UNIVERSITY of CALIFORNIA
Santa Barbara

Coherent THz Manipulation of Ensembles of Hydrogenic Electrons in GaAs

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Physics

by

Matthew F. Doty

Committee in charge:

Professor Mark Sherwin, Chair
Professor David D. Awschalom
Professor Walter Kohn

December 2004

The dissertation of Matthew F. Doty is approved:

Professor David D. Awschalom

Professor Walter Kohn

Professor Mark Sherwin, Chair

September 2004

Coherent THz Manipulation of Ensembles of Hydrogenic Electrons in GaAs

by

Matthew F. Doty

Acknowledgements

I have already begun to feel nostalgic about the six years I've spent at UCSB. The frustrations and challenges are being overshadowed in my memory by the joys, which are due almost entirely to my great friends and coworkers. It is easy to go a long time without acknowledging everyone who has had such a positive impact on my life, but it's time to convey some well-earned thanks.

First, to Mark Sherwin, who has displayed remarkable patience and optimism over the last four years. He has been a wonderful example of an academic with a well-balanced life, who pursues challenging science without losing sight of the fact that research is never the only, or even the most important, thing in life. I hope I can treat my future students or employees as well as he has treated me.

To Tom, Kohl, Victoria, Xomalin, Sam, Nathan, Brendan, Dan, Jing, Sangwoo and the other members of the FEL or Sherwin group that I've had the opportunity to work with over the years. The spirit of cooperation is alive and well, and I truly appreciate how willing everyone has always been to go out of their way to help me with whatever I need. In particular I want to thank Tom, who taught me an awful lot about lasers, and Victoria, Xomalin, and Kohl, who have patiently and supportively listened to far more than their fair share of my complaining and my ideas.

To David Enyeart, without whom the FEL would not function and none of this research would have been possible. His astonishing capacity to fix things pales in comparison to his capacity to remain positive when I break things. I hope I have the privilege to work with such a genuinely good-hearted person in the future.

To Chris, Heather, Robbie, Craig, Jess and Bübbles, who were my core family of friends from the very beginning. Without their help, I would have done rather poorly in any number of classes the first few years. More importantly, without their friendship I wouldn't have so many wonderful memories.

To Steve, Brian, and everyone else at the Santa Barbara Judo Club, which has been a vital outlet for me throughout graduate school. I have learned so much from everyone at the club, and only some of it has been about Judo.

To Jess, Robbie, Julie, Tobias, Chris and Helen, who have been intimately involved in and supportive of my life for the past few years. We have shared a lot of good times, and I trust that they are but a small fraction of the good times to come.

Last, to Heather, who has been my best friend from the beginning. She has challenged me and supported me through everything, and constantly inspires me to be a better person. It is a great comfort to know that whatever we will face in the future we will face together, and I dedicate this thesis to her.

Curriculum Vitæ

Matthew F. Doty

Education

1998–2004	Ph.D. Physics University of California, Santa Barbara Santa Barbara, California
1994–1998	B.S. Physics The Pennsylvania State University State College, Pennsylvania
1991–1994	Emmaus High School Emmaus, Pennsylvania

Employment

6/00-10/04	Graduate Student Researcher, University of California, Santa Barbara. Advisor: Prof. Mark S. Sherwin
7/99-6/00	Head Teaching Assistant, Department of Physics, University of California, Santa Barbara
10/98-6/99	Teaching Assistant, University of California, Santa Barbara
1/97-5/98	Teaching Assistant, The Pennsylvania State University
1/95-5/98	Undergraduate Researcher, The Pennsylvania State University. Supervisor: Prof. John Yeazell

Publications

2004 M. F. Doty, B. E. Cole, B. T. King and M. S. Sherwin Wavelength-specific laser-activated switches for improved contrast ratio in generation of short THz pulses. Rev. Sci. Inst. **75** (September 2004, in press).

2004 M. F. Doty, B. T. King, M. S. Sherwin and C. R. Stanley Selective coherent manipulations of hydrogenic transitions in n-GaAs. in preparation.

Abstract

Coherent THz Manipulation of Ensembles of Hydrogenic Electrons in GaAs

by

Matthew F. Doty

Electrons bound to shallow donors in GaAs have energy levels that approximate those of the hydrogen atom. The effective mass of the electron and the dielectric constant of GaAs combine to rescale the energies of the orbital state transitions to the THz frequency range. These states are of interest for quantum information schemes, and I wish to explore the dynamics of coherent manipulations of the states.

After providing a theoretical description of the hydrogenic model, and an overview of previous related experiments, I proceed to describe experimental techniques and experiments. The UCSB Free Electron Laser is used to generate the required THz radiation. I use a magnetic field to tune the states into resonance with the available THz frequencies and read out the final state of the electrons using photoconductivity.

By varying the polarization of the THz radiation, I first examine the selection rules for the hydrogenic transitions. Verification of these selection rules is important for the design of THz photonic cavities that may interact with hydro-

genic transitions and for selective addressing of state transitions. I then drive Rabi oscillations between hydrogenic states by applying short intense pulses of THz radiation 10-100 ps in length. These short pulses are generated by using laser-activated semiconductor switches to "slice" short pulses from the output of the UCSB Free Electron Laser.

I model the behavior of these hydrogenic electrons using the density matrix equations of motion. Using this model to fit to experimental data, I extract information on the dephasing time of the ensemble of electrons. For the 1s to $2p^+$ transition at 2.52 THz, the dephasing time is of order 20 ps and appears to be limited by the presence of the excited state in the conduction band. With numerical calculations using the model, I explore the range of parameter space in which Rabi oscillations can be observed.

Contents

List of Figures

Chapter 1

Introduction and Motivation

The existence of shallow donors in semiconductors has been known for a long time; the first comprehensive theoretical treatment of these states was by Kohn in 1957 [5]. This theory describes the energy levels of electrons bound to shallow donors by a hydrogenic model. In particular, we will be interested in the states of donors in GaAs, where the crystal symmetry simplifies the model. There have been extensive spectroscopic studies of the states that show excellent agreement with the theory. The transition energy between these hydrogenic states lies in the THz frequency regime, and so until relatively recently it has been impossible to reach the strong coupling regime for the purposes of studying coherent manipulations of the orbital states.

Since the discovery of Shor's algorithm [6], there has been great interest in the potential computational power of quantum information processing. There are many publications on the subject; see Awschalom, Loss and Samarth [7] or Ben-

nett and DiVincenzo [8] for reviews. A variety of schemes for implementing quantum information processing exist, but a fundamental component of any scheme is a qubit: a two-level quantum system that can be prepared in an arbitrary superposition of its two states and made to undergo conditional interactions with other qubits. A workable qubit would have a decoherence time that is long compared to the time necessary to perform a conditional operation. There has therefore been extensive work carried out to identify and characterize the decoherence times of two-level systems that might serve as qubits.

The hydrogen atom is one of the simplest and most exhaustively studied quantum systems. Hydrogenic donors may be a way to combine this simple quantum system with the wealth of solid state processing techniques that have been developed for integrated circuit technology. In addition, there are schemes for quantum information processing [9] in which qubits are coupled via solid state cavity Quantum Electrodynamics (QED). One challenge of such a scheme is to selectively address the individual qubit states. If a cavity (for instance a photonic band gap structure) were prepared with a cavity resonance at optical or near-infrared (NIR) frequencies, the small size of the cavity (of order 1 μm) would be comparable to the wavelength of light and would make it challenging to optically address separate qubits within the cavity. A THz-frequency cavity would be of order 90 μm, simplifying the selective addressing of individual qubits with visible radiation. Since they occur naturally in GaAs growth, have transitions in the THz regime and are extremely uniform, hydrogenic donors are an interesting candidate

for use as the qubits in such a scheme. The THz cavity quantum computation scheme is discussed in more detail in Sherwin, et al. [10]

These potential applications in quantum information processing have led us to undertake a study of the hydrogenic donors in GaAs. In particular, we are interested in studying the selection rules for the orbital transitions and investigating coherent state manipulations. Understanding the selection rules is necessary for performing selective excitation of orbital states and also for coupling such state transitions to cavity modes. Coherent manipulations provide two important pieces of information. First, we can determine the required time to perform a single qubit operation, i.e. a single Rabi oscillation between the two states. Second, we can determine the dephasing time of an ensemble of hydrogenic donors, which provides a lower bound on the decoherence time of a single donor.

Chapter 2

The Hydrogenic Model

2.1 Shallow Donors

A substitutional impurity in a semiconductor arises when a normal constituent element of the semiconductor is replaced by a different element. These impurities arise naturally from contamination of the source materials when semiconductors are grown, for example by molecular beam epitaxy (MBE). Impurities can also be introduced intentionally. Substitutional impurities in semiconductors can be broken into two categories: donors and acceptors. Acceptors arise if the substituting atom contains fewer valence electrons than the atom it replaces. In this case, the impurity can bind an electron from the valence band of the semiconductor, leaving a free hole. This results in p-type doping of the semiconductor. Donors arise if the substituting atom contains more valence electrons than the atom it replaces. With sufficient energy, these additional electrons can be ionized into the

conduction band of the semiconductor, resulting in free electrons; this is n-type doping. In the current work, I will consider only donor impurities. In this section, I follow the work of Kohn [5] to introduce a theoretical description of donor-bound electrons.

Donor impurities can be further subdivided into two categories: shallow and deep donors. Deep donors (for example gold substituting on an arsenic site in GaAs) have an ionization energy of about 0.5 eV. This relatively large energy results in an electron that is relatively tightly bound to the donor site, which contains additional positive charge relative to the neighboring atoms. The electron orbits the donor site with a radius of a few Å. Shallow donors, on the other hand, have a binding energy of ≤ 0.1 eV. The orbit of the electron around the donor site is of order 100 Å. In this case, the electron orbit overlaps $\sim 10^5$ neighboring lattice sites. The particular interactions with individual lattice sites are essentially averaged out, and the slowly varying nature of the wave function makes it reasonable to make the effective mass approximation. Thus, the energy level of the electron can be described simply by the Coulomb interaction with the donor site, with a few small corrections that will be described in Section 2.3. We are only interested in shallow donors in this work.

In the case of an indirect-gap semiconductor (Si) the Hamiltonian for an electron bound to a shallow donor at a given band gap minimum is given by

$$\left(\frac{\hbar^2}{2m_l}\frac{\partial^2}{\partial z^2} - \frac{\hbar^2}{2m_t}\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right) - \frac{e^2}{r\kappa}\right)\Phi(r) = E\Phi(r), \tag{2.1}$$

where m_l is the longitudinal effective mass, m_t is the transverse effective mass, e is the electron charge and κ is the dielectric constant of the material. The total wavefunction of the electron is then given by

$$\Psi(r) = \Sigma_{j=minima} a_j \Phi_j(r) \vartheta_j(r), \tag{2.2}$$

where the sum is taken over all minima of the band gap with a_j a coefficient, $\Phi_j(r)$ the hydrogenic envelope function of that particular minima and $\vartheta_j(r)$ the Bloch wave at that minima.

In this work, we will only consider GaAs, so we need only work with the equations for the case of a direct-gap semiconductor, where the Hamiltonian is substantially simpler. It is given by

$$\left(-\frac{\hbar^2\nabla^2}{2m^*} - \frac{e^2}{r\kappa}\right)\Psi(r) = E\Psi(r), \tag{2.3}$$

with m^* the isotropic effective mass and κ the dielectric constant. This is just the Hamiltonian of the hydrogen atom, and the solutions are well known. The resulting quantized energy levels of the electron are

$$E_n = \frac{1}{n^2}\frac{(\frac{e^2}{\kappa})^2}{2\hbar m^*} \tag{2.4}$$

for integer values of n. From Cohen-Tannoudji [11], the wave functions of the electron for n = 1 and n = 2 are given by

$$\begin{aligned} \Psi(r)_{n=1,l=0,m=0} &= \frac{1}{\sqrt{\pi a^{*3}}}\exp(-r/a^*) \\ \Psi(r)_{n=2,l=0,m=0} &= \frac{1}{\sqrt{8\pi a^{*3}}}\left(1-\frac{r}{2a^*}\right)\exp(-r/2a^*) \end{aligned} \tag{2.5}$$

$$\begin{aligned}
\Psi(r)_{n=2,l=1,m=1} &= \frac{-1}{8\sqrt{\pi a^{*3}}}\frac{r}{a^*}\exp(-r/2a^*)\sin(\Theta)e^{i\phi} \\
\Psi(r)_{n=2,l=1,m=0} &= \frac{1}{4\sqrt{2\pi a^{*3}}}\frac{r}{a^*}\exp(-r/2a^*)\cos(\Theta) \\
\Psi(r)_{n=2,l=1,m=-1} &= \frac{1}{8\sqrt{\pi a^{*3}}}\frac{r}{a^*}\exp(-r/2a^*)\sin(\Theta)e^{-i\phi},
\end{aligned}$$

where a^* is an effective Bohr radius given by

$$a^* = \frac{\hbar^2\kappa}{m^* e^2}. \tag{2.6}$$

When appropriate values for GaAs are inserted for m^* and κ ($m^* = 0.0665\ m_e$ with m_e the mass of the electron, $\kappa = 12.56$) the effective Bohr radius is found to be 10 nm. As shown in Equation 2.5, the hydrogenic states can be defined by three quantum numbers: n, l and m, where n is the principal quantum number, l is the angular momentum quantum number and m is the magnetic quantum number. The wavefunctions given in Equation 2.5 correspond to the electron states traditionally notated as $1s^0$, $2s^0$, $2p^+$, $2p^0$ and $2p^-$, respectively. There is also a spin quantum number, s, but the spin states are not resolved in any of the experiments in this work, and I will ignore them from here forward. The only contribution from spin states would be to double the number of electrons that could occupy any state.

2.2 Shallow donors in a magnetic field

If we now place our semiconductor in a magnetic field oriented in the $\overrightarrow{z}$ direction, the full Hamiltonian becomes

$$\left(-\frac{\hbar^2\nabla^2}{2m^*} - \frac{e^2}{r\kappa} + \frac{1}{2}\hbar\omega_c l_z + \frac{1}{8}m^*\omega_c^2\left(x^2+y^2\right)\right)\Psi(r) = E\Psi(r). \tag{2.7}$$

Here ω_c is the cyclotron resonance frequency and l_z is the z component of the angular momentum. The states of the electron are again quantified by three quantum numbers: N, m and ν. N is the Landau level, m is the z component of the angular momentum and ν is the number of nodes of the wavefunction in the z direction. Figure 2.1 depicts all of the states, which can be divided into three categories: free electron states, metastable states and bound states.

The first category of states, free electrons, arises from the splitting of the conduction band under a magnetic field into Landau levels. There is a continuum of states available at each Landau level, where the energy of each Landau level is given by

$$E_N = \left(N - \frac{1}{2}\right)\frac{\hbar eB}{m^*}, \tag{2.8}$$

where B is the magnetic field in Tesla [12].

Metastable states, the second category, are bound states that have quantum numbers $N \geq 1$ and $m < N$. This designation of quantum numbers is not strictly correct because there are matrix elements that mix states of equal m but different N, and so N is not a unique quantum number. However, as described in Klaassen, et al. [1], for gyromagnetic ratio $\gamma \geq 0.1$, which is the case in this work, one value

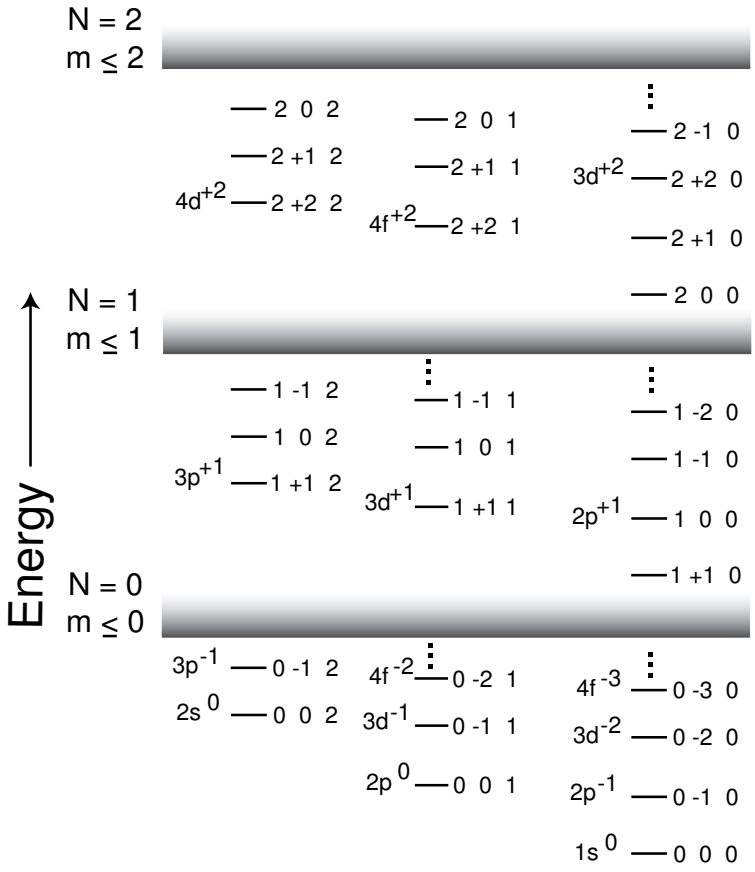

Figure 2.1: Energy levels of shallow donor electrons in a magnetic field. The bound and metastable states are denoted by the quantum numbers N, m and ν. The bound states are also labelled with their conventional hydrogen state notation. This figure is taken from Klaassen, et al. [1]

of N dominates and we will continue to use it as a label. These states are called metastable because they can undergo a radiationless, energy- and momentum-conserving (i.e. $\Delta m = 0$) transition to free Landau level states.

The third category of states, bound states, are those with N=0 and/or m=N. For states with these quantum numbers, no free state with the same m and same energy exists, making radiationless transitions impossible. As the magnetic field, B, goes to zero, these states extrapolate to hydrogen-like states with wavefunctions given by Equation 2.5. I will therefore continue to refer to these state with their hydrogenic notation. I call electrons in these bound states hydrogenic electrons, and it is the transitions between the hydrogenic states that are the subject of this thesis.

2.3 Hydrogenic states in strong magnetic fields

The energy levels of a hydrogen atom in a strong magnetic field are of interest both for semiconductor systems, as I have been describing, but also for astrophysical systems, in which neutron stars may have magnetic fields up to $10^7 - 10^9$ T [13]. A strong magnetic field is defined by Klaassen, et al. [1] as a field whose "action results in a sizable modification of the electronic structure." The strength of the magnetic field is generally given by the gyromagnetic ratio, $\gamma = \frac{1}{2}\hbar\omega_c/R^*$ where $\omega_c = eB/m^*c$ is the cyclotron frequency and R^* is the effective Rydberg [13]. In these dimensionless units, a strong magnetic field is one where $\gamma \geq 1$.

With the rescaling of energy levels described above (for the electron effective mass and GaAs dielectric constant), $\gamma = 1$ occurs for a hydrogenic electron in GaAs at B=6.57 T [1]. The experiments we perform therefore approach the strong-field conditions.

Substantial work has been done over the years to calculate the energy levels of a hydrogenic electron in the presence of a strong magnetic field. Makado and McGill used a variational approach to calculate the energy levels and in their paper they provide a comprehensive description of many of the preceding attempts [13]. I have chosen to use the results of Makado and McGill since their method appears to be somewhat more accurate for excited state energy levels and because they provide a higher density of calculated energies in the range of magnetic fields applicable to this work. Using the calculations of Makado and McGill, Figure 2.2 shows the energy levels of the lowest few states as a function of magnetic field. I have used values of $R^* = 46.1$ cm^{-1} and $\gamma = B(T)/6.57$ to convert their dimensionless data. These values were taken from Klaassen [1]. Also included in Figure 2.2 are the two lowest Landau levels of the semiconductor.

As mentioned in Section 2.1, the energy levels of shallow donor electrons in a semiconductor require some small corrections from the purely re-scaled energy levels of the bare hydrogen atom. The major correction is called the central cell correction, and accounts for the electron's interaction with the donor site itself. The zero-field magnitude of this correction depends on the actual donor species. Fetterman has shown that the magnitude of the correction scales with magnetic

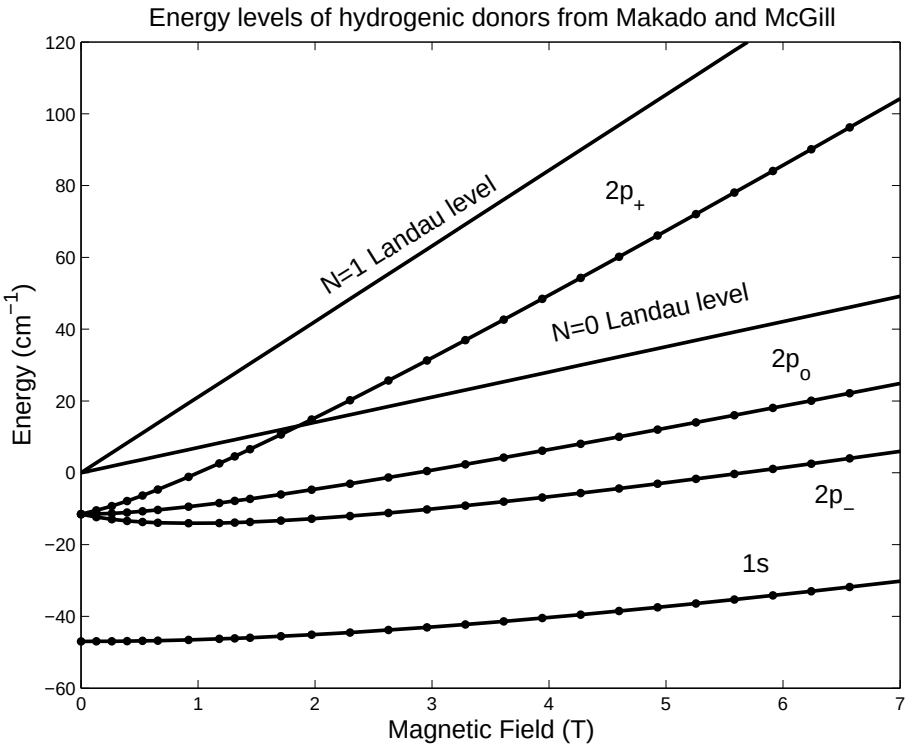

Figure 2.2: The energy levels of the 1s and 2p states as a function of magnetic field. Lines are drawn through points from the Makado and McGill calculation. The lowest two Landau levels are included, and the central cell shift is included in the 1s state energy. The energy of the lowest Landau level at zero field is taken to be zero.

field according to

$$\Delta E_j = \kappa_j |\Phi(0)|^2, \tag{2.9}$$

where κ_j is the zero field correction for a particular donor and $\Phi(0)$ is the envelope function of the hydrogenic ground state at a given field [14]. Values for $|\Phi(0)|^2/|\Phi(0)|^2_{B=0}$ have been calculated by Cabib, Fabri and Fiorio [15]. From these values, a theoretical calculation of the scaling of the central cell correction can be made. This theoretical scaling has been confirmed experimentally by Heron, et al. [14]. It is also possible to accurately determine the zero-field central cell correction for S and Si donors in GaAs from Heron's work.

Using Heron's value for the zero-field central cell correction for S (0.110 meV), Figure 2.3 shows the energy of the 1s ground state as calculated by the bare hydrogen model and as modified by the central cell correction. Of the states considered in this work, only the 1s state needs to be modified, as all 2p states have zero probability of finding the electron at the donor site where the central cell effects occur. This central cell correction is included in the energy of the 1s state in Figure 2.2.

There are two other corrections that I should also mention: polaron effects and band nonparabolicity. Both can be ignored in this work. Polaron effects only become significant at energies of roughly 300 cm^{-1}, well above the energies of these experiments (84 cm^{-1} and below) [14]. Band nonparabolicity also becomes more significant at higher energies. Since the effects of band nonparabolicity are

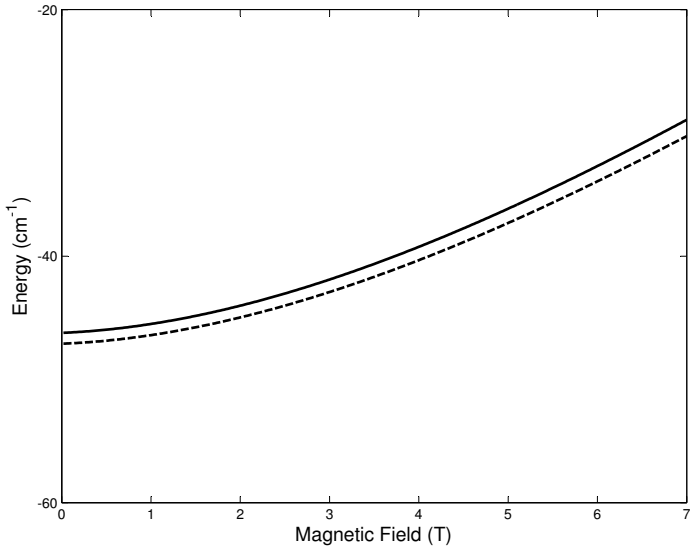

Figure 2.3: The energy of the 1s state as calculated by the model of Makado and McGill (solid) and as modified by the central cell correction (dashed).

small in comparison to the central cell correction, we can safely neglect them as well.

Since the central cell correction was included, Figure 2.2 shows the energies of the 1s and 2p states of a S donor in GaAs. It should be noted that the central cell corrections measured by Heron, et al. were obtained from samples grown by Colin Stanley of the University of Glasgow, who also provided the samples used in our experiments. While the material is not from the same growth, the growth conditions are nearly identical. The energy levels measured by Heron, et al. are therefore expected to match the energy levels of our samples.

2.4 Experimental Studies of Hydrogenic Electrons in GaAs

The hydrogenic states have been extensively studied by a large number of experimental groups. In this section, I will provide a brief history of experiments that are relevant to this thesis. I have chosen five experiments, all performed on n-type GaAs: interferometric magnetospectroscopy, fixed-illumination-frequency magnetospectroscopy, measurement of the excited state lifetime, time-resolved dynamics and measurements of dephasing time. The experiments I undertook in this thesis are, in many ways, extensions of experiments like those described in this section. However, the UCSB FEL generates sufficient FIR intensity that we are able to extend experiments on shallow donors into the regime of non-linear phenomena.

Several of the experiments described in this section use photoconductivity to measure the electron orbital state. The higher orbital states ionize into the conduction band more easily than the ground state, and therefore an increase in the conductivity of the sample indicates the presence of the electrons in the higher state. This is the same method of measuring excited state population that is used in the experiments of this thesis, and more details can be found in Section 5.4.

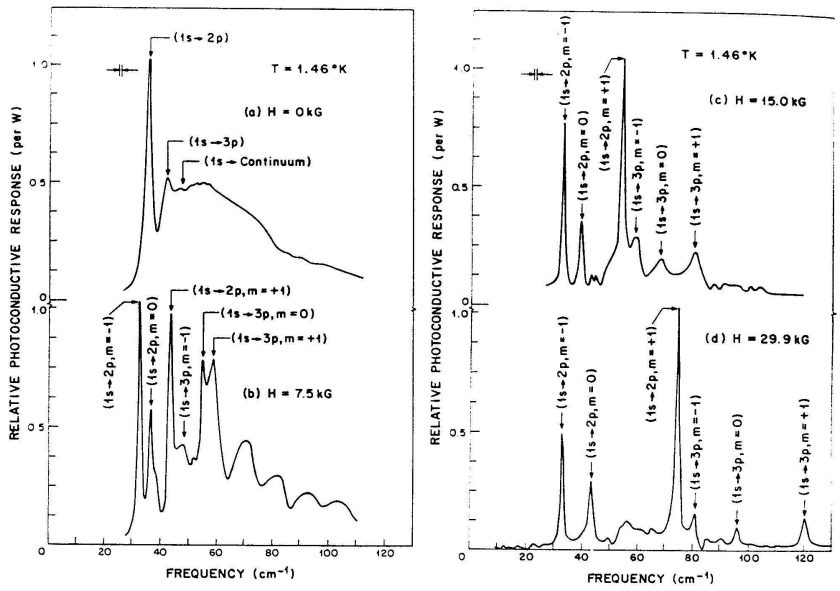

Figure 2.4: Spectroscopy of hydrogenic electrons by Stillman, et al.[2] at several fixed magnetic fields.

2.4.1 Interferometric magnetospectroscopy

Stillman, et al. studied the transitions between donor-bound electron energy levels by placing a sample of GaAs in a light pipe at the center of a superconducting solenoid [2]. Using an interferometer, they measured photoconductivity as a function of illumination frequency. Figure 2.4 is from the paper of Stillman, et al. showing spectroscopy at several fixed magnetic fields. The authors use the hydrogenic notation to identify the orbital transitions that correspond to the various peaks. The specific donor species responsible for the transitions was not determined. In Figure 2.5, also from the paper by Stillman, et al., the transition energy for the 1s to 2p and 1s to 3p transitions are plotted as a function of magnetic field.

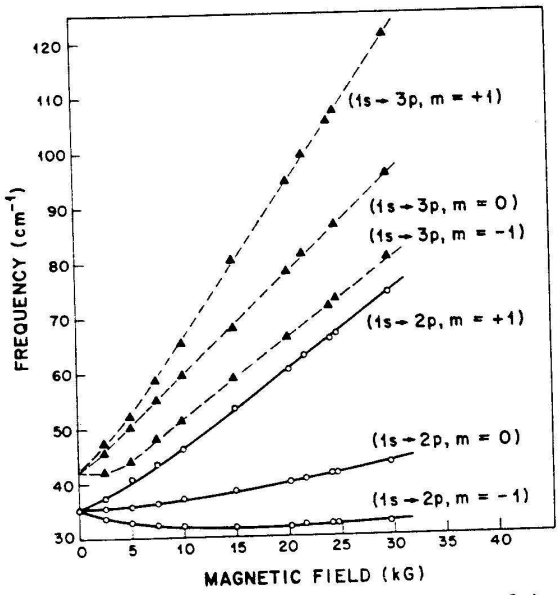

Figure 2.5: A plot of donor-bound electron orbital transition energies versus magnetic field from Stillman, et al.[2]

2.4.2 Fixed-illumination-frequency magnetospectroscopy

A common method of magneto-optical spectroscopy experiment is to illuminate an n-type GaAs sample with a fixed frequency of THz radiation and sweep the applied magnetic field. Some examples are the work of Klaassen, et al. [1], and Heron, et al. [14]. When the magnetic field value brings a hydrogenic transition into resonance with the THz field, the electrons are excited into a higher state, and an increase in photoconductivity is observed.

