cube of the matrix size). We could not increase the problem size larger than this on the XT3, because the memory per node was just 1GB at the time. Results for different discretization sizes need to be scaled to the limit where the discretization size goes to zero to achieve reliable results. We could show, by comparing to exact sum rules [7], that our results improved in accuracy from an average error of a few percent to less than one-tenth of a percent when we could scale to the zero discretization limit. Our CAP project used about 600,000 CPU hours for the scaling analysis and production runs over a period of about 6 weeks. We are currently preparing these results for publications in peer-reviewed journals.

References:

[1] W. Metzner and D. Vollhardt, "Correlated Lattice Fermions in d=∞ Dimensions," Phys. Rev. Lett. **62**, 324 (1989).

[2] A. Georges, G. Kotliar, W. Krauth, M. J. Rozenberg, "Dynamical mean-field theory of strongly correlated Fermion systems and the limit of infinite dimensions," Rev. Mod. Phys. **68**, 13 (1996).

[3] L. P. Kadanoff and G. Baym, *Quantum Statistical Mechanics*, (New York, Benjamin, 1962).

[4] L. V. Keldysh, "Diagram Technique for Nonequilibrium Processes," J. Exptl. Theoret. Phys. (USSR) **47**, 1515 (1964) [Sov. Phys. JETP **20**, 1018 (1965)].

[5] L. M. Falicov and J. C. Kimball, "Simple model for semiconductor-metal transitions: SmB_6 and transition-metal oxides," Phys. Rev. Lett. **22**, 997 (1969).

[6] V. Turkowski and J. K. Freericks, "Nonlinear response of Bloch electrons in infinite dimensions," Phys. Rev. B **71**, 085104 (2005).

[7] V. O. Turkowski and J. K. Freericks, "Spectral moment sum rules for strongly correlated electrons in time-dependent electric fields," Phys. Rev. B **73,** 075108—1-15 (2006).