

Nonlinear Response of Strongly Correlated Materials to Large Electric Fields

J.K. Freericks and V. Turkowski
Department of Physics, Georgetown University,
Washington, DC
{freericks, turk}@physics.georgetown.edu

V. Zlatić
Institute of Physics of the University of Zagreb,
Zagreb, Croatia
zlatic@ifs.hr

Abstract

Strongly correlated electron materials are materials where the electron-electron interaction is so strong it is of primary importance in determining how the materials behave. These materials may be tuned to pass through a metal-insulator transition as a function of chemical doping, pressure, or temperature. This tunability of their properties makes them good candidates for so-called “smart materials” that may change their properties to respond to the particular needs of a device. We employ a massively parallel algorithm to solve exactly for the response of these strongly correlated materials to the presence of a large electric field, including all nonlinear and non-equilibrium effects. The work was performed as a CAP Phase II project on the ERDC XT3 computer in the Winter of 2006. The code displayed good scalability up to approximately 3000 processors although most production runs employed about 1500 processors, because of machine load.

1. Introduction

In this work, we examine the problem of what happens to a strongly correlated material when it is exposed to a large electric field. Ohm’s law^[1] describes the response of materials in the linear region, where the current is proportional to the field; here we are interested in what happens as the field is made larger—so large that the linear-response approach breaks down. While this may seem like a formidable problem, the underlying theory and formalism for treating the nonequilibrium quantum statistical mechanics problem was developed over 40 years ago.^[2,3] Unfortunately, little progress has been made since then on exactly solving the resulting equations due to their complexity, although numerous approximate calculations have been performed. In this work, we report on a Capabilities Application Project that utilized the Engineer Research and Development Center (ERDC) Cray XT3 machine (which has 4096 processors) to solve these equations for a range of different types of

materials. In total, about 600,000 CPU hours were used on the project over about a six week period.

The reasons we are interested in the nonlinear behavior of the material or device are practical. Nonlinear behavior is often employed in devices to achieve specific goals, such as amplification of a signal, fast switching, sensing external fields or chemical species, etc. As we move toward developing so-called “smart electronics”, we need to understand how the device properties may be tuned, and what effect that has on the nonlinear behavior. Strongly correlated materials may play an important role in such tunable devices, because they often can show high tunability of their properties. In addition, we are confronted with larger and larger fields across a sample as electronics move into the nanoscale. When a feature size is on the order of 100 nm, a potential difference of 1 V produces an electric field of $E \sim 10^7$ V/cm over the feature area. Such a large field is expected to cause nonlinear effects. Finally, the military is interested in the robustness of electronics devices to large pulsed fields that can arise from natural sources like lightning, or from man-made sources like those employed in electronic warfare. These high energy-density short-time pulsed fields may be difficult to filter out of a device and may cause the device to “burn out.” These effects can also be treated within our formalism, but in this initial study, we concentrate on a constant electric field turned on at a specific time.

Strongly correlated electron materials require the electron-electron repulsion to be taken into account in determining how the material responds to external perturbations. In conventional metals, insulators, and semiconductors, it is adequate to ignore the mutual electron-electron repulsion, and treat all of the electrons as independent, moving in an average field created by the other electrons. This is the regime where band-theory holds and has been highly successful in describing real-materials properties. But as the electron-electron interactions are made stronger relative to the kinetic energy of the electrons, the electron correlations need to be treated properly. Hence electrons are not treated in an average way, but instead we need to take into account

where the electrons are and how they move as every other electron moves. This adds significant complexity to the quantum-mechanical treatment of the problem. It also opens the door to new and exciting physical results, such as the existence of a metal-insulator transition as the interaction between the electrons grows; this transition is called the Mott-Hubbard transition and the insulating phase is a Mott insulator (to distinguish it from the ordinary band insulator).

The simplest model which takes into account strong electron-electron correlations is the Falicov-Kimball model.^[4] This model has two kinds of electrons: itinerant electrons and localized electrons. They interact by a Coulomb repulsion when they both occupy the same unit cell of the lattice. If the number of itinerant electrons plus the number of localized electrons is equal to the number of lattice sites, then the system will undergo a metal-insulator transition as the Coulomb repulsion is increased. This is the model we examine.