Figure 2.6 shows results from spectroscopy of donor-bound electron states by Klaassen, et al. [1] Using the N, m, ν notation, the authors have identified each of the transitions they observed. The transition lines observed are inhomogeneously broadened. One might expect that the transition for different donors of the same

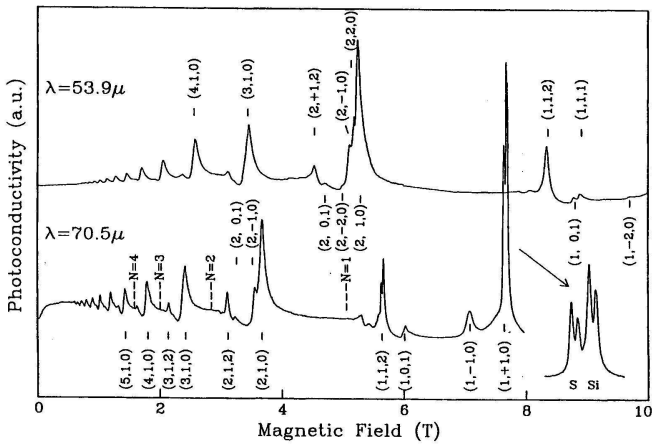

Figure 2.6: Magneto-optical spectroscopy of donors in n-type GaAs by Klaassen, et al.[1] Note the resolution of the different donor species and spin states of the (1,+1,0) peak.

species should be homogeneously broadened. In fact, the random distribution of lattice defects and compensating impurities leads to a random distribution of local electric field for each of the hydrogenic electrons and thus to the inhomogeneous broadening. This inhomogeneous broadening was explained in a theory by Larsen and also applies to the work of this thesis [16].

2.4.3 Excited state lifetime measurement

Kalkman, et al. undertook a pump-probe experiment of the $2p^-$ state of a donor-bound electron in GaAs using FIR of wavelength 292.14 μm and a magnetic field of 5.1 T [17]. In this experiment, a pump beam approximately 300 ns long was used to excite the 2p- state. A time-delayed probe beam of about 1 ns was derived from the same pulsed FIR source by laser-activated semiconductor switches. This probe beam passed through the sample and was reflected to a fast Schottky diode

detector. By measuring the absorption of the probe pulse as a function of delay between the pump and probe pulses, the lifetime of the excited state was explored. Although interference effects made an exact measurement difficult, a lifetime of approximately 350 ns was determined. This long lifetime suggests that the 2p-state is well isolated and may therefore also have a long dephasing time. It is thus of particular interest for the experiments of this thesis.

2.4.4 Time-resolved dynamics

Burghoorn, et al. did an experiment to measure the ionization time of electrons excited into the $2p^+$ state [3]. They used a short pulse of FIR radiation (118.8 μm) switched from a CW molecular-gas laser by laser-activated semiconductor switches. A magnetic field of 3.6 T was used to tune the state transition into resonance with the THz field. Figure 2.7 shows a plot from their paper showing photoconductivity as a function of time for four different laser intensities. Using a three-level rate equation model, the authors fit to the data to determine that the minimum ionization rate of the $2p^+$ state was 5×10^9 s^{-1}. The authors also study the recapture rate of the electrons by the ionized donors, shown in Figure 2.8, and find that electrons typically spend about 12 ns in the conduction band. This is comparable to the decay of photoconductivity observed in this thesis. Some differences from varying donor concentration are expected.

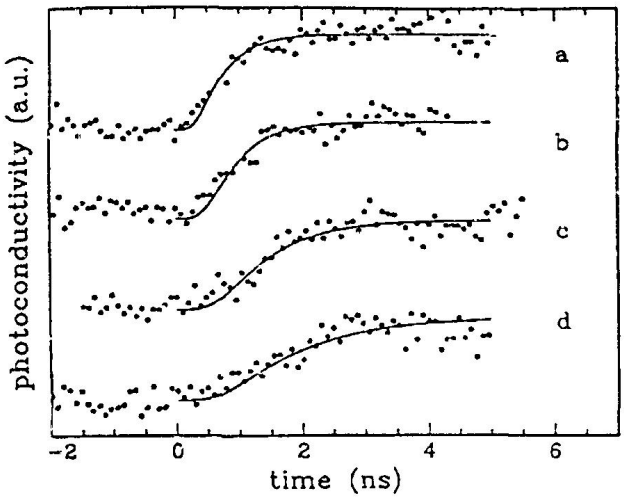

Figure 2.7: Photoconductivity as a function of time for laser intensities of (a) 17, (b) 5.1, (c) 1.3 and (d) 0.68 W cm^{-2}. The data is from Burghoorn, et al.[3] The solid lines are fits using a three-level rate equation model.

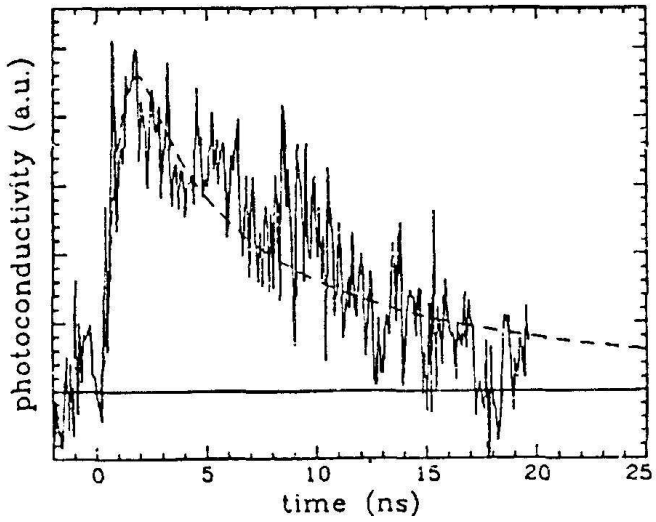

Figure 2.8: A plot from Burghoorn, et al.[3] of photoconductivity as a function of time in response to a 1.3-ns FIR pulse with a peak power of 9 W. The decay of the photoconductivity gives the recapture rate of electrons from the conduction band.

2.4.5 Measuring dephasing times

Planken, et al. performed an experiment to measure the dephasing time of electrons in the $2p^+$ state [18]. They used high-intensity short pulses (about 3 ps) from FELIX, a Dutch mode-locked FEL. This short pulse was split into two pulses, one delayed relative to the other. The first pulse began to drive electrons from the 1s to the $2p^+$ state. The second pulse either continued to drive the electron into the upper state, or drove it back down to the lower state, depending on the phase relationship between the two pulses (controlled by the delay). The phase memory time of the ensemble of hydrogenic electrons is the dephasing time of the ensemble. Therefore, by observing the decay time of the interference effects between the two pulses, the authors were able to determine the dephasing time of the ensemble to be about 18 ps. Unlike the experiments described in this thesis, the experiment of Planken, et al. does not drive electrons through a full cycle of the quantum states.

Chapter 3

TeraHertz pulse shaping

In this section I will first describe the UCSB FEL, our source for high-intensity THz radiation, and our methods for varying the polarization of the THz radiation. For many of the experiments described in this thesis, I require extremely short pulses of THz radiation. The bulk of this chapter is therefore devoted to the experimental methods by which I generate these short pulses.

3.1 THz source and technology

In recent years there has been substantial development of THz sources, as well as increased effort to develop optical components to support use of these sources [19]. In particular, there have recently been tremendous improvements in Quantum Cascade Lasers that promise reasonably sized solid state THz lasing. At present, however, only free electron lasers can achieve the pulse intensity nec-

essary to reach the nonlinear coupling regime in which coherent manipulations occur. The THz source for these experiments was the UCSB FEL. The theory and operation of the FEL has been described elsewhere, for example by Ramian [20], but a short overview is in order.

The FEL operates based on the principle that oscillating electric charges emit radiation. The frequency of the emitted radiation is determined by the period of the oscillations. In the UCSB FEL, a beam of relativistic electrons is generated by a Van De Graaf generator coupled to a linear accelerator. This beam of electrons is directed through a series of electromagnets that steer and shape the beam. The beam is then passed through an undulator that consists of a fixed sequence of permanent magnets of alternating polarity. As the electrons pass through the first magnetic field, their path bends slightly in one direction. The next magnetic field bends the electron beam back toward its original direction. The series of these magnets therefore "wiggles" the electron beam back and forth. The period of the electron oscillations is determined by the fixed spacing of the permanent magnets and the speed of the electrons. The electron speed is determined by the energy of the accelerator, and thus the frequency of the emitted radiation can be controlled by varying the voltage across the accelerator.

The undulator in the UCSB FEL is implanted in a waveguide that is capped at either end by mirrors. Thus, the THz radiation spontaneously emitted by the oscillating electrons reflects back and forth along the length of the undulator. This radiation stimulates emission of more photons, hence the lasing action. The

output radiation is coupled through a small hole in one end mirror and is then distributed to experimental stations through an evacuated transport system that prevents losses due to atmospheric absorbtion. The peak power output of the FEL can be as high as 1 kW at 2.52 THz, but decreases by roughly an order of magnitude as the frequency is lowered to 1 THz. The output is limited to relatively short pulses (a few μs) due to the maximum length of the electron pulse. The repetition rate of the laser is roughly 1 Hz.

It is also appropriate to add a short word on the optical components and techniques used when working at THz frequencies, since they differ somewhat from conventional visible and near-infrared (NIR) wavelengths. I should note that THz frequencies lie outside the visible range, and no commercially available viewers exist. This means that it is extremely difficult to precisely locate a THz beam. To facilitate creating and aligning an optical system in the THz, I use visible lasers that are set to follow the same optical path as the THz. The first such "alignment" beam passes through the lasing cavity of the FEL and so follows the same output path as the THz. Additional alignment beams can be introduced later by defining an optical axis along which both the visible and THz beams pass. The alignment of the THz beam to such an optical axis can be verified by using a pyroelectric detector to trace out the position of the FIR beam. More information on this can be found in Section C.3.

The wavelength of THz frequency beams is quite large (1 THz = 300 μm). For beams with such large wavelengths, it is necessary to use wide beams to avoid

diffraction when propagating the light over large distances (1 m or more). For the UCSB FEL, a beam width of up to 10 cm is used. This means that all optics used must have a diameter that exceeds this dimension, and visible alignment beams are expanded to similar dimensions. Gold coated glass mirrors of 4 $\times$ 6 inches are typically used. The depth of the gold coating is normally of order 1 μm since the skin depth of THz radiation incident on a metal can be quite large.

To build a focusing element, one would like to look for a material that has a reasonable index of refraction in the THz frequency regime. Silicon, with an index of 3.42, is one candidate. However, silicon's opacity at optical frequencies precludes the use of visible alignment beams. One must therefore find a material with the same index of refraction for THz and visible frequencies. The only material satisfying this requirement that I am aware of is TPX plastic. TPX lenses are used in the transport system of the UCSB FEL to compensate for beam diffraction over the long path lengths. However, TPX plastic also absorbs THz radiation, and so lenses of this material are lossy. To avoid these losses, focusing is normally done with an off-axis parabolic mirror (OAP). These OAPs are generally cut from brass following the surface of a solid paraboloid, as in Figure 3.1. Since geometry dictates that any ray initially directed toward the directorix of a parabola will reflect toward a single focal point, these can be used in place of lenses. The focal length of the "lens" is determined by the curvature of the parabola. The brass is then coated with gold; since all reflection occurs on the front surface, the properties of the OAPs are not wavelength specific.

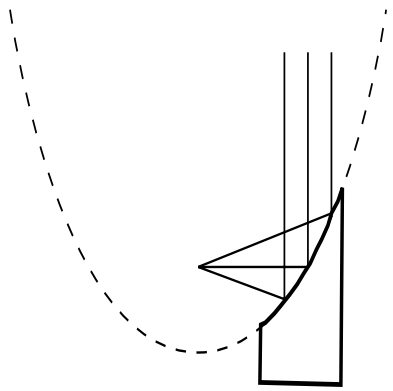

Figure 3.1: An outline trace of an off-axis parabolic mirror. The dotted line indicates the parabola that defines the surface curvature. The reflection of three example rays towards a single focal point is shown.

Visible wavelength polarizers are often made with some kind of oriented polymer where the dimensions of the polymer strands and their separation is comparable to the wavelength of the light. For the very long wavelengths of the THz frequency regime, a polarizer is made of a closely spaced grid of thin metal wires, as show in Figure 3.2. When the electric field of the incident radiation is parallel to the wires, currents are driven along the wires and the beam is reflected and absorbed. When the electric field is perpendicular to the wires, no currents can be driven and the beam propagates through the polarizer. The efficiency of the polarizers depends somewhat on the relative size of the incident wavelength and the wire grid spacing. The polarizers used throughout this work are most efficient at low frequencies. At 1 THz they have been measured to transmit 100% of light polarized along the transmission axis and 0.5% of light polarized perpendicular to the transmission axis. For comparison, at a frequency of 2.52 THz the polarizers transmit 3% of light polarized perpendicular to the transmission axis

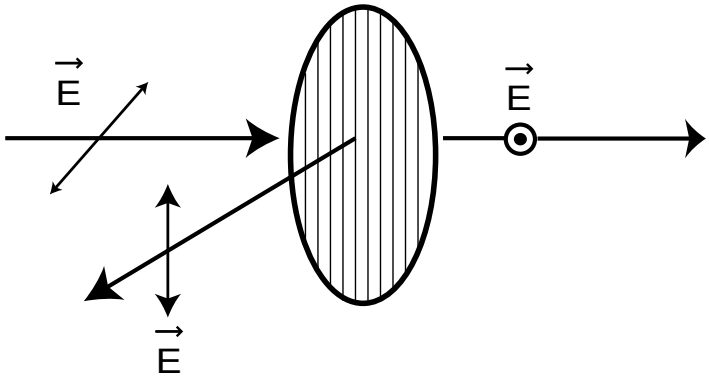

Figure 3.2: A schematic of a wire grid polarizer showing the polarization of the transmitted and reflected light.

As mentioned above, the THz output of the UCSB FEL is transported in vacuum to avoid atmospheric absorption. This absorption arises primarily from water vapor in the air, as the vibrational modes of the tri-atomic water molecule lie in the THz spectrum. Although this evacuated transport system permits transport of the THz to experimental stations with minimal loss, for practical reasons the experiments themselves are not done under vacuum. Since the optical path length of the THz system used in the particular experiments described here (Section 3.3) is of order 6 m, absorption would still be a significant effect. For this reason the experiments are generally all undertaken in "water windows:" frequency ranges where the atmospheric absorbtion is quite low.

3.2 Variable THz Polarization

One set of experiments that I will describe here involves probing the selection rules for transitions between the hydrogenic states. I will then take advantage of

these selection rules to excite selective manipulations. To study and then use the selection rules, it is necessary to vary the polarization state of the THz radiation. For visible frequencies, this is easily accomplished by using quarter- and/or half-wave plates. Quarter- and half-wave plates are manufactured from birefringent materials. Birefringent materials have a different index of refraction for light that is polarized along different crystal axes. These indices of refraction are typically labelled as n_o and n_e for ordinary and extraordinary index, respectively. The crystal is oriented so that the electric field of the incident beam projects equally onto the two different crystal axes. One component of the total electric field passes through the crystal with a speed that is determine by n_o and the other component has a speed determined by n_e. At the output of the crystal, the two components of the electric field recombine with a phase delay between them. This phase delay depends on the relative values of n_o and n_e and on the length of the material through which the beam has passed. A quarter-wave plate is designed to introduce a phase delay of $\pi/2$, resulting in a circularly-polarized output beam. A half-wave plate creates a phase delay of π and thus rotates the polarization by 90^o.

For THz frequencies, there are no commercially available birefringent phase shifters. However, I can accomplish the same phase delay effect using a mirror and a polarizer, as shown in Figure 3.3a. A linearly polarized THz beam hits a polarizer that is set at 45^o to the polarization of the incident beam. This polarizer reflects one component of the polarization and transmits the orthogonal

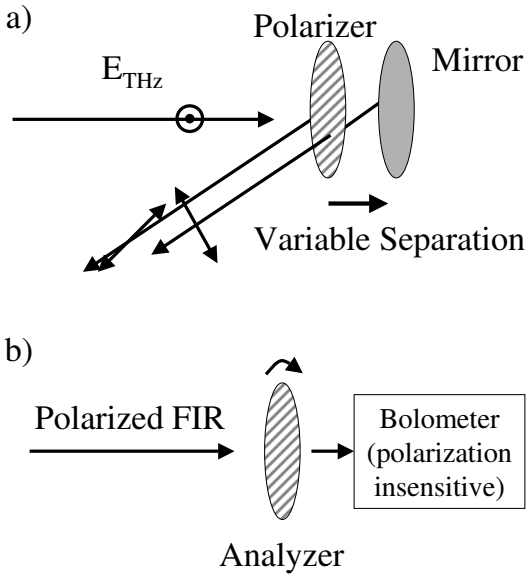

Figure 3.3: a) A diagram of a variable polarizer used to control the polarization state of THz radiation. b) The experimental setup used to probe the polarization state of THz radiation.

component (of equal intensity, since the polarizer is set to 45^o). The transmitted beam reflects off a mirror and returns through the polarizer. The polarizer and mirror are aligned to be parallel to one another, so the output of the combined polarizer-mirror pair is the sum of the two electric fields. These are just the two orthogonal electric field components of the incident beam, with a phase delay introduced by the spacing between the mirror and the polarizer. By varying this spacing, an arbitrary polarization can be generated.

The polarization is measured with a polarization analyzer, which is simply a polarizer that can be rotated to transmit light polarized along a single axis. The polarization state is quantified using the Stokes parameters, S_0, S_1, S_2 and S_3 [21]. S_0 gives the total intensity of the light: the sum of the intensity passing

through an analyzer set at 0^o and the intensity passing through an analyzer at 90^o. S_1 measures the degree of polarization along the 0^o axis and is given by the difference between intensity transmitted through the analyzer at 0^o and at 90^o. S_2 measures the degree of polarization along the 45^o axis, i.e. the difference in intensity passing through the analyzer at 45^o and at 135^o. Finally, S_3 measures the difference in intensity between right- and left-circularly-polarized light. This is measured by inserting a quarter-wave plate that delays the 90^o-axis component of polarization by $\pi/2$ relative to the 0^o-axis component. S_3 is then given by the difference in transmission through an analyzer placed after the quarter-wave plate when the analyzer is at 45^o versus 135^o. Knowing the four Stokes parameters is sufficient to fully characterize the polarization state of a beam of light.

As discussed above, in the case of THz measurements I do not have access to a quarter-wave plate, which makes direct measurements of S_3 impossible. However, the Stokes parameters are related by

$$S_0^2 = S_1^2 + S_2^2 + S_3^2. \tag{3.10}$$

So if I can measure S_0, S_1 and S_2 I can easily determine the value of S_3^2. The experimental setup for making these measurements is quite simple and is shown in Figure 3.3b. I place an analyzer in front of a polarization insensitive detector, in our case an InSb hot-electron bolometer. The bolometer is polarization insensitive because it works by detecting changes in conductivity induced by the heating from

the THz radiation. By measuring the transmitted intensity at analyzer angles of 0, 45, 90 and 135^o I can calculate the value of S_0, S_1 and S_2. In the experiment, I vary the spacing of the mirror and polarizer and measure the Stokes parameters at various values of this spacing.

Finally, the sign of S_3 can be determined with a little additional knowledge about the orientation of the variable polarizer in our system. I know that the polarizer is set to reflect light along the 45^o axis and transmit light along the -45^o axis. When the output light is polarized along the 0^o axis, there must be zero phase delay between the electric field pointing to $\pm 45^o$. If I slightly increase the spacing between the mirror and the polarizer, the -45^o component is delayed, and therefore the light must be right hand elliptically (σ^+) polarized. Thus I can manually reinsert the sign of S_3 by comparing the phase delay at the particular mirror-polarizer spacing with the phase delay when the output light was polarized along the 0^o axis.

3.3 The Pulse Slicer

For the purpose of generating and time resolving coherent manipulation of orbital states, I must have a control pulse with a temporal width that is comparable to the dephasing time of the ensemble of donors. This is expected to be in the picosecond to nanosecond range for the low-impurity-density GaAs I study. The output of the FEL is typically a few μs long. The width is somewhat controllable

on the μs scale, but can not be reduced to below about 2 μs. This lower bound exists for two reasons. First, the FEL does not produce a stable electron-beam pulse with such short duration. Second, the lasing buildup within the FEL occurs over the first few microseconds and the pulse amplitude does not reach a steady state value until 1-2 μs after the onset of THz output. Thus to perform experiments with short pulses of THz radiation it is necessary to extract a ps-ns width pulse from the steady state section of the pulse envelope of the FEL output. The "pulse slicer" is our name for the system that has been built to do this pulse extraction.

The pulse slicer was originally built by Frank Hegmann, a postdoc, in 1996. It is based on the principle of laser-activated switching. Laser-activated semiconductor switches have been used in the mid and far infrared since 1975 [22]. For a more detailed history of their development and use, see Hegmann, et al. ([4]). The switches used in our system are made of silicon, which is transparent to THz radiation until it is activated by illumination with near-infrared (NIR) radiation across the silicon band gap. This NIR generates an electron-hole plasma on the front surface of the silicon wafer. The plasma reflects the THz radiation, but the reflectivity is a function of both plasma density and wavelength of the incident THz, as shown in Figure 3.4.

Figure 3.4 shows that generation of sufficient plasma density for reflection of FIR pulses with wavelengths as short as 100 μm requires plasma densities of order 10^{19} cm^{-3}. This requires NIR pulse energies of at least 1 mJ per switch [4]. In order to have the switching time short in comparison to the temporal width of the

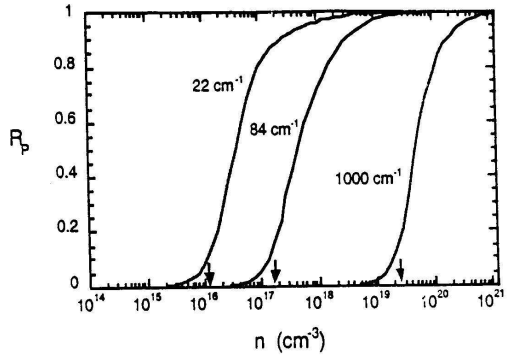

Figure 3.4: The reflectivity of an electron-hole plasma as a function of plasma density. Curves are shown for different incident THz frequencies. The plot is taken from Hegmann, et al.[4]

sliced pulse (ps), this 1 mJ must arrive in a pulse shorter than 1 ps. It requires special regenerative amplification techniques to generate short NIR pulses with such high energy densities; these techniques are discussed in Section 3.4.

The pulse slicer is built by using these silicon switches in two different geometries. In the first geometry, the FIR is focused onto a silicon switch using an OAP (off-axis parabolic mirror). When the NIR pulse activates the wafer, the reflected radiation is kept in the optical path, defining the leading edge of the pulse. This is called a "reflection switch." For the second geometry, the radiation transmitted through the switch is kept in the optical path and the reflection is dumped. This "transmission switch" defines the trailing edge of the pulse. Figure 3.5 illustrates how pulse widths can be controlled simply by varying the delay between the arrival of the near-infrared pulses at each switch. A diagram of the full UCSB pulse slicer is shown in Figure 3.6; two reflection switches are used in series to increase contrast ratio.

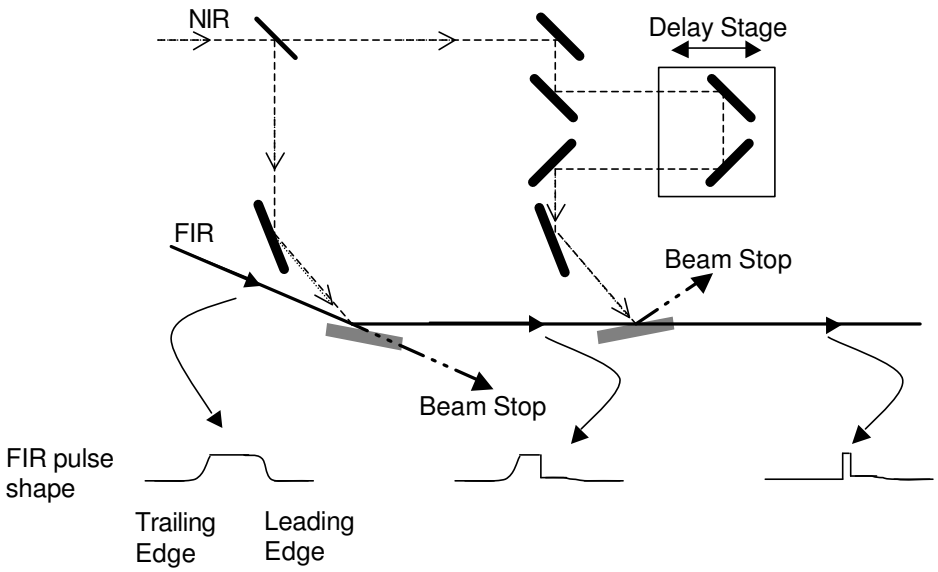

Figure 3.5: A schematic of the UCSB Pulse Slicer.

The contrast ratio of sliced pulses is the ratio between the peak power of the short pulse and the power of the background FIR reflecting off the switches prior to activation. The peak power of the short pulse is determined by the reflectivity of an activated wafer, which has been measured to be $\sim 60\%$. The background FIR can be suppressed by p-polarizing (electric field parallel to the plane of incidence) the incident beam and setting the wafers at Brewster's angle. While it is necessary to use small $f\backslash$ optics to create a FIR spot size small enough to be illuminated with reasonable near-infrared pulse energy, the small $f\backslash$ has the negative side effect of broadening the range of incident angles for the FIR. Since suppression of reflection near Brewster's angle depends strongly on the angle, this broadening increases pre-activation FIR reflection. There are several methods that can be

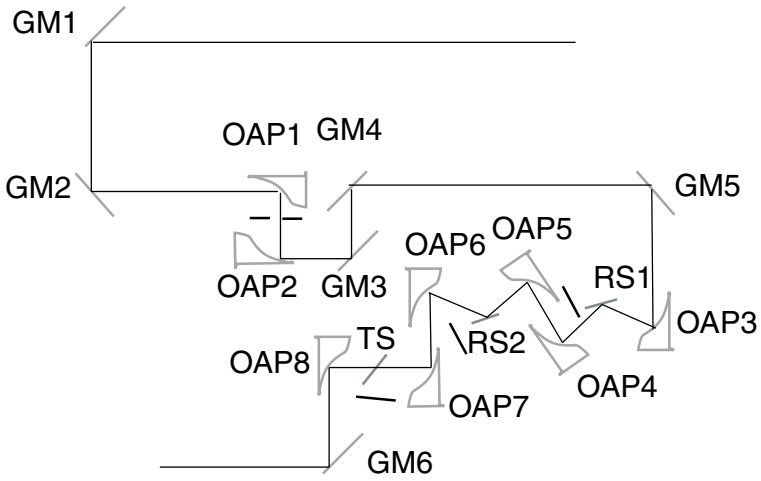

Figure 3.6: A diagram of the UCSB Pulse Slicer. The incident THz enters at the upper right, reflects of gold Mirrors (GM), is focused through a spatial filter and then is focused in turn onto two reflection switches (RS1 and RS2) and finally onto a transmission switch (TS) before exiting at the lower left. The short black segments near each switch indicate the location of beam stops used to absorb THz transmitted through the reflection switches before activation and reflected from the transmission switch after reflection. The focusing elements are off-axis parabolic mirrors (OAP).

used to increase the contrast ratio. One method is to tune the thickness of the silicon switch wafers so that FIR reflections from the front and back surface of the wafers interfere destructively, suppressing the pre-activation reflectivity. This method is discussed in detail in Chapter 4.

Another technique used to improve contrast ratio is to aperture down the FIR pulse, reducing the spread of angles incident on a switch. Although this also results in a decrease in the total throughput power, it can increase the contrast ratio and is a useful tool. Beam apertures are discussed in more detail in Appendix C.4. Another method of reducing the spread of incident angles is to make the FIR beam diameter as narrow as possible. Maintaining a narrow beam can be challenging due to the rapid diffraction at long wavelengths. Using a narrow beam also increases the effective $f\backslash$ of the OAPs used to focus the beam, which results in a wider focal spot and lower FIR intensity on a sample. However, if these effects can be tolerated, reducing the beam diameter can increase the contrast ratio significantly. Section 4.3 contains data on measurements of contrast ratio with varying wafer thickness and beam diameter.

3.4 Near-Infrared regenerative amplification

As mentioned in Section 3.3, activation of the silicon switches requires illumination by a laser pulse of energy larger than the silicon band gap, around 1.1 eV at room temperature. It requires at least 1 mJ of energy per switch and there

are three such switches in the UCSB pulse slicer. To generate a top-hat pulse shape (a pulse with rising and falling edges that are nearly step functions) the 1 mJ must arrive within a time that is short compared to the width of the pulse. Since the pulse slicer is intended to generate pulses with a width as short as only a few ps, the activating laser pulse must be substantially shorter than a ps. I use a regenerative amplifier to produce pulses of approximately 150 fs duration centered at 795 nm with energies of up to 10 mJ per pulse. This section is intended to provide an overview of the regenerative amplification system. Details on the alignment and operation of the UCSB regen can be found in Appendix B.

A regenerative amplifier works by taking the output of a mode-locked laser, which is spectrally broad but temporally short, and chirping the pulse. A pulse is said to be chirped when the lowest spectral components arrive at the beginning of the pulse and the highest at the end of the pulse, rather than simultaneously. This chirped pulse is used to stimulate emission from a second laser cavity, whose output has the same temporal and spectral characteristics as the original seed beam, but is significantly more intense. This output is further amplified and finally compressed by reversing the chirp procedure. The final output is once again temporally short, but now contains significantly more energy per pulse than the original beam. Amplification of a sub-ps pulse to such high energies requires regenerative amplification to avoid exceeding the energy density damage threshold of the lasing gain crystals.

The initial seed beam for our system is the output of a mode-locked Tita-

nium:Sapphire laser (Spectra-Physics Tsunami) pumped by a CW diode laser operating at 532 nm (Spectra-Physics Millenia V). By mode-locked, I mean that the lasing cavity is modulated at its free spectral range to creates a periodic disturbance in the cavity gain. This forces all of the separate cavity modes to arrive at the output coupler at the same time. The output of the laser is thus a train of pulses. I passively mode lock the Tsunami and find output pulses of 150-fs duration and energy of order 10 nJ per pulse. This pulse energy is calculated from the laser repetition rate of 76 MHz and a maximum output power of 730 mW. The bandwidth is approximately 10 nm and centered on 795 nm.

This output pulse is then passed into the stretcher portion of the regen cavity and reflected off of a diffraction grating, which spatially disperses the frequency components of the pulse. This spatially broad pulse is reflected through an optical path that requires a longer time of flight for one part of the spatial (and thus spectral) distribution and finally recombined on the grating. After this recombination on the grating, the pulse is once again spatially narrow but is now temporally spread out by the difference in time of flight for the different components. This is a chirped pulse: the longest wavelength is at the leading edge of the pulse and the shortest at the trailing edge. The width of the chirped pulse is approximately 200 ps.

The chirped pulse is then injected into a lasing cavity (called the oscillator) that is pumped by a cavity-dumped frequency-doubled Nd:Yag laser operating at 10 Hz (Spectra-Physics Quanta-Ray). The Nd:Yag lases at 1064 nm, so the

resulting frequency-doubled output is at 532 nm, precisely like the CW diode laser mentioned above. The Nd:Yag, however, is flashlamp pumped and the energy in the lasing cavity is allowed to reach its maximum value before being Q-switched out in a single short pulse. This results in pump pulses for the regenerative amplifier that are very intense, though at a very low repetition rate.

