2001

Measuring Electron Gas Relaxation in Gold through Second Harmonic Generation

Paul SanGiorgio
Harvey Mudd College

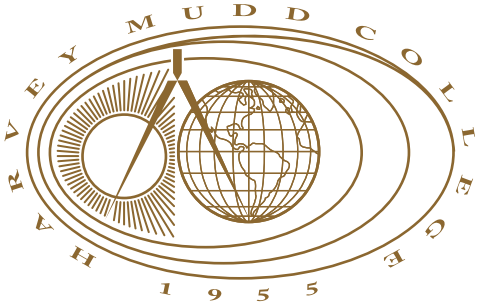

Modeling Electron Thermalization in Metals

by
Paul SanGiorgio
Tom Donnelly, Advisor

Advisor: _____

Second Reader: _____

(Lisette DePillis)

May 2001

Department of Mathematics

HARVEY MUDD
COLLEGE

Abstract

Modeling Electron Thermalization in Metals

by Paul SanGiorgio

May 2001

In a thermally equilibrated system, electron behavior in a metal is described by the Fermi-Dirac equation. With ultrafast lasers, electrons can be excited into temporary distributions which are not described by the Fermi-Dirac equation and are therefore not at a well-defined temperature. These nonthermal distributions quickly equilibrate through two primary processes: electron-electron scattering and electron-phonon scattering. In most situations, these effects are unnoticeable, since they are completed within $5\ ps$. A probabilistic numerical model for electron-electron scattering is presented. The model is robust, scaleable, and requires only one parameter. The success of the model suggests future work on a similar electron-phonon scattering model, which would provide a complete description of the elctron distribution during thermalization. Once complete, this model can be tested by measuring the amount of second harmonic light generated by an ultrafast laser in a pump-probe experiment.

Table of Contents

List of Figures

Acknowledgments

I would like to thank Kevin Moore for his original work on this problem, Nick Breznay for his assistance in the laboratory, and of course Professor Tom Donnelly, without whom, none of this would be possible.

Chapter 1

Introduction

Thermalization is the process by which a group of objects interact to reach a common statistical temperature. In general, many different processes are responsible for thermalization, and in practical experience the time scale of these interactions can be anywhere from months (temperature changes on a global scale) to milliseconds (burning a finger on a hot plate). Yet, thermalization on an atomic scale differs from practical experience in two important ways: first, the time scale involved is much shorter, and second, in many cases, an object is not at a definable temperature during the entire thermalization process. This is possible because in order for a system to be at a "temperature," it must be internally equilibrated. In atomic systems, we have the ability to perturb particles so violently and quickly that they become completely unequilibrated. Thus, easily perturbed atomic systems are our best laboratories for studying thermalization on extremely short time scales.

1.1 Nonthermal Distributions in Metals

Metals are excellent materials for studying thermalization for a number of reasons. First, it is relatively easy to induce non-thermal distributions using ultrafast laser pulses. Second, the heat capacity of the metal lattice is sufficiently larger than the heat capacity of the electrons so that the lattice absorbs little energy and we can consider only the interaction of the pulse with the electrons [2]. Finally, the be-

havior of the electrons and the processes by which they equilibrate are reasonably simple, and can be approximated using electron gas theory [13].

1.2 The Fermi-Dirac Distribution

Due to the lattice structure of metals, the electrons are not localized to particular atoms, so we imagine them as a sort of gas that is dispersed throughout the metal. We say that a particular electron state is characterized by its energy, ε, and, from the Pauli Exclusion Principle for fermions, each available state can either be occupied or unoccupied. We call the total number of available states between ε and $\varepsilon + \delta\varepsilon$ the density of states, $\mathcal{D}(\varepsilon)d\varepsilon$. It is a simple exercise of statistical mechanics [7] to show that fractional occupancy of a particular energy is given by the Fermi-Dirac distribution:

$$f(\varepsilon) = \frac{1}{1 + e^{\left(\frac{\varepsilon - E_F}{kT}\right)}} \tag{1.1}$$

E_F is the "Fermi energy" of the metal, and, along with the density of states, must be determined experimentally; k is Boltzmann's constant, and T is the temperature in Kelvin. As $T \to 0$, $f(\varepsilon)$ approaches a step function. In fact, the metal must be at a very high temperature indeed to have a sizeable curvature, as shown in Figure (1.1). With this and the density of states, we can calculate the number of electrons with energy between ε and $\varepsilon + \delta\varepsilon$:

$$N(\varepsilon)d\varepsilon = f(\varepsilon)\mathcal{D}(\varepsilon)d\varepsilon \tag{1.2}$$

This simple equation is qualitatively responsible for many metallic properties, and is the foundation for much of our further discussion.

1.3 Techniques for Probing Nonthermal Distributions

Much work has been done on the subject of electron thermalization in metals, as it has important applications to: the interaction of molecules with surfaces, sur-

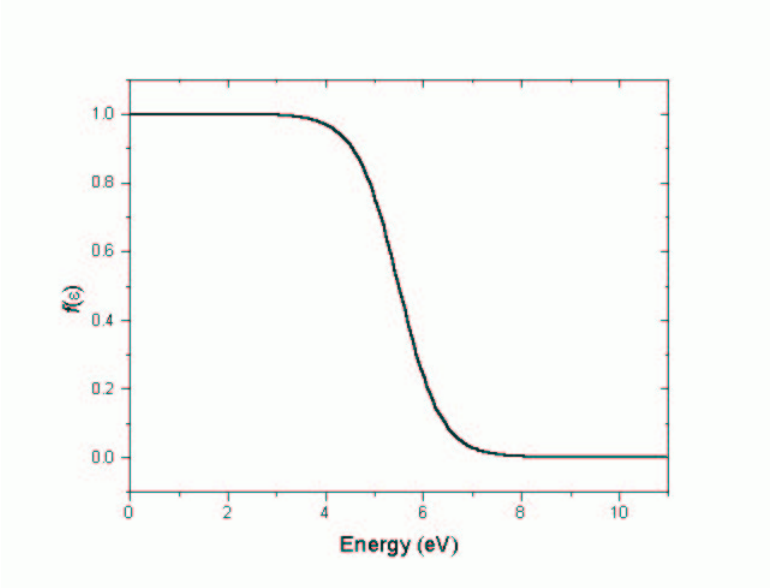

Figure 1.1: Only at high temperatures does the distribution differ noticeably from a step function. Here, $T = 5000^o\ K$, roughly the temperature of the surface of the sun.

face photochemistry, optical processing of surfaces, magneto-optical recording, and electrical charge and heat transport [12]. Since the time scale of thermalization in this case is on the order of 100 femtoseconds [2][3] ($10^{-13}s$), research in this area has only been made possible recently with the advent of the ultrafast laser. An ultrafast laser produces light pulses with durations as small as $4.5\ fs$ [10]. For a concise introduction to the physics of ultrafast lasers, see Donnelly *et al.* [1].

Ultrafast lasers are an excellent tool for perturbing electron systems, as they can deliver a large amount of energy to a very small area in a very short time. Also, there are a number of techniques which can be used to analyze the electrons once they have been perturbed. Fann *et al.* [2][3] use pump-probe electron spectroscopy to directly measure the electron distribution. In this setup, the individual photon energy ($\hbar\omega$) is large enough to eject the electron from the metal through the photoelectric effect, and one can infer the initial energy of the electron simply

by subtracting the fixed energy of the photons from the measured kinetic energy. This process, however, requires that the target be under high vacuum since, once ejected, the electrons have a very short lifetime in air. Thus, it can be prohibitively expensive. Petek *et al.* [12] use time-resolved two-photon photoemission (TR-2PP), which in addition to being in high vacuum is also done in a cryostat, yet offers unmatched temporal resolution. Other researchers have also used more mundane techniques such as transmissivity/reflectivity measurements [15] and thermal resistivity [6]. Recently, Moore *et al.* [9] used a technique called Second Harmonic Generation (SHG) which is discussed in detail below.

Chapter 2

Second Harmonic Generation

SHG is an optical effect in which a system absorbs photons with energy $\hbar\omega$ and emits photons with energy $2\hbar\omega$. As one would expect, this is a nonlinear process, and in most regimes it is insignificant compared to first-order effects (transmission, reflection, and absorption). Thus, it was not until the invention of the laser that it was demonstrated (Francken, 1961) and, even then, it was not until the advent of ultrafast physics that it became possible to produce intense enough pulses to make SHG a common laboratory tool.

Quantum mechanically, we imagine an SHG event as such: an electron absorbs one, then another photon of energy $\hbar\omega$ before emitting a single photon of energy $2\hbar\omega$.