2. Formalism and Numerical Algorithm

We consider the Falicov-Kimball (FK) model in the presence of an external electric field that is spatially uniform, but can be time-dependent, have an arbitrarily large amplitude. The FK model has two kinds of electrons: itinerant electrons with creation and annihilation operators c_i^\dagger and c_i for conduction electrons at site i and localized electrons with the corresponding operators f_i^\dagger and f_i . The FK Hamiltonian is

$$\mathcal{H} = -\sum_{ij} t_{ij} c_i^\dagger c_j + U \sum_i c_i^\dagger c_i f_i^\dagger f_i - \mu \sum_i c_i^\dagger c_i, \quad (1)$$

where t_{ij} is the nearest-neighbor hopping matrix, U is the on-site repulsion between c and f -electrons, and is the chemical potential of the conduction electrons. The f -electrons are localized and do not move. We choose them to occupy half of the lattice sites. We also choose the conduction electrons to occupy half of the lattice sites ($\mu = U/2$). This case is called the case of half-filling, and it can lead to a metal-insulator transition if U is large enough, because the conduction electrons will then avoid the lattice sites occupied by the f -electrons, and since there will be no empty sites left over, the entire system will be “frozen” and cannot conduct electricity well.

The uniform electric field $\mathbf{E}(t)$ is described by a vector potential $\mathbf{A}(t)$ in the Hamiltonian gauge where the scalar potential vanishes:

$$\mathbf{E}(t) = -\frac{1}{c} \frac{\partial \mathbf{A}(t)}{\partial t}. \quad (2)$$

We assume that the vector potential $\mathbf{A}(t)$ is smooth enough, that we can neglect the magnetic field associated

with the time-varying electric field, which is normally required by Maxwell’s equations.

The electric field is introduced into the Hamiltonian in Eq. (1) by the so-called Peierls’ substitution^[5,6] where we neglect interband transitions because we are considering only a single band model:

$$t_{ij} \rightarrow t_{ij} \exp \left[-\frac{ie}{\hbar c} \mathbf{A}(t) \cdot (\mathbf{R}_j - \mathbf{R}_i) \right]. \quad (3)$$

In this work $t_{ij} \neq 0$ only for nearest-neighbor sites i and j .

For simplicity we shall study the case of a d -dimensional hypercubic lattice in the limit of large spatial dimensions $d \rightarrow \infty$. In this limit, the electron self-energy becomes local, which simplifies both the formalism and the numerical calculations. This approximation corresponds to the dynamical mean-field theory (DMFT) limit.^[7] The simplest electric field is one that lies along the unit cell diagonal^[8]:

$$\mathbf{A}(t) = A(t)(1, 1, \dots, 1). \quad (4)$$

After the Peierls’ substitution, the “band-structure” in the electric field becomes

$$\epsilon_{\mathbf{k}} = -2t \sum_i \cos \left[a \left(\mathbf{k}_i - \frac{e\mathbf{A}_i(t)}{\hbar c} \right) \right], \quad (5)$$

with a the lattice spacing which we will take to be one. With our choice for the electric field along the diagonal, this becomes

$$\epsilon_{\mathbf{k}} = \cos \left(\frac{eA(t)}{\hbar c} \right) \epsilon_{\mathbf{k}} + \sin \left(\frac{eA(t)}{\hbar c} \right) \bar{\epsilon}_{\mathbf{k}}, \quad (6)$$

with

$$\epsilon_{\mathbf{k}} = -\frac{t^*}{\sqrt{d}} \sum_i \cos k_i \quad (7)$$

and

$$\bar{\epsilon}_{\mathbf{k}} = -\frac{t^*}{\sqrt{d}} \sum_i \sin k_i. \quad (8)$$

being generalized energy functions and t^* ’s a renormalized hopping parameter: $t = t^*/2\sqrt{d}$ in the limit $d \rightarrow \infty$,^[7] t^* will be used as our energy unit.

We find that many quantities we want to determine involve a summation over momenta of functions of ϵ and $\bar{\epsilon}$. These summations can be performed more easily by determining a joint density of states for the two energies in Eqs. (7) and (8); the result in the limit of the infinite dimensions^[9] becomes:

$$\rho_2(\epsilon, \bar{\epsilon}) = \frac{1}{\pi t^{*2}} \exp \left[-\frac{\epsilon^2}{t^{*2}} - \frac{\bar{\epsilon}^2}{t^{*2}} \right].$$

Hence, a summation over an infinite-dimensional Brillouin zone may be re-expressed as a two-dimensional Gaussian integral (the integrand turns out to be *matrix-valued*).

The Green's function describes what happens as we create an electronic excitation at a particular time and destroy it at another time. In nonequilibrium problems, there are two Green's functions that need to be considered: (i) the retarded Green's function, which determines the quantum mechanical states, and how they are distributed in energy, and (ii) the lesser Green's function, which determines how those quantum states are occupied. Both Green's functions may be determined from the so-called contour-ordered Green's function, which is defined on the Kadanoff-Baym-Keldysh contour shown in Figure 1. The definition of the contour-ordered Green's function is

$$G_{ij}^c(t, t') = -\frac{i}{\mathcal{Z}} \text{Tr} e^{-\beta \mathcal{H}(0)} \mathcal{T}_i^c c_i(t) c_j^\dagger(t'), \quad (9)$$

where $\mathcal{Z} = \text{Tr} \exp[-\beta \mathcal{H}(0)]$ is the partition function, the Fermionic operators are expressed in the time-dependent Heisenberg picture, the symbol \mathcal{T}_i^c denotes time ordering along the Kadanoff-Baym-Keldysh contour, and $\beta = 1/T$ is the inverse temperature. In this case, the lesser Green's function is found when t lies on the lower real branch and t' lies on the upper branch. The retarded Green's function can also be determined, but is somewhat more complicated to write out. The procedure to numerically calculate the Green's function is complicated, and has been sketched in References 10 and 11. Here we will focus on some of the mathematical issues and on the computational algorithm that needs to be employed.

The system of the equations is solved by iteration, starting from some initial guess for the self-energy (usually chosen to be the equilibrium self-energy). The most technically challenging piece is to calculate the local Green's function by summing the momentum-dependent Green's function over all momentum. As mentioned above, the infinite-dimensional integral can be converted to a two-dimensional integral, but the integrand is matrix valued because the momentum-dependent Green's function is a continuous matrix-valued operator. The integrand can be found from the noninteracting Green's function (in the presence of the field) and the self-energy, but it requires one matrix inversion and one matrix multiplication to be calculated. Once the local Green's function has been found, we solve an impurity problem by first determining the dynamical mean field. Next we solve for the impurity Green's function in the dynamical mean field, and finally extracting the impurity self-energy. We equate this self-energy with that of the lattice, and repeat the iteration until the self-energy stops

changing. We call this iterative solution approach the DMFT algorithm.




Figure 1. Kadanoff-Baym-Keldysh integration contour for the time variables. The time-domain cutoffs are symmetric at $\pm t_{\max}$. The direction for the integration of the line integral is indicated by the arrows. The dashed line schematically shows where we typically turn on the electric field, as represented by the vector potential; it is commonly turned on when the time is equal to $-t_{\max} + 5$. Note that for the lesser functions, we choose the first time argument on the upper real time branch, and the second time argument on the lower real branch. When the contour is discretized, we use a step spacing of Δt along the real axis, and a step size of 0.1 along the imaginary axis. All calculations presented here have $\beta = 10$ corresponding to one hundred steps along the imaginary axis.

The first numerical issue that we need to resolve is that we need to find a way to discretize the continuous matrix operators so they can be represented as discrete matrices and we can use linear algebra packages to perform operations on them. We do this by choosing a step size Δt on the top two branches of the contour, and a step size $\Delta \tau = 0.1$ on the imaginary axis. The imaginary branch always has a fixed step size of 0.1, while we vary the step size on the upper and lower branches from 0.1 to 0.025. We typically take the time cutoff at $t_{\max} = 20$. The matrices are general complex matrices, which range in size from 900×900 up to about 2200×2200 . The computational time will grow like the cube of the linear dimension of these matrices.

Next we need to evaluate efficiently the two-dimensional Gaussian-weighted matrix-valued integral. This is handled by using Gaussian integration along each of the dimensions. We take the average of the integral with 100 points and 101 points in each dimension. This results in $100 \times 100 + 101 \times 101 = 20,201$ quadrature points for each iteration (see Reference 10 for a discussion of the accuracy and limitations of the Gaussian integration scheme).

Finally, we use LAPACK and BLAS routines for the required matrix inversions and matrix multiply operations. The matrix inversions are the most time-consuming parts of the algorithm. We need 20,201 inversions to determine the local Green's function

for each iteration, which is followed by six more inversions to solve the impurity problem, which completes the DMFT iteration; typically between 15 and 40 iterations are needed to achieve a self-consistent solution.

As can be quickly seen, this algorithm has a clear parallelizable part (the matrix inversions for each quadrature point) and a clear serial part (solving the impurity problem) since the inversions for the impurity-problem solver must all take place sequentially. Hence, we expect the algorithm to be the combination of a fixed time plus a time that scales inversely with the number of processors. The algorithm is also easily parallelized in the master-slave format.

Once the DMFT algorithm has been solved, then we may determine the properties of the system as a function of time. One interesting quantity is the current density that is driven by the external electric field:

$$\mathbf{j}_l(T) = -i \frac{et^*}{\sqrt{d}} \sum_{\mathbf{k}} \sin\left(\mathbf{k}_l - \frac{e\mathbf{A}_l(T)}{\hbar c}\right) G^c(\varepsilon_{\mathbf{k}}, \bar{\varepsilon}_{\mathbf{k}}, t, t), \quad (10)$$

expressed in terms of the momentum-dependent lesser Green's function, with each vector component identical when the electric field lies along the diagonal.

It is well known that the current response to an electric field is counterintuitive for a perfect conductor that has no electron scattering; there is an *ac* response to a *dc* field due to the lattice periodicity, which does not allow the momentum of the electron to get too large before an umklapp scattering event occurs with the lattice and changes the sign of the momentum. This phenomenon is called a Bloch oscillation,^[12,13,14] and it should be seen in any material that is free enough of defects and other sources of scattering. No conventional metal has ever been grown that has small enough scattering to exhibit Bloch oscillations. Instead, the scattering occurs so rapidly, that the steady-state current is a constant, which increases linearly with the electric field until nonlinear effects take over. Bloch oscillations have been seen in semiconducting heterostructures.^[15]

Bloch oscillations are also seen in DMFT, with a time-independent electric field (E constant and $A(t) = -Ect$ for $t > 0$)^[8]:

$$j(T) \sim \sin\left(\frac{eA(t)}{\hbar c}\right) \int d\varepsilon \frac{df(\varepsilon - \mu)}{d\varepsilon} \rho(\varepsilon), \quad (11)$$

producing an oscillating current density [$\rho(\varepsilon)$ is the noninteracting density of states, which is equal to the integral of ρ_2 over $\bar{\varepsilon}$ and $f(\varepsilon) = 1/\{1 + \exp(\beta\varepsilon)\}$ is the Fermi-Dirac distribution]. The frequency of the oscillation is $\omega_{\text{Bloch}} = eE/\hbar$ and is called the Bloch oscillation frequency. We expect these oscillations to survive in the presence of scattering if the field is large enough that the relaxation time due to scattering is

significantly larger than the Bloch oscillation period. The frequency of oscillation is undoubtedly too high for the Bloch oscillations to be directly observed ($\omega_{\text{Bloch}} > 10^{12}$ Hz).

3. Scaling to Large Numbers of Processors

Because the algorithm has such a clear parallelization—sending each quadrature point to a slave node, and then accumulating the results on the master node before solving the impurity problem—it was thought that the code would easily scale inversely with the number of processors as long as the number of inversions on each of the slave nodes was larger than those needed by the impurity solver on the master node. But when we tested the scaling on the XT3, we found that once we reached a moderate number of nodes, the codes stopped scaling with the number of processors (see Figure 2). It immediately became obvious where the bottleneck was. After calculating the integrand matrix at each quadrature point, we sent the matrix to the master node before having the slave node calculate the next quadrature point. This many-to-one transfer of data caused problems, because each node completed their matrix inversions and matrix multiplications at about the same time, and everyone tried to send the data to the master node at once. Since the data is only needed on the master node after all quadrature points have been evaluated, we devised a method to work around this bottleneck, which we call the recursive binary gather operation. The procedure was to store the intermediate results on each node, accumulating the matrix as each quadrature point for the node is generated, then once all quadrature points have been calculated, the nodes are divided in half, with half of the nodes sending their data to the other half, where they are accumulated with the data that already was present on that node. Then we repeat this procedure for all active nodes until all of the data has been transferred to one slave node. That node then transfers its data to the master node. Once we implemented the recursive binary gather, we found the code to scale precisely as it should for a code with a parallelizable and a serial piece to it. The scaling fit the predicted behavior perfectly (see Figure 2).

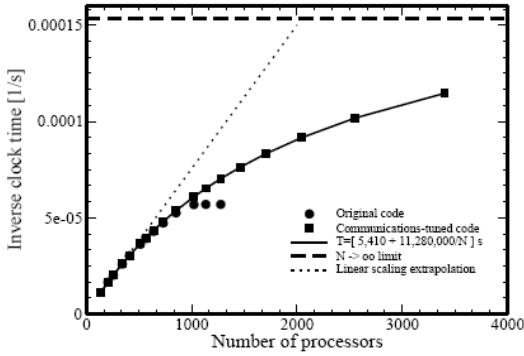


Figure 2. Scaling plot of the inverse of the total computational time as a function of the number of processors used on the XT3. The problem size is moderate (1700×1700 matrices) and we perform six iterations of the DMFT loop. If we had linear scale up of the algorithm, the inverse computational time would scale linearly, as indicated by the dotted line. But because there is a serial part of the code, the fastest the computation could be performed in theory corresponds to the dashed line at the top of the figure. The actual data was poor before we tuned the communications using the recursive binary gather, as indicated by the solid circles, which saturate at about 900 processors due to the communications bottleneck. After implementing the recursive binary gather operation, we find good scaling of the parallel part of the code (square symbols), as it agrees with the theoretical prediction (solid curve). Running at about 1500 processors provides a nice balance between overall efficiency and real-time speed. It runs at better than 70% efficiency of the linearly scaling limit (dotted line).

4. Benchmarking the Nonequilibrium Results

It is important to try to find ways to benchmark the nonequilibrium results to determine measures for how accurate the numerical solutions are. It is well known, for the case of half-filling that the imaginary part of the equal-time retarded Green's function $\text{Im}G_{ii}^R(t, t^+)$ is equal to 1, and the imaginary part of the equal-time lesser Green's function $\text{Im}G_{ii}^<(t, t)$ is $\rho_e = 1/2$. These relations follow from the definition of the Green's functions and the anticommutation properties of the Fermionic operators. Similarly, the first two derivatives of the Green's functions, which are equivalent to the first two moments of the spectral function, (i.e., multiplying the density of states by one or two powers of frequency and then integrating over frequency) also can be determined exactly^[11]. We use these results as a benchmark of the accuracy of the calculation. In Figure 3, we plot the zeroth moment for a strongly scattering normal metal with $U = 0.5$, $E = 1$, and $T = 0.1$. This system has a mean-free path of a few lattice spacings. Note how the moment is least accurate in the equilibrium case, and how the accuracy improves with decreasing step size; the

extrapolated results (pointwise quadratic extrapolation using a Lagrange interpolation formula) are almost indistinguishable from the exact result; the errors are largest at the extremal times, due, in part, to an artificial boundary effect. The calculations for this metallic phase appear to be under good numerical control.




Figure 3. Plot of the zeroth moment of the local retarded Green's function at different times with $U = 0.5$. Note that the electric field ($E = 1$) is turned on at the origin in this plot. The exact moment is equal to 1, and the curves with different symbols correspond to different values of Δt . The solid line is the extrapolated result using a quadratic Lagrange interpolation formula for the extrapolation. One can see that as the step size is made smaller, the moment comes closer to its true value, and the extrapolated result is quite accurate.

In Figure 4, we plot the second moment versus its exact result for the same strong scattering metal (the first moment exactly vanishes). Here we also see good control of the numerics, with the extrapolated results agreeing well with the exact results and being least accurate at small and large times. The lesser Green's function moments are similar, so we do not include them.




Figure 4. Plot of minus the second moment of the local retarded Green's function at different times. Note that the electric field is turned on at the origin in this plot. The exact moment is equal to -0.5625 , and the curves with different symbols correspond to different values of Δt . The solid line is the extrapolated result using a quadratic Lagrange interpolation formula for the extrapolation. One can see that as the step size is made smaller, the moment comes closer to its true value, and the extrapolated result is quite accurate.

We have results for calculations on four different types of systems: (i) a strongly scattering metal ($U = 0.5$);

(ii) an anomalous metal ($U = 1.0$, which has a resistivity that decreases as T increases); (iii) a near critical Mott insulator ($U = 1.5$, with an almost vanishing gap region); and (iv) a small gap Mott insulator ($U = 2$). We did not have enough resources to complete results for a larger gap Mott insulator, as described below. In general, the accuracy of the calculations is reduced as U increases. We found that we could, nevertheless, still perform scaling analyses for all cases except the largest U value, because the dependence on the step size was regular. In the Mott insulator phase, Δt needs to be significantly smaller to reach the scaling realm, and we ran into memory issues in being able to reach that regime (during CAP testing, the nodes all had 1 Gb of memory per node; now it has been increased to 2 Gb per node). In Table 1, we summarize the errors for the cases where the results can be scaled to $\Delta t \rightarrow 0$. The data shows the maximal percent error for the calculation, and the maximal error after the field has been turned on, because the error is usually higher in the equilibrium region, before the field is turned on. One can see the errors are under excellent control, but they do increase as the interaction strength increases, and we enter the insulating phase.

Table 1. Maximum error for the zeroth and second moments of the retarded green's function for different values of U . The first number for each moment is the maximum error over the full time range, the second is the maximal error over a range where the field has been turned on.

Percent Error in Moment	$U = 0.5$	$U = 1.0$	$U = 1.5$
Zeroth (maximal)	0.3%	1.5%	1.7%
Zeroth (max., field on)	0.1%	0.3%	1.7%
Second (maximal)	0.3%	1.7%	2.5%
Second (max., field on)	0.2%	0.16%	2.0%

We may examine the errors in the small-gap Mott insulator as well. Because the system has not yet reached the scaling regime, we are not able to show scaled results. Also, because we are limited in the size of the matrices, we need to reduce the maximal time values as Δt is made smaller. The errors in this case are quite large for large Δt . They show the right trend as Δt is made smaller, but the dependence on Δt is not yet regular enough to perform a proper scaling analysis. In Figure 5, we plot the results for the zeroth moment and a variety of different Δt values. We do not show any scaling results. Note that because the larger U results require more iterations before they converge, these calculations become increasingly more difficult to carry out.




Figure 5. Plot of the zeroth moment of the local retarded Green's function at different times with $U = 2$. Note that the electric field ($E = 1$) is turned on at the origin in this plot. The exact moment is equal to 1, and the curves with different symbols correspond to different values of Δt . One can see that as the step size is made smaller, the moment comes closer to its true value, but these results cannot be extrapolated yet.




Figure 6. Plot of minus the second moment of the local retarded Green's function at different times with $U = 2$. Note that the electric field ($E = 1$) is turned on at the origin in this plot. The exact moment is equal to -1.5 , and the curves with different symbols correspond to different values of Δt . One can see that as the step size is made smaller, the moment comes closer to its true value, but these results cannot be extrapolated yet.

In Figure 6, we show the second moment for the case with $U = 2$. The results are similar to the zeroth moment. There is an improving trend towards the exact value as Δt is made smaller, but one cannot yet successfully extrapolate the results to the $\Delta t \rightarrow 0$ limit. At the moment, these results are at the limit of acceptable accuracy for the available computational time, and for the size of the matrices that can fit within the memory.

In this section, we have shown that there is a systematic way to understand the errors in the calculations as a function of the discretization size. As Δt is reduced, the calculations improve, and if U is small enough, we can scale the results to the $\Delta t = 0$ limit. We find excellent results for this procedure when the scaling hypothesis is valid. As U increases to 2 and beyond, the calculations become more difficult to carry out to the same level of

accuracy, but we can still see the same systematic behavior.

5. Bloch Oscillations and the Nonequilibrium Distribution of Electrons

The most important measurement to make for this system is the expectation value of the current density as a function of time. The current must vanish until the field is turned on, and then, if the field is large enough, we expect to see Bloch oscillations that will decay until we reach the steady state which may still be oscillating, or it may approach a constant value.

The first issue that we address is the issue of convergence, or extrapolation to $\Delta t = 0$, for the strongly scattering metal with $U = 0.5$. In Figure 7, we plot the current for different values of Δt and for the extrapolated result (pointwise, using a quadratic Lagrange interpolation formula). Notice how the raw data has current at negative times, but that structure is vastly reduced in the extrapolated results. This is a strong indication of the fidelity of the data for this case. In addition, the results do not change too much from one value of Δt to another, which also indicates that the system is converging nicely with $\Delta t \rightarrow 0$. The data also show Bloch oscillations that decay, but we cannot tell whether they decay altogether, or survive in the steady state with the time cutoff that we have.

The data for the other values of U also extrapolates well for the anomalous metal $U = 1$ and the near-critical Mott insulator ($U = 1.5$); this data will be presented and discussed elsewhere.

In Figure 8, we show similar results for the small-gap Mott insulator ($U = 2$), but we are not able to extrapolate the data successfully. Note how the anomalous current for negative times is much larger here, and how it systematically decreases as Δt is reduced, as expected. The Mott insulator does not appear to have any remnant of the conventional Bloch oscillations remaining, but it does show irregular oscillatory behavior.

Another interesting issue to examine is how the electrons are distributed over the Brillouin zone. The distribution function depends only on the two band energies ϵ and $\bar{\epsilon}$, so it may be plotted on a two-dimensional contour plot using false color to indicate the height of the distribution function. This is done at a sequence of time steps in Figures 9–12. We examine the anomalous metal with $U = 1$, and the same parameters as before ($E = 1$ and $T = 0.1$). The distribution function here, is represented by the equal time lesser Green's function in the Brillouin zone. Because there are numerical errors in the zeroth moment sum rule for both the retarded and the lesser Green's function, we calculate this distribution function by taking the ratio of the lesser

Green's function to the retarded Green's function, which appears to partially compensate for the numerical errors.

In the equilibrium limit, the distribution function can be shown to be independent of $\bar{\epsilon}$, so it should be represented by horizontal lines. This is approximately true in our data in Figure 9. The errors should be solely numerical ones. It turns out that there are many choices to make for what Green's function should be examined. In all of our calculations, we have been working with Green's functions evaluated in the vector-potential only gauge. It is possible to make a transformation of the Green's function to a so-called gauge-invariant Green's function, but we do not do so here. The gauge-invariant Green's functions involve a time-dependent shift of the momentum wavevector, and may be constructed from our results, if desired. In the Hamiltonian gauge, the noninteracting distribution function is unchanged as the field is turned on, because it remains equal to $1/\{1 + \exp[\beta(\epsilon - \mu)]\}$, for all time.

In Figure 10, we see that in the interacting case, the distribution function begins to evolve as the field is turned on. As time increases further, the occupation becomes more and more mixed together in a swirling pattern (Figures 11 and 12). If we look at more images in turn, we can see that the distribution function appears to have a clockwise swirling motion to it, as the states become more and more evenly spread out over the Brillouin zone. It is difficult to determine just what the steady-state distribution will be, for the same reasons as why we cannot determine the steady-state current, but we definitely see the distribution function becoming more uniform around $1/2$ as the field remains on.




Figure 7. Current as a function of time for the strongly scattering metal ($U = 0.5$) with $E = 1$ and $T = 0.1$. We plot the raw current as a function of the discretization size (black, red, and green colors) and the results scaled to $\Delta t = 0$ (blue). Note how the raw results actually show some current before the field is turned on, and how that current is pushed to zero in the scaled results. The Bloch oscillations shown here decay as time increases, but we cannot tell whether they continue to decay, or survive at a lower amplitude, in the steady state.




Figure 8. Current as a function of time for the small-gap Mott insulator ($U = 2$) with $E = 1$ and $T = 0.1$. We plot the raw current as a function of the discretization size, but do not extrapolate any results. Note how the raw results show substantial current before the field is turned on, but that current does systematically decrease as Δt is reduced. The Bloch oscillations shown here have a very different character in the insulating phase, and do not give any indication of surviving into the steady-state limit.

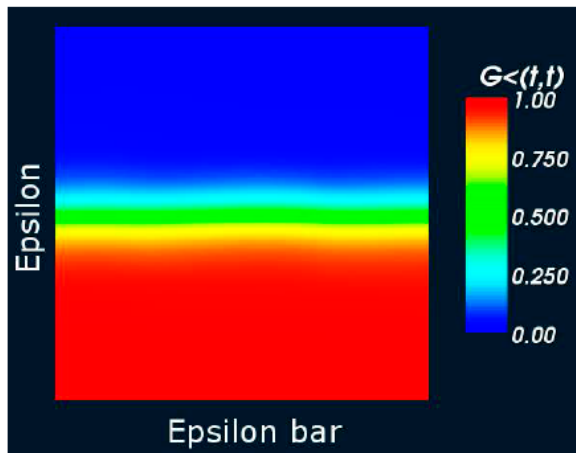


Figure 9. Distribution function for $U = 1$, $E = 1$, $T = 0.1$ and $t = -3$




Figure 10. Distribution function for $U = 1$, $E = 1$, $T = 0.1$ and $t = 1$

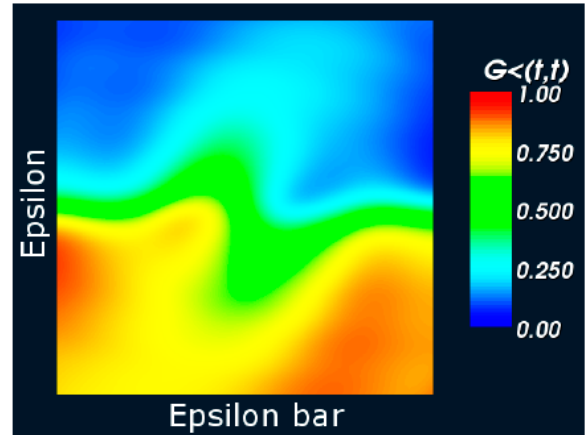


Figure 11. Distribution function for $U = 1$, $E = 1$, $T = 0.1$ and $t = 3$




Figure 12. Distribution function for $U = 1$, $E = 1$, $T = 0.1$ and $t = 5$. In these series of images (Figures 9–12), we show the distribution of the electrons in false-color plots for four different times, with the field turned on at $t = 0$. In equilibrium ($t = -3$), the distribution function has no $\bar{\epsilon}$ dependence, and hence should appear as horizontal lines. As the field is turned on, the distribution function evolves in a quite complex fashion. The range of the horizontal and vertical axes is $-4 \leq \epsilon, \bar{\epsilon} \leq 4$.

6. Conclusions

In this work, we have solved one of the longest standing problems in strongly correlated electron physics—the problem of the nonlinear response of a correlated Mott insulator to a large electric field. The solution is numerically exact, and may often be scaled systematically to zero discretization size, improving the overall quality of the data significantly.

The computational scheme involved a massively parallel algorithm which showed good scaling with the number of processors once we were able to identify and resolve the communications bottleneck. We found the XT3 performed quite well on this problem, and the ability

to use more than one thousand processors per run allowed the calculations to be completed very rapidly in real time.

The results that we found illustrated many interesting physical effects. We could see how the Bloch oscillations set in, and how they decay in time, but we could not resolve whether the decay is complete in the steady state, or whether the oscillations can survive with a smaller amplitude in the large time limit. For the insulator, we found the calculations to be much more difficult to scale, and ended up using the smallest Δt value that we calculated rather than scaling the results. The current has very interesting behavior—showing irregular oscillations, that decay, but do not have the regularity of the Bloch oscillations.

We also examined the distribution functions at different instants of time. Here we found interesting results as well—the distribution showed a swirling motion that tended to distribute the electrons more evenly over the Brillouin zone as time increased.

Overall, the CAP project was a success. We were able to accomplish solving a significant problem with an intense period of large-scale production runs and we were able to locate and resolve a communications bottleneck to accomplish near linear scaling of the code. We will continue to analyze our data and evaluate other parameters and regimes to build up a more coherent picture of the nonlinear response of strongly correlated materials, which was initiated by the success of this project.

Acknowledgments

We acknowledge support of the Office of Naval Research under grant number N00014-05-1-0078 and from the National Science Foundation under grant number DMR-0210717. Supercomputer resources (Cray XT3) were provided by the Engineering Research and Development Center (ERDC) under a DoD High Performance Computing Modernization Program CAP Phase II grant.

References

1. Ohm, G.S., *The Galvanic current investigated mathematically*, J.G.F. Kniestädt, Berlin, 1827.
2. Kadanoff, L.P. and G. Baym, *Quantum Statistical Mechanics*, W.A. Benjamin, Inc., New York, 1962.
3. Keldysh, L.V., “Diagram technique for nonequilibrium processes.” *J. Exptl. Theoret. Phys.*, 47, pp. 1515–1527, 1964, [*Sov. Phys. JETP* 20, pp. 1018–1026, 1965].
4. Falicov, L.M. and J.C. Kimball, “Simple model for semiconductor-metal transitions: SmB_6 and transition-metal oxides.” *Phys. Rev. Lett.*, 22, pp. 997–999, 1969.
5. Peierls, R.E., “Theory of diamagnetism of conduction electrons,” *Z. Phys.*, 80, pp. 763–791, 1933.
6. Jauho, A.P. and J.W. Wilkins, “Theory of high-electric-field quantum transport for electron-resonant impurity systems.” *Phys. Rev. B*, 29, pp. 1919–1938, 1984.
7. Metzner, W. and D. Vollhardt, “Correlated lattice fermions in $d = \infty$ Dimensions,” *Phys. Rev. Lett.*, 62, 324–327, 1989.
8. Turkowski, V. and J.K. Freericks, “Nonlinear response of Bloch electrons in infinite dimensions.” *Phys. Rev. B*, 71, 085104-1-11, 2005.
9. Schmidt, P., “Time-dependent dynamical mean-field theory”, *Diplome thesis*, University of Bonn, 1999.
10. Freericks, J.K., V.M. Turkowski, and V. Zlatić, “Real-time formalism for studying the nonlinear response of ‘smart’ materials to an electric field.” *Proceedings of the HPCMP Users Group Conference 2005, Nashville, TN, June 28–30, 2005*, edited by D.E. Post (IEEE Computer Society, Los Alamitos, CA, 2005), pp. 25–34.
11. Turkowski, V.M. and J.K. Freericks, “Spectral moment sum rules for strongly correlated electrons in time-dependent electric fields.” *Phys. Rev. B*, 73, 075108--15, 2006; Erratum, *Phys. Rev. B*, 73, 209902(E)-1–1, 2006.
12. Bloch, F., “Quantum mechanics of electrons in crystals.” *Z. Phys.*, 52, pp. 555–599, 1928.
13. Zener, C., “A theory of the electrical breakdown of solid dielectrics.” *Proc. R. Soc. (London) Ser. A*, 145, pp. 523–529, 1934.
14. Aschcroft, N.W. and N.D. Mermin, *Solid state physics*, Holt, Rinehart, and Winston, Philadelphia, 1976.
15. Waschke, C., H.G. Roskos, R. Schwedler, K. Leo, H. Kurz, H. and K. Köhler, “Coherent submillimeter-wave emission from Bloch oscillations in a semiconductor superlattice.” *Phys. Rev. Lett.* 70, pp. 3319–3322, 1993.