The pump pulse from the Nd:Yag laser provides the population inversion for a Titanium:Sapphire crystal in the oscillator cavity. The emission of this crystal is stimulated by the chirped pulse. This results in output pulses that have the same temporal and spectral profile as the seed beam, but higher intensity. The oscillator cavity is then dumped and the output passed through a single-pass amplifier.

Finally, the output pulse is passed off of another diffraction grating and then along another path where the different spatial (spectral) modes have different path lengths. This path is designed to exactly invert the chirping induced in the stretcher section and so the final output of the compressor regains the temporal profile ($\sim$150 fs) of the input pulse, but now contains up to 10 mJ per pulse. This final output is divided into three pieces and sent to the silicon switches used in the pulse slicer.

3.5 A note on synchronization

There are five different lasers used in this experiment: the Millenia, Tsunami, Quanta-Ray (Nd:YAG), regenerative amplifier and FEL. Successfully slicing a

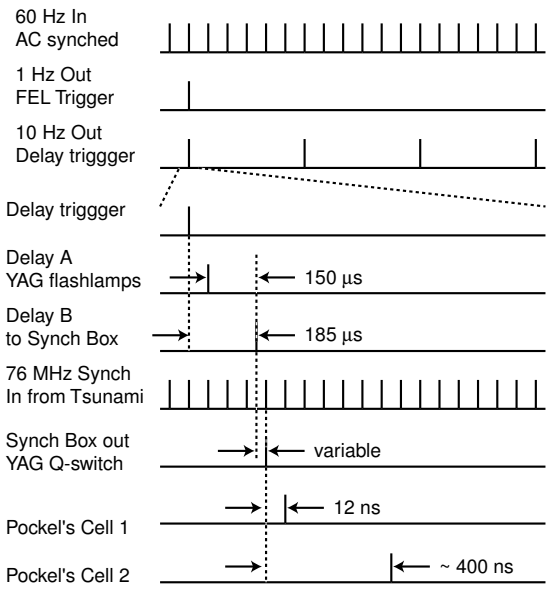

Figure 3.7: A timing diagram of the pulse slicer components.

pulse requires that the pulses generated by all these different lasers overlap both spatially and temporally. The alignment procedures for the lasers and pulse slicer are described in Appendices B and C.3, respectively. This alignment achieves the spatial overlap.

The temporal overlap of the pulses is achieved by combining some commercial pulse delay and synchronization boxes with some homemade electronics. The homemade electronics were originally built by Bryan Cole and were then modified by Tom King. I have attempted to diagram the timing relationship between the various events in Figure 3.7.

The fundamental clock for our experiment comes from the FEL. David Enyeart has built a clock box that generates a 60-Hz TTL pulse that is synchronized to

the AC power lines. This is important because the FEL runs most stably when each pulse is generated at the same phase point in the AC signal. This 60-Hz TTL pulse is sent to our lab and enters the first homemade timing box. In this box, the 60-Hz signal is divided down into a 10-Hz signal and a 1-Hz signal that remain synchronized to the 60-Hz input. Actually, the 1 Hz can be varied over a range of frequencies from about 0.25 Hz up to about 5 Hz, but I will refer to the signal as 1 Hz throughout this discussion.

The 1-Hz signal is sent back out to the FEL and provides the trigger for firing the electron pulse that generates the FIR output. The additional, controllable, delay built into the FEL software should be set to 0.21 ms. This value compensates for the time of flight of signals and the delay between the trigger pulse and actual photon output.

The 10-Hz signal is then used as the trigger for a Stanford Delay generator that sends out two time-delayed signals. I will call these Delay A and Delay B. Delay A is sent to the Quanta-Ray Nd:YAG and determines the time at which the flashlamps fire. This starts the lasing build up in the YAG cavity. Delay B is sent to the Q-switch that dumps the light out of the YAG cavity. Delay B thus determines when pump pulses actually arrive in the Regen. The output pulse of the regen is timed relative to this pump beam, so the delay between the FEL trigger and this Delay B sets the exact time in the FEL photon pulse where slicing occurs. The nominal value of this is 185 μs, though it requires adjustment depending on where you want to slice and the build-up time of the

photon pulse at different wavenumbers. The time difference between Delay A and Delay B determines how much build up has occurred in the Quanta-Ray lasing cavity before the dump. This should be set for maximum output power, which I find is about 150 μs. Using the Stanford delay box, I define Delay B to be 185 μs after the 10-Hz input trigger and Delay A to be 150 μs **before** Delay B. This allows us to adjust Delay B, changing the relative position of a slice within the FIR photon pulse, without changing the relative time between the flashlamps and Q-switch in the Quanta-Ray (i.e. maintaining maximum power output from the Nd:YAG).

Before Delay B is actually sent to the Quanta-Ray Q-switch it passes through a synchronization box provided by Spectra-Physics. This synchronization box receives two inputs: the Trigger In (Delay B) and a Synch In signal that comes from the photodiode monitoring the mode-locking behavior of the Tsunami. The pulse repetition rate of the mode-locked Tsunami is 76 MHz. The synchronization box delays the Trigger In (Delay B) until the arrival of the next pulse from the Synch In line and then sends it out on Trigger Out. This delay makes sure that there is always a constant time interval between the arrival of the pump beam (defined by the Quanta-Ray Q-switch) and the Tsunami pulse in the regen cavity. Since the Tsunami operates at 76 MHz, the delay added at the synchronization box is never more than $\sim$ 13 ns, which does not matter for purposes of getting the regen output at the right point of a several-μs-long THz pulse. This delay does affect the timing of our data sampling, which will be addressed below. When the

Trigger Out pulse is sent from the synchronization box, it also starts a counter for several other delays: regen Pockel's Cell 1, regen Pockel's Cell 2 and the monitoring oscilloscope.

As described up until now, the timing electronics make sure that the pump beam from the Quanta-Ray arrives in the regen cavity with a fixed time relationship to a pulse from the Tsunami and synchronized to the FIR pulse from the FEL. In passive mode, each pulse from the Tsunami enters the regen oscillator cavity, makes one round trip around the cavity and then leaves the cavity. When Pockel's Cell 1 is fired, any pulse in the cavity is trapped and all subsequent pulses are prevented from entering. Thus the timing of Pockel's Cell 1 is chosen to trap a single seed pulse in the oscillator cavity. This delay time is fixed by the geometry of the setup to be 12 ns.

When Pockel's Cell 2 is fired, it changes the polarization of the light in the oscillator cavity, which results in dumping the cavity into the amplifier and compressor stages. The delay on Pockel's Cell 2 is chosen to dump the oscillator at the point of maximum gain in order to achieve maximum power output. This time varies over the course of any given day from 285 ns to 400 ns as thermal drifts change the build-up time of the oscillator cavity. Pockel's Cell 2 ultimately determines the output time of the regenerative amplifier.

There are two delays that vary over the course of any given day. First, the exact timing of Pockel's Cell 2 is periodically adjusted. Second, there is a shot-to-shot variation in the delay induced by the Spectra-Physics synchronization box.

These two variable delays make it impossible to use a fixed delay to trigger an oscilloscope measuring a sample's response to the sliced pulse. Fortunately, the LeCroy oscilloscope I use has a smart trigger option that allows us to solve this problem. In smart-trigger mode, the oscilloscope arms the trigger when a TTL pulse arrives from the FEL, indicating that it is about to fire. The trigger actually fires when a signal arrives on one of the scope inputs. By connecting this input to a photodiode that sees some scattered light from the output of the regen, the scope trigger can be exactly synchronized to the NIR pulse and thus to the actual sliced THz pulse, regardless of any variable delays.

Chapter 4

Tuned Wafers

As described in Section 3.3, tuned wafers can be used to dramatically improve the contrast ratio of sliced pulses. Bryan Cole first used this technique, but manufactured his wafers by trial and error. He thinned his wafers slightly, then measured them with the FEL and repeated this process until he found a local minimum in total reflectivity. Since I wanted to extend the pulse slicer to operation at a variety of different wavelengths, it was important to find a reliable way to calculate the desired wafer thickness for a given wavelength, and then manufacture wafers to the desired specifications.

A tuned wafer is illustrated in Figure 4.1. Since the electron-hole plasma is generated only on the front surface, tuning the wafer thickness should have no effect on the energy reflected after the near-infrared pulse activates the switch. Experimental tests verify that upon activation a tuned wafer reflects about 60% of the incident FIR, identical to an untuned, unprocessed wafer.

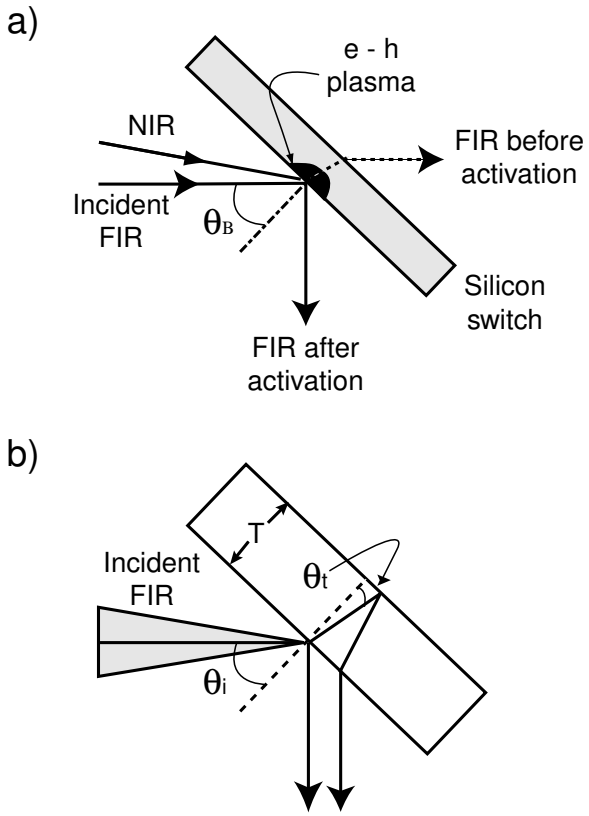

Figure 4.1: (a) Silicon switches transmit incident far-infrared radiation (FIR) before they are activated by a near-infrared (NIR) laser pulse that generates an electron-hole plasma on the surface, reflecting the FIR. Although separated here for clarity, the NIR and FIR are actually collinear. (b)The FIR is focused onto the switch, leading to a range of incident angles and some pre-activation reflection. The wafer thickness, T, is tuned to achieve destructive interference between the front and back surface reflections.

Although the technique described here was designed specifically for the UCSB pulse slicer, it can be applied to any laser-activated switching system in the mid to far infrared. Laser-activated switches have been used, for example, to dump pulses from passive resonators [3]. The switch-tuning method described here could be used to increase the Q of such resonant cavities by reducing pre-activation loss from the switch itself. Another method of reducing pre-activation reflectivity would be to illuminate the switch at Brewster's angle with a collimated beam, and thus reduce the spread of incident angles. However, the near-infrared energy required to activate a sufficiently large area for reflection of an unfocused beam would be prohibitive. An incident Gaussian of diameter 2.5 cm, for example, would require NIR pulses 150 fs long with energies in excess of 100 mJ, achievable only with extremely complex laser systems. Any switching method in which the FIR is brought to a focus on the switch must broaden the range of incident angles and can thus benefit from this technique.

4.1 Calculating the necessary wafer thickness

The switch thickness necessary for destructive interference can be calculated simply from geometric arguments if the incident wavelength and angle are known. Figure 4.1b indicates the relevant angles; it is important to remember that a π phase shift occurs upon reflection from the second interface.

The thickness of the wafer is given by

$$T = \frac{m\lambda}{2(\frac{n_{Si}}{\cos\Theta_t} - n_{air}\tan\Theta_t \sin\Theta_i)}, \tag{4.11}$$

where T is the thickness of the wafer, λ the wavelength of the incident light, n_{Si} the index of refraction of silicon, n_{air} the index of refraction of air, Θ_i the angle of incidence and Θ_t the angle of the transmitted beam in the wafer calculated from Snell's law. The order of the interference is given by m, which can be chosen to provide a reasonable wafer thickness. The index of refraction of silicon in the far infrared is roughly 3.42 and Brewster's angle is therefore 73.7 degrees. For wavelengths in the range of 100-300 μm, the first-order destructive interference corresponds to wafer thicknesses of $15.24 - 45.77$ μm.

If the FIR is incident precisely at Brewster's angle, Equation 4.11 can be simplified to

$$T = \frac{m\lambda}{2}\frac{(n_{Si}^2 + n_{air}^2)^{1/2}}{n_{Si}^2}. \tag{4.12}$$

In practice, however, experimental setups sometimes require variations in incident angle and Equation 4.11 may be more useful.

The pre-activation reflectivity of a switch arises primarily from the outer portions of the focused beam, which are incident at an angle different from Brewster's angle. The total reflectivity of a switch can be calculated by assuming a Gaussian beam of diameter 5 cm incident on a focusing optic of $f\backslash$ 2.4, with the center of the beam from this optic incident on the switch at Brewster's angle. Integrating the product of the surface reflectivity and the amplitude of the switch etalon effect

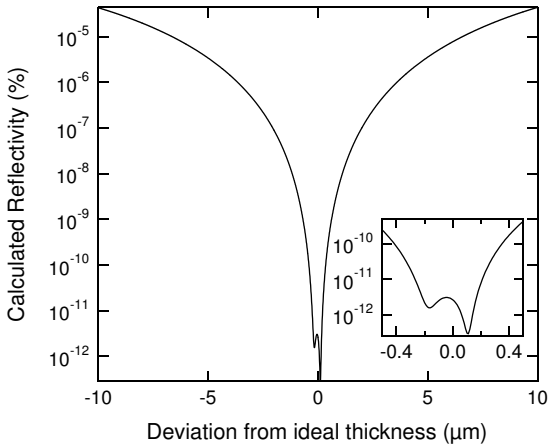

Figure 4.2: Theoretical calculation of the reflectivity of a wafer as a function of the variation in wafer thickness from ideal. The calculation is done with a Gaussian 5-cm-diameter beam of wavelength 294 μm focused by an optic of $f\backslash$ 2.4 onto a wafer of 44.87 μm at Brewster's angle. The calculation includes both reflection from the surface and interference effects. The inset shows an enlargement of the area around the minimum value.

over the distribution of incident angles yields a measure of the total reflectivity. Figure 4.2 shows a plot of this calculated reflectivity under illumination by 294-μm radiation as a function of the deviation from a switch thickness of 44.87 μm. The etalon-effect calculation assumes no loss in the switch, equal reflection at each surface and perfectly parallel surfaces. Although 44.87 μm is the ideal switch thickness for suppression of 294-μm radiation incident exactly at Brewster's angle, Figure 4.2 shows that a slightly thicker wafer is ideal for suppression of the total radiation incident at a distribution of angles. Since maximum contrast ratios require minimizing total reflectivity, $\pm$0.15 μm was chosen as an acceptable margin of error on the wafer thickness.

We chose silicon over GaAs as a switch material for longer electron-hole re-

combination times, which are necessary for generation of "flat-top" pulse shapes. High-resistivity silicon ($\sim 10,000\ \Omega cm$) is used to minimize transmission losses and thus enhance the etalon effect, though I have not investigated the performance of switches with lower resistivity. Surface roughness will also impact the performance of the switches as mirrors. A typical acceptable surface roughness is $\lambda/10$, or $\sim 10\ \mu$m for FIR radiation. However, the above calculation indicates that such a variation would have a severe impact on the efficiency of etalon suppression of reflectivity. To remain within the margin of error for wafer thickness, the surface roughness must be less than 150 nm. The use of visible lasers for alignment also requires surface roughness of 100 nm or less.

4.2 Wafer Manufacturing

To manufacture reflection switches, a high-resistivity silicon wafer must be thinned down to within 0.15 μm of the target thickness. The thinning process must maintain an optical-quality surface finish and minimize thickness variations across the wafer. A focused beam of 294-μm wavelength has a diffraction-limited spot diameter of only 0.86 mm, but in practice the steep incident angles and need to eliminate FIR scattering sources require a workable switch area of at least 1 cm^2, with only minor thickness variations or surface roughness across this area.

There are many ways to etch silicon. I investigated chemical etching and reactive ion etching but found them to be unsuitable for reasons I will detail

below. I ultimately selected a cyclical process of oxidation followed by oxide etching. Wafers were purchased from University Wafer; a typical test wafer had an initial thickness of approximately 50 μm and an initial rms surface roughness of $\sim$ 50 nm over an area of 1 mm^2. Surface roughness measurements were made with a WYCO NT1100 Optical Profiling System manufactured by Veeco. Measurements of wafer thickness were made with FTIR, to be described in more detail below.

Chemical etching using a solution of 90% Nitric acid, 5% Acetic acid and 5% Hydrofluoric (HF) acid was limited by the production of hydrogen bubbles on the surface of the silicon. These bubbles limited the flow of fresh etchant to the surface, resulting in a random distribution of small mesas, an rms surface roughness of 180 nm and a peak-to-peak roughness of $\sim$ 900 nm. Under agitation, the bubbles masked the downstream area of the wafer, leading to enlarged mesa size. Non-uniformity of the etch across the wafer due to unequal flow of fresh etchant to various areas was observed. Chemical etching using KOH was also investigated, but found to create pits in the $\langle 100 \rangle$ silicon surface after only a few minutes. This may be due to the anisotropic nature of the KOH etch, which etches the $\langle 111 \rangle$ silicon plane more slowly than other crystal planes.

Reactive Ion Etching (RIE) was investigated using etch plasmas of chlorine gas, sulfur hexafluoride (SF_6) and combinations of SF_6, Ar and O_2. Both Bosch and continuous processes were investigated. Typical switches after RIE processing had an rms surface roughness of 118-230 nm. A subsequent oxidation and oxide etch reduced this roughness to $\sim$ 150 nm with a peak-to-peak roughness of $\sim$ 890

nm. More troubling than surface roughness was the large non-uniformity in the etch across the wafer surface. This non-uniformity could be attributed to either deformation of the RF fields by the wafer itself, or loading of the etch plasma due to the relatively large area of silicon being etched.

The method ultimately selected for use was a cyclical process of growing an oxide layer into the silicon surface and subsequently etching away the oxide layer with HF acid. I oxidized 2-inch wafers in a Tystar wet oxidation furnace. The growth rate of oxides in such a furnace is well described by the Deal - Grove model. As described in Gandhi [23], growth of an oxide of thickness t μm consumes 0.44 t μm of silicon. A representative oxidation at 1050^o C for a duration of 15 hours 19 minutes 48 seconds grew a 2.39-μm oxide layer on each side of the wafer, consuming a total of 2.11 μm of silicon. The wafers were then submerged in HF acid, which etched away the silicon dioxide and stopped at the silicon surface.

Growth of oxides under wet thermal oxidation is limited by diffusion of water vapor to the silicon - silicon dioxide interface, which takes increasingly longer times as the oxide layer gets thicker. It is thus impractical to etch a switch wafer to its target thickness with a single oxide growth. There is no difficulty, however, in making several sequential oxidations and etches. Placement of the wafers such that the oxide layer is grown on both faces can also be used to accelerate the process. The rms surface roughness of wafers manufactured with this process was found to be 26 nm with a peak-to-peak roughness of $\sim$ 190 nm.

To determine the required oxidation time and characterize the manufactured

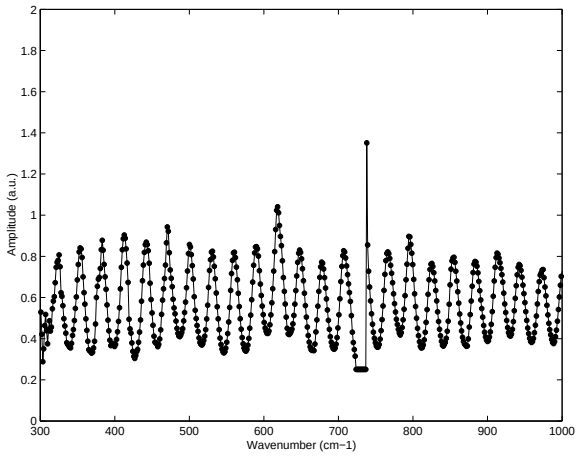

Figure 4.3: A characteristic scan from the FTIR for mid-infrared radiation incident normally on a thin silicon wafer. A fit to this data yields the period of the interference fringes, which can be used to calculate the wafer thickness.

switches, it is necessary to have a method for measuring the wafer thickness. I chose to use an interferometric measurement based on FTIR. Broadband mid-infrared light from an interferometer was sent through a wafer at normal incidence. By fitting to the resulting Fabry-Perot interference fringes, the thickness of the wafer can be calculated; a representative fringe pattern is shown in Figure 4.3. Initial variations in wafer thickness and slight non-uniformity across the wafer during the thinning procedure made it advantageous to measure many locations on a single wafer to locate the area with best wafer thickness before mounting the switch.

Once manufactured, I mounted the reflection switches on the front face of "C"-shaped frames machined from stainless steel. Since the FIR is incident at a very steep angle, the frames were positioned so that the FIR transmitted through the wafer passed through the open side of the "C," reducing the scattered FIR.

The frames attached kinematically to a positioning system, allowing fine control of the portion of the switch face actually used for switching. This frame system has dramatically increased the reliability of switching operations by eliminating switch breakage problems from directly handling thin silicon wafers.

4.3 Measurements of Wafer Efficiency

Wafers were manufactured according to the process above for wavenumbers of 34 cm^{-1} (λ = 294 μm, wafer = 44.9 μm), 51 cm^{-1} (λ = 196 μm, wafer = 29.9 μm) and 84 cm^{-1} (λ = 119 μm, wafer = 36.3 μm). The second-order interference was chosen for the wafer manufactured for 84 cm^{-1} for increased structural stability over the first-order wafer (18.1 μm).

The pre-activation reflectivity of a single wafer of 44.9 μm is plotted in Figure 4.4a for several different wavenumbers, along with a calculated plot of expected reflectivity. The agreement between the calculation and experimental data is surprisingly good, considering that the calculation assumes an exact wafer thickness, no loss and perfectly flat and parallel wafer surfaces. Figure 4.4b shows data and calculated expected reflectivity for reflection off two sequential wafers each of thickness 29.9 μm. Although the total reflectivity for the two-wafer case is considerably lower than that for the single-wafer case, it is evident that the reflectivity does not approach the low values calculated for ideal conditions.

It is important to remember that the focusing of the FIR onto the switch

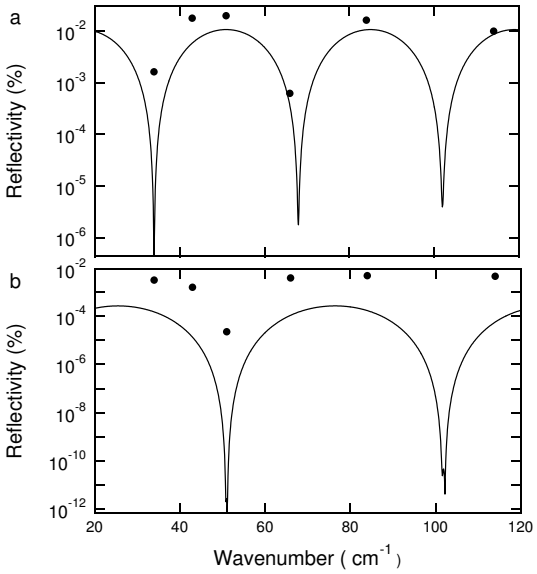

Figure 4.4: a) Theoretical calculation of the reflectivity of a single 44.9-μm wafer as a function of the wavenumber of the incident FIR. Symbols indicate data obtained with a wafer near this thickness for several different incident frequencies. b) Calculated reflectivity and data for two sequential wafers of 29.9 μm.

leads to a range of incident angles not only in the plane of reflection shown in Figure 4.1 but also in the orthogonal plane. These orthogonal rays also affect the pre-activation reflectivity of the switches, but I did not consider them in this analysis.

Direct measurements of the contrast ratio were undertaken in the two-switch setup of the UCSB Pulse Slicer using a FIR beam of approximately 4 cm. The average contrast ratio with a pair of off-resonant wafers was 60. With a pair of wafers tuned to the ideal thickness for the illumination frequency, contrast ratios as high as 4.5×10^4 were achieved, with an average contrast ratio of 2.5×10^4. The contrast ratio of tuned wafers was on average 400 times larger than that

obtained with untuned wafers. The measured increases ranged from a factor of 100 to a factor of 740 with an uncertainty of 10%. This large range with relatively small uncertainty indicates that variations from the wafer processing still play a significant role.

When the FIR beam was reduced to approximately 1.8 cm, contrast ratios as large as 7.4×10^4 were measured. The pre-activation reflectivity was measured to be 1800 times smaller, on average, than the pre-activation reflectivity of an untuned wafer. Although further work to reduce the variation from wafer to wafer can be undertaken, the techniques presented here provide a reliable method for manufacturing laser-activated switches for use at specific FIR wavelengths that can improve switching contrast ratio by a factor of up to 1800.

Chapter 5

Probing the Hydrogenic Donors

In this section I will describe the growth and preparation of the samples used in the experiment and the experimental methods I used to analyze the orbital state of the donors. I did essentially two different types of measurements: short and long pulse. Long-pulse measurements were used to probe selection rules and perform some magneto-spectroscopy of the state transitions. These experiments were done using full-length FEL output pulses ($\sim$ 4-5 μs). Short-pulse experiments used the pulse slicer, described in Section 3.3, to generate THz pulses with widths between 0 and 3.5 ns. For both types of experiments, I used photoconductive detection as described in this chapter.

5.1 Samples

I would like to study electrons bound to donors in GaAs in the limit in which the electrons are well isolated from one another. Since the Bohr radius of a single donor-bound electron is roughly 10 nm, this requires a low donor density. It turns out to be sufficient to grow GaAs without any intentional doping and rely only on the nonzero flux of impurities to provide the donors. Since the impurity density is quite low, the depletion length on the surface of the sample will be quite long. This requires reasonably thick samples to make sure that there are plenty of donors with electrons still present. Thick samples are also advantageous because they allow us to increase the total number of donors studied, and therefore the amplitude of our signals, without reducing the spacing between donors.

The particular samples I used were grown by Molecular Beam Epitaxy by Colin Stanley of the University of Glasgow. A 15 μm layer of GaAs is grown with no intentional doping on top of a semi-insulating GaAs substrate. A 20-nm capping layer of GaAs doped with silicon (1×10^{18} cm^{-3}) is grown on top. This capping layer facilitates Ohmic contacts to the sample. Hall studies performed by the Glasgow group determined that the donor density of the sample I used (B434) was 2.8×10^{14} cm^{-3} and the mobility was 138,000 at 77 K. Sulfur is the dominant hydrogenic donor.

The patterning of the contacts was done with conventional optical photolithography, e-beam metal deposition and chemical etches. The recipes for these

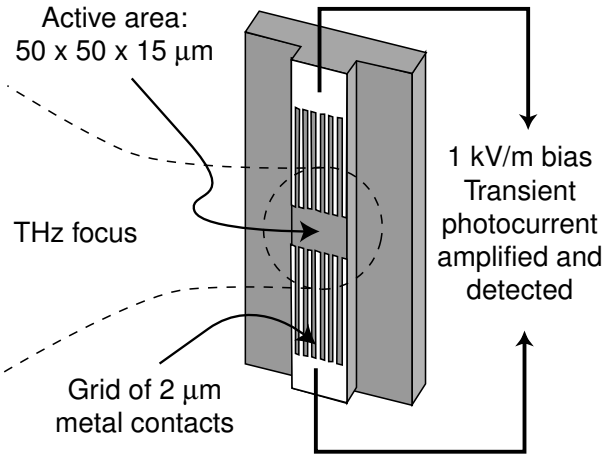

Figure 5.1: A schematic of the samples used in the experiment.

processes can be found in Appendix A. A diagram of the sample is shown in Figure 5.1. The Ohmic contacts are made with a "finger contact" pattern, originally designed by Bryan Cole, to minimize the distortion of the THz electric field. If the contact was a solid strip of metal, THz electric fields along the direction of the metal edge could drive currents in the metal; the gaps in the metal prevent currents from being driven along the edge while still allowing the application of a uniform electric field across the sample. The grid consists of 10 periods of a 2-μm strip of metal alternating with a 2-μm gap. The Ohmic contacts are made with NiAuGe and are then annealed. After the anneal, the 20-nm capping layer is etched off the remaining area of the sample, which ensures that there are no extra silicon dopants in the active area of the sample.

After the finger contact pattern is deposited, a second round of photolithography and e-beam deposition is used to pattern TiAu contact pads. These pads

overlap the finger contact metal and then expand away from the sample into an area large enough to make physical contact with the wires that will be used to make electrical contact with the outside world. This electrical contact is made by pressing a thin ($25 - 50$ μm) gold wire into indium that has been pressed onto the contact pad. Once the samples have been mounted, this gold wire is run to a Samtek pin on the sample mount. Finally, a mesa 15 μm deep is etched around the finger contact area to isolate the active area from everything else.

As shown schematically in Figure 5.1, the focal spot of the THz radiation was kept larger than the active area of the sample. This ensured that the field across the active area was uniform in intensity. This requirement is in fact very easy to achieve since diffraction makes focusing the long-wavelength THz radiation to a tight spot difficult.

5.2 Lens-Mounted Samples

One of the major challenges faced in this experiment was coupling sufficient THz intensity into the sample. To address this problem, I developed a method by which a sample could be mounted on the back side of a silicon solid-immersion lens. A schematic of this is shown in Figure 5.2. The solid-immersion lens is predicted to increase the intensity at the active area by a factor of the index of refraction squared ($n_{Si} = 3.42, n_{Si}^2 \approx 10$). Although it was not possible to measure the actual intensity at the active area, a substantial improvement in

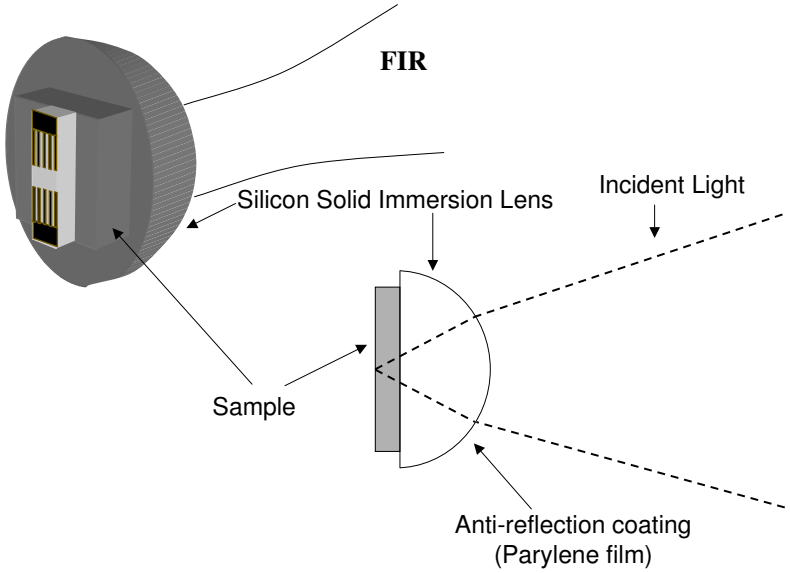

Figure 5.2: A schematic of a lens-mounted sample.

coupling efficiency was observed.