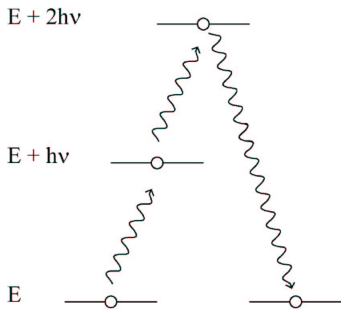

Figure 2.1: An electron absorbs one, then another photon of frequency ω before emitting one of frequency 2ω.

Since the state $\varepsilon + \hbar\omega$ is not the equilibrium state of the electron, it should decay back to a lower state quickly, thus we expect that this process will depend heavily

on the time between the first and second photon (i.e. if the second photon comes too late, then the electron will have already decayed back to its original energy level). Thus, examining the amount of second harmonic light generated as a function of time should give us some measure of the frequency with which the excited electrons are decaying. This is the fundamental idea behind our experimental efforts, and will be explored later.

2.1 First-Order Approximation

It is straightforward to show (either classically [14] or quantum mechanically [4]) that to first-order, the intensity of SHG light ($I_{2\omega}$) depends quadratically on the incident light intensity.

$$I_{2\omega} \propto I_{\omega}^{2} \tag{2.1}$$

Naively, we could have guessed this relation, if we consider that the probability that a particle will absorb a photon is proportional to the number of photons present (the incident intensity), and since SHG is a two photon process, then the probability of the event should go like the number of photons squared. Although not particularly exact, this type of probabilistic thinking will be useful later in developing a semi-classical model of electron thermalization.

2.2 Energy Distribution Effects

At low energies, the quadratic law for SHG at metal surfaces is well-known and studied. Yet, at higher intensities, a supra-quadratic behavior has been observed [11]. In order to understand the reason for this behavior, we must first look at how electron energies change during a SHG event, and how the Pauli Exclusion Principle affects the process.

As discussed in the previous section, in a SHG event, an electron absorbs two photons, then emits one. From the Pauli Exclusion Principle, we know that in or-

der for the electron to do this, there must be available excited states at both $\varepsilon + \hbar\omega$ and $\varepsilon + 2\hbar\omega$, as well as a space for it to fall back into at ϵ once it has emitted the photon. Once again, we argue that the total number of SHG events should be proportional to I_ω^2, yet we must also account for the amount of available space at the appropriate energy levels. But, why is it that this factor only becomes important at high intensity levels? To understand this, we must look at how the electrons interact with the laser.

The majority of light incident on the metal surface will be reflected. In reflection, electrons in the metal make virtual transitions to excited energy levels and then emit photons coherently. Yet, a fraction of the incident light will be absorbed by an electron that will be excited to a real energy level, and remain there without emitting a photon immediately. At high enough intensities, the laser can make a significant change to the energy distribution of the electrons; electrons within $\hbar\omega$ of the Fermi edge can absorb a photon and occupy one of the many empty spaces outside the Fermi sea. Figure (2.2) shows this effect in a typical electron distribution.

As mentioned before, in order for SHG to occur, there must be sufficient available states at the right energy levels. At $T = 0$, even though there are a large number of available states at high energy levels, an electron could not then fall back into the Fermi sea (easily), since it is almost entirely full. At some finite T, we expect that those electrons that are at partially filled energy levels will be most likely to undergo SHG since these events are proportional to the number of electrons $(f(\varepsilon)\mathcal{D}(\varepsilon)d\varepsilon)$ and the number of avialable states $((1 - f(\varepsilon))\mathcal{D}(\varepsilon)d\varepsilon)$, which is nonzero only when ε is partially filled.

Now, let's consider the effects of a high intensity laser. All of a sudden, there are a lot more electrons at partially filled energy levels due to the absorbed energy. Thus, we increase the number of electrons that can undergo SHG, which is the principle cause of supra-quadratic behavior.

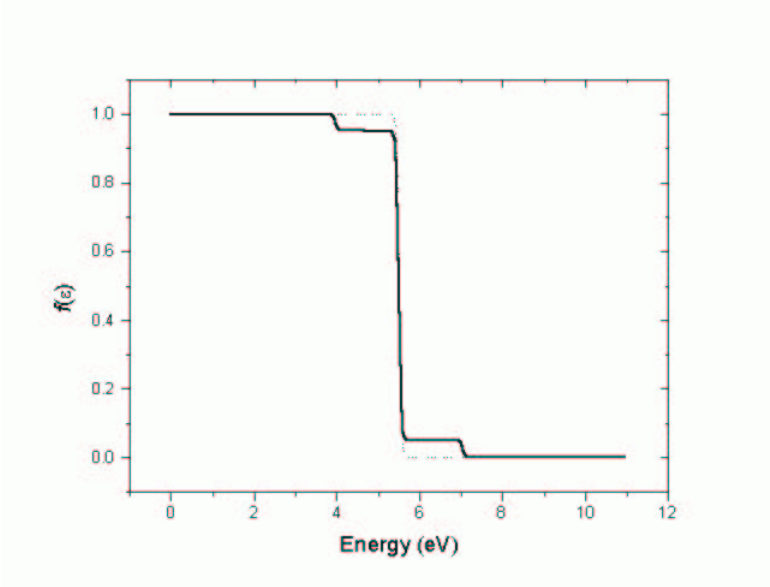

Figure 2.2: Before interacting with the laser, the electrons are in thermal equilibrium (dotted line). The light pulse moves a "block" of electrons over the Fermi edge, disturbing the equilibrium. The effect has been exaggerated for clarity by a factor of 5.

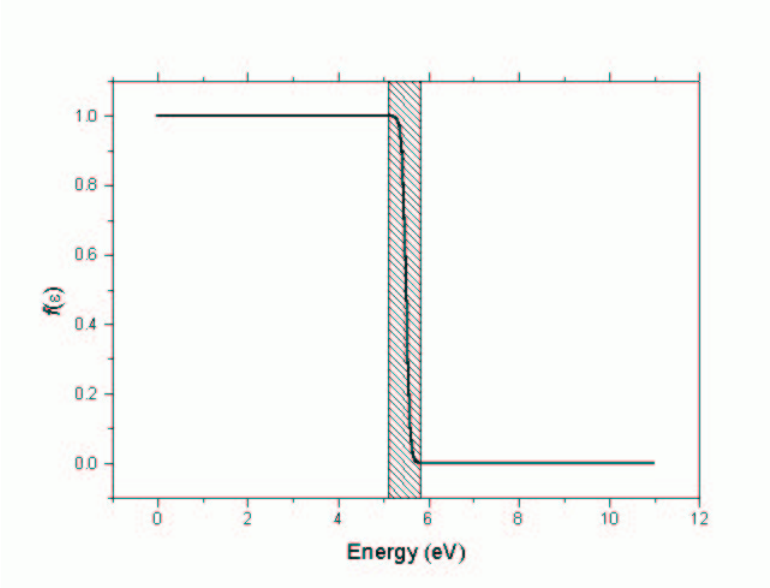

Figure 2.3: At low intensities, the distribution is not very different from an undisturbed Fermi-Dirac distribution, so onnly the shaded region has partially filled energy levels, and therefore the number of electrons that could contribute to SHG is very small.

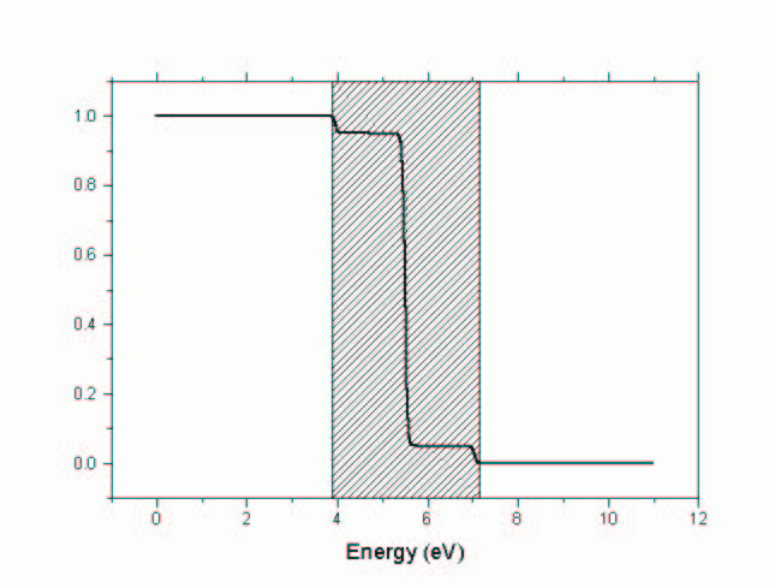

Figure 2.4: By moving a group of electrons past the Fermi edge, the laser increases the number of electrons that can undergo SHG. The magnitude of the shift has been exaggerated for clarity, but the width is realistic.

Chapter 3

Electron Thermalization

The model worked out in Chapter 2, developed by K Moore [8], adequately describes second-harmonic generated by $25\ fs$ ultrafast pulses. Yet, we know from statistical mechanics that as soon as the electrons are disturbed from their equilibrium positions, they will equilibrate, thus changing the distribution. K Moore concludes that the time scale of equilibration is much larger than the length of the laser pulse, and therefore no significant amount of equilibration occurs (this analysis is consistent with the results of Fann *et al.* [2][3]). But, as was discussed in Chapter 1, SHG should be an ideal tool for measuring thermalization processes due to its dependency on the electron distribution. Thus, it is only logical to extend K Moore's model to include thermalization processes.

3.1 Physical Processes

Petek gives a complete review of all the dynamical processes that occur at a metal surface on the femtosecond time scale [12], here I shall only attempt to review the two most important causes of thermalization: electron-electron scattering and electron-phonon scattering.

After a laser has interacted with the surface of the metal, the electron distribution resembles figure (2.2). It is clear that this does not correspond to a Fermi-Dirac distribution, and is not therefore at a definable temperature. We also know from statistical mechanics that the Fermi-Dirac distribution represents the maximal entropy of the system, so we should guess that the quickest thermalization process

will be the one that restores the Fermi-Dirac distribution. This happens through electron-electron (e-e) collisions. Essentially, there are too many "hot" electrons (i.e. with $\varepsilon > E_F$), and they will scatter elastically with "cold" electrons until a Fermi-Dirac distribution has been reached. One may be inclined to think that since the hot electrons are giving up energy to the cold electrons that this process is dissipative, but due to the elastic nature of the scatterings, the energy remains in the form of a Fermi-Dirac distribution at a higher temperature than before the excitation. Using electron-spectroscopy, Fann *et al.* [3] found that the relaxation time of excited electrons was given by:

$$\tau = \tau_0 \left(\frac{E_F}{\delta E} \right)^2 \tag{3.1}$$

as expected from Fermi Liquid Theory [13]. Here δE is $\varepsilon - E_F$,and τ_0 was experimentally determined to be $5fs$. This means that at the highest excitation levels ($\varepsilon = E_F + \hbar\omega$), $\hbar\omega = 1.55eV$, $\tau \approx 77fs$ in gold, while some of the excited states with ε closer to E_F have relaxation times on the order of picoseconds. Since there are generally not many nonthermal electrons with δE near E_F, we find that the electrons reach a well-defined temperature within $400fs$ in gold at typical intensities [2][3][12].

In addition to scattering with other electrons to reach a statistical temperature, the hot electrons (which have absorbed almost all of the laser energy) will interact with the cold lattice (which absorbs very little of the laser energy). In this process, electrons absorb and emit phonons, which are quanta of vibrations in the lattice. The excess electron energy is transformed into lattice vibrations which quickly dissipate in the bulk, and the electrons finally become fully thermalized as they cool to lattice (room) temperature. Once again, Fann *et al.* [2][3] have used electron spectroscopy to measure the dynamics of this process, and have found that the characteristic time scale of electron-phonon scattering is $\approx 2 - 4\ ps$.

3.2 A Numerical model of e-e scattering

We will approach e-e scattering in the same probabilistic manner that we addressed supra-quadratic SHG. We will say that our scattering electrons start with energies ε_i and ε_j. After the collision, they are scattered into states with energies ε_k and ε_l, respectively. Since these collisions are inelastic [12], we know that energy must

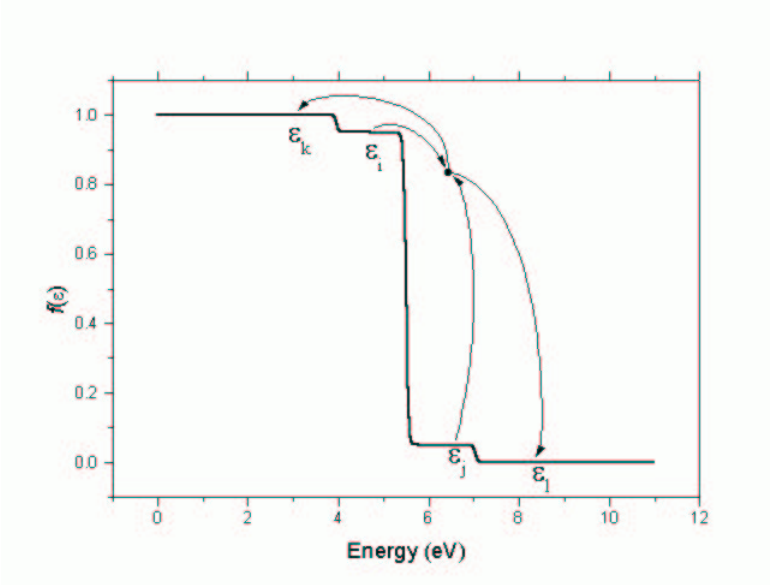

Figure 3.1: Two electrons collide elastically and are scattered into different energy levels.

be conserved (more on that in chapter 4), and of course, total electron number must be conserved. Now, the question become, how likely is this event to occur? Probabilistically, we expect that the process should be proportional to four things:

- The number of electrons with energy ε_i
- The number of electrons with energy ε_j
- The number of vacant states with energy ε_k

- The number of vacant states with energy ε_l

Motivated by this, we define the transition probability from energies ε_i and ε_j to ε_k and ε_l, $T_{ij\rightarrow kl}$, as:

$$T_{ij\rightarrow kl} = g * N(\varepsilon_i) * N(\varepsilon_j) * V(\varepsilon k) * V(\varepsilon l) \tag{3.2}$$

where $V(\varepsilon)$, the vacancy function, is defined by:

$$V(\varepsilon)d\varepsilon = (\mathcal{D}(\varepsilon) - N(\varepsilon))d\varepsilon \tag{3.3}$$

and g is a parameter with the dimensions of s^{-1} (we will look more closely at g in Chapter 4). To characterize the behavior of all the electrons, we sum $T_{ij\rightarrow kl}$ over all possible values of i, j, k, and l.

3.3 A Numerical model of e-p scattering

Most models that people have developed for understanding how the hot electron gas interacts with the cold lattice have used the simplifying assumption that the electrons have already reached a temperature, T, (i.e. that e-e scattering has already run its course) and therefore are not valid on the femtosecond time scale [2][3][5][15]. We are motivated by the simplicity of the e-e scattering model to develop a similar model for e-p scattering, i.e. one that treats each electron individually thus allowing greater flexibility.

Phonons are the vibrational modes of the metal lattice, much like photons are the vibrational modes of the electromagnetic field. Knowing this allows us to visualize the process of hot electrons bouncing off the cold lattice as an exchange of energetic particles. Following similar reasoning as above, if the phonon has an energy, $\varepsilon_p = \hbar\omega_p$, then we expect that the probability for an electron with energy ε to absorb a phonon to be given by:

$$P_{absorb} = g_a N(\varepsilon) V(\varepsilon + \hbar\omega_p) \tag{3.4}$$

Similarly, the probability to emit a phonon should be:

$$P_{emit} = g_e N(\varepsilon) V(\varepsilon - \hbar\omega_p) \tag{3.5}$$

These equations bring up a few key questions:

- Is $g_a = g_e$? Should absorption be as common as emission?
- How many different phonon frequencies, w_p, contribute to the process? Is g a function of frequency?
- Are g_a and g_e functions of ε (i.e. are different energy electrons more or less likely to interact with the lattice)?
- How do the frequencies of the phonons relate to the temperature of the lattice?

Unfortunately, I have not been able to complete this model to any degree of satisfaction, so the resolution of these issues will be left up to my eventual successor. We shall now focus our attention on the actual implementation of the e-e scattering model.

Chapter 4

Numerical Results

As is the case with any computational model, the first order of business is to discretize our system. This is done by dividing the electron distribution into N energy bins of width ΔE. Each bin will be labeled by an energy, E_i, and will contain a number of electrons, N_i. Referring to equation (1.2), it is straightforward to see that:

$$N_i = f(E_i)\mathcal{D}(E_i)\Delta E \tag{4.1}$$

Here, $f(E_i)$ represents the filling fraction of the energy bins – it does not have to be the Fermi-Dirac distribution. Since conservation of energy is a requirement, it is instructive to consider the discrete energy distribution associated with our system. First, let's consider a simple distribution that is broken up into only five bins, $f(E_i)$, with a constant density of states function, $\mathcal{D}(E) = D$. If this were the real (continuous) distribution, it would have the associated energy function $Ef(E)D$, which would look like this: It is clear that in energy space, each bin is trapezoidal, not rectangular, so the actual amount of energy in each bin is not $E_i f(E_i) D\Delta E$, but rather:

$$Energy(E_i) = (E_i + \frac{1}{2}\Delta E) f(E_i) D\Delta E \tag{4.2}$$

4.1 e-e scattering

We shall begin with equation (3.2), where we hypothesize that the transition rate from bins i and j to bins k and l is given by:

$$T_{ij\to kl} = gN(E_i)N(E_j)V(E_k)V(E_l) \tag{4.3}$$

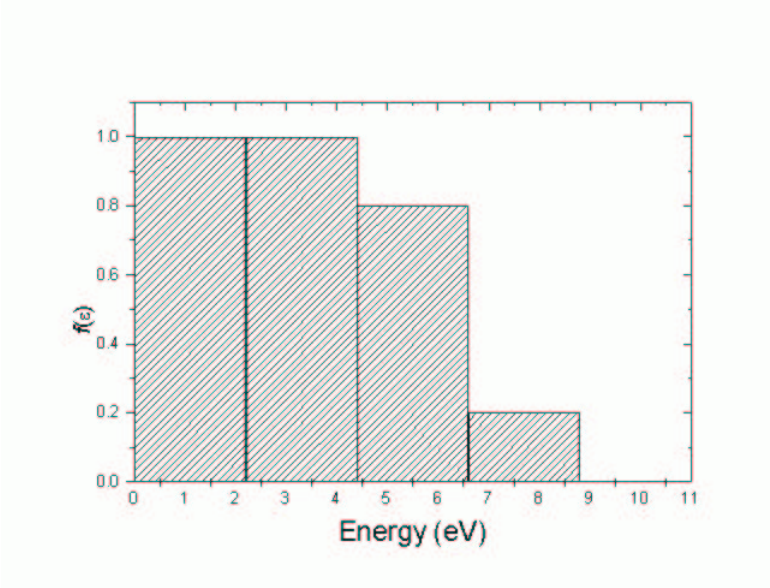

Figure 4.1: This simple distribution has two full bins, two partially filled bins, and one empty bin.

Replacing $N(E_i)$ with our definition from above, we see that:

$$T_{ij\to kl} = gf(E_i)f(E_j)(1-f(E_k))(1-f(E_l))\mathcal{D}(E_i)\mathcal{D}(E_j)\mathcal{D}(E_k)\mathcal{D}(E_l)(\Delta E)^4 \quad (4.4)$$

As it stands, this is just about enough for us to begin a numerical simulation. But, we must first understand the nature of g. We want the total number of particles that will be moving from the bins to be on the order of the number of particles in that bin. Since we are multiplying together four "numbers of electrons," it makes sense that g should carry with it a $\frac{1}{\mathcal{N}^3}$, where $\mathcal{N}$ is the total number of electrons in our system (not to be confused with N, the number of bins). Still, we are not done, as we need to make sure that the program behaves the same *regardless of the number of bins.*

$T_{ij\to kl}$ goes like ΔE^4, so as of now, it is not scalable in the number of bins. But, $T_{ij\to kl}$ refers only to a single scattering event, and the total number of events in each calculation will depend on the number of bins, so we must consider how

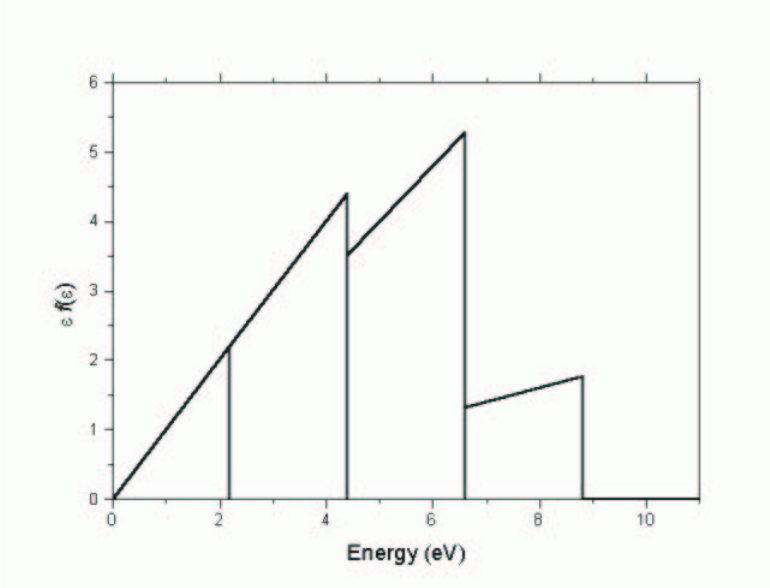

Figure 4.2: To calculate the energy in a distribution, we multiply by E. This is the resulting function.

$T_{ij\to kl}$ is implemented before we can understand how it scales. We claimed earlier that all was needed to do was sum T over all possible values of i, j, k, and l to simulate a scattering event. From symmetry between i and j and also between k and l, we know that we only have to consider $i = 0 \to N-1$, $j = i \to N-1$ and $k = 0 \to N-1$, $l = k \to N-1$. Also, each transition must conserve energy, so that the total amount of energy leaving bins i and j should equal the energy entering bins k and l.

Bin	Before Transition	After Transition
i	$(E_i + \frac{1}{2}\Delta E) f(E_i)\mathcal{D}(E_i)\Delta E$	$(E_i + \frac{1}{2}\Delta E)(f(E_i)\mathcal{D}(E_i) - T_{ij\to kl})\Delta E$
j	$(E_j + \frac{1}{2}\Delta E) f(E_j)\mathcal{D}(E_j)\Delta E$	$(E_j + \frac{1}{2}\Delta E)(f(E_j)\mathcal{D}(E_j) - T_{ij\to kl})\Delta E$
k	$(E_k + \frac{1}{2}\Delta E) f(E_k)\mathcal{D}(E_k)\Delta E$	$(E_k + \frac{1}{2}\Delta E)(f(E_k)\mathcal{D}(E_k) + T_{ij\to kl})\Delta E$
l	$(E_l + \frac{1}{2}\Delta E) f(E_l)\mathcal{D}(E_l)\Delta E$	$(E_l + \frac{1}{2}\Delta E)(f(E_l)\mathcal{D}(E_l) + T_{ij\to kl})\Delta E$

After summing up both columns, setting them equal to each other, and canceling terms, this simplifies to:

$$E_i + E_j = E_k + E_l \tag{4.5}$$

Since we discretized our system linearly, $E_i = \Delta E i$, so this further simplifies to:

$$i + j = k + l \tag{4.6}$$

Computationally, this means that once we have specified i, j, and k, l is automatically determined, so we have three sums instead of four. Now, we can return to the issue of choosing g so that it scales properly.

Consider the total flux through a single bin in a time step. This can be calculated using the balance law:

$$\frac{\Delta N(E_i)}{\Delta t} = \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} (T_{jk \to i(j+k-i)} - T_{ij \to k(i+j-k)}) \tag{4.7}$$

This equation, although not entirely true (to be precise, we need to adjust our summation limits to ignore negative indices), but it has the right basic form. By the way that we have defined everything, the only variables that are dependent on the bin width are ΔE and N, so if we can arrange for all of those to cancel out, we will be OK. On the left side of equation (4.7), we have one ΔE (in the $N(E_i)$), while on the right, we have four (all of which are in T). One might then conclude that we are haplessly left with three unwanted ΔE's on that side that need to be canceled. But, we also have N's in the limits of the summation. Since this is a double sum, the number of terms will go like N^2. And, since ΔE goes like N^{-1}, this effectively cancels out two more of our ΔE's, and we are left with one ΔE. This implies that as we make our approximation better and better (i.e. $\Delta E \to 0$), the transitions would eventually stop. One could have predicted that we would be in this situation from the beginning, since by conservation of energy, we are choosing exactly which bin

the second scattered electron goes to, but as we go from continuous to discrete, the number of states at that energy approaches 0. One approach to solve this would be to assume that energy is not exactly conserved and, from an uncertainty principle argument, distribute the scattered electron over a group of bins with width ΔE_l given by:

$$\Delta E_l \Delta t \geq \frac{\hbar}{2} \tag{4.8}$$

where Δt is the average scattering time. This approach adds unnecessary complications to implementing the model, and it is much more efficient to adjust g so that there is a factor of ΔE in the denominator of T to cancel this out. This is most easily done by a factor of N in the numerator. So, the final form of g is:

$$g = \frac{N}{(\mathcal{N})^3} \tag{4.9}$$

up to some constant which, if we have adjusted everything properly, only carries the time information and should not have to be changed as we adjust the discretization.

The model is essentially complete. All that remains is to test it.

4.2 Relaxation to a Fermi-Dirac Distribution

Since the purpose of this program is to relax a group of non-thermal electrons into a thermalized Fermi-Dirac distribution, this should be the principle measure of its value. We will test our model with a simple "staircase" distribution with a constant density of states (figure (4.3)). This distribution is a somewhat exaggerated example of a laser-perturbed electron distribution. If we choose a reasonable value for g – here we mean one that is calculationally convenient – the distribution will relax within approximately 200 time steps to figure (4.4). The results are compared to a best-fit Fermi-Dirac distribution, and if we look at the residuals, we see that the points closest to E_F relax the slowest, but in general, the distribution has relaxed

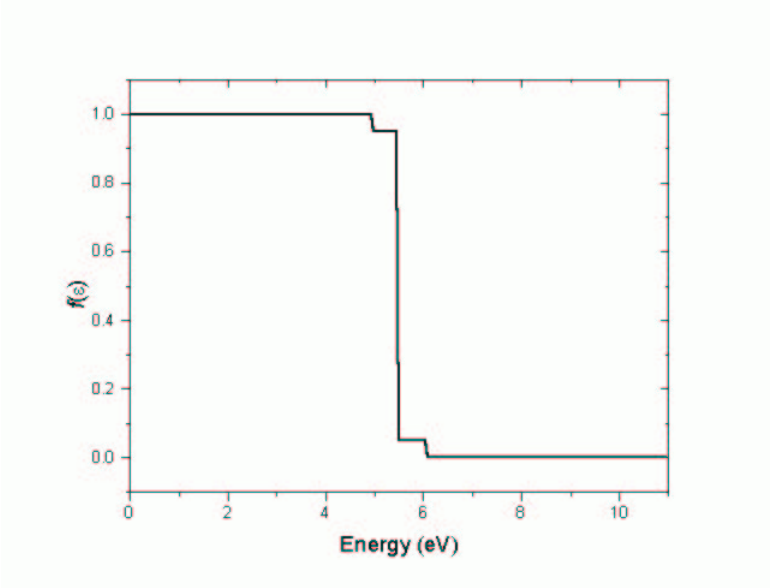

Figure 4.3: An idealized laser-disturbed electron distribution.

beautifully. This example alone is not enough to demonstrate the effectiveness of the model. Let's begin with a much more abnormal distribution, with an equally abnormal density of states (figure (4.5)). After running the relaxation program for enough time (the program is stopped when the changes in the distribution are too small to be noticed) and then fitting to a Fermi-Dirac distribution, we get figure (4.6). The model is robust enough to deal with any possible initial distribution. But, have we really made the proper adjustments to g so that we can discretize the system as we please?

4.3 Scalability of the Method

Given that the model runs in time $O(N^3)$, where N is the number of bins, it is very important indeed to determine what the minimum value of N is that retains sufficient detail. Clearly, this will depend on the exact circumstances of the experiment,

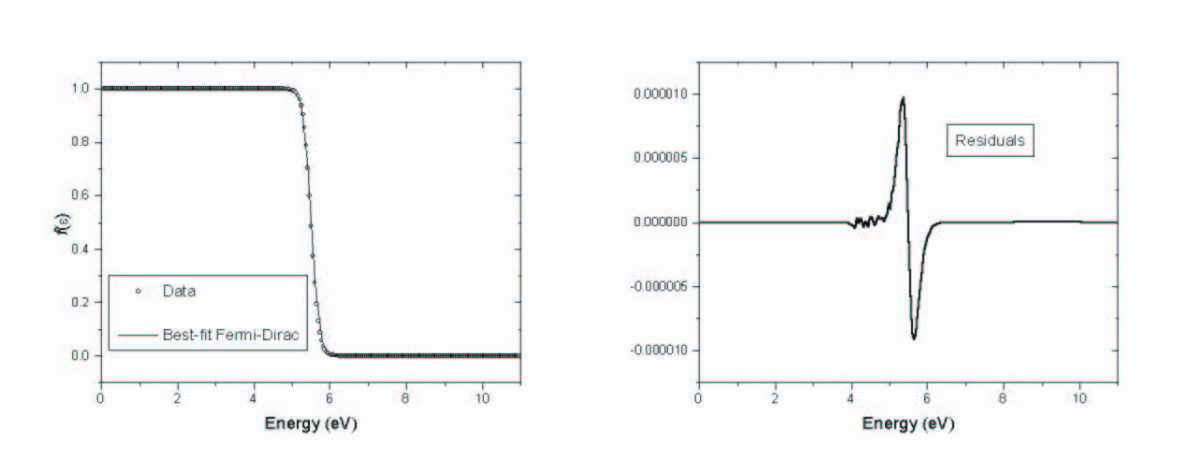

Figure 4.4: With only the Pauli Exclusion Principle as a guide, the nonthermal distribution relaxes to a Fermi-Dirac distribution within any degree of accuracy. Residuals are largest near E_F, as predicted by Fermi Liquid Theory.

so in our case, we will use a simulated laser-perturbed distribution, with a constant density of states (figure (4.3)) for the remainder of these runs.

We begin by dividing the distribution into 200, 400, 600, 800, and 1000 bins. As it stands, my program absorbs the factor of ΔE into the density of states function, so that the "200" distribution has a density of states, $\mathcal{D}(E_i)$ which is 5 times greater than the density of states of the "1000" distribution. This is due more to poor programming than anything else, and should ultimately be changed. Nevertheless, the first thing to check is that the different representations actually relax to the same Fermi-Dirac distribution. To test this, the program was run for a long enough time that the change in each distribution was on the order of one part in ten million. Then, the distributions were fitted to Fermi-Dirac distributions, with the only free parameters being the temperature, kT, and the Fermi energy, E_F. The results are summarized in figure (4.7).

The results are a little bit disturbing. First, we note that the Fermi energies are consistent, which is an indication that we are conserving the particle number

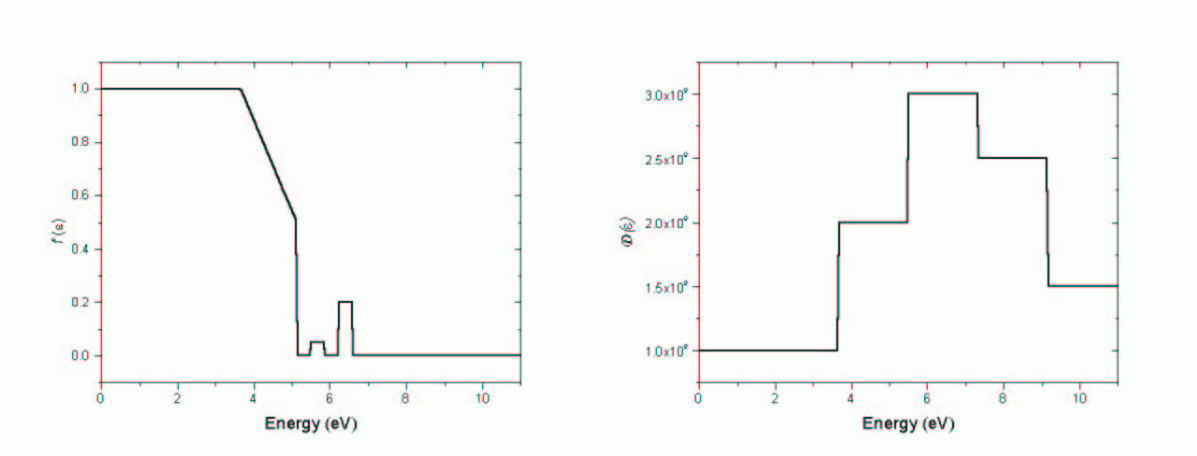

Figure 4.5: Although this does not correspond to any physical system, this should represent an extreme test of the robustness of the model.

well, but why do the temperatures seem to approach a limit as $N \rightarrow \infty$? The temperature of the system is the best measure of the amount of energy deposited, so one is lead to believe that the finer binings have somehow created more energy than the coarser ones. However, if we look at the amount of energy in the systems, we find that it is equal across the board, both before and after relaxation (assuming of course that we are using the proper equation for finding the energy!). So, why do the coarser binings *appear* to have less energy? This turns out to be a simple consequence of the discretization itself. If we consider a continuous Fermi-Dirac distribution, and then discretize it into N bins, figure (4.8) shows the relationship between total energy and number of bins.

The energy sum increases as a function of N, finally levelling off at a a rate of $\approx \frac{1}{N}$. Why does this happen? Well, because the distribution has a negative slope to it, and using this simple method of discretization, the fewer bins there are, the less of the negatively curved region you will be able to capture. In the initial distribution, which has no curvature, there is no error. But, the more bins one uses, the higher an energy they will calculate for a given Fermi-Dirac distribution, so,

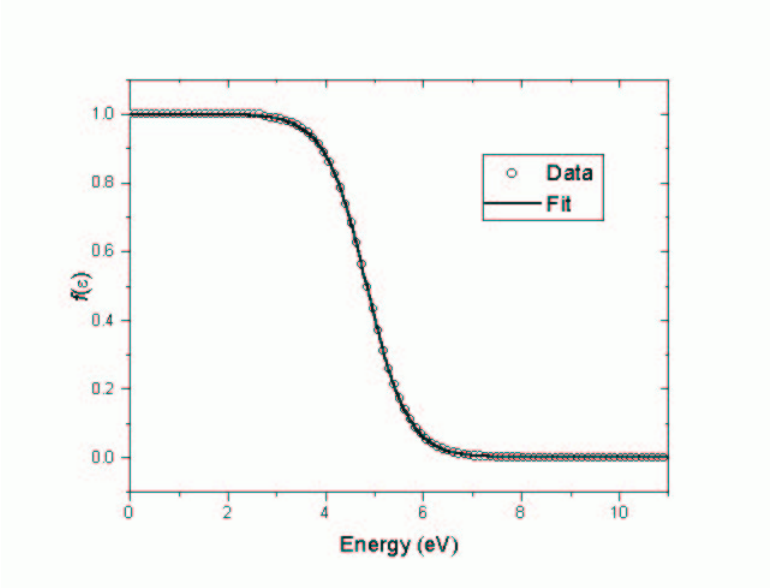

Figure 4.6: Despite its bizarre origins, the distribution relaxed smoothly back to a Fermi-Dirac distribution. The slight kink at $E = 2.75\ eV$ is an artifact of a calculational short-cut, described in section 4.4.

since my program conserves energy exactly, it is impossible for all of the binings to relax to the same Fermi-Dirac distribution, because that would mean that they had different energies. Therefore, the coarsest bining relaxes to the lowest (apparent) energy, and the finest the highest (apparent) energy (as measured by kT). Finally, though, none of this matters too much, as another glance at figure (4.7) tells us that the difference in temperature between the "200" and "1000" binings is a paltry 4 degrees Kelvin. This is well within the level of accuracy that we expect from our measurements, so we now move on to the time evolution of the systems.

If g is adjusted properly, each different bining should relax at the same rate toward the final Fermi-Dirac distribution. For the following set of runs, the distribution was deemd "relaxed" if the maximum fractional change in electron number in a bin was less than one part in ten million – this roughly corresponds to changes of single electrons. The amount of time each distribution took to achieve this is

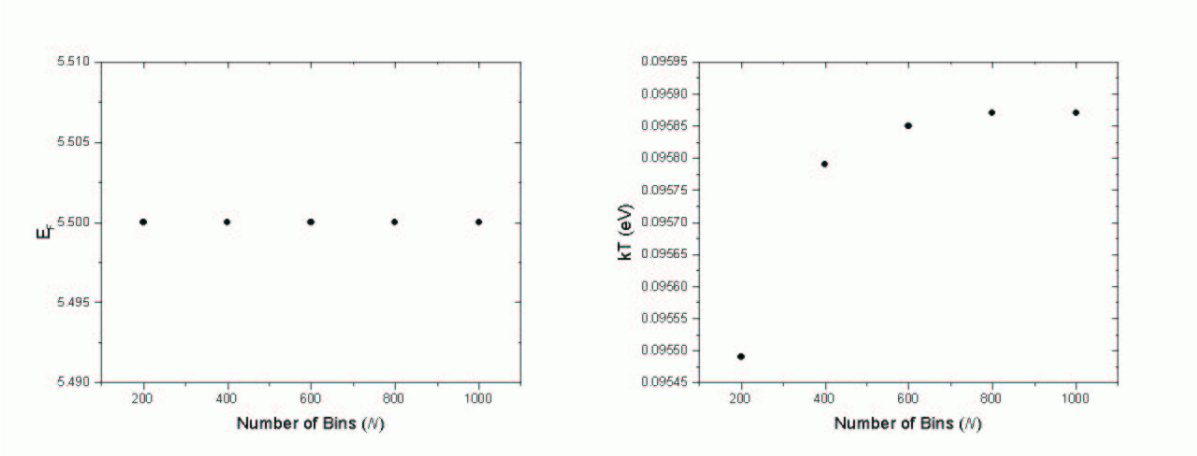

Figure 4.7: The value of E_F is constant for each bining, but kT increases as a function of the number of bins due to discretization.

plotted in figure (4.9).

Once again, we see a similar "leveling off" behavior. First, we note that we must be choosing g properly just because these numbers are all roughly the same; the factors that we have inserted into g are on the order of 10^9 (total number of electrons) and 100 (number of bins). If we had the wrong functional form, we would expect the relaxation time to be off by similar orders of magnitude. But, then, why are they different at all? This is a remnant of equation (3.1), which says that the electrons nearest the Fermi edge will relax the slowest. Due to the nature of discretization, as we increase the number of bins, we will necessarily have more electrons closer to E_F than in coarser binings. And, if we look at the residuals, we see that it is the single electron bin which is closest that requires the most time to relax (I have not included graphs of the residuals, since they very closely resemble figure 4.4. Thus, this small difference in relaxation time is not something that will go away, but at least we are satisfied that the program is functions properly.

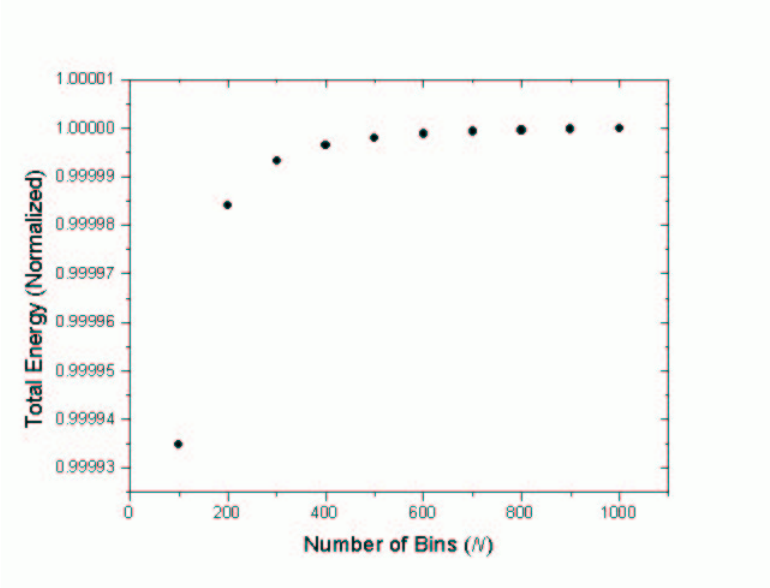

Figure 4.8: Due to the nature of the discretization, an approximation of the total energy will be too small, but approaches the correct value as $N \to \infty$.

4.4 Implementation

Now that we have established the scalability of the model, it begs the question: how much should we scale it? How many bins are sufficient? Since there is some amount of "transient" behavior that goes away as we increase the number of bins (i.e. the anomalously low temperature, the fast relaxation times), it makes sense that we want enough bins that we are realistically modeling a continuous distribution, but not too many. As mentioned before, the run time of this program goes like N^3, which means that while a 200 bin distribution executes a time step very quickly (less than half a second), the 1000 bin distribution takes nearly 125 times as long (more than a minute). The first thing we can do to combat this problem is focus our program more effectively. Since only the electrons near E_F are going to move appreciably, we can limit our calculation of T to just those bins. This was accomplished by "cutting off" the first and final quarter of the distribution,

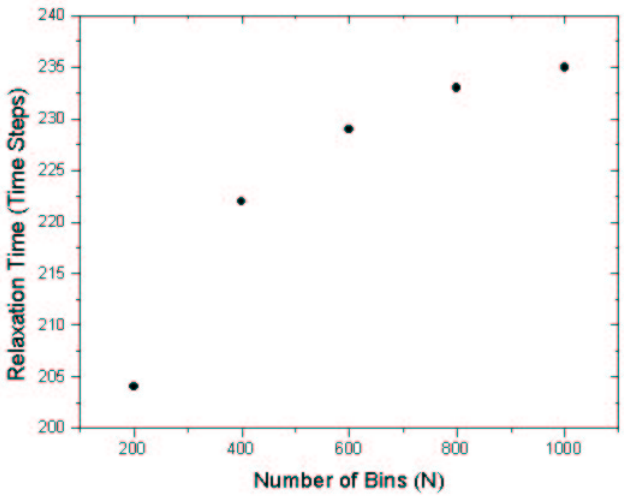

Figure 4.9: Although the relaxation times are not identical for different binings, they are close enough to suggest that the difference is due to the very nature of discretizing the problem, and not a systematic error in g.

and appending them unchanged after the calculation had been made. This alone increases the speed of the program eightfold, but since a typical relaxation time is ≈ 400 steps, this still means that the 1000 bin distribution will take 50 minutes to run. One could further increase the speed by decreasing the number of bins; how coarse a bining one would choose would depend on how sensitive the model needs to be to with respect to the experimental data.

Further development of the model might include implementation on a parallel processing cluster, such as the Harvey Mudd Beowulf cluster. Each transition probability can be calculated independently, so this calculation should be ideal for parallel computing.

Chapter 5

Experimental Attempts

The purpose of the modeling is, of course, to predict or verify an experimental result. Unfortunately, without a working e-p scattering model, the e-e scattering model cannot currently make relevant predictions. Nonetheless, a significant effort has been made in the lab to obtain time-resolved SHG data.

5.1 The Pump-Probe Setup

In contrast to the work done by K Moore *et al.* [8], we are using a two pulse pump-probe setup to gather SHG data. A transform-limited $800\ nm$ pulse $15 \sim 50\ fs$ in duration is split by a 50/50 beam splitter, with the pump arm focused at a nearly perpendicular angle, and the probe arm incident at 45^o. A precision translation stage on the probe arm changes the delay time, τ, in steps of $1\ \mu m$, which corresponds to a temporal difference of $6.6\ fs$. The laser is p-polarized, and is focused to spot sizes of $\sim 5\ \mu m$ by $8\ \mu m$. Pulse duration is adjusted through a prism-pair compressor and is measured by an interferometric autocorrelator. Temporal and spatial beam overlap is found using a KDP doubling crystal. The target, a $100\ nm$ thick polycrystalline gold surface, was chosen due to its relatively low cost, its high damage threshold, and its ability to withstand atmospheric conditions without oxidizing. At $800\ nm$, gold has a fundamental reflectivity of 0.986 and a skin depth of $20\ nm$ [16].

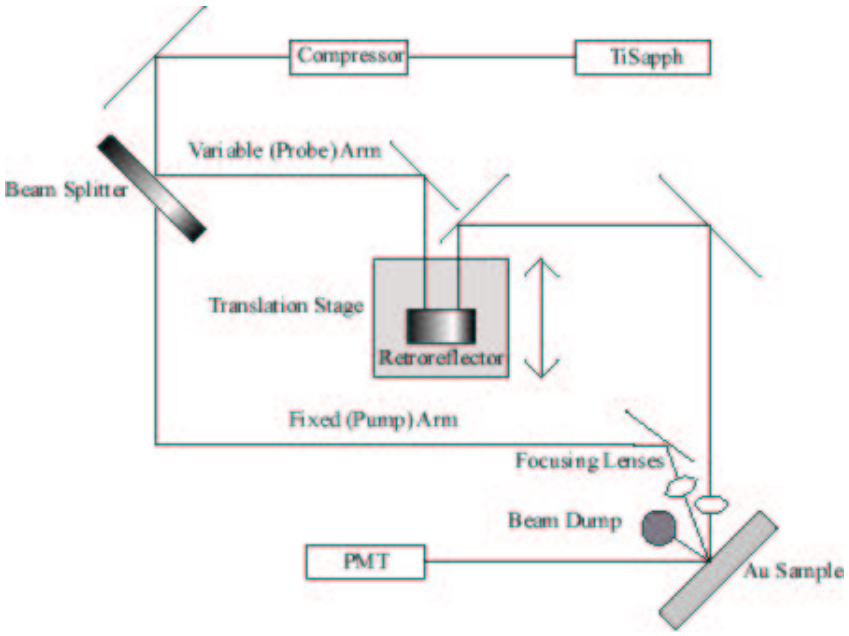

Figure 5.1: The Experimental Setup

5.2 Experimental Difficulties

Presently, we have not obtained any (interesting) time-resolved data; there seems to be no change in SHG output versus τ. There are a number of possible reasons for this. In our setup, only the second harmonic light generated by the probe beam is being captured, and it may be that we are in the regime in which the SHG light is not highly dependent on the electron distribution. So, despite the fact that the pump beam "warps" the distribution, and then the probe beam comes along and warps it again, we may be at such a low intensity that the presence or absence of the pump beam may not mean very much. In order to test this, only a minor modification needs to be made to K Moore's "femtoSHG" program. If we assume that even after two pulses we are not in a regime where thermalization will play a key role, we can compare the SHG output for when the pump and probe come one after the other and when the pump energy has been completely dissipated by the

time of the probe beam.

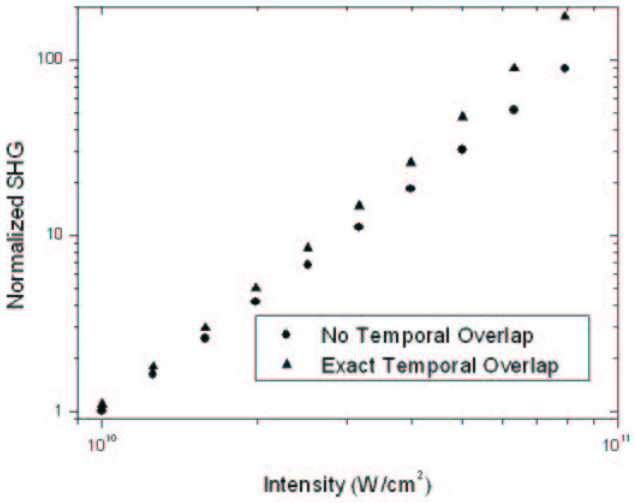

Figure 5.2: Even at relatively low intensity levels ($< 30\ GW/cm^2$) the change in SHG should be on the order of a few percent. At high levels, the difference between overlapped and non overlapped pulses is nearly an order of magnitude

Even at relatively low intensities ($10\ GW/cm^2$), there should still be an $\sim 7\%$ enhancement to the SHG. Currently, we are able to get $35\ GW/cm^2$ on target, so this should correspond to a difference of nearly 25%. This is well within the range of our sensitivity, and the fact that we are not seeing this suggests that our two beams are not well overlapped. Although our overlapping method seems to be robust and stable, it is not clear that we are maintaining the required spatial overlap at the gold surface; this will be explored in further work.

Chapter 6

Conclusion

There is still much work to be done on this experiment, both technically and theoretically. Most importantly, the process for overlapping the pump and probe beams must either be verified or corrected; and if the overlap is good– but our intensity is still too low – then the intensity of the laser must be increased. Using a standard ultrafast amplifier, pulses with durations of $50\ fs$ can be generated with enough intensity to cause time-delay dependent effects. Theoretically, the first order of business is the creation of a viable electron-phonon scattering model. Even a crude model will be enough to generate meaningful predictions about the electron distribution which could guide laboratory work. Also, once complete, the model could be easily extended to calculating third harmonic generation intensities (a similar effect), or possibly any other distribution-dependent effect in metals.

The model described in this paper is a simple tool for understanding the time evolution of electron distributions. It is scalable to any realistic level of accuracy, it requires only a single input parameter (τ_0), and it can be used without modification to model e-e scattering in any metal. I am confident that further work in this field will yield a simple, accurate model of electron thermalization that will correctly predict SHG at metal surfaces.

Appendix A

Source Code

`relaxx.cpp` is the main code and can be compiled for unix machines using the command `g++ -g relaxx.cpp -orelaxx`. The program is designed to be flexible; the code can be changed to output whatever data is desired with little adjustment.

```
/*
Electron Relaxation Program : Version X
Filename:  relaxx.cpp
Author:  Paul SanGiorgio
Date:  26 April 2001
The puprose of this program is to simulate electron-electrons
scattering using probabilistic methods.  The input should be
of the form:
E[0] f[0] D[0]
E[1] f[1] D[1]
 :    :    :
Where E[i] is the energy of bin i, f[i] is initial filling fraction,
and D[i] is the total number of states with energy E[i].  The program
will run for a fixed amount of timesteps, or until the changes in the
distribution are less than a chosen value.  The output of the file
is of the form:
E[0] f[0] f'[0]
E[1] f[1] f'[1]
 :    :    :
E and f are the same as the input file, and f'[i] is the final
filling fraction.  This can easily be changed to output whatever
neccessary.  The "main" section of this code is at the bottom, with
the functions listed first.
Note:  I have found that a good value of g (i.e. one that does not
overflow the bins) is usually around 10 - 100.  The program will run
most quickly when g is chosen to be as large as possible without
overfilling the bins.
*/
#include <iostream.h>
#include <fstream.h>
#include <string.h>
#include <math.h>
//variable initialization.
int tempnumbins, numbins, numsteps;
double g, normalizer;
double temp1 = 0, temp2 = 0;
// This function calculates the absolute value of a number
double absvalue(double number) {
  if (number > 0)
    return number;
  else
    return -number;
}
// This function clears an array.
void cleararray(double trans[]) {
```

```
  for (int i = 0 ; i < numbins ; i++) {
    trans[i] = 0;
  }
}
// This calculates the transition number for
// electrons in bins i and j going into bins k and l.
double transition(double num[], double numstates[],
  int i, int j, int k, int l) {
  temp1 = num[i]*num[j]*(numstates[k] - num[k])*(numstates[l] - num[l]);
  temp1 /= normalizer;
  temp1 *= g;
  return temp1;
}
/* This evolves the program one timestep.  The transition from i,j to
   k,l is calculated, then put into a transition array (it is added to
   k and l, and subtracted from i and j).  After all the probabilities
   have been calculated, the transition array is added to the original
   number array, thus advancing the program one time step.
*/
double evolve(double num[], double trans[], double numstates[]) {
  double avgdif(0);
  cleararray(trans);
  for (int i = 0 ; i < numbins ; i++) {
    for (int j = i ; j < numbins ; j++) {
      for (int k = 0 ; k < numbins ; k++) {
if ( (i + j - k >= 0) && (i + j - k < numbins)
     && (i + j - k >= k)) {
  temp2 = transition(num, numstates, i, j , k , i + j -k);
  trans[i] -= temp2;
  trans[j] -= temp2;
  trans[k] += temp2;
  trans[i + j - k] += temp2;
}
      }
    }
  }
  // This calculates the average change in each bin.  This will be
  // used to measure how much change is taking place.
  for (int w = 0; w < numbins; w++) {
      avgdif += absvalue(trans[w]/numstates[w]);
  }
  avgdif /= numbins;
  // Here, we add in the transition array into our number array.  Checks
  // need to be made to see if we are overfilling or underfilling any
  // bins.  This will happen if g is to large.
  for (int z = 0; z < numbins ; z++) {
    num[z] += trans[z];
    if (num[z] > numstates[z] || num[z] < 0) {
      if (num[z] > numstates[z])
cout << "Overfilled bin! " << z << " " << trans[z]
     << " " << num[z] << " " << numstates[z] << "\n";
      if (num[z] < 0)
cout << "Negative bin! " << z << " " << trans[z]
     << " " << num[z] << "\n";
      exit (-1);
    }
  }
  return avgdif;
}
main() {
  char infilename[25], outfilename[25];
  cout << "Name of data file: ";
  cin >> infilename;
  cout << "Name of output file: ";
  cin >> outfilename;
  cout << "Enter g: ";
  cin >> g;
  cout << "Enter maximum number of steps: ";
  cin >> numsteps;
  cout << "Enter number of bins: ";
  cin >> tempnumbins;
  cout << "Reading File" << "\n";
  ifstream data_file;
  data_file.open(infilename);
  ofstream out_file;
  out_file.open(outfilename);
```

```
  if (data_file.bad()) {
    cerr << "Error: Could not open " << infilename << "\n";
    exit (8);
  }
  cout << "\n";
  // We only want to deal with the "middle half" of the distribution
  // since the edges are not going to be involved in this proccess.
  // Thus, the tempxxxxx variables refer to the original (full)
  // distribution, while its counter part xxxxx refers to the halved
  // distribution.  Some arrays (energy[]) do not need counterparts,
  // since they are not involved in the calculation at all.
  numbins = tempnumbins/2;
  double energy[tempnumbins];  // Contains "x" values of all datapoints
  double numstates[numbins];   // the number of states in a bin
  double tempnumstates[tempnumbins];  // temp. number of states in a bin
  double num[numbins];  // the number of electrons in a bin
  double tempnum[tempnumbins];  // the temp. number of electrons in a bin
  double trans[numbins];  // the transition array
  double init[tempnumbins];  // the initial distribution
  // this reads in the input into the proper array
  for (int i = 0; i < tempnumbins ; i++) {
    data_file >> energy[i];
    data_file >> init[i];
    data_file >> tempnumstates[i];
    tempnum[i] = init[i]*tempnumstates[i];
  }
  // This calculates the total number of electrons.
  double totalnumber = 0;
  for (int i = 0; i < tempnumbins ; i++ ) {
    totalnumber += tempnum[i];
  }
  // This removes the wings of the distribution, since they
  // will not thermalize noticeably.
  for (int i = 0; i < numbins; i++ ) {
    num[i] = tempnum[i+ numbins/2];
    numstates[i] = tempnumstates[i+numbins/2];
  }
  // This calculates the normalizer properly.
  normalizer = totalnumber*totalnumber*totalnumber/numbins;

  double maxdif = 1;
  // This is the guts of the program, which evolves it until
  // either the time limit has been reached, or the changes in
  // the distribution become too small to be important.
  for (int t = 0; t < numsteps ; t++) {
    if (maxdif > 0.000001) {
      cout << "t = " << t << "\t"
   << "maxdif = " << maxdif << "\n";
      maxdif = evolve(num, trans, numstates);
    }
  }
  // This section outputs the results.  It is a little bit
  // complicated due to the fact that we need to reinsert the
  // wings that we cut off earlier.
  for (int i = 0; i < numbins/2; i++) {
    out_file << energy[i] << "\t" << init[i]
     << "\t" << init[i] << "\n";
  }
  for (int i = numbins/2; i < (numbins*3)/2; i++) {
    out_file << energy[i] << "\t" << init[i]
     << "\t" << num[i- numbins/2]/numstates[i-numbins/2]
     << "\n";
  }
  for (int i = (numbins*3)/2; i < tempnumbins; i++) {
    out_file << energy[i] << "\t" << init[i]
     << "\t" << init[i] << "\n";
  }
  return (0);
}
```

Bibliography

[1] T.D. Donnelly and Carl Grossman. Ultrafast phenomena: A laboratory experiment for undergraduates. *American Journal of Physics*, 66(8):677–685, AUG 1998.

[2] W.S. Fann, R. Storz, H.W.K. Tom, and J. Bokor. Direct measurement of nonequilibrium electron-energy distributions in subpicosecond laser-heated gold films. *Physical Review Letters*, 68(18):2834–2837, MAY 1992.

[3] W.S. Fann, R. Storz, H.W.K. Tom, and J. Bokor. Electron thermalization in gold. *Physical Review B*, 46(20):592–595, NOV 1992.

[4] A.T. Georges. Theory of multiphoton photoelectric effect: a stepwise excitation process. *Physical Review B*, 15, 1995.

[5] J. Hohlfeld, U. Conrad, and E. Matthias. Does femtosecond time-resolved second-harmonic generation probe electron temperatures at surfaces? *Physics D and Optics*, B63(5):541–544, 1996.

[6] M. Kaveh and H. Wiser. Electron-electron scattering in conducting materials. *Adv. Phys.*, 33(4):257–372, 1984.

[7] Charles Kittel and Herbert Kroemer. *Thermal Physics*. W.H. Freeman and Company, second edition, 1980.

[8] K. Moore. An experimental and numerical examination of nonequilibrium electron distributions in gold using second harmonic generation. Technical report, Harvey Mudd College, 1999.

[9] K.L. Moore and T.D. Donnelly. Probing nonequilibrium electron distributions in gold by use of second-harmonic generation. *Optics Letters*, 24(14):990–992, JUL 1999.

[10] M. Nisoli, S. Stagira, S. DiSilvestri, O. Svelto, S. Sartania, Z. Cheng, M. Lenzner, Ch. Spielmann, and F. Krausz. A novel-high energy pulse compression system: generation of multigigawatt sub-5-fs pulses. *Applied Physics B*, 65(189), 1997.

[11] N.A. Papadogiannis and S.D. Moustaizis. Nonlinear enhancement of the efficiency of the second harmonic radiation produced by ultrashort laser pulses on a gold surface. *Optics Communication*, 137(1–3):174–180, 1997.

[12] H. Petek and S. Ogawa. Femtosecond time-resolved two-photon photoemission studies of electron dynamics in metals. *Progress in Surface Sciences*, 56(4):239–311, 1997.

[13] David Pines and Philippe Nozières. *The Theory of Quantum Liquids*, volume 1. Advanced Book Classics, 1994.

[14] Y.R. Shen. *The Principles of Nonlinear Optics*. John Wiley & Sons, 1980.

[15] C.K. Sun, F. Vallée, L.H. Acioli, E.P. Ippen, and J.G. Fujimoto. Femtosecond-tunable measurement of electron thermalization in gold. *Physical Review B*, 50(20):337–348, NOV 1964.

[16] Robert C. Weast, editor. *Handbook of Chemistry and Physics*. CRC Press, 51st edition, 1971.