This increase in intensity arises from a tighter focal spot, which arises because of two effects. First, the incident rays are bent toward the center of the lens as they cross the surface. The change in incident angle is given by Snell's Law, so

$$\Theta_r = n_{Si} \sin(\Theta_i), \tag{5.13}$$

where Θ_r is the refracted angle inside the lens and Θ_i is the initial incident angle. This change in angle brings the light to a tighter focus, and accounts for one factor of n_{Si}. The second factor arises from the wavelength of the light, which is modified from λ to λ/n_{Si} in the lens. This decrease in the wavelength also allows

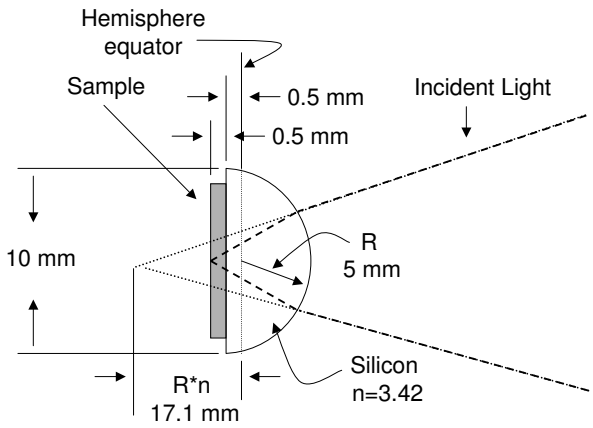

Figure 5.3: The dimensions of a hyper-hemispherical solid-immersion lens and required focal distances.

for a tighter focus and accounts for the second factor of n_{Si}.

The lenses used to accomplish this focusing are manufactured from silicon into a hyper-hemispherical shape. This is simply a sphere that is cut off along a plane that is offset from the equator. I used lenses with a radius of 5 mm cut off along a plane 0.5 mm from the equator. This is detailed in Figure 5.3. The sample, which is also 0.5 mm thick, is mounted to the backside of the lens. To focus properly on the sample surface after the refraction at the lens surface, light is focused by an external element (in our case an OAP) to a point that is a distance R*n from the backside of the lens.

As mentioned above, our experiments in coherent control require that the focus of the THz be larger than the active area of the sample. The diffraction-limited spot size of the THz after focusing by the lens is roughly 50 μm, the same size as the sample. However, it is nearly impossible to achieve a diffraction-limited spot

with THz radiation, and so the broadening of the spot size makes it sufficiently wide to overlap the sample area.

Before attaching the sample, the silicon lenses were coated with a layer of Parylene D to create an anti-reflection coating. Parylene D is a polymer with index of refraction 1.62 at THz frequencies. This is quite close to the square root of the index of silicon ($\sqrt{3.42} = 1.85$), which is ideal for an anti-reflection coating. A layer of Parylene D 45.4 μm thick was deposited on the curved surface of the lens by Specialty Coating Systems. This thickness was chosen to cause destructive interference of reflections at 34 cm^{-1} and therefore increase the coupling of THz radiation into the sample. The transmission window for a film of this thickness was wide enough that the coating did not negatively affect the coupling of radiation at 51 or 84 cm^{-1}.

Perhaps the most challenging task is to mount the sample so that the 50-μm active area is positioned at the exact focal point of the lens. To do this, I built an alignment jig, detailed in Figure 5.4. This aligner illuminates the lens with 1.3-μm light (below the silicon band gap) that is originally focused past the silicon lens surface as described above. The sample is placed on the back side of the lens and the light then emerges from the sample surface. By using an infrared viewer, it is possible to see both the sample features that define the active area and the light emerging from the sample that indicates the focal point. The sample can then be positioned over the focus using micro-manipulator arms. Finally, the sample is glued to the lens surface using cyano-acrylate glue (superglue), which is

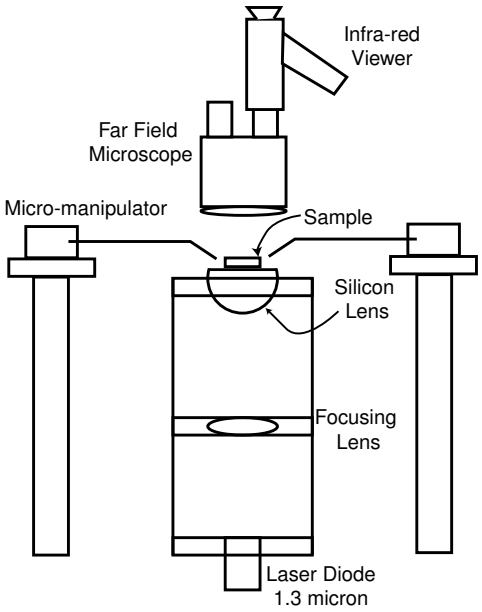

Figure 5.4: A schematic of the jig used to align the sample to the focus of the solid immersion lens.

transparent to THz frequencies and holds at low temperatures.

5.3 Magneto-optical Cryostat

Since shallow donors have ionization energies of only a few meV, it is necessary to perform experiments at low temperatures in order to have electrons bound to the donors. The ionization energy of a hydrogenic donor at zero magnetic field is just over 4 meV, which corresponds to a temperature of about 10 K, so a liquid helium bath (4 K) is necessary. This is provided using a magneto-optical cryostat. A diagram of the cryostat is shown in Figure 5.5.

The helium bath provides sufficiently low temperatures for the experiment, but there is a constant flow of bubbles as the helium boils off. These bubbles act

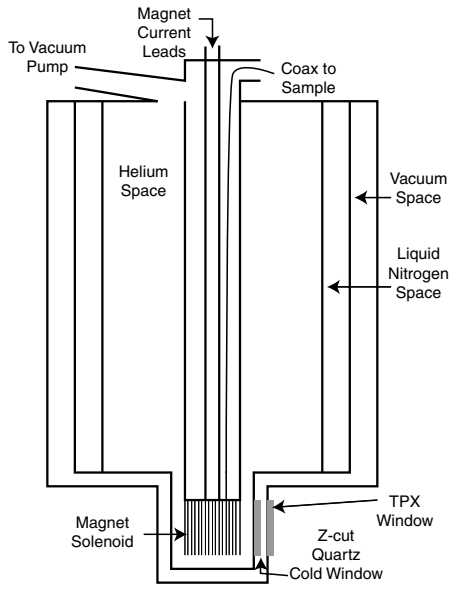

Figure 5.5: A diagram of the magneto-optical cryostat used in the experiment.

as scattering centers for the incident THz radiation. To eliminate these bubbles, I pump on the helium space to drive the helium into the superfluid phase. In this phase, all evaporation comes from helium at the top of the bath and all bubbles in the bath are eliminated.

At the base of the cryostat is a superconducting magnet used to apply the magnetic field to the sample, which is mounted in the bore of the magnet near the center of the solenoid for a uniform field. The entire magnet is relatively small, and is a single, not split, coil. The distance from the outer window to the center of the magnet bore is short (about 3 inches), which permits the use of small $f\backslash$ optics to obtain a tight focus and high THz intensity on the sample. Relatively high magnetic fields (up to 7 T) can be generated. The magnet is powered by a

HP 6260B power supply controlled by an American Magnetics magnet controller. The current passing through the magnet is monitored using a 1-mΩ resistor in series with the magnet, and the current-to-field conversion for the magnet is well calibrated. The THz radiation is incident on the sample along the direction of the magnetic field (Faraday geometry). The inner (cold) window of the cryostat is made of z-cut quartz and the outer (warm) window is made from TPX plastic.

5.4 Photoconductive Detection Electronics

I detect the presence of electrons in the excited states by using photoconductive detection, which takes advantage of the more rapid ionization of the electron into the conduction band from the excited states rather than the ground state. To measure the presence of the electrons in the conduction band, I apply a DC bias to the sample and measure the transient change in the conductivity that occurs when the additional electrons are ionized. This transient change in conductivity is measured as a voltage drop across the resistance of the sample leads.

I determine the number of electrons that will be excited to the conduction band from the sample parameters. The sample is 15 μm thick, with a doping density of 2.8×10^{14} cm^{-3}. With this doping density, the depletion length from the surface is roughly 1 μm. I therefore take a sample volume of $50 \times 50 \times 14$ μm, multiplied by the donor density, to determine that there are $\sim 1 \times 10^7$ donors in the sample. This limits the magnitude of the signal I can measure.

Figure 5.6 shows a schematic of the electronic components in the system. I apply a bias voltage of 50 mV across the 50-μm sample area using a 9-V battery. The resulting field (1 kVm^{-1}) is well below the threshold for impact ionization of the donor bound electrons. I couple into the sample using a bias-tee (Mini-Circuits ZFBT-4R2G), which separates the RF and DC components of the signal. Thus the transient photocurrent that is generated when electrons are excited into the conduction band is coupled out through the RF channel and then amplified using fast amplifiers (Mini-Circuits ZFL-1000LN, 0.1-1000 MHz bandwidth). This amplified signal is then coupled into a 1-GHz digitizing oscilloscope (Lecroy LC584AM) and the transient photocurrent can be measured. A representative scope trace of a sample responding to a sliced pulse is shown in Figure 5.7.

The trace from the oscilloscope is then read by a computer. To assign a numerical value to the measured photosignal, I numerically integrate over the regions indicated in Figure 5.7. I subtract the integrated value over region 1 from the value for region 2 to eliminate any background offset. This final value (in arbitrary units) is the measured photosignal.

5.5 RF shielding

The transient photocurrent measured in this experiment had a rise time on the order of a few ns, which puts the signals in the radio-frequency (RF) regime. RF noise was thus a major problem for the experiment. The RF noise can be

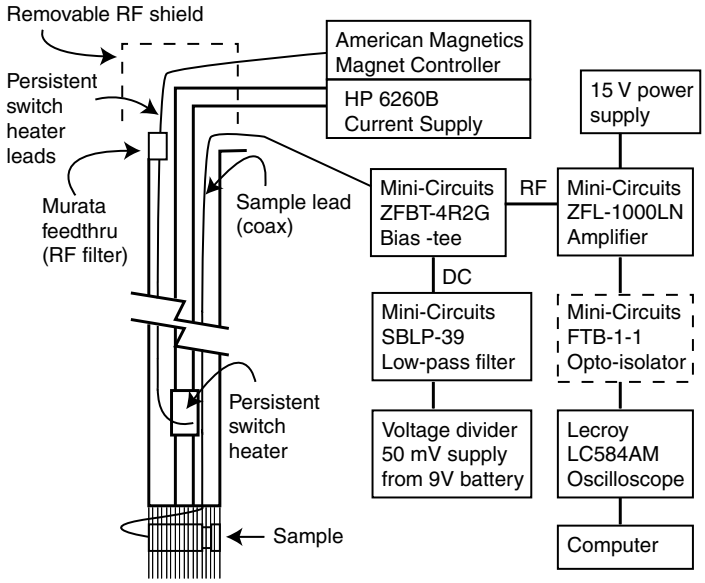

Figure 5.6: A schematic of the electronic components used to apply a bias voltage and read out the transient photosignal. The dashed box on the right encloses a Mini-Circuits high-speed optical isolation transformer that was sometimes used to reduce ground loop noise.

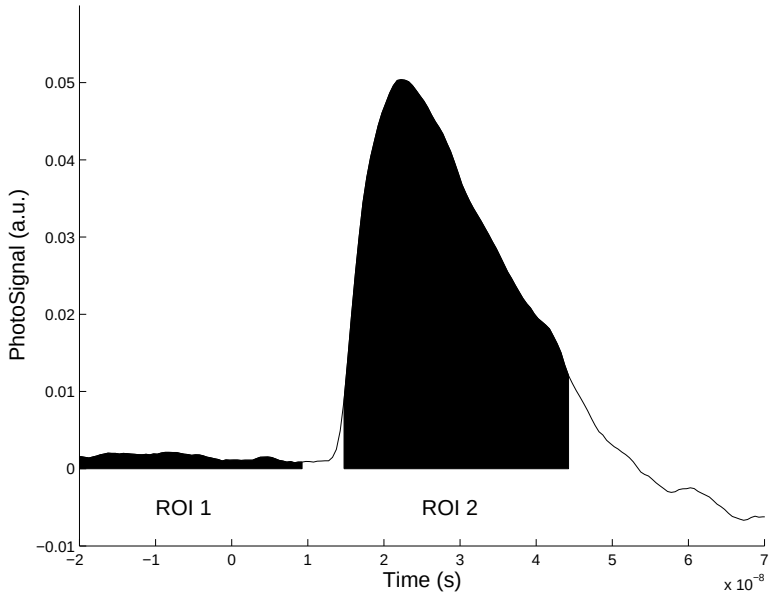

Figure 5.7: A representative photosignal with the regions of intrest (ROI) for numerical integration indicated.

broken down into two categories: random noise and synchronized noise. Random noise sources include AM radio stations and possibly RF sources from the nearby Santa Barbara Airport. Since the noise is random, the signal-to-noise ratio could be improved by averaging the signal. Synchronized noise arises almost exclusively from the Pockel's Cells used in the regenerative amplifier to rapidly switch the polarization state of the light and dump it out of the cavity. These Pockel's cells define the arrival time of the NIR light at the pulse slicer and thus the creation of the short THz pulse. The Pockel's Cells operate by switching a high voltage across a birefringent crystal. This rapid high-voltage switching generates RF noise that is synchronized to the THz pulse and thus to the transient photocurrent I am trying to measure. Averaging over traces, therefore, does not significantly improve the signal-to-noise ratio. Attenuation and filtering of RF noise is therefore quite important to the experiment.

To reduce the noise coupled out of the regenerative amplifier box, Bryan Cole replaced the plastic lids on the cavity with metal lids contacting a conductive strip around their edges to provide a maximally continuous conductive enclosure. I found it necessary to make a number of improvements to the RF shielding of the cryostat itself. Most of these improvements are shown in Figure 5.6. One of the largest mechanisms for coupling in RF noise was the magnet current leads, which essentially act like a large antenna.

A first step to reducing the noise picked up by the magnet leads was to shield the current leads with a "lossy line" material. The lossy line material is a hollow

rubbery tube that is conductive but has very high electrical loss. RF signals therefore penetrate the lossy line filter but are dramatically attenuated before they reach the inner hollow where the magnet current leads are run. While the lossy line shielding did reduce the RF noise, it was not sufficient.

The next step in RF shielding was to install a persistent switch on the magnet current leads just above the magnet itself. A persistent switch is superconducting at helium temperatures and provides a continuous path for the current flowing through the magnet. Thus when the switch is engaged, the magnet current simply flows in a closed path through the magnet solenoid, around through the persistent switch and back into the solenoid. The current leads that connect the magnet to the outside world are not necessary to maintain field in this configuration and so they can be completely disconnected from the system. This allowed us to set the magnetic field using a power supply, engage the persistent switch, disconnect the magnet leads from the cryostat and place a RF shield (closed metal surface) over the magnet current lead inputs into the cryostat. The persistent switch is turned on and off simply by applying a voltage to a resistive wire wound through the switch, thus heating the switch and destroying the superconducting shunt. This switch-heater lead was also shielded by placing Murata feedthrus (RF filters) at the entrance to the cryostat. The persistent switch could be used to reduce RF noise whenever operating at constant magnetic field values, but could not be used when scanning the magnetic field.

The electronics used to detect the transient photoconductivity were also sensi-

tive to RF noise. To reduce the RF pickup, the bias voltage applied to the sample was generated from a 9-V battery and then passed through a low-pass filter before the bias tee. The output of the bias tee was channelled directly into the fast amplifiers. The power supply for the fast amplifiers was also passed through a low-pass filter before being applied to the amplifiers. The entire battery bias and amplifier assembly was encased in a continuous metal box. Last, an optical isolation transformer was sometimes used before coupling the signal into the oscilloscope to eliminate any potential ground-loop noise.

Chapter 6

Experimental Results

In this chapter I will present results from experiments on ensembles of hydrogenic electrons using the experimental techniques described in the previous chapters. I organized this chapter so that the results of each section build on the techniques or results of previous sections. However, this organization is not chronological. Sometimes a particular section was motivated by preliminary results from a subsequent section. I have summarized the motivation for each section independently. In Section 6.2 I present spectroscopic data showing photoconductivity as a function of magnetic field under illumination by a fixed frequency of THz light. In Section 6.3 I present data on the verification of the selection rules for hydrogenic transitions. In Section 6.4 I present data on Rabi oscillations driven between the 1s and $2p^+$ states with a transition energy of 84 cm^{-1}. Last, in Section 6.5, I present results on attempts to extend the Rabi oscillation experiment to other frequencies and transitions.

All experiments described in this Chapter were performed on unintentionally-doped GaAs as described in Section 5.1. Except for experiments with the lens-mounted sample described in Section 6.5, the sample was mounted so that the incident radiation was focused directly onto the 50 $\times$ 50-μm active area on the front side of the wafer. For all experiments, the samples were mounted in the magneto-optical cryostat described in Section 5.5.

6.1 Saturation

For the spectroscopy experiments described in Section 6.2 and studies of the selection rules described in Section 6.3, it is important to operate as close to the regime of linear response as possible. The samples I studied are highly sensitive photoconductors that are easily saturated by the high power necessary for the nonlinear experiments described in Section 6.4. In this section, I investigate the intensity-dependent photoconductive response to few-μs-long pulses from the UCSB FEL. I determined that tremendous attenuation is necessary to approach the linear regime while maintaining a sufficient signal-to-noise ratio.

In Figure 6.1a, I show the photoconductivity from the sample versus relative intensity of THz radiation. The sample is illuminated with a pulse of radiation at 52.3 cm^{-1} about 5 μs long and the intensity modulation is done with calibrated plastic attenuators. It was not possible to calibrate the exact intensity of the maximum signal. Looking at the blue trace, it is clear that the photoconductiv-

ity displays saturation behavior as the intensity rises. In the other traces I have included additional un-calibrated attenuation to reduce the overall intensity and determine whether the saturation behavior appears only at high absolute intensities. The saturation behavior remains evident even at low absolute intensities. In Figure 6.1b, I show the same plot for illumination with a $\sim$ 5-μs pulse at 35.8 cm^{-1}. In this case, the intensity modulation is done with crossed polarizers, but the same saturation effect is seen when calibrated plastic attenuators are used. The non-zero minimum value for the blue trace arises from the imperfections of the polarizer discussed in Section 3.1. Again, with the addition of un-calibrated attenuation, the saturation behavior is still apparent.

The saturation may be due to the finite number of donors in the sample. The impurity density is quite low (2.8×10^{14} cm^{-3}); under high-intensity illumination the electrons are rapidly promoted to higher orbital states and ionized into the conduction band. Once the ground state has been depopulated, continued illumination will not increase the number of carriers. At best, an equilibrium number of carriers is sustained by continually re-ionizing electrons that are recaptured by the donors. Continued illumination of electrons ionized to the conduction band may also result in carrier heating, which could affect the photoconductivity by changing the mobility. I suspect that the saturation effect arises from a combination of the depopulation and carrier heating (possibly reaching a steady-state temperature).

In Sections 6.2 and 6.3 I describe magnetospectroscopy of the state transitions

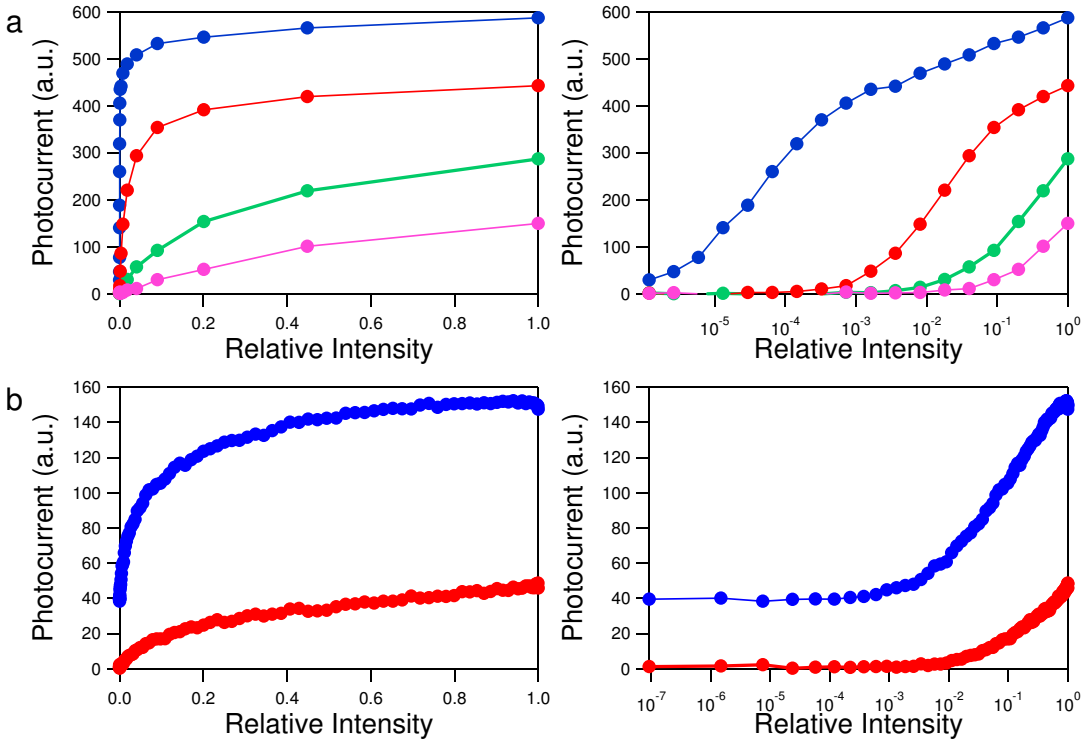

Figure 6.1: Photoconductivity versus relative intensity for (a) illumination with THz radiation at 52.3 cm^{-1} with intensity modulation using calibrated attenuators and (b) illumination with THz radiation at 35.8 cm^{-1} with intensity modulation using crossed polarizers. The left and right graphs show identical data, but use linear and log scales, respectively, for clarity. The different traces in each plot are taken with a different amount of uncalibrated attenuation, which reduces the absolute intensity. The lower traces in each plot show that a saturation effect is still present at low absolute intensities.

and verification of the selection rules. These experiments were both performed using long pulses ($\sim 5\ \mu s$) of THz radiation, and so the effect of saturation on the photoconductivity is particularly significant. To minimize the effects, the experiments were performed with sufficient attenuation to be in the region where the slope of the photoconductivity versus intensity was large. However, as shown in Figure 6.1, even with significant attenuation there is still some saturation effect, which means that a small measured change in photoconductivity corresponds to a significant change in the intensity of light coupling in to the transition being studied. In the spectroscopy and selection-rules experiments, the measured photoconductivity is taken as the numerical integration of the photosignal over the first $\sim 1\ \mu$s of the response to the FEL. This window is chosen to avoid roll-off from the fast amplifiers, which have a bandwidth of 0.1-1000 MHz. No difference in behavior was observed when measurements were made using a slow amplifier and integrating over the entire 5-μs pulse.

6.2 Spectroscopy

The spectroscopy experiments were primarily motivated by preliminary results of the study of selection rules, where I observed several confusing effects. Spectroscopic identification of the state transitions that were resonant at particular magnetic fields and frequencies eliminated some of this confusion. One particular concern was to determine whether two-photon excitation was influencing the

photoconductivity signal observed from the $2p^-$ transition. This could occur if the THz radiation were close to resonance for both the 1s to $2p^-$ transition and the $2p^-$ to conduction band (N = 0) transition.

Since the sample is cooled to $\sim$ 1.5 K and recombination to the ground state (1s) is much faster than the repetition rate of the laser, the electrons are expected to be in the ground state before every pulse of the THz radiation. I perform spectroscopy of the state transitions by illuminating with a fixed THz frequency and sweeping the magnetic field. I monitor the photoconductivity signal that arises when the transition from the 1s state to an excited state comes into resonance. The data should be compared to the theoretical transition energy: the difference between the electron state energies shown in Figure 2.2. In Figure 6.2 I show a plot of the transition energy as a function of magnetic field. The transition between the N = 0 and N = 1 Landau levels is called the cyclotron transition. The $2p^-$ to N = 0 transition is included in Figure 6.2 because of the concern about a two-photon transition.

Spectroscopic experiments were performed with pulses of radiation approximately 5 μs long. The calibration of the illumination frequencies is based on measurements of side-band generation performed by Sam Carter. In sideband generation, the THz frequency modulates an NIR beam in a quantum-well structure to generate light at the sum and difference of the NIR and THz frequencies. The frequency of the initial NIR frequency and sum (or difference) frequency can be measured with a high degree of accuracy to determine the THz frequency. In

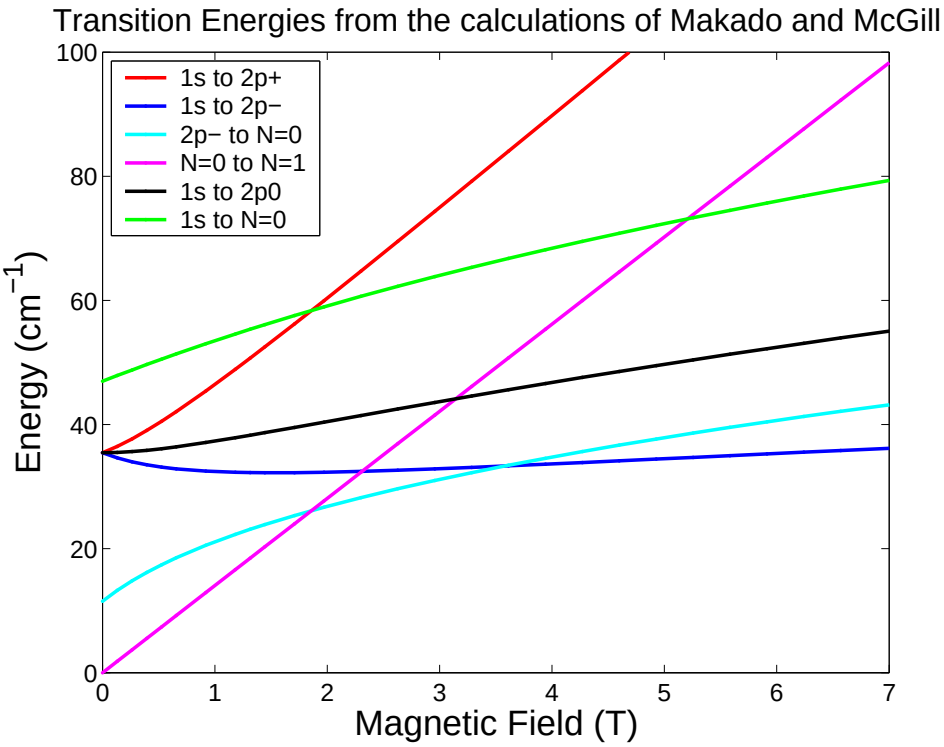

Figure 6.2: The energy of transitions between the various electron states shown in Figure 2.2 as a function of magnetic field.

the range of 32.5 to 35.8 cm^{-1}, side-band measurements of the THz frequency at a few specific values of the FEL terminal voltage were used to calibrate the terminal voltage-to-frequency conversion over the range.

For some of the spectroscopy experiments, the polarization of the incident radiation was varied to help distinguish between transitions with different selection rules. For example, the cyclotron transition requires σ^+ light, while the $2p^-$ transition should require σ^-. A more extensive study of the selection rules is described in Section 6.3. The different polarizations were generated using the variable polarizer described in Section 3.2. At each wavenumber, measurement of the Stokes parameters was used to determine the micrometer setting that corresponded to horizontally-polarized light (π_x), left-hand-circularly-polarized light (σ^-) and right-hand-circularly-polarized light (σ^+). The spectroscopy was then repeated at each of these polarizations. The handedness of polarization (σ^+ versus σ^-) that should couple to a particular transition depends on the direction of the magnetic field (+z versus -z). Unfortunately, I did not always use the same direction of magnetic field. For clarity of the discussion, I have corrected all of the data in this section to be consistent; for data taken with the magnetic field in the -z direction I have switched the designation of σ^+ and σ^-. Any spectroscopy that shows only a single trace was performed with π_x-polarized light.

Since the THz emitted by the FEL is quite intense (up to 1 kW at around 2.5 THz), long pulses of radiation can generate a photoconductivity signal even when not resonant with a particular state transition. This was described in more

detail in Section 6.1. To avoid the saturation effects, the spectroscopy is always performed with some attenuation. The exact amount of attenuation for a given wavenumber is chosen to provide a clear signal with good signal-to-noise ratio.

6.2.1 Magnetospectroscopy of the $2p^+$ state

In Figure 6.3 I show the results of magnetospectroscopy at 84.5 cm^{-1} (2.53 THz). The theoretical transition energies (Figure 6.2) are plotted with the data, but the range of transition energies shown is chosen to zoom in around the actual illumination frequency. At low magnetic fields, a very large signal appears. This signal decays away as the field is increased to reveal the state transitions. Although I do not fully understand the origin of this low-field signal, it appears at all frequencies. I believe it is related to the convergence of the energy levels of many higher-order states near zero field. The 1s to $2p^+$ state transition and the cyclotron transition appear in the spectroscopy data and agree well with the theoretical transition energy. I have added cut lines through the data in green to indicate the illumination frequency and the magnetic field at which the peaks are observed. If theory and experiment agree perfectly, the theory line for a particular transition should go through the crossing of the two cut lines. At least one additional peak appears around 2 T, which presumably corresponds to a transition to a higher excited state not included in the theoretical calculation. The good agreement between the cyclotron frequency and the observed cyclotron resonance

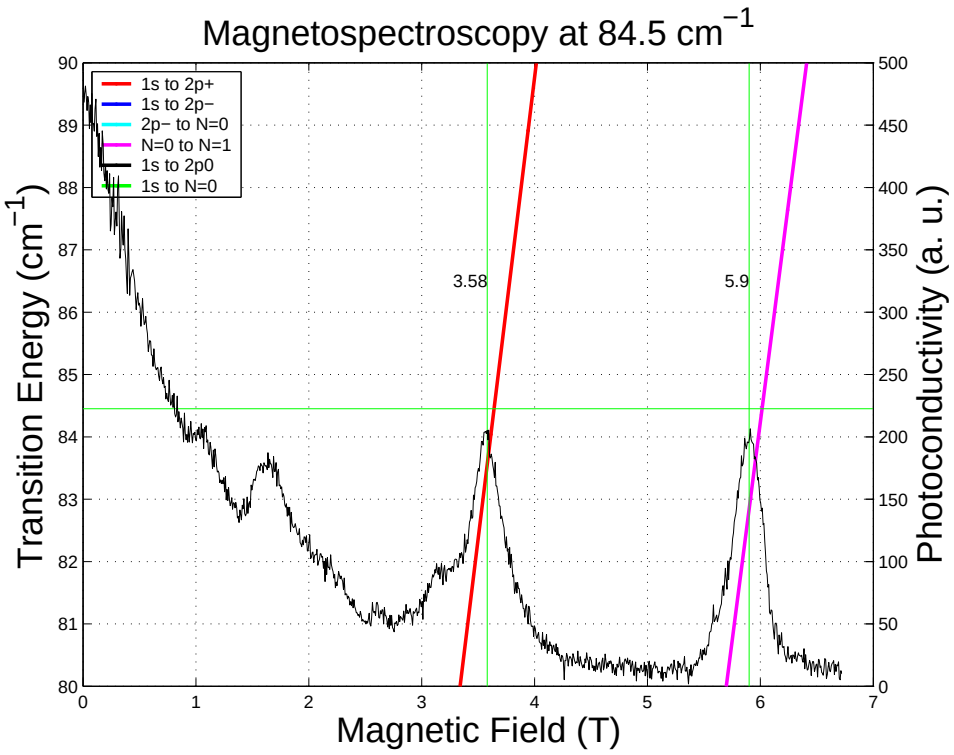

Figure 6.3: Photoconductivity as a function of magnetic field under illumination at 84.5 cm^{-1}. Theoretical transition energies near the illumination frequency are shown.

confirms that the magnetic field is well calibrated. The observation of the cyclotron transition implies that electrons are present in the N = 0 Landau level before THz illumination arrives. However, as described by Bluyssen, et al. [24], the N = 0 population does not have to be large.

In Figure 6.4 I again show the results of magnetospectroscopy overlayed with theoretical transition energies, this time for illumination at 52.3 cm^{-1} (1.57 THz). The 1s to 2p^{+} transition and the cyclotron transition are again evident and agree with the theoretical transition energy. Also evident is a small peak at 5.35 T. This resonance is about 1.5 cm^{-1} (or, equivalently, 0.6 T) from the calculated

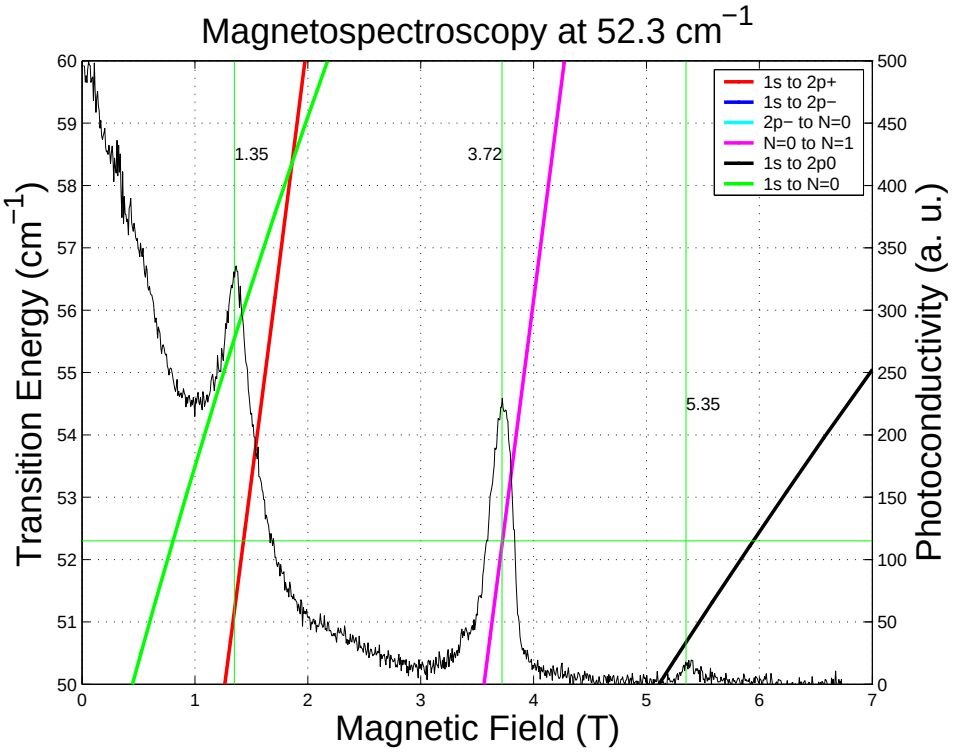

Figure 6.4: Photoconductivity as a function of magnetic field under illumination at 52.3 cm^{-1}. Theoretical transition energies near the illumination frequency are shown.

1s to 2p^0 transition energy. Although this transition is forbidden in the Faraday geometry, it appears due to a small component of the THz electric field that is not perpendicular to the sample face.

6.2.2 Magnetospectroscopy of the 2p$^-$ state

In Figure 6.5 the results of magnetospectroscopy are shown for illumination at 35.8 cm^{-1} (1.07 THz). The cyclotron transition appears at 2.5 T; for clarity I have expanded the region around 2.5 T in Figure 6.6. It is clear that σ^+ light couples into the cyclotron transition much more strongly than σ^- light. The 2p$^-$

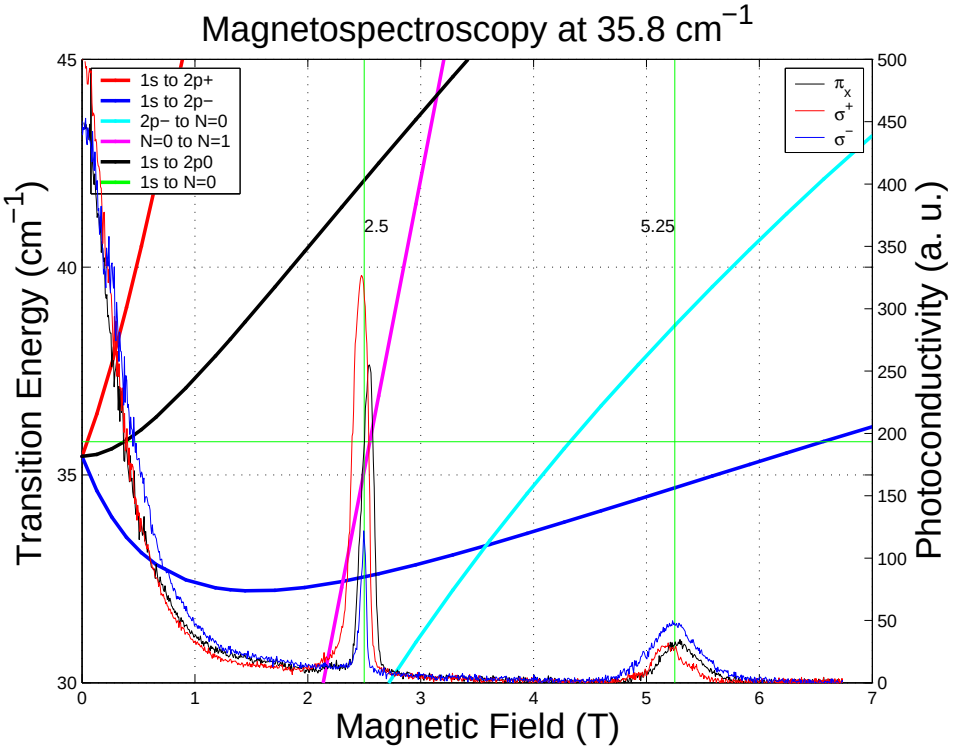

Figure 6.5: Magnetospectroscopy with illumination at 35.8 cm^{-1}.

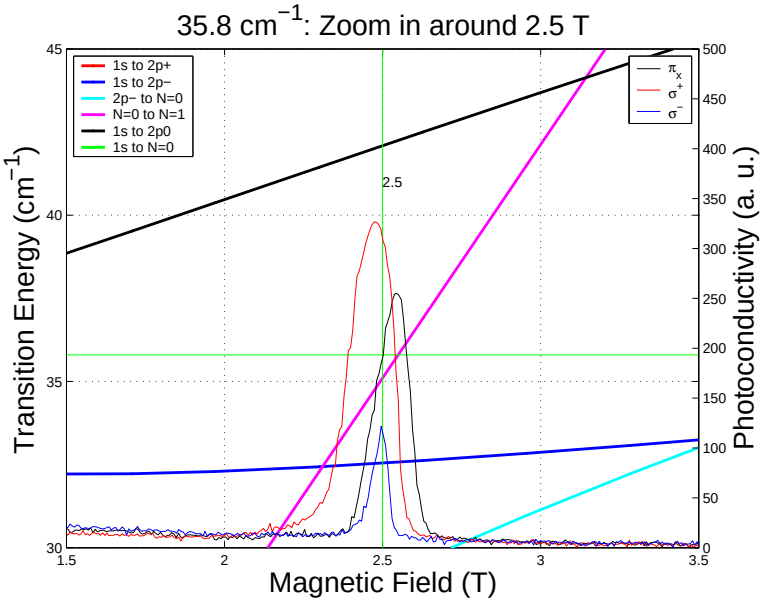

Figure 6.6: Magnetospectroscopy with illumination at 35.8 cm^{-1}, zoomed in around 2.5 T for clarity.

transition is at 5.25 T and σ^- light couples into this transition more strongly than σ^+ light. The location of the $2p^-$ peak appears to be about 1 T lower than the expected location of the peak based on the theoretical calculation.

Figure 6.7 shows magnetospectroscopy at 35.0 cm^{-1} (1.05 THz). The cyclotron transition appears at 2.47 T and the $2p^-$ transition at 4.4 T. As with illumination at 35.8 cm^{-1}, this is below the theoretical magnetic resonance value. Figure 6.8 shows results for illumination at 33.6 cm^{-1} (1.01 THz). The cyclotron peak is clear at 2.43 T and is largest for σ^+, as expected. According to the theoretical curves, $2p^-$ peaks should appear at about 0.4 T and 4 T. A peak in the photoconductivity signal using σ^- light does appear at low field, and it is clear that σ^- light induces a larger signal up to about 5 T, with the exception of the cyclotron peak. However, no clear peak is visible near 4 T. This is consistent with the trend in the magnetospectroscopy data that the $2p^-$ resonance appears at lower fields than calculated.

Figure 6.9 shows data obtained with illumination at 33.4 cm^{-1} (1.00 THz). This frequency crosses the theoretical two-photon resonance point. The cyclotron peak is clear at 2.38 T, but there is no clear $2p^-$ peak. There does appear to be a large shoulder on the high-magnetic-field side of the cyclotron peak, which is consistent with the $2p^-$ transition appearing below the theoretical resonance value. This offset of the observed $2p^-$ transition from the calculated field value explains why we do not observe a resonance at the two-photon point.

Figure 6.10 shows magnetospectroscopy at 33.1 cm^{-1} (0.99 THz). The data

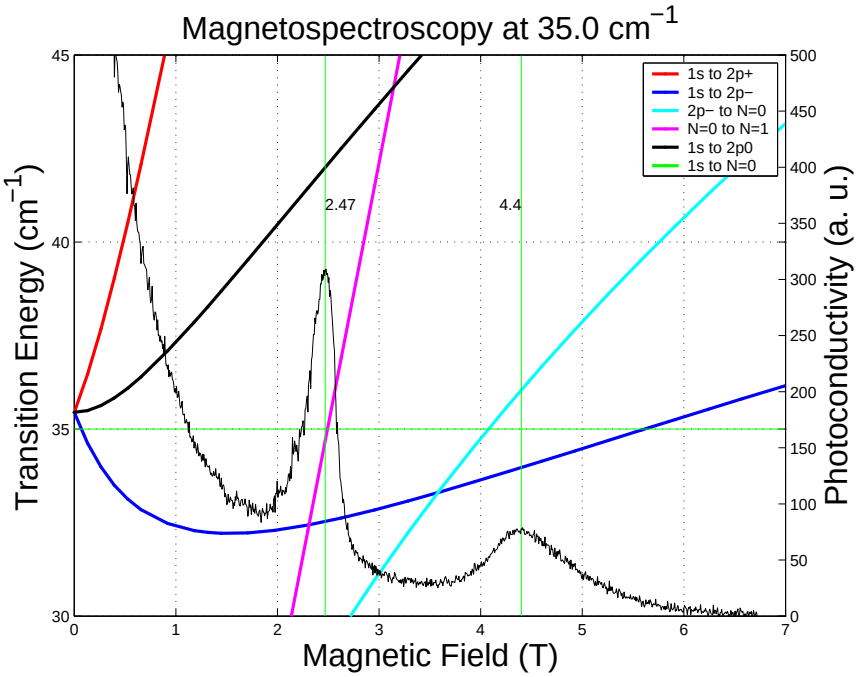

Figure 6.7: Magnetospectroscopy with illumination at 35.0 cm^{-1}.

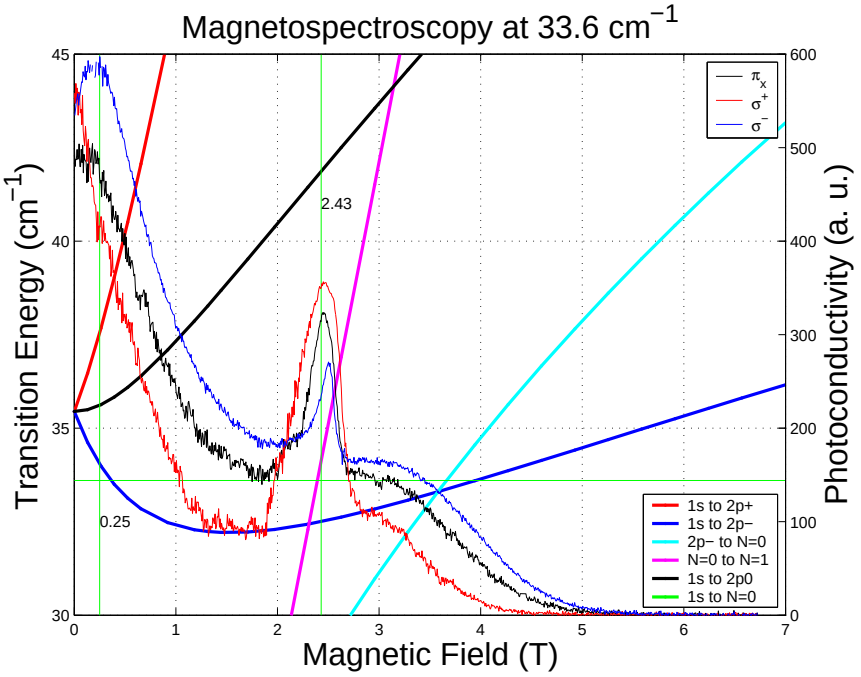

Figure 6.8: Magnetospectroscopy with illumination at 33.6 cm^{-1}.

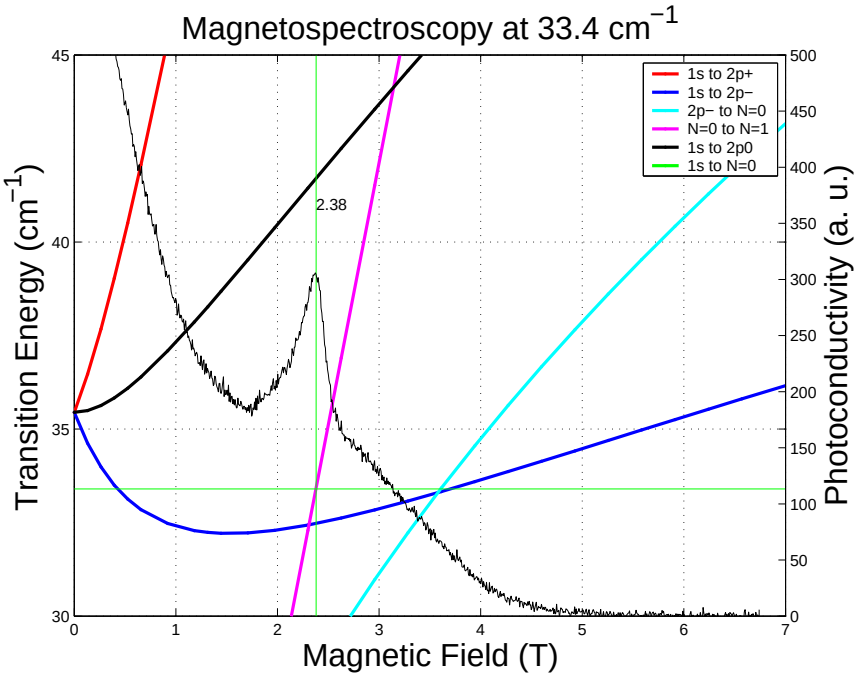

Figure 6.9: Magnetospectroscopy with illumination at 33.4 cm^{-1}.

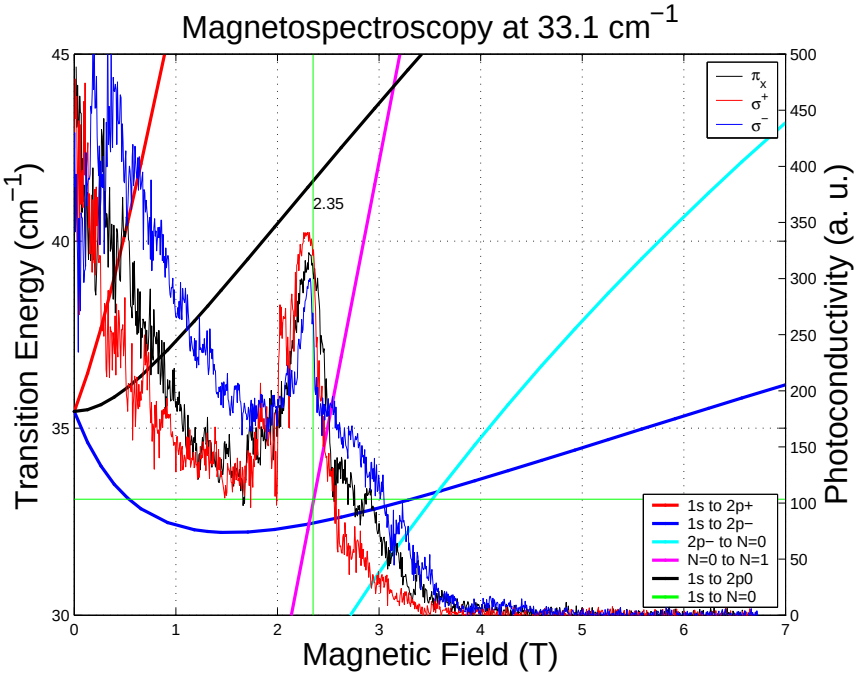

Figure 6.10: Magnetospectroscopy with illumination at 33.1 cm^{-1}.

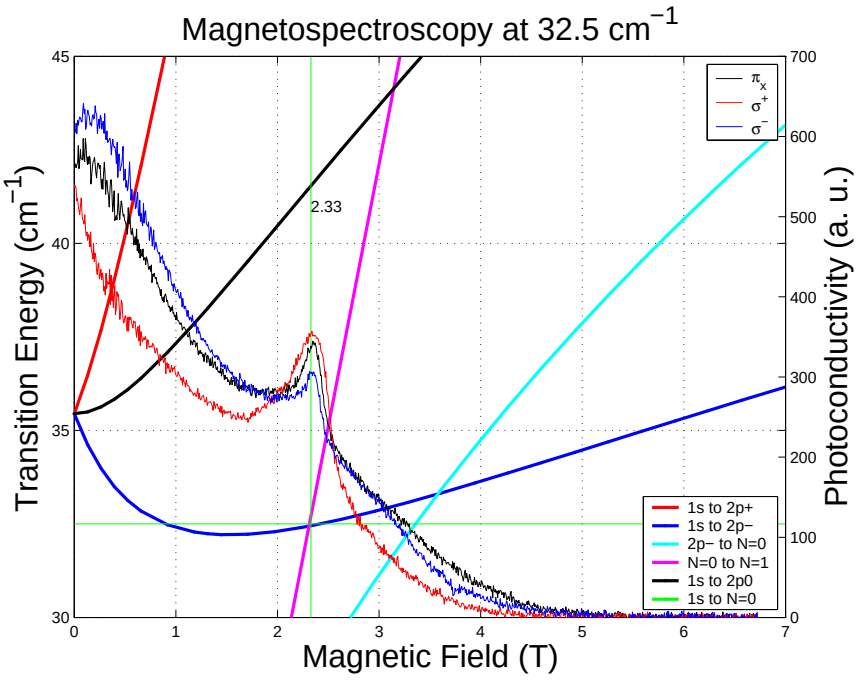

Figure 6.11: Magnetospectroscopy with illumination at 32.5 cm^{-1}.

is particularly noisy because of an atmospheric-water-absorbtion line at this frequency. The cyclotron transition appears at 2.35 T. No clear peak for the 2p$^-$ state is visible. However, it is clear that σ^- light generates a larger photoconductivity signal for magnetic fields between 0.5 and 3.5 T, where the energy of the 2p$^-$ transition is relatively flat and close to the illumination frequency. The exception is the cyclotron transition, where σ^+ light generates a larger signal, as expected. Finally, Figure 6.11 shows data at 32.5 cm^{-1} (0.98 THz). Again, σ^- light generates the larger photosignal, except on resonance with the cyclotron transition.

6.2.3 Discussion of spectroscopy data

The results of magnetospectroscopy of the $2p^+$ state are not surprising. We observe the state transition to occur very close to the magnetic field value calculated from the Makado and McGill state energies. The cyclotron transition also appears at the expected magnetic field. For the $2p^-$, however, the results are more confusing. In every case where a clear $2p^-$ resonance peak can be identified, the peak occurs at a significantly lower magnetic field than calculated. The difference between the observed and calculated peak position is much larger than the margin of error in magnetic field calibration that can be estimated from the $2p^+$ results. However, the slope of the $2p^-$ transition energy is more shallow than the $2p^+$. This means that a small offset in energy (of either the illumination frequency or the actual state transition) would correspond to a large change in magnetic field for the resonance.

While investigating possible sources of the offset in the $2p^-$ transition energies, I contacted both Tjeerd Klaassen and Rachel Heron, who generously provided their data on spectroscopy of the $2p^-$ state. The data of Klaassen was taken for silicon donors in GaAs, and matches the calculated curves with the silicon central cell correction quite well. Heron has data taken with both sulfur and silicon donors in GaAs. I extracted numerical values from her spectroscopy of sulfur donors using the program DataThief. In Figure 6.12 I show her data along with the theoretical transition energies. Heron's data is offset from the theoretical

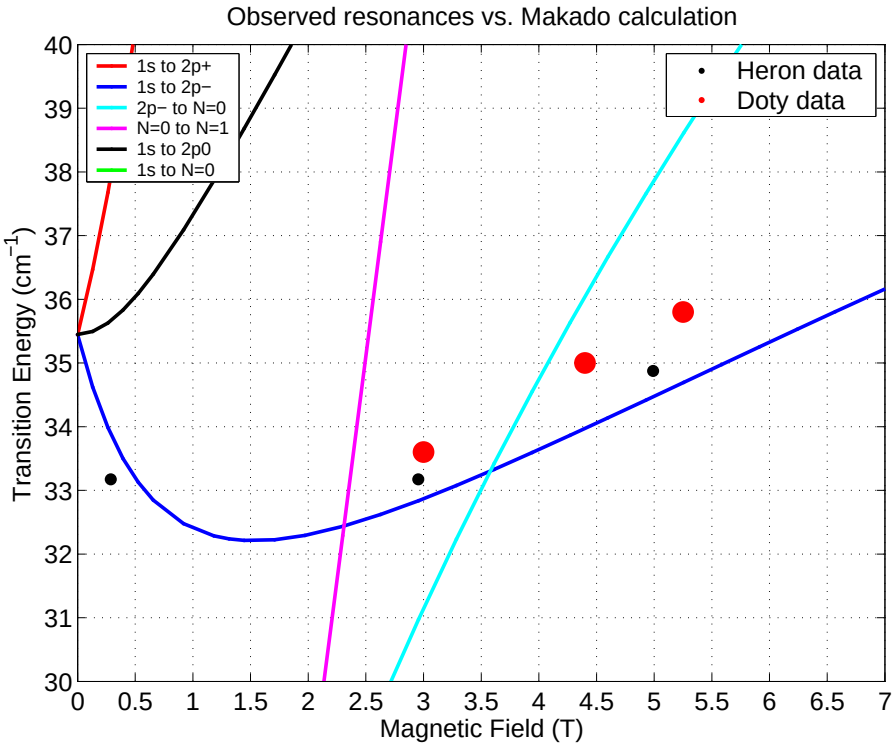

Figure 6.12: Data of observed resonances for the 1s to $2p^-$ transition by Heron (black) and me (red). The size of the markers for my data is chosen to represent the error bars. The solid curves show the theoretical values of the transition energy.

curve by approximately 0.5 T. In Figure 6.12 I also plot the transition resonances from my spectroscopy at 33.6, 35.0 and 35.8 cm^{-1}. I use only three points because it is difficult to extract a resonant field value from spectroscopy in the region of the cyclotron transition. The size of the markers is chosen to represent the error bars. The resonance values that I observe are offset from the theoretical curve by about 1 T.

In Figure 6.13 I plot my observed values of the 1s to $2p^-$ transition (red dots, connected by dashed blue line). I calculate and plot the expected energy of the $2p^-$ to $N = 0$ transition (dashed cyan line) by subtracting the observed 1s to $2p^-$

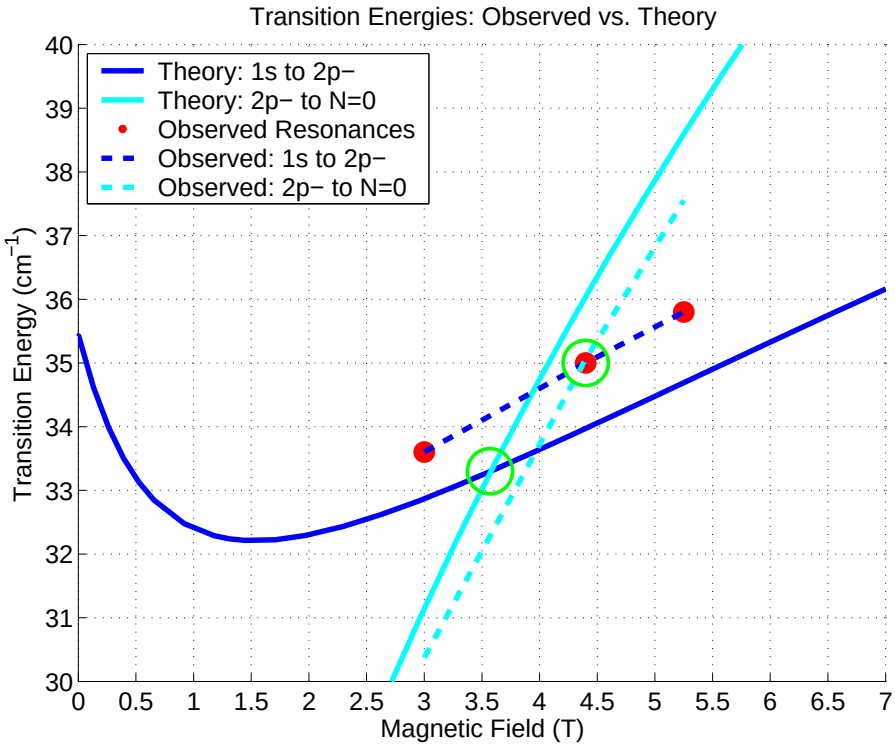

Figure 6.13: A comparison of the theoretical and observed transition energies of the 1s to $2p^-$ and $2p^-$ to N = 0 transitions. The green circles are added to emphasize the crossing points.

energy from the calculated 1s to N = 0 energy. I plot the theoretical transition energies from the Makado and McGill calculation with the solid lines. According to the theory, the 1s to $2p^-$ transition crosses the $2p^-$ to N = 0 transition at about 3.5 T. However, the crossing point of the dashed lines, which are taken from the observed resonances, is at about 4.4 T. The crossing point of these two curves indicates the position at which we would expect to observe any effects from two-photon absorption. The spectroscopy data indicates that the crossing is near 4.4 T. I will discuss the two-photon absorption further after presenting additional data in Section 6.3.

6.3 Selection Rules

The selection rules for orbital state transitions of the hydrogen atom are well known, and I refer to them in the previous section when, for example, I say that the 1s to $2p^+$ transition requires a σ^+ polarized photon. Since the electron effective mass in GaAs is isotropic, the selection rules for the bare hydrogen atom might be expected to hold. However, symmetry in a semiconductor can easily be lost. In quantum dots, for example, elongation of the dot along specific crystal planes can break the dot symmetry. The broken symmetry leads to a splitting in energy between polarization-specific excitons and relaxation between the polarization eigenstates. [25] In this section I will describe experiments to verify that the hydrogen atom selection rules hold, which is important for two reasons. First, we would like to take advantage of the selection rules to drive selective manipulations of quantum states. Second, the scheme for THz cavity quantum information processing that I discussed in Chapter 1 [10] imagines coupling THz photons in a cavity to a state transition. Verification of the selection rules for the transition will be important to designing the polarization of the electric field in the cavity.

To verify the selection rules, I illuminate the sample with a fixed frequency of THz radiation and tune the state transition into resonance using the magnetic field. I then vary the spacing between the mirror and polarizer in the variable polarizer described in Section 3.2. This alters the polarization of the radiation, and I monitor the photoconductivity as a function of the mirror-polarizer spacing.

I then measure the Stokes parameters at various fixed values of the spacing to determine the polarization state of the light. I will denote the polarization state of the light in the data by the ratio S_3/S_0, which is +1 for σ^+ light and -1 for σ^-.

If the selection rules are obeyed, I would expect to see a sine curve with a peak corresponding to the circular-polarization handedness that couples into the transition and a valley for the opposite handedness. To verify that the changes in photoconductivity as a function of polarization are due only to the selection rules and not to some other effect, I perform the same experiment twice with anti-parallel directions of the magnetic field (+z and -z). As discussed in Section 6.2, reversing the direction of the magnetic field reverses the handedness of light that couples into a transition. I expect that for the anti-parallel magnetic field (-z) we will see a sine curve with a π phase shift relative to the +z data.

As discussed in Section 6.1, data taken with long pulses always shows a saturation-like behavior as a function of intensity. I took all data on selection rules with sufficient attenuation to provide a clear signal. Some of the curves still show an asymmetric shape that arises from the saturation effect, which also reduces the apparent contrast between the high and low values of the photoconductivity.

As an initial verification and calibration, I checked the cyclotron transition (N = 0 to N = 1 Landau level), which requires σ^+-polarized light for momentum conservation. The results are shown in Figure 6.14. The red (+z) and blue (-z) symbols show the data for the two anti-parallel directions of the magnetic field. As

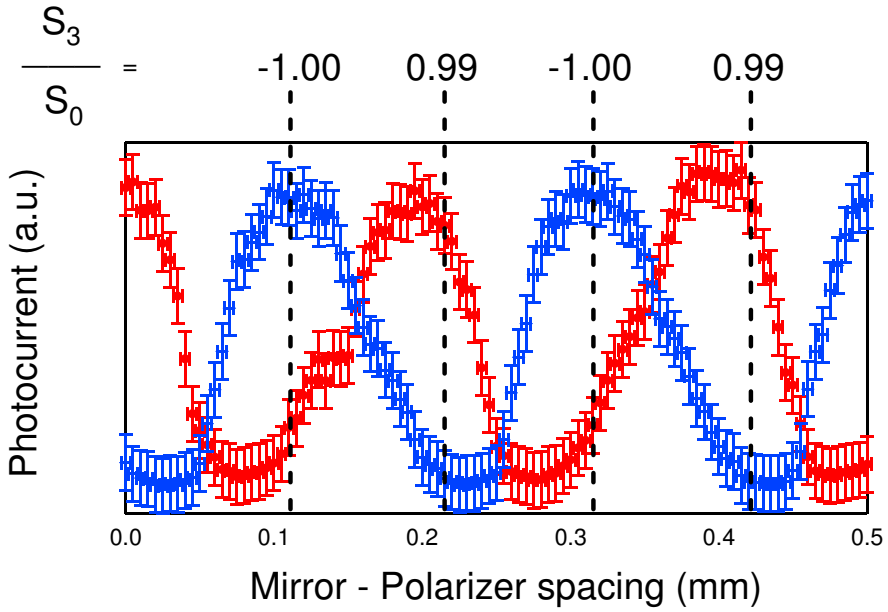

Figure 6.14: Photoconductivity versus mirror-polarizer spacing under illumination at 35.0 cm^{-1} at a magnetic field of 2.5 T, resonant with the cyclotron transition. The blue and red symbols correspond to anti-parallel directions of the magnetic field. The cut lines indicate spacings at which the Stokes parameters were measured, with the value of these parameters given in the form S_3/S_0.

expected, the largest photocurrent is observed for σ^+ radiation when the magnetic field is oriented in the +z direction (red). For the -z direction of the magnetic field (blue) the peaks are shifted by approximately π. The dashed lines indicate the mirror-polarizer spacing at which the Stokes parameters were measured, with the measured value of S_3/S_0 placed above each line.

In Figure 6.15 I show the results of this experiment performed at 84.5 cm^{-1} and 3.5 T. It is clear that the peaks of the red data (+z) line up with σ^+ polarized light and the valleys with σ^-. This agrees with the expected selection rules. When the magnetic field is reversed (blue data), we see that the signal has a phase shift of about π. The black lines through the data are fits using a sine curve. The

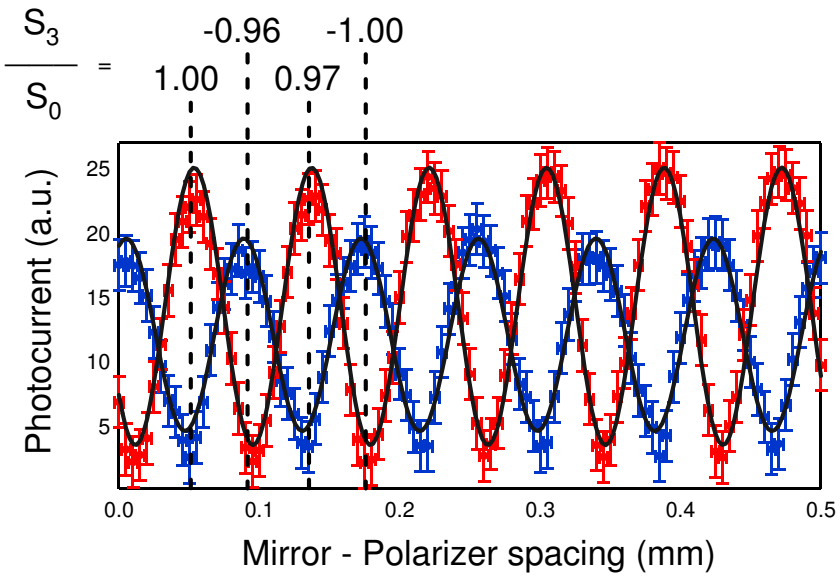

Figure 6.15: Photoconductivity versus mirror-polarizer spacing under illumination at 84.5 cm^{-1} at a magnetic field of 3.5 T, resonant with the 1s to $2p^+$ transition.

phase difference between the sine curves fit to the anti-parallel directions is 3.63 radians, a 15.5% difference from π.

Figure 6.16 shows data for the experiment at 52.3 cm^{-1} and 1.35 T. As before, the peaks of the red data line up with σ^+ polarized light and the valleys with σ^-, with the reverse for the blue data. The phase difference between the fit curves is 3.13 radians, a 0.5% difference from π.

In Figure 6.17 I show data on the selection rules studied at 35.8 cm^{-1} and 5.25 T, on resonance with the 1s to $2p^-$ transition. The peaks for the red trace (+z) now line up with σ^--polarized light and the blue data (-z) is again reversed. This is consistent with the expected selection rules for the 1s to $2p^-$ transition. The phase difference between the curves is 3.57 radians, a 13.6% difference from

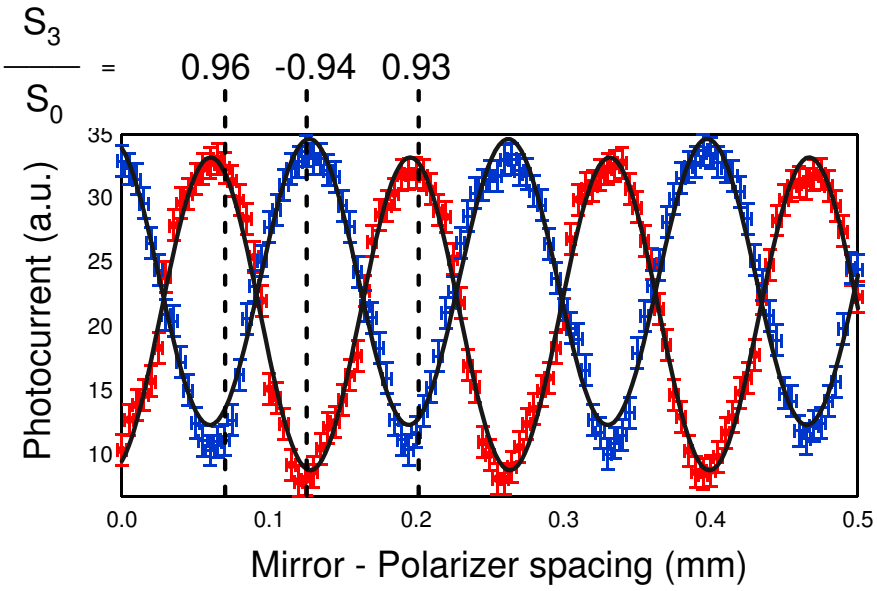

Figure 6.16: Photoconductivity versus mirror-polarizer spacing under illumination at 52.3 cm^{-1} at a magnetic field of 1.35 T, resonant with the 1s to $2p^+$ transition.

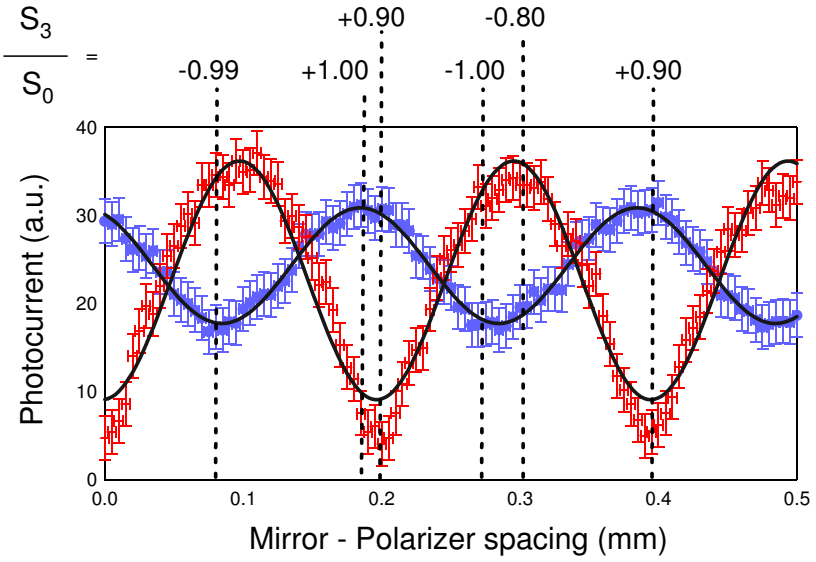

Figure 6.17: Photoconductivity versus mirror-polarizer spacing under illumination at 35.8 cm^{-1} at a magnetic field of 5.25 T, resonant with the 1s to $2p^-$ transition.

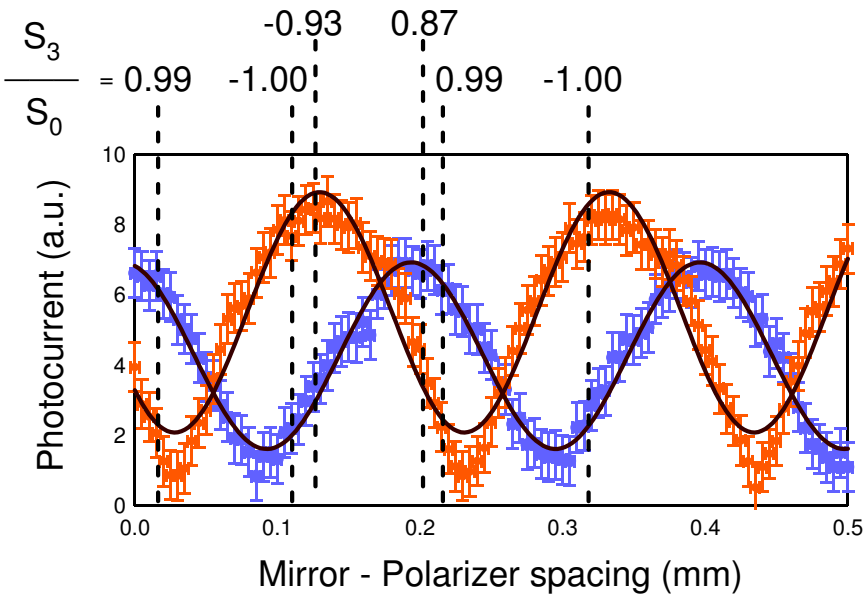

Figure 6.18: Photoconductivity versus mirror-polarizer spacing under illumination at 35.0 cm^{-1} at a magnetic field of 4.47 T, resonant with the 1s to $2p^-$ transition.

π. There is a significant amplitude difference between the two directions of the field, which appears at several different frequencies exciting the $2p^-$ transition. The +z direction data appears to be consistently larger, but I do not understand the origin of the difference.

With an illumination frequency of 35.0 cm^{-1}, the $2p^-$ transition comes into resonance at 4.47 T. The data is presented in Figure 6.18. The phase difference between the +z and -z data is surprisingly large: 4.28 radians, a 36.3% difference from π. The peaks of the photoconductivity signal do not seem to correspond to σ^+ and σ^- polarized light, but rather seem to be shifted toward the mirror-polarizer spacing that corresponds to horizontally-polarized light. This shift suggests some distortion of the selection rules, and I will come back to it

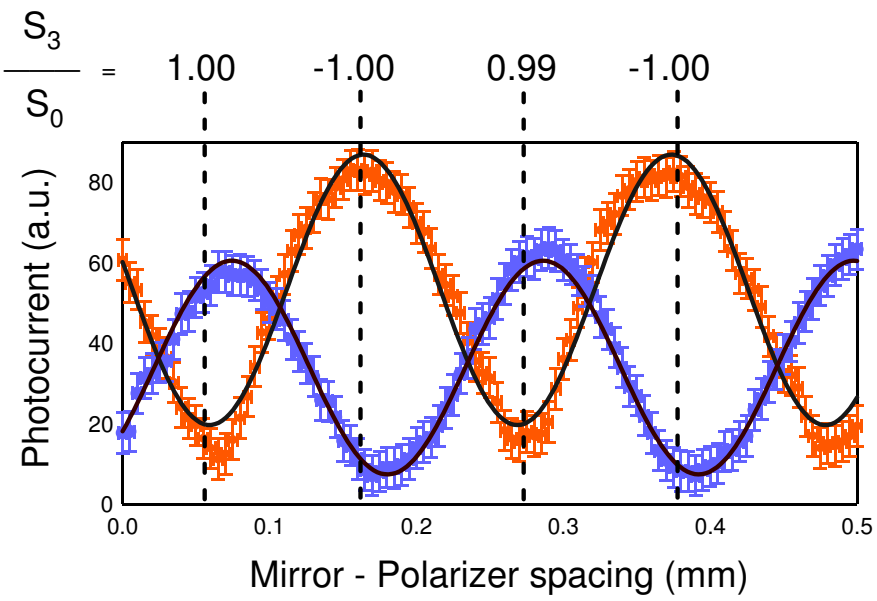

Figure 6.19: Photoconductivity versus mirror-polarizer spacing under illumination at 33.6 cm^{-1} at a magnetic field of 3.50 T, resonant with the 1s to $2p^-$ transition.

later.

At an illumination frequency of 33.6 cm^{-1}, the $2p^-$ transition appears at 3.50 T. Figure 6.19 shows that the distortion of the selection rules that was evident at 35.0 cm^{-1} is gone. The phase difference between the +z and -z data is 3.57 radians, a 13.8% difference from π.

Finally, data taken at an illumination frequency of 32.5 cm^{-1} and field of 1.30 T appears in Figure 6.20. The phase difference is 2.45 radians, a 22.1% difference from π. At this magnetic field, it appears that the -z data lines up well with the points of maximal circular polarization, while the +z data is shifted.

The sine curves fit to the selection-rules data have a period that is directly related to the wavelength of the incident light. The fit value of the period can

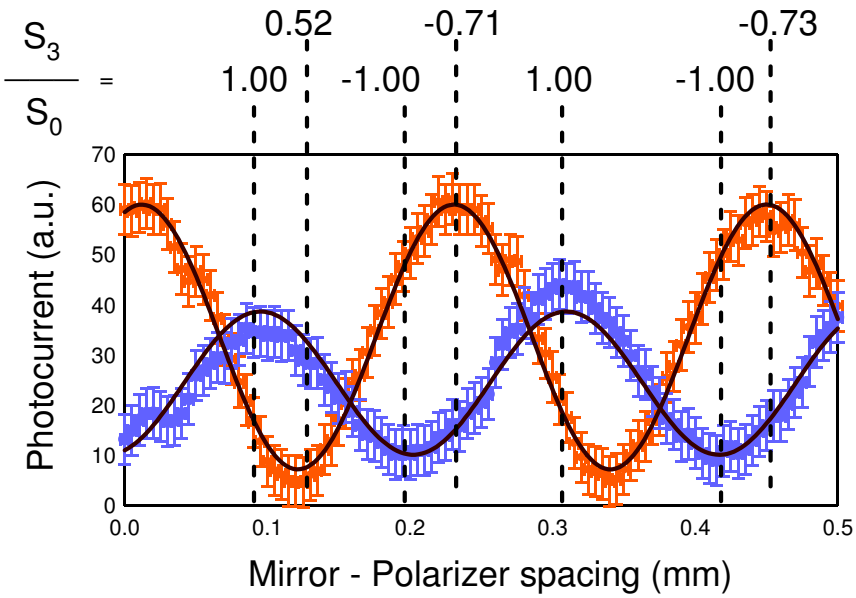

Figure 6.20: Photoconductivity versus mirror-polarizer spacing under illumination at 32.5 cm^{-1} at a magnetic field of 1.30 T, resonant with the 1s to $2p^-$ transition.

therefore be used as an additional verification of the frequency of the incident light. For example, at 84.5 cm^{-1} the period of the sine curves is 83.7 μm. This should correspond to a one-wavelength path-length difference between the beams reflected by the polarizer and the mirror. When the 45^o angle of incidence is taken into account, a spacing of 83.7 μm corresponds to a path-length difference of 118.37 μm. This is negligibly different from the actual wavelength of 118.34 μm. In Table 6.1 I show the value of the spacing calculated from the period of the fitted sine curves, the corresponding change in path length, the actual wavelength and the percent discrepancy.

For each frequency, I can extract a phase difference between the curves for the two anti-parallel directions of the field, which are summarized in Table 6.2. The

Wavenumber	Spacing	Δ path length	Wavelength	% discrepancy
84.5	83.70	118.37	118.34	0.03
52.3	135.13	191.11	191.20	0.05
35.8	198.08	280.12	279.33	0.28
35.0	202.55	286.44	285.71	0.26
33.6	209.52	296.31	297.62	0.44
32.5	219.05	309.78	307.69	0.68

Table 6.1: A comparison of fit values and actual wavelengths. Lengths are given in μm and wavenumber in cm^{-1}.

Wavenumber	Resonant Field (T)	Δ Phase (rad)	% discrepancy
84.5	3.50	3.63	15.5
52.3	1.35	3.13	0.50
35.8	5.25	3.57	13.6
35.0	4.47	4.28	36.3
33.6	3.50	3.57	13.8
32.5	1.30	2.95	22.1

Table 6.2: A summary of the data from the selection-rules experiments at several different wavenumbers.

percent discrepancy from a π phase shift is 15% or lower, with two exceptions.

At 32.5 cm^{-1} the phase discrepancy is 22.1%. One possibility is that this shift arises from the broad background photoconductivity signal seen at low magnetic fields and discussed in Section 6.2. Looking back to Figure 6.11, we also see that the observed 1s to $2p^-$ transition is quite close to the cyclotron transition. Near-resonant excitation of the cyclotron transition is probably contributing to the photoconductivity signal. Since the cyclotron transition is excited by the opposite handedness of polarization than the 1s to $2p^-$ transition, this is probably the source of the larger-than-usual phase discrepancy.

At 35.0 cm^{-1} the phase discrepancy is 36.3%. The cyclotron transition is far

from resonance and can not explain this shift. A possible source for the distortion would be a parametric two-photon absorption from the 1s to $2p^-$ and then $2p^-$ to N = 0 levels. By parametric I mean that the transition from the $2p^-$ to N = 0 level (which requires a σ^+-polarized photon) is not independent of the excitation from the 1s to $2p^-$ state. Rather, a σ^- photon is first required to excite electrons to the $2p^-$ state. From the $2p^-$ state, electrons can arrive in the conduction band either by phonon ionization or by subsequent absorbtion of a σ^+ photon. This parametric two-photon absorption would be expected to contribute most strongly at the point where the energies of the 1s to $2p^-$ and $2p^-$ to N = 0 transitions are degenerate.

According to the calculations of the transition energies, the $2p^-$ to N = 0 transition should cross the 1s to $2p^-$ transition at a field of about 3.6 T with a frequency near 33.6 cm^{-1}. However, as shown in Figure 6.13, the spectroscopy data presented above indicate that the transitions actually cross at about 4.4 T. The data on the selection rules support the conclusion that the the crossing is at 4.4 T: no unusual phase shift is observed for illumination at 33.6 cm^{-1} and 3.6 T, while a large phase shift is seen at 35.0 cm^{-1} and 4.4 T. The combination of the spectroscopy and selection rules data strongly suggests that a two-photon parametric absorption is contributing to the photoconductivity signal near 4.4 T. The shift of the photoconductivity peaks toward horizontally-polarized light in Figure 6.18 is probably due to this parametric two-photon absorption.

At all frequencies, there is a non-zero photocurrent for the incorrect handedness

of polarization. However, the significance of this non-zero signal is exaggerated by the saturation behavior described in Section 6.1. Additionally, the measurement of the Stokes parameters is done outside the cryostat. To illuminate the sample, the radiation is focused by an off-axis parabolic mirror (which introduces some astigmatism) and passes through two windows (TPX and z-cut quartz). The sample contacts are designed to minimize distortions of the incident polarization, but may not do so perfectly. The combination of these elements probably introduces some distortion of the polarization that is not reflected in the Stokes parameters. This distortion is the probable source of the non-zero photoconductivity observed for the incorrect handedness of circular polarization.

The results of these experiments verify that the selection rules for the hydrogen atom hold for hydrogenic electrons in GaAs: σ^+-polarized light couples most strongly into the $2p^+$ hydrogenic state transition for a +z-oriented magnetic field, while σ^--polarized light couples most strongly into the $2p^-$ transition. The polarization that excites a transition is switched for an anti-parallel direction of the magnetic field, verifying that the effect is due only to the selection rules. However, the photoconductive response to polarized radiation can be distorted by the proximity of the cyclotron transition or a parametric two-photon absorption.

6.4 Rabi oscillations at 84 cm^{-1}

I want to drive Rabi oscillations between electron orbital states to demonstrate the feasibility of coherent manipulations with applications in quantum information processing and to study the dynamics of these manipulations. Driving Rabi oscillations requires THz pulses that are short with respect to the dephasing time of the ensemble of electrons. I discussed the techniques for generating these short pulses in Section 3.3, so here I will only present an overview of the actual experiment.

I tune the magnetic field to bring the 1s to $2p^+$ transition into resonance with radiation at 84.5 cm^{-1} (2.54 THz). No attenuation is used because I want strong coupling in order to reach the nonlinear regime in which Rabi oscillations occur. The strong field shifts the energy of the transition slightly (Stark shifting) and so the resonant field must be adjusted slightly from the values observed in the spectroscopy data presented above. Because of the saturation effects described in Section 6.1, finding this field value with a long pulse and no attenuation is difficult. I find it most effective to identify the resonant magnetic field using a pulse of about 10 ps, which is short enough to avoid the saturation problems.

The measurement of the magnetic field resonance is complicated by the pre-activation reflectivity of the reflection switches. I call the sample's response to the pre-activation reflectivity the "pre-pulse." The magnitude of the pre-pulse depends both on the amplitude of the reflection from the wafers and on the detuning of

the magnetic field from the resonance for the transition. When the magnetic field is scanned, the pre-pulse will get larger at the resonant field. This reduces the number of electrons still in the ground state at the time of arrival of the sliced pulse, and therefore the amplitude of the response to the 10-ps pulse can actually be smaller on resonance than off resonance. With careful optimization of the contrast ratio, this negative response can usually be eliminated.

Using the pulse slicer, I generate FIR pulses with a controllable temporal width. For a relatively long pulse (for example, 100 ps) the electrons undergo Rabi oscillations between the 1s and $2p^+$ states until they dephase and are ultimately all ionized into the conduction band. A shorter pulse width cuts off the driving field at a mid-point of the cycle and I can then measure the population of the states by the photoconductivity signal, described in Section 5.4. I plot the photoconductivity as a function of the pulse width to get a map of the electron-state population as a function of time. This is often called a "stroboscopic" measurement.

Bryan Cole did the first generation of experiments on Rabi oscillations at 84 cm^{-1} and his results are given in his paper [26]. Figure 6.21 shows a similar set of results from my experiments. The traces are offset vertically for clarity. The top trace is done with no attenuation, and each successively lower trace is taken with an additional attenuator that decreases the THz intensity by a factor of about 4. The important feature is the decrease in photoconductivity that appears in the top three traces at 20-40 ps. The decrease in excited-state population as the total pulse energy increases is the signature of a Rabi oscillation. As the

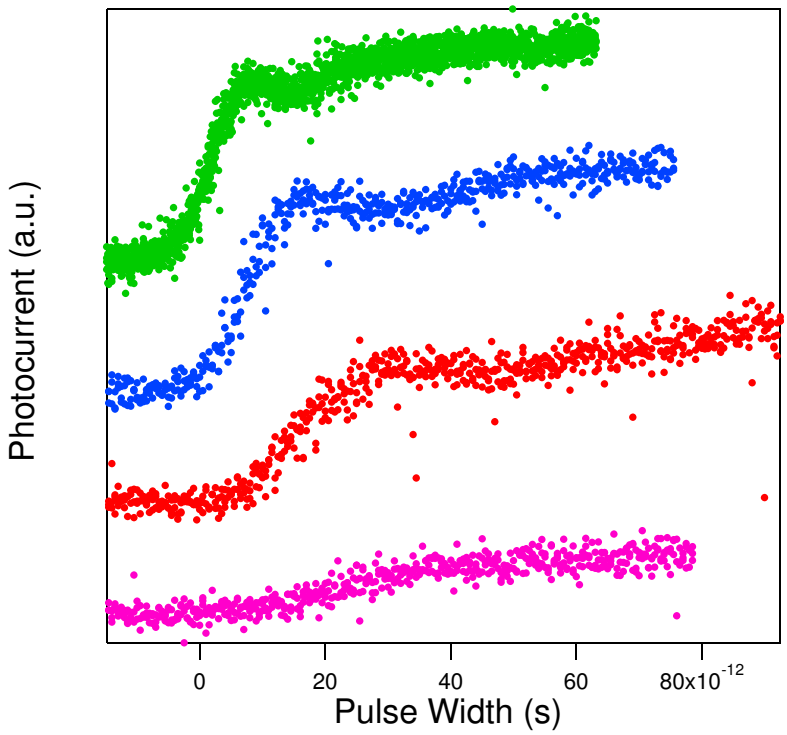

Figure 6.21: Photoconductivity versus pulse width for illumination at 84.5 cm^{-1} and magnetic field of 3.62 T. The decrease in photoconductivity (at about 15 ps for the green trace) is the signature of Rabi oscillations, demonstrating coherent manipulation of the states. The lower traces are each taken with attenuation that lowers the intensity by a factor of about 4 for each successive trace.

intensity decreases, the period of the Rabi oscillation (the time it takes to get to the dip) increases. This is consistent with the theory of Rabi oscillations, where the frequency of the oscillation is known to depend on intensity. For the lowest curve in Figure 6.21, no oscillation is visible at all.

I extended the Rabi oscillation experiment by using the variable polarizer to set the polarization of the incident light. Experimental results are shown in Figure 6.22. For σ^+-polarized light, I observe a Rabi oscillation with a period of about 15 ps. For horizontally-polarized (π_x) and vertically-polarized (π_y) light, a Rabi

oscillation is still visible, but the period of the oscillation is 20-25 ps. Since the 1s to $2p^+$ transition requires σ^+-polarized light, only half of the linearly-polarized light is actually coupling in to the transition. This causes a decrease in the effective intensity and an increase in oscillation period similar to that shown in Figure 6.21. I suspect that small errors in the mirror-polarizer spacing caused the difference in oscillation period between the scans taken with horizontally- and vertically-polarized light. Theoretically there should be no difference between the two. When I illuminate the sample with σ^--polarized light, no Rabi oscillation is evident. There is a shallow increase in photoconductivity, which I attribute to a non-zero σ^+ component of the polarization arising from the coupling in to the cryostat, as described in Section 6.3.

The absence of a Rabi oscillation when the hydrogenic electrons are illuminated with high-intensity but incorrectly-polarized radiation demonstrates a selective state manipulation. Selective state manipulations have been used to perform complicated manipulations of exciton states [27] to demonstrate quantum logic operations. The demonstration of selective manipulation of hydrogenic states is a first step toward implementing more complicated schemes using these donor-bound electrons.

A great deal of information about the dynamics of the electrons can be extracted from fitting a theoretical model to the data. The modelling, fitting and resulting information about the dynamics are all described in Chapter 7. Looking at the raw data, one can see that the minima of the dip in photoconductivity dur-

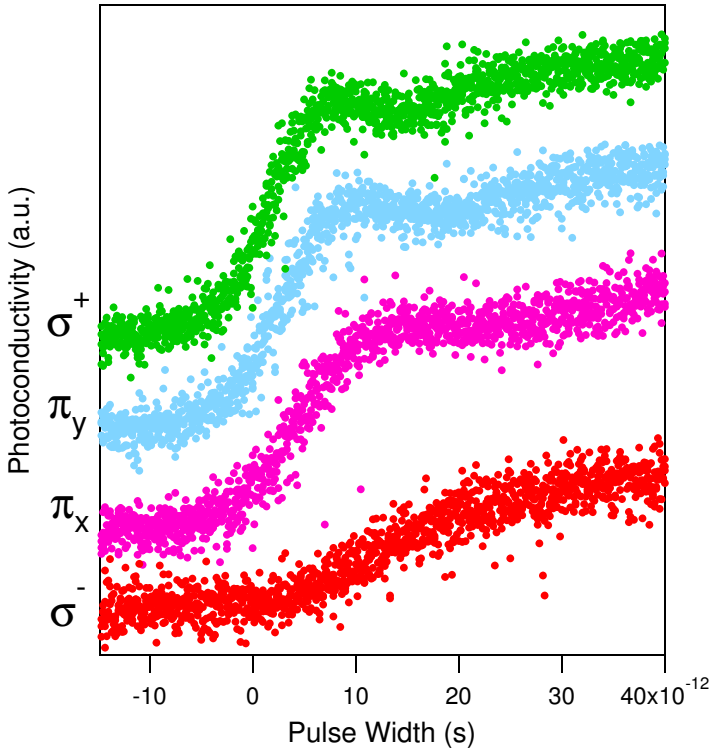

Figure 6.22: Photoconductivity versus pulse width using different polarizations. The illumination frequency is 84.5 cm^{-1} and the magnetic field is 3.62 T.

ing the Rabi oscillation is not particularly low; the electron population does not fully return to the ground state. This happens because the excited state is located above the conduction band and the electrons ionize rapidly into the conduction band, leaving the driven two-state system. The modelling and fitting in Chapter 7 will provide a quantitative analysis of the ionization and dephasing rates that cause this loss.

To avoid this loss to the conduction band, I would like to drive oscillations between orbital states where the excited state is more isolated. This would allow me to observe a deeper dip in the excited-state population and also to observe more cycles of Rabi oscillations before ionization or dephasing destroys the coherent signal. From the point of view of quantum information applications, an isolated excited state is critical to storage of information in the quantum bit. My attempts to observe oscillations with a more isolated excited state are described in Section 6.5.

6.5 Rabi oscillations at other frequencies

To find more isolated excited states, I looked back at the energy level diagram in Figure 2.2 to find ranges of magnetic field where either the $2p^+$ or $2p^-$ state is below the conduction band. I chose to look for Rabi oscillations between the 1s and $2p^+$ states at a magnetic field of about 1.3 T, where the excited state is just below the conduction band. The illumination frequency for this transition is 52.3

cm^{-1}. For the 1s to $2p^-$ transition, I chose to look at a magnetic field of 4.7 T with illumination at 35.0 cm^{-1}.

One criterion for choosing 35.0 cm^{-1} was that the energy required to excite from the $2p^-$ state to the N = 0 level be larger than the energy of the 1s to $2p^-$ transition to avoid two-photon excitation to the conduction band. I chose this frequency from the theoretical energy levels before performing the spectroscopy and selection rules experiments described above. The data from those experiments indicate that the two-photon absorption is at a higher magnetic field and frequency than expected from the theoretical energy levels. In retrospect, it would have been better to choose a higher frequency and magnetic field for the Rabi oscillation experiment so that the excited state was farther below the conduction band and two-photon effects were avoided. However, the modelling presented in Chapter 7 indicates that the proximity to the two-photon absorption was probably not the limiting factor.

The experimental procedure to drive Rabi oscillations at 52.3 or 35.0 cm^{-1} is exactly the same as described in Section 6.4: the photoconductivity is plotted as a function of pulse width under illumination at a fixed frequency and with the magnetic field set to the resonant value for the transition. However, I have never seen a shifting of the resonant magnetic field (as observed at 84 cm^{-1}) that would indicate that I am in the strong-coupling regime. The total power output of the FEL drops with decreasing wavenumber, and is approximately a factor of ten smaller at 35 cm^{-1} than at 84 cm^{-1}.

As shown in Figure 6.21, the rise time of the photoconductivity signal depends on the intensity of the light being coupled in to the sample. If the time it takes for the photoconductivity to rise to saturation is longer than the dephasing time of the ensemble of electrons, no coherent phenomena can be observed. Thus, in the absence of a Rabi oscillation, the primary feedback I have for optimizing the system is the pulse width required to saturate the transition. I attempt to minimize the rise-to-saturation time by maximizing the intensity of the light at the sample without generating an unacceptably large pre-pulse.

In Figure 6.23a I show a plot of the best rise-to-saturation time I have observed at 52.3 cm^{-1}. The zero-delay point of the x axis is somewhat arbitrary. The signal rises from zero to saturation in about 45 ps, and there is no sign of a dip in photoconductivity indicating a Rabi oscillation. The best rise to saturation I have observed at 35.0 cm^{-1} is shown in Figure 6.23b. In this case, the rise is drastically slower and takes almost 3.5 ns to saturate. The maximum intensity of the FEL is lower at 35.0 cm^{-1} than 52.3 cm^{-1}, but the drastic difference in rise time may also be related to the ionization rate from the $2p^-$ to the conduction band. This will be discussed in more detail in Chapter 7.

To improve the rise-to-saturation time at these frequencies, I needed to increase the intensity of the light at the active area of the sample. I did this by mounting a sample of the same material, with the same finger-contact pattern, on the back side of a solid-immersion lens. The details of this technique are given in Section 5.2. This tighter focus should lead to an increase in intensity at the sample by

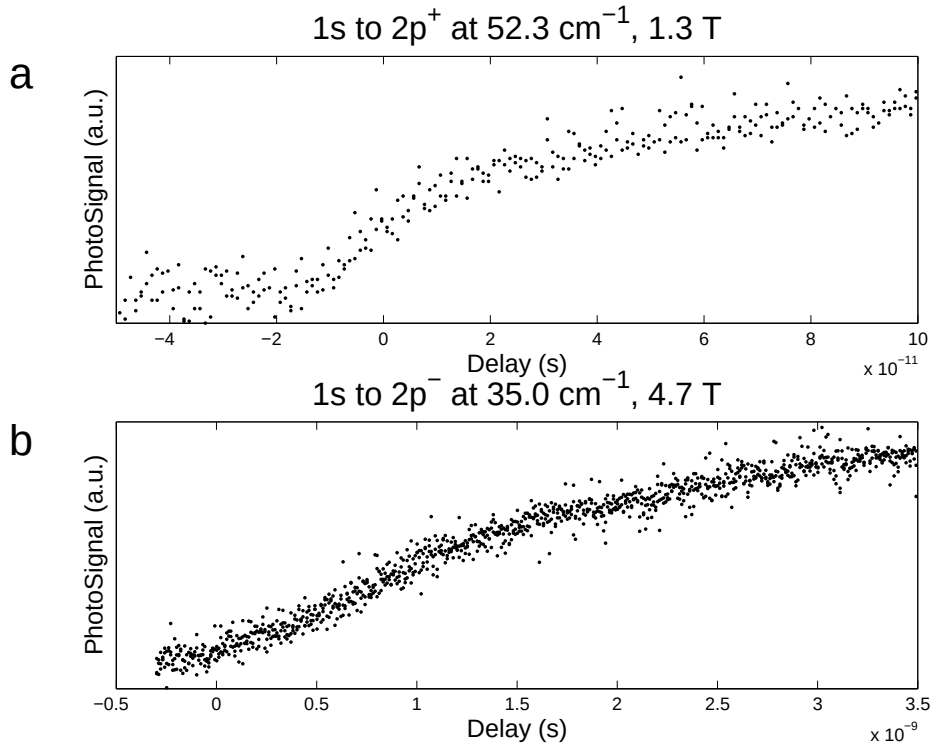

Figure 6.23: Photoconductivity versus pulse width for (a) illumination at 52.3 cm^{-1} and magnetic field of 1.3 T and (b) illumination at 35.0 cm^{-1} and magnetic field of 4.7 T. These show the fastest rise to saturation observed at these frequencies when using the regular sample.

a factor of about 10, and the antireflection coating should lead to an increase of 15-20%. The overall intensity should therefore increase by approximately a factor of 12.

In Figure 6.24 I show characteristic photoconductivity using the lens-mounted sample with sufficient attenuation to be slightly below the intensity that was incident on the non-lens-mounted sample. At 52.3 cm^{-1}, a sharp rise in the photoconductivity is seen from 0 to 0.5 ns, followed by a shallow rise from there on. The time to saturation, taken at the peak of the sharp rise, is approximately 500 ps, far slower than for the non-lens-mounted sample. For illumination at 35.0 cm^{-1} no rapid rise is evident. Rather, there appears to be only an exponential-like rise with pulse width.

If I had a totally isolated two-level system with perfect contrast ratio, the rise-to-saturation time shown in Figure 6.24a should decrease as I increase the intensity coupled in to the sample. In Figure 6.25 I show curves taken with the lens-mounted sample under illumination by 52.3 cm^{-1} radiation at a number of different intensities. The intensity is calibrated relative to the maximum intensity (I_0) by using calibrated attenuators. Since the attenuators reduce the intensity by a non-integer amount, the fractional intensities listed are not exact but give an approximate value of the relative intensity. I have added a constant horizontal offset to each curve to clarify the display. As the intensity gets larger (moving to the left of the figure), the sharp rise gets faster. Before the rise time approaches the rise-to-saturation time observed with the non-lens-mounted sample (45 ps),

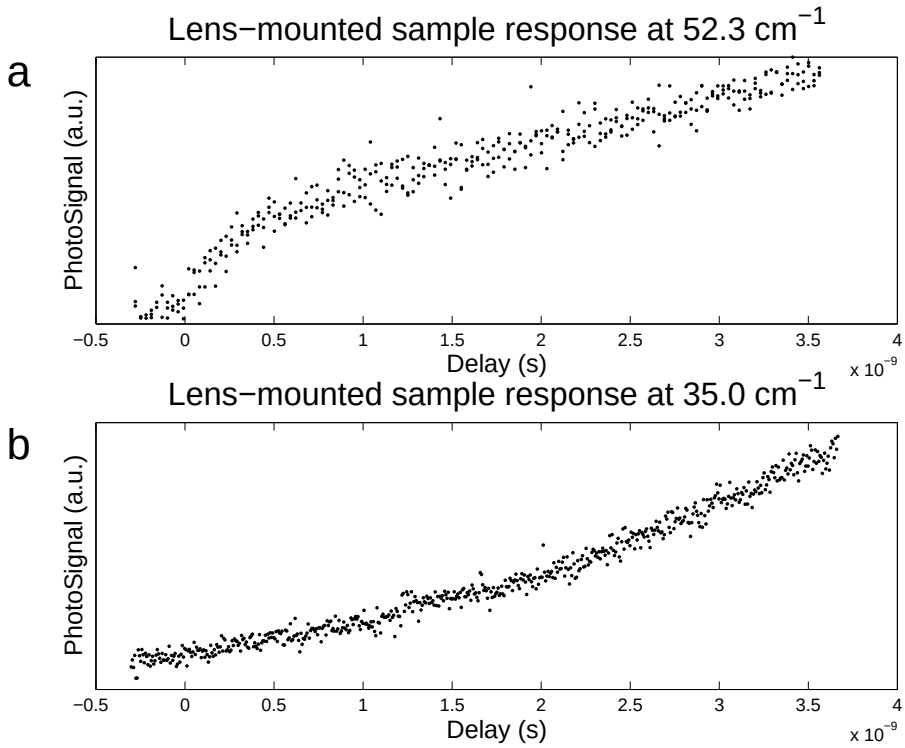

Figure 6.24: Photoconductivity versus pulse width for (a) illumination at 52.3 cm^{-1} and magnetic field of 1.3 T and (b) illumination at 35.0 cm^{-1} and magnetic field of 4.7 T. These traces are examples of the characteristic behavior using a lens-mounted sample. The base of the graphs corresponds to zero photoconductivity signal.

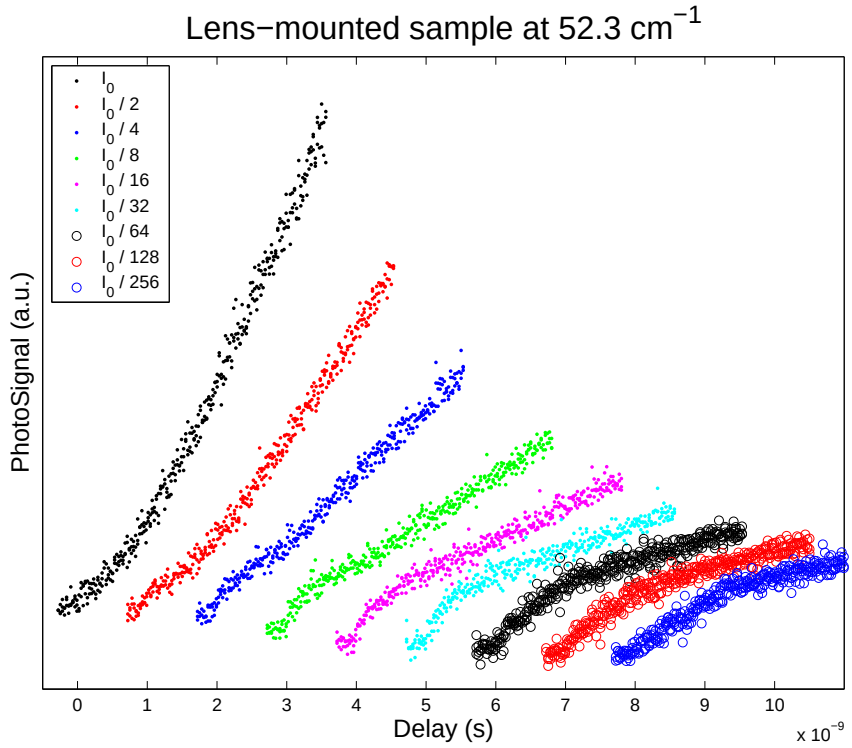

Figure 6.25: Photoconductivity versus pulse width for illumination at 52.3 cm^{-1} using several different intensities. A horizontal offset is added between traces.

the second, exponential-like, rise dominates. Also of note is the small vertical offset between the traces. This is not added to the data, but rather reflects a change in the background photoconductivity of the sample (before the rising edge of the short THz pulse).

In Figure 6.26 I show similar curves taken with the lens-mounted sample under illumination by 35.0 cm^{-1} radiation. Again, the fractional intensities listed are only approximate and I have added a horizontal offset to each curve. No sharp rise is apparent at any intensity and the exponential rise grows at higher intensities. The vertical offset between the traces is more pronounced at this frequency than at 52.3 cm^{-1}.

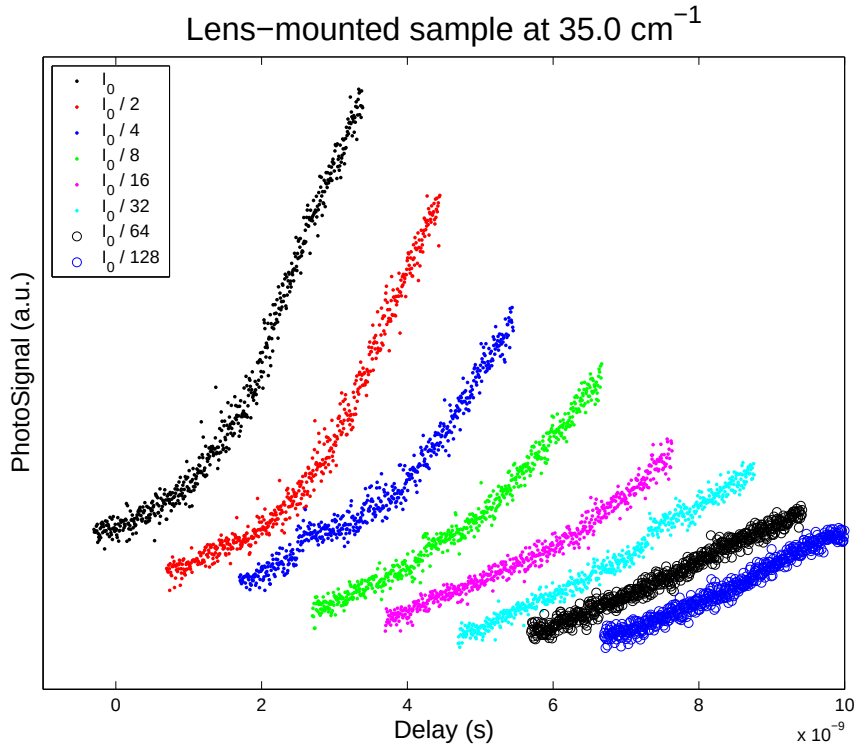

Figure 6.26: Photoconductivity versus pulse width for illumination at 35.0 cm^{-1} using several different intensities. A horizontal offset is added between traces.

The vertical offset appearing at higher intensities when using the lens-mounted sample is probably due to the increased efficiency with which the pre-pulse is coupled into the sample. When doing the experiments, I observed that it was extremely difficult to reduce the pre-pulse response of the sample to an adequate level. A large pre-pulse may preclude observation of coherent phenomena by dephasing the system before a time-resolved pulse arrives. This possibility is considered in the model of Chapter 7.

The exponential-like rise in photoconductivity at higher intensities is not a feature of the two-level Rabi oscillation system. It must arise from other physics that I have not considered in this work. For example, the intense THz field may be heating the sample or (more likely) the conducting electrons by promoting them to higher Landau levels. These effects would change the conductivity of the sample. The photoconductive detection mechanism does not allow us to distinguish these changes in conductivity from changes in the population of the conduction band due to excitation of the higher orbital states.

Chapter 7

Modelling and Fitting Rabi Oscillations

In Chapter 6 I presented data demonstrating Rabi oscillations between the 1s and $2p^+$ state of a donor bound electron in GaAs. To extract useful information from this data, I create a model of the system that contains relevant parameters like the dephasing rate, ionization rate and actual electric field. This model is then fit to the data and from the fit I can determine the values of these parameters for our system.

To model the system, I use the density matrix equations of motion, based on those described by Boyd [28]. In Section 7.1 I will describe the two-level open-system model. This was originally used by Jon Williams to fit Bryan Cole's data [26]. Using this model, I re-derived the fitting routine. I compared the fits made with my fitting routine to Jon's and found a maximum deviation from his values

of 2%. Although I don't know why, my derivation proved capable of fitting to a larger number of data sets and all fits and parameters described in this chapter are done using my routine. I will close this section by describing the limitations of the two-level open-system model.

In Section 7.2 I will describe a three-level model. This is a more appropriate model to use for the 1s to $2p^-$ transition. Although the model is too complex to use in a fitting routine, it does allow me to extract a great deal of information about the expected behavior of our system and the experimental limitations.

7.1 Two-level open-system model

7.1.1 Deriving the density-matrix equations of motion

In Figure 7.1 I show a diagram of a two-level open system. In the case of the hydrogenic electrons I want to model, state 1 is the ground state (1s) and state 2 is the excited state ($2p^+$). Throughout this section I will simply refer to them as states 1 and 2. The model contains a driving term (from the electric field of a laser), V, a lifetime for the excited state, $1/\Gamma_1$, and a dephasing rate, γ_2. The model is called an open system because it also contains a rate of ionization out of state 2, γ_3. The number of particles in the system is therefore not strictly conserved.

I will first derive the density matrix equations of motion for the two-level

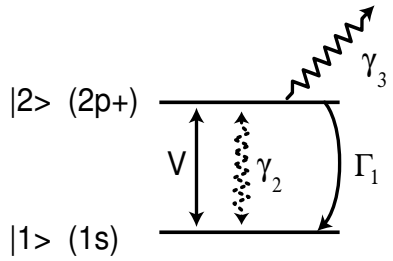

Figure 7.1: A diagram of the states for the two-level open-system density-matrix model. V is a driving term between states 1 and 2, Γ_1 is the inverse lifetime of the excited state, γ_2 the dephasing rate and γ_3 the ionization rate out of state 2.

system, neglecting the damping terms Γ_1, γ_2 and γ_3. Once I have derived the equations of motion for the two-level system in the presence of the driving term, V, I can then add in the damping terms.

I start with the Hamiltonian of the system,

$$\widehat{H} = \widehat{H}_o + \widehat{V}(t), \tag{7.14}$$

where $\widehat{H}_o$ is the Hamiltonian for the two-level system and $\widehat{V}(t)$ is the term for the interaction of the two-level system with the laser field. I know that the unperturbed two-level system Hamiltonian has energies given by

$$\begin{aligned} \widehat{H}_o|1\rangle &= \hbar\omega_1|1\rangle \\ \widehat{H}_o|2\rangle &= \hbar\omega_2|2\rangle. \end{aligned} \tag{7.15}$$

If I assume that the interaction with the laser field can be described in the electric-dipole approximation, I can write

$$V_{nm}(t) = eE(t)z_{nm}, \tag{7.16}$$

where e is the charge of the electron and z is the dipole matrix element. I know from dipole selection rules that $V_{11} = V_{22} = 0$ and that $V_{12} = V_{21}^*$.

I now describe the two-level system using the density matrix

$$\widehat{\rho} = \begin{bmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{bmatrix}, \tag{7.17}$$

where ρ_{11} and ρ_{22} give the populations of the 1 and 2 states, respectively, and I know that $\rho_{12} = \rho_{21}^*$.

Neglecting damping for the moment, I can calculate the evolution of the density matrix

$$\begin{aligned}
\dot{\rho}_{nm} &= -\frac{i}{\hbar}\left[\widehat{H},\widehat{\rho}\right]_{nm} \qquad (7.18)\\
&= -\frac{i}{\hbar}\left[(\widehat{H}\widehat{\rho})_{nm} - (\widehat{\rho}\widehat{H})_{nm}\right]\\
&= -\frac{i}{\hbar}\sum_{\nu}\left[H_{n\nu}\rho_{\nu m} - \rho_{n\nu}H_{\nu m}\right]\\
&= -\frac{i}{\hbar}\sum_{\nu}\left[(H_{o,n\nu} + V_{n\nu}(t))\rho_{\nu m} - \rho_{n\nu}(H_{o,\nu m} + V_{\nu m}(t))\right]\\
&= -\frac{i}{\hbar}\sum_{\nu}\left[H_{o,n\nu}\rho_{\nu m} + V_{n\nu}(t)\rho_{\nu m} - \rho_{n\nu}H_{o,\nu m} - \rho_{n\nu}V_{\nu m}(t)\right]\\
&= -\frac{i}{\hbar}\sum_{\nu}\left[H_{o,n\nu}\rho_{\nu m} - \rho_{n\nu}H_{o,\nu m} + V_{n\nu}(t)\rho_{\nu m} - \rho_{n\nu}V_{\nu m}(t)\right]\\
&= -\frac{i}{\hbar}\left(H_{o,nn}\rho_{nm} - \rho_{nm}H_{o,mm} + \sum_{\nu}\left[V_{n\nu}(t)\rho_{\nu m} - \rho_{n\nu}V_{\nu m}(t)\right]\right)\\
&= -\frac{i}{\hbar}(\hbar\omega_n\rho_{nm} - \rho_{nm}\hbar\omega_m) - \frac{i}{\hbar}\sum_{\nu}\left[V_{n\nu}(t)\rho_{\nu m} - \rho_{n\nu}V_{\nu m}(t)\right]\\
&= -i\omega_n\rho_{nm} + i\rho_{nm}\omega_m - \frac{i}{\hbar}\sum_{\nu}\left[V_{n\nu}(t)\rho_{\nu m} - \rho_{n\nu}V_{\nu m}(t)\right],
\end{aligned}$$

where I have used Equations 7.14 and 7.15 in the derivation. For the two-level

model I have only states 1 and 2 and I know that $V_{11} = V_{22} = 0$ and that $V_{12} = V_{21}^*$. I define $\omega_0 = \omega_2 - \omega_1$ and write out the components of the density matrix

$$\begin{aligned}
\dot{\rho_{11}} &= -\frac{i}{\hbar}\left[V_{12}(t)\rho_{21} - \rho_{12}V_{21}(t)\right] \\
\dot{\rho_{22}} &= -\frac{i}{\hbar}\left[V_{21}(t)\rho_{12} - \rho_{21}V_{12}(t)\right] \\
\dot{\rho_{12}} &= i\omega_0\rho_{12} - \frac{i}{\hbar}\left[V_{12}(t)\rho_{22} - \rho_{11}V_{12}(t)\right] \\
\dot{\rho_{21}} &= -i\omega_0\rho_{21} - \frac{i}{\hbar}\left[V_{21}(t)\rho_{11} - \rho_{22}V_{21}(t)\right].
\end{aligned} \tag{7.19}$$

7.1.2 The Rotating Wave approximation

If I now consider that the driving electric field has the form

$$E(t) = E_0\cos(\omega t), \tag{7.20}$$

with E_0 the electric field magnitude, I can use Equation 7.16 to write

$$\begin{aligned}
V_{nm} &= eE_0\cos(\omega t)z_{nm} \\
&= eE_0 z_{nm}(e^{-i\omega t} + e^{i\omega t}).
\end{aligned} \tag{7.21}$$

In general, the density matrix equations of motion cannot be solved exactly with a driving term of this form. However, I can simplify the driving term by making the rotating wave approximation.

To motivate the rotating wave approximation, notice that if I set $V = 0$ I get

$$\begin{aligned}
\dot{\rho_{11}} &= 0 \\
\dot{\rho_{22}} &= 0
\end{aligned} \tag{7.22}$$

$$\begin{aligned}\dot{\rho_{12}} &= i\omega_0\rho_{12} \\ \dot{\rho_{21}} &= -i\omega_0\rho_{21},\end{aligned}$$

which has a solution of $\rho_{11} = \rho_{22} = 0$ and $\rho_{12} = \rho_{12}(0)e^{i\omega_0 t}$ (and similarly for $\rho_{21} = \rho_{12}^*$). Thus when the driving frequency, ω, is close to the resonant frequency, ω_0, the part of V_{12} that goes as $e^{i\omega t}$ will be a much more effective driving term than the part that goes as $e^{-i\omega t}$. The rotating wave approximation is simply keeping only the effective driving term. So I get

$$\begin{aligned}V_{12} &= eE_0 z_{12} e^{i\omega t} \\ V_{21} &= eE_0 z_{21} e^{-i\omega t}.\end{aligned} \tag{7.23}$$

Inserting this function into the components of the density-matrix equations of motion, I get

$$\begin{aligned}\dot{\rho_{11}} &= -\frac{i}{\hbar}\left[eE_0 z_{12} e^{i\omega t}\rho_{21} - \rho_{12} eE_0 z_{21} e^{-i\omega t}\right] \\ \dot{\rho_{22}} &= -\frac{i}{\hbar}\left[eE_0 z_{21} e^{-i\omega t}\rho_{12} - \rho_{21} eE_0 z_{12} e^{i\omega t}\right] \\ \dot{\rho_{12}} &= i\omega_0\rho_{12} - \frac{i}{\hbar}\left[eE_0 z_{12} e^{i\omega t}\rho_{22} - \rho_{11} eE_0 z_{12} e^{i\omega t}\right] \\ \dot{\rho_{21}} &= -i\omega_0\rho_{21} - \frac{i}{\hbar}\left[eE_0 z_{21} e^{-i\omega t}\rho_{11} - \rho_{22} eE_0 z_{21} e^{-i\omega t}\right].\end{aligned} \tag{7.24}$$

Now, finally, I want to calculate steady-state solutions to these equations: solutions that hold after the transients from the turn-on of the electric field have decayed away. To do this, I introduce a slowly varying quantity

$$\begin{aligned}\sigma_{12} &= \rho_{12} e^{-i\omega t} \\ \sigma_{21} &= \rho_{21} e^{i\omega t}.\end{aligned} \tag{7.25}$$

From straightforward differentiation, I get

$$\begin{aligned}
\dot{\sigma}_{12} &= \dot{\rho}_{12}e^{-i\omega t} - i\omega\rho_{12}e^{-i\omega t} \qquad (7.26)\\
&= \dot{\rho}_{12}e^{-i\omega t} - i\omega\sigma_{12}\\
&= \left[i\omega_0\rho_{12} - \frac{i}{\hbar}\left[eE_0z_{12}e^{i\omega t}\rho_{22} - \rho_{11}eE_0z_{12}e^{i\omega t}\right]\right]e^{-i\omega t} - i\omega\sigma_{12}\\
&= i\omega_0\sigma_{12} - \frac{i}{\hbar}\left[eE_0z_{12}\rho_{22} - \rho_{11}eE_0z_{12}\right] - i\omega\sigma_{12}\\
&= -i(\omega-\omega_0)\sigma_{12} - \frac{i}{\hbar}\left[eE_0z_{12}\rho_{22} - \rho_{11}eE_0z_{12}\right]
\end{aligned}$$

and similarly for σ_{21}. Inserting the substitution into Equations 7.24 I get

$$\begin{aligned}
\dot{\rho_{11}} &= -\frac{i}{\hbar}\left[eE_0z_{12}\sigma_{21} - \sigma_{12}eE_0z_{21}\right] \qquad (7.27)\\
\dot{\rho_{22}} &= -\frac{i}{\hbar}\left[eE_0z_{21}\sigma_{12} - \sigma_{21}eE_0z_{12}\right]\\
\dot{\sigma_{12}} &= -i(\omega-\omega_0)\sigma_{12} + \frac{i}{\hbar}\left[\rho_{11}eE_0z_{12} - eE_0z_{12}\rho_{22}\right]\\
\dot{\sigma_{21}} &= i(\omega-\omega_0)\sigma_{21} - \frac{i}{\hbar}\left[eE_0z_{21}\rho_{11} - \rho_{22}eE_0z_{21}\right].
\end{aligned}$$

Because the dipole matrix element is simply calculated from the overlap of the state wavefunctions, $z_{12} = z_{21}$. This allows me to define

$$\begin{aligned}
V &\equiv eE_0z_{12}/\hbar \qquad (7.28)\\
&= eE_0z_{21}/\hbar.
\end{aligned}$$

With this definition, the equations of motion for the two-level system, in the absence of damping, are given by

$$\dot{\rho_{11}} = \frac{1}{2}iV\sigma_{12} - \frac{1}{2}iV\sigma_{21} \qquad (7.29)$$

$$\begin{aligned}
\dot{\rho_{22}} &= -\frac{1}{2}iV\sigma_{12} + \frac{1}{2}iV\sigma_{21} \\
\dot{\sigma_{12}} &= -i(\omega - \omega_0)\sigma_{12} + \frac{1}{2}iV(\rho_{11} - \rho_{22}) \\
\dot{\sigma_{21}} &= i(\omega - \omega_0)\sigma_{21} - \frac{1}{2}iV(\rho_{11} - \rho_{22}).
\end{aligned}$$

I can now return to Figure 7.1, where I graphically represented three damping terms. The first, Γ_1, is the rate of decay from state 2 to state 1. The second, γ_2, gives the dephasing between states 1 and 2. The third, γ_3, gives the rate of ionization out of state 2. I add these terms to the equations of motion in the following way:

$$\begin{aligned}
\dot{\rho_{11}} &= \dot{\rho_{11}}(no\ damping) + \Gamma_1\rho_{22} \\
\dot{\rho_{22}} &= \dot{\rho_{22}}(no\ damping) - \Gamma_1\rho_{22} - \gamma_3\rho_{22} \\
\dot{\sigma_{12}} &= \dot{\sigma_{12}}(no\ damping) - \gamma_2\sigma_{12} \\
\dot{\sigma_{21}} &= \dot{\sigma_{21}}(no\ damping) - \gamma_2\sigma_{21}.
\end{aligned} \tag{7.30}$$

Adding in all the damping terms I arrive at the full equations of motion for the two-level open system

$$\begin{aligned}
\dot{\rho_{11}} &= \frac{1}{2}iV\sigma_{12} - \frac{1}{2}iV\sigma_{21} + \Gamma_1\rho_{22} \\
\dot{\rho_{22}} &= -\frac{1}{2}iV\sigma_{12} + \frac{1}{2}iV\sigma_{21} - \Gamma_1\rho_{22} - \gamma_3\rho_{22} \\
\dot{\sigma_{12}} &= -i(\omega - \omega_0)\sigma_{12} + \frac{1}{2}iV(\rho_{11} - \rho_{22}) - \gamma_2\sigma_{12} \\
\dot{\sigma_{21}} &= i(\omega - \omega_0)\sigma_{21} - \frac{1}{2}iV(\rho_{11} - \rho_{22}) - \gamma_2\sigma_{21}.
\end{aligned} \tag{7.31}$$

7.1.3 Analytic solutions to the density-matrix equations of motion

To get an analytic solution to the equations of motion, it is easiest to put the equations into matrix form

$$\begin{pmatrix} \dot{\rho_{11}} \\ \dot{\rho_{22}} \\ \dot{\sigma_{12}} \\ \dot{\sigma_{21}} \end{pmatrix} = \begin{pmatrix} 0 & \Gamma_1 & \frac{1}{2}iV & -\frac{1}{2}iV \\ 0 & -\Gamma_1 - \gamma_3 & -\frac{1}{2}iV & \frac{1}{2}iV \\ \frac{1}{2}iV & -\frac{1}{2}iV & -i(\omega - \omega_0) - \gamma_2 & 0 \\ -\frac{1}{2}iV & \frac{1}{2}iV & 0 & -i(\omega - \omega_0) - \gamma_2 \end{pmatrix} \begin{pmatrix} \rho_{11} \\ \rho_{22} \\ \sigma_{12} \\ \sigma_{21} \end{pmatrix}. \tag{7.32}$$

I can summarize this matrix equation as

$$\dot{\widehat{X}} = \widehat{M}\widehat{X}, \tag{7.33}$$

where $\widehat{X}$ is the vector of components of the density matrix and $\widehat{M}$ is the matrix of coefficients. A differential equation of this form has a solution given by

$$\widehat{X}(t) = \sum_{i=1}^{4} a_i \widehat{E}_{\lambda_i} e^{\lambda_i t}, \tag{7.34}$$

where λ_i are the eigenvalues of the matrix of coefficients, $\widehat{E}_{\lambda_i}$ are the corresponding eigenvectors and a_i are coefficients. As an example, the solution for ρ_{11} is given by

$$\rho_{11} = a_1\widehat{E_{\lambda_1 1}}\exp(\lambda_1 t) + a_2\widehat{E_{\lambda_2 1}}\exp(\lambda_2 t) + a_3\widehat{E_{\lambda_3 1}}\exp(\lambda_3 t) + a_4\widehat{E_{\lambda_4 1}}\exp(\lambda_4 t). \tag{7.35}$$

Using this form, I find the analytic solution by calculating the eigenvalues and eigenvectors of the matrix of coefficients, expressed in terms of the parameters of the equations of motion (V, Γ_1, etc.) I calculated the eigenvalues and eigenvectors using Mathematica. The problem is simplified somewhat by the fact that one eigenvector component is zero. The algebra is not very enlightening, so I will not include it here. All of the algebraic manipulations are contained, with documentation, in three Mathematica files entitled Rabi_Model, Rabi_Model_Simplify and Rabi_Inversion_Check.

Once the eigenvalues and eigenvectors have been determined, I solve for the coefficients a_i by using the initial conditions

$$\begin{aligned} \rho_{11}(0) &= 1 \\ \rho_{22}(0) &= 0 \\ \sigma_{12}(0) &= 0 \\ \sigma_{21}(0) &= 0, \end{aligned} \tag{7.36}$$

which correspond to the initial population being entirely in the 1 state and the driving term being turned on instantaneously to full strength at time t = 0. It is not so complicated to write the coefficients as a function of the eigenvalues and V. However, for the purpose of using the analytic solutions for a fitting routine, I need the analytic solution as a function only of the eigenvalues. This is achievable, since V is in fact a function of the eigenvalues. Unfortunately, the algebra for expressing V, and then the coefficients, in terms of the eigenvalues is extremely complicated

(to put it mildly), which is the primary reason I used Mathematica to do the manipulations.

Once the algebra is completed, I have an analytic solution to the density-matrix equations of motion that can be used in a fitting routine. I take the photoconductivity to be proportional to $1-\rho_{11}$. This is a reasonable model to use for the $2p^+$ state at a magnetic field of approximately 3.5 T because the $2p^+$ state lies within the first Landau level. Since the photoconductivity signal I observe persists for approximately 30 ns after the sliced pulse of THz radiation, it is safe to assume that all electrons that are excited out of the ground state end up in the conduction band within our probing window. I also simplify matters by setting $\Gamma_1 = 0$, which is a safe approximation for the 1s to $2p^+$ experiment since the lifetime represented by Γ_1 is long in comparison to the other processes.

Finally, the analytic solution can be fit to existing data by varying only the eigenvalues to achieve the best fit. I did this using Igor Pro, a commercial data analysis program. Once the eigenvalues are determined from the fit, the algebraic relationships are used to calculate the corresponding values of V, γ_2 and γ_3.

7.1.4 Fitting results

For the 1s to 2p transition at zero magnetic field the dipole matrix element, z, is 52.6 Å. However, as the magnetic field increases the electrons are more tightly bound to the donor site; the reduced orbital radius leads to a larger overlap of

wavefunctions and a larger matrix element for the orbital transition. For the purpose of getting numerical results, I take an approximate dipole matrix element of 100 Å. The fitting routine is executed with six parameters. The first parameter varies the point in the data that is identified as t = 0, which allows for uncertainty in the zero-path-difference point of the delay stage in the pulse slicer. The second parameter is a constant offset and the third parameter is an overall scale factor. These convert the analytic function, which has a maximum range of 0 to 1, to the range and base value of the data. The last three parameters are the eigenvalues of the analytic solution. The remaining terms in Equation 7.35 are all expressed in terms of the eigenvalues, and so introduce no additional free parameters.

In Figure 7.2 I show data of excitation of the two-level system using THz radiation of varying polarizations. This data was presented in Chapter 6. Here, I have added fits to the data using the model. From the fits, I see that, for the most effective coupling (σ^+), the dephasing rate, γ_2, is $1.98 \times 10^{11}\ s^{-1}$, which corresponds to a dephasing time, T_2^*, of 5 ps. The ionization rate into the conduction band, γ_3, is $1.22 \times 10^{11}\ s^{-1}$, which corresponds to a lifetime of the electrons in the $2p^+$ state of about 8 ps. For less effective coupling into the conduction band, the fits indicate that both the dephasing and ionization rates decrease.

In Figure 7.3 I show fits to data from Chapter 6 on the response to illumination with varying intensities of THz radiation. From these fits, I see that the dephasing and ionization rates decrease with intensity. When using one attenuator (the blue

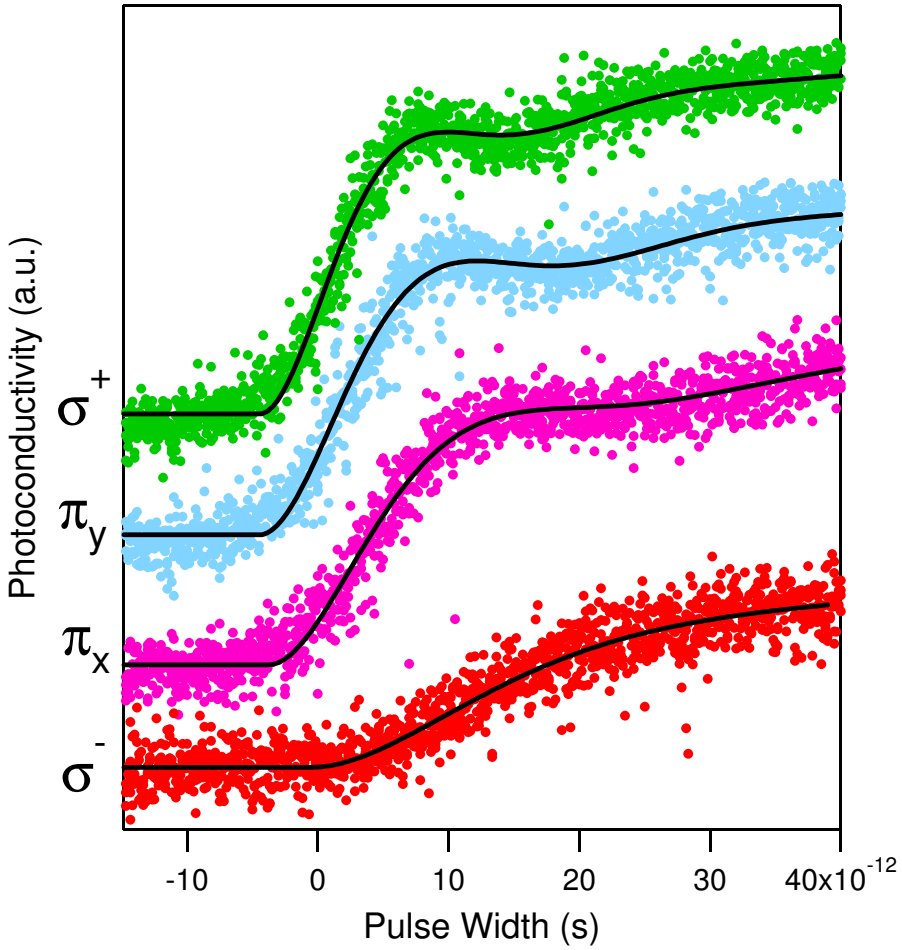

Figure 7.2: A plot of data demonstrating the response of the two-level system to different polarizations of incident light. The solid black lines are fits using the two-level open-system density matrix model. The relevant model parameters extracted from each fit are shown.

trace in Figure 7.3), Rabi oscillations are still evident. The fits indicate that γ_2 is $1.07 \times 10^{11}\ s^{-1}$, which corresponds to a dephasing time, T_2^*, of about 9 ps. The ionization rate into the conduction band, γ_3, is $6.54 \times 10^{10}\ s^{-1}$, which corresponds to a lifetime of the electrons in the $2p^+$ state of about 15 ps. With two attenuators (red trace) the intensity of the incident THz radiation is considerably smaller than with no attenuators (green trace) but Rabi oscillations are still weakly evident. γ_2 is $1.23 \times 10^{10}\ s^{-1}$, corresponding to a T_2^* of about 80 ps, and γ_3 is $4.01 \times 10^{10}\ s^{-1}$, corresponding to a lifetime of the electrons in the $2p^+$ state of about 25 ps.

From data and fits such as those shown in Figure 7.3 and those by Cole, et al. [26], we conclude that the observed dephasing times are limited by the rapid ionization from the $2p^+$ state into the conduction band. This conclusion is supported by the work of Brandi, Latge and Oliveira [29], who pursued more complicated fits and modelling of the data presented by Cole, et al. [26]. This limitation is well grounded conceptually, since the excited state ($2p^+$) lies within the conduction band and interactions with conduction-band states may provide the primary mechanism for decoherence and dephasing. At reduced intensities, the slower excitation to the excited state translates into a longer time before this dephasing mechanism contributes. It is this limitation that primarily motivated the extension of the Rabi oscillations to new wavelengths where the excited state would not lie within the conduction band and would therefore, hopefully, have a longer dephasing time.

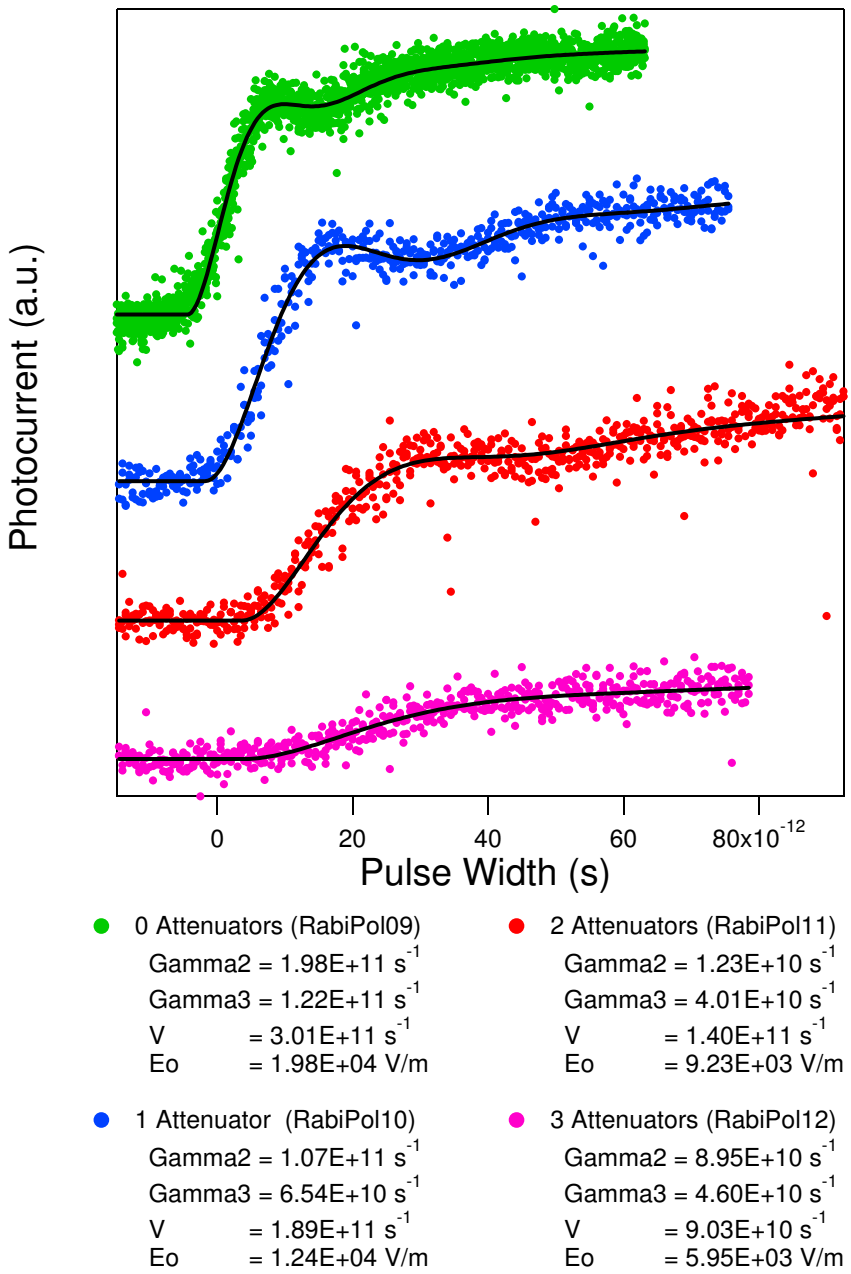

Figure 7.3: A plot of data demonstrating the response of the two-level system to different intensities of σ^+ polarized incident light. Each attenuator reduces the intensity by a factor of about 4. The solid black lines are fits using the two-level open-system density matrix model. The relevant model parameters extracted from each fit are shown.

7.1.5 Limitations of the two-level open-system model

As mentioned above, the two-level open system is an appropriate model to use for the experiment where Rabi oscillations are driven between the 1s and $2p^+$ state at a magnetic field of approximately 3.5 T. The model, however, makes two primary assumptions. First, it assumes that all electrons excited out of the ground state end up in the conduction band within the sampling window. Second, the model is integrated only over the duration of the sliced pulse. It does not consider an imperfect contrast ratio where the illuminating electric-field strength is low but non-zero for a long time before stepping up to its maximum value.

The first assumption is false if the excited state of the Rabi oscillations does not lie within the conduction band and the ionization rate from the excited state is not fast compared to the sampling window. The second assumption will begin to cause problems as the contrast ratio decreases or the lifetime of the excited (2p) state increases, thus allowing a larger population to build up in this state before illumination at the full electric-field strength initiates Rabi oscillations. In the case of Rabi oscillations between the 1s and $2p^-$ states at magnetic fields of roughly 4.5 T, the excited state lies well below the conduction band ($\sim$40 cm^{-1}) and has a measured lifetime of 350 ns [17]. The two-level open-system model is therefore not sufficient to model this case.

7.2 Three-level model

7.2.1 Extending the two-level open-system model

To properly model Rabi oscillations between the 1s and $2p^-$ states and to understand conditions under which Rabi oscillations might be observed, I extend to a three-level model that separately tracks the population of the conduction band state. In addition, I make the amplitude of the electric field a function of time. This allows me to integrate the equations of motion for a long time with low-level illumination to accurately represent the imperfect contrast ratio and then turn the electric field to full strength to calculate the system response to the sliced pulse.

Since I am separating the conduction band from the excited state and anticipating a slower ionization rate into the conduction band, it is no longer sufficient to take the photoconductivity to be proportional to $1 - \rho_{11}$. Rather, I must now take the photoconductivity to be proportional to the population of the conduction band. In addition, since the sampling window persists for a substantial time after the sliced pulse has ended, I must allow for the possibility of continued ionization into the conduction band after the conclusion of the sliced pulse. This will be done by terminating the integration of the equations of motion at the end of the sliced pulse and reading out the populations of each state. These values will then be used as the starting point for a second integration of the equations of motion in which the driving term is set to zero and the excited state continues to ionize to

the conduction band. This integration will be allowed to run until the population of the conduction band has reached an equilibrium value. The photoconductivity will be taken to be proportional to this equilibrium value.

With these additions, it will no longer be sufficient to run a single integration of the equations of motion and plot the state populations as a function of time within the sliced pulse. Now it will be necessary to integrate the equations of motion for a pulse of specified width and then calculate the equilibrium value of the conduction band population. By repeating this process for various values of the pulse width, the population as a function of pulse duration can be mapped out.

7.2.2 The density-matrix equations of motion

In Figure 7.4 I show a diagram of a three-level system. State 1 remains the ground state (1s) and state 2 is the excited state ($2p^-$). The driving term (from the electric field of a laser), V, remains, though it will become a function of time. Γ_{21} is the inverse lifetime for the excited state. γ_{21} is the dephasing rate for the oscillations between states 1 and 2. γ_{2c} is the ionization rate from state 2 to state c (the conduction band). Γ_{c2} is the lifetime of the conduction band state before capture by an ionized donor. Γ_{c1} is the lifetime of the conduction band state before returning to the ground state through a separate channel.

Fortunately for the reader, I do not need to re-derive the density matrix equa-

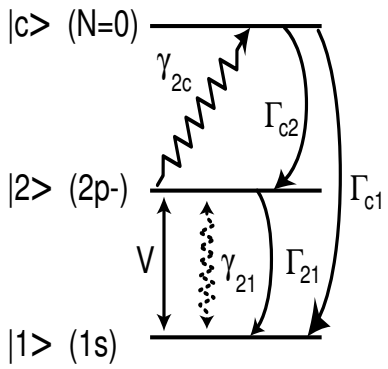

Figure 7.4: A diagram of the states for the three-level density matrix model. V is a driving term between states 1 and 2, Γ_{21} is the lifetime of the excited state, γ_{21} the dephasing rate, γ_{2c} the ionization rate from state 2 to state c (the conduction band), Γ_{c2} the lifetime of the conduction band state before capture by an ionized donor and Γ_{c1} the lifetime of the conduction band state before returning to the ground state through a separate channel.

tions of motion for this system. Our derivation above for the two-level system was done in the absence of damping, with the damping terms added by hand at the end. For the three-level model, I have not changed any of the driving terms, only added additional damping mechanisms. See Boyd [28] for additional information on the choice of these damping terms. I therefore take Equations 7.29 and add in the damping terms for the three-level system to get

$$\begin{aligned}
\rho'_{cc}(t) &= \Gamma_{2c}\rho_{22}(t) - \Gamma_{c2}\rho_{cc}(t) - \Gamma_{c1}\rho_{cc}(t) \qquad (7.37)\\
\rho'_{11}(t) &= \frac{1}{2}iV\sigma_{12}(t) - \frac{1}{2}iV\sigma_{21}(t) + \Gamma_{21}\rho_{22}(t) + \Gamma_{c1}\rho_{cc}(t)\\
\rho'_{22}(t) &= -\frac{1}{2}iV\sigma_{12}(t) + \frac{1}{2}iV\sigma_{21}(t) - \Gamma_{21}\rho_{22}(t) - \gamma_{2c}\rho_{22}(t) + \Gamma_{c2}\rho_{cc}(t)\\
\sigma'_{12}(t) &= -i(\omega-\omega_0)\sigma_{12}(t) + iV(\rho_{11}(t) - \rho_{22}(t)) - \gamma_{21}\sigma_{12}(t)\\
\sigma'_{21}(t) &= i(\omega-\omega_0)\sigma_{21}(t) - iV(\rho_{11}(t) - \rho_{22}(t)) - \gamma_{21}\sigma_{21}(t),
\end{aligned}$$

where ρ_{cc} is the population of the c state (conduction band).

As initial conditions I take

$$\begin{aligned} \rho_{cc}(0) &= 0 \\ \rho_{11}(0) &= 1 \\ \rho_{22}(0) &= 0 \\ \sigma_{12}(0) &= 0 \\ \sigma_{21}(0) &= 0, \end{aligned} \tag{7.38}$$

which again corresponds to the initial population being entirely in the 1 state.

To simulate an imperfect contrast ratio, I take V to be given by

$$\begin{aligned} V &= \frac{eE_0(t)z_{12}}{\hbar} \\ E_0(t) &= \sqrt{E_{max}^2/CR} \qquad t < t_0 \\ E_0(t) &= E_{max} \qquad\qquad t \geq t_0, \end{aligned} \tag{7.39}$$

where CR is the contrast ratio of the sliced pulse (the ratio of the peak intensity to the intensity before activation of the reflection switches) and E_{max} is the maximum electric-field amplitude during the sliced pulse.

7.2.3 Numerical modelling

Since the three-level model contains substantially more parameters than the two-level open-system model, as well as a step function, the algebra necessary to find an analytic solution to the model is prohibitively complex. It is also not clear that it would be appropriate to use this model to fit data taken for the 1s

to $2p^-$ transition, since no Rabi oscillations have been observed between these states. Instead, I will use this model to generate numerical results that map out the expected response of the system to various conditions.

To map out the parameter space, I need to assign numerical values to the parameters and then vary these values. Given the large number of parameters, it is not instructive to attempt to map out the full parameter space. I will therefore take fixed values for those parameters that can be determined with some confidence. The values of these fixed parameters and relevant constants are

$$\begin{aligned}
\hbar &= 6.63 \times 10^{-34}/(2\pi)\ Js \qquad (7.40)\\
e &= 1.602 \times 10^{-19}\ C\\
m_o &= 9.109 \times 10^{-31}\ kg\\
a^* &= 10 \times 10^{-9}\ m\\
z_{12} &= 52.6 \times 10^{-10}\ \mathring{A}\\
\omega - \omega_0 &= 0\\
\Gamma_{21} &= 2.85 \times 10^{6}\ s^{-1}\\
\Gamma_{c1} &= 6.66 \times 10^{6}\ s^{-1}\\
\Gamma_{c2} &= 2 \times 10^{7}\ s^{-1}\\
t_0 &= 1 \times 10^{-6}\ s\\
t_s &= 50 \times 10^{-12}\ s.
\end{aligned}$$

The choice of these values requires some explanation. a^* is the effective Bohr

radius of the electron orbit, determined from the hydrogenic model to be 100 Å. The zero-field dipole matrix element, $z_{12} = 52.6$ Å, is used since a reliable calculation of how this element scales with magnetic field is not available. The condition $\omega - \omega_0 = 0$ specifies that the THz radiation is exactly resonant with the energy difference between states 1 and 2. Γ_{21} is the inverse lifetime of the $2p^-$ state before decaying to the 1s state. This is taken to be 350 ns from measurements by Kalkman, et al. [17], giving a rate of $2.85 \times 10^6\ s^{-1}$. From calculations using the model, I find that setting this rate to zero makes no qualitative difference, further validating the use of the $\Gamma_1 = 0$ assumption above. Γ_{c1} is the decay rate of conduction band electrons to the ground state. This is taken to be $6.66 \times 10^6\ s^{-1}$ from Klaassen, et al. [1], who state that conduction-band electrons return to the ground (1s) state in about 150 ns. Γ_{c2} is the rate for conduction-band electrons to be captured by an ionized impurity. This is taken to be $2 \times 10^7\ s^{-1}$ from Klaassen et al., who indicate that the mean time for recapture is 50 ns. t_0 is the time at which the electric field switches from the prepulse value to full strength. This is taken to be 1×10^{-6} s, in rough agreement with experimental conditions. Finally, t_s is the time allowed for electrons in state 2 to ionize to the conduction band before reading out the conduction band population. This is taken to be 1×10^{-12} s from tests of the model that indicate that the conduction band population saturates within this time for an ionization rate (γ_{2c}) of $1 \times 10^{10}\ s^{-1}$.

This leaves four parameters that will be varied to map out the behavior of the three-level system: γ_{2c}, γ_{21}, E_{max} and CR. To study the dependence on each

parameter, I will fix the other three and run the numerical calculations to generate a family of curves showing the system behavior for various values of the parameter in question. The code used to generate these numerical results is written in Mathematica and is contained in a file named HydrogenicElecStates_v3. The fixed values for the parameters are taken from the fits to Rabi oscillation data at 84 cm^{-1} and the maximum CR observed by Cole [26].

These calculations only plot the population of the conduction band as a function of illuminating pulse duration. They are intended to aid in understanding the limits of the observation of coherent phenomena using transient photoconductivity. In some cases, coherent oscillations between states 1 and 2 may still be occurring, though they are not visible in the conduction band population.

Varying contrast ratio

I will begin by varying the contrast ratio of the sliced pulse used to illuminate the sample. I fix the following parameters:

$$\begin{aligned} \gamma_{2c} &= 1 \times 10^{11}\ s^{-1} \\ \gamma_{21} &= 1.98 \times 10^{11}\ s^{-1} \\ E_{max} &= 1.98 \times 10^{4}\ V/m. \end{aligned} \tag{7.41}$$

I vary the contrast ratio from 10^6, the value quoted by Cole [26], to 10^3, well below my measured CR values. The results are shown in Figure 7.5 and indicate that Rabi oscillations are only clearly visible for contrast ratios of 5×10^4 or larger.

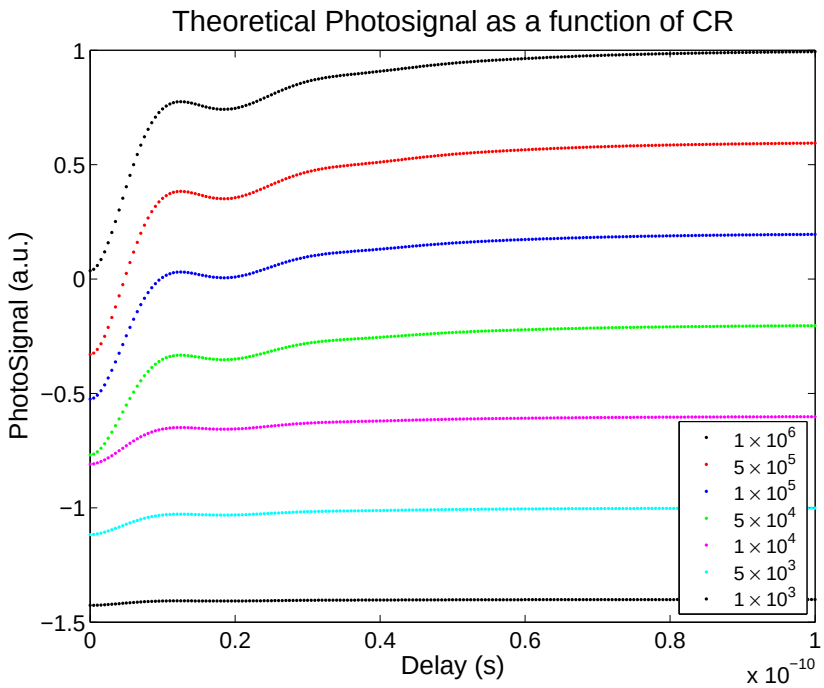

Figure 7.5: Calculations of the theoretical photocurrent as a function of varying contrast ratio.

At smaller contrast ratios, the preactivation radiation is sufficiently intense that a large excited state population is created before the intense sliced pulse arrives. If the populations of the ground and excited state are nearly equal, oscillations between the two states do not result in a significant change in the conduction band population.

Varying E_{max}

Next I will vary the value of the maximum electric field. I fix the following parameters:

$$\gamma_{2c} \;=\; 1 \;\times\; 10^{11}\ s^{-1} \tag{7.42}$$

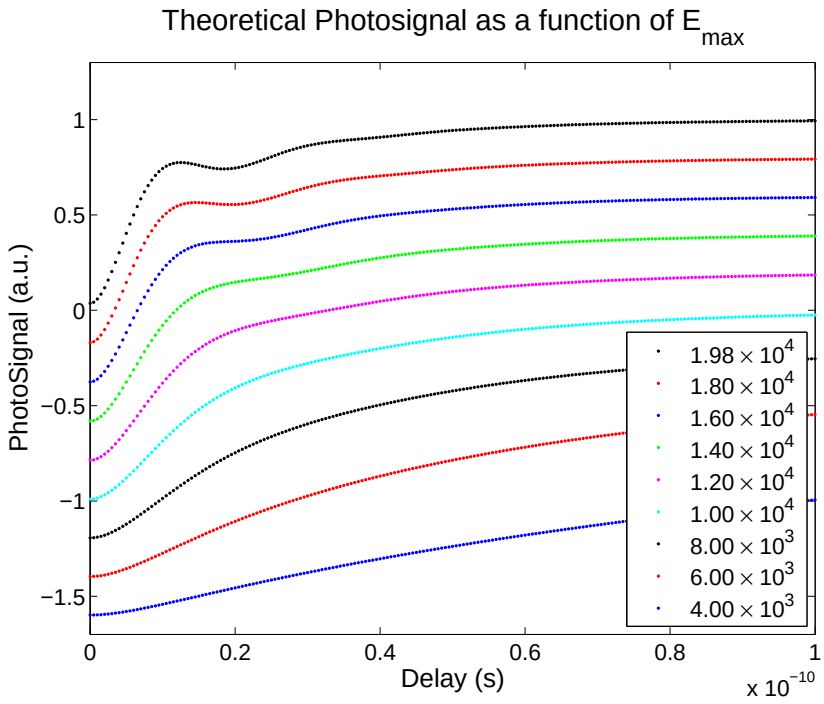

Figure 7.6: Calculations of the theoretical photocurrent as a function of varying maximum electric field.

$$\gamma_{21} = 1.98 \times 10^{11}\ s^{-1}$$

$$CR = 10^{6}.$$

E_{max} is varied from $1.98 \times 10^4\ V/m$ to $4.00 \times 10^3\ V/m$. The upper value is the fitted value for observed Rabi oscillations at 84 cm^{-1}, where the FEL intensity is highest; the lower value corresponds to an intensity roughly 25 times weaker, lower than the expected FEL intensity near 34 cm^{-1}.

The results are shown in Figure 7.6 and indicate that Rabi oscillations are only clearly visible for a maximum electric field of $1.4 \times 10^4\ V/m$ or greater. At smaller electric fields, the oscillation between states 1 and 2 is slow compared to the damping terms and no coherent phenomena can be observed in the conduction-

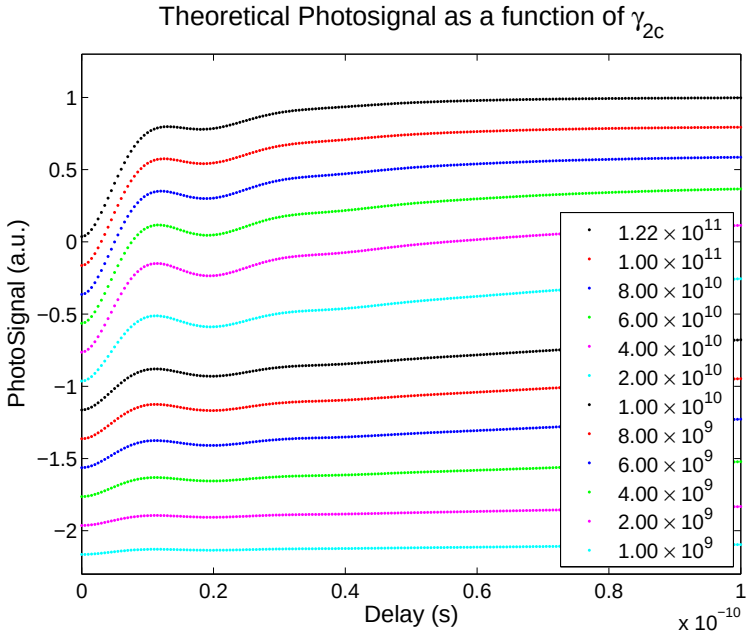

Figure 7.7: Calculations of the theoretical photocurrent as a function of varying γ_{2c}.

band population.

Varying γ_{2c}

Next I will vary the value of γ_{2c}, the ionization rate into the conduction band. I fix the following parameters:

$$\begin{aligned} \gamma_{21} &= 1.98 \times 10^{11}\ s^{-1} \\ CR &= 10^6 \\ E_{max} &= 1.98 \times 10^4\ V/m. \end{aligned} \tag{7.43}$$

γ_{2c} is varied from $1.22 \times 10^{11}\ s^{-1}$ to $1.00 \times 10^9\ s^{-1}$. The upper value corresponds to the fitted value for observed Rabi oscillations at 84 cm^{-1} (an ionization time of 8 ps); the lower value corresponds to an ionization time of 1 ns, well longer

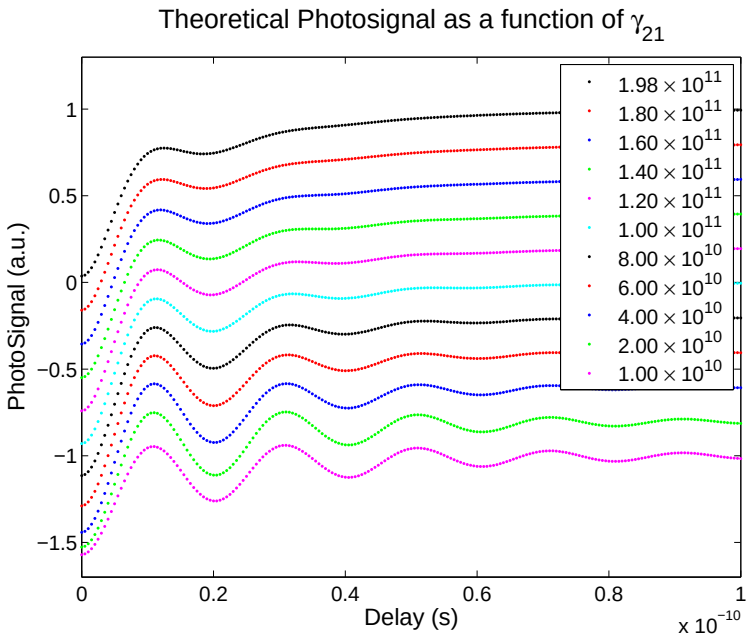

Figure 7.8: Calculations of the theoretical photocurrent as a function of varyingγ_{21}.

than the 100 ps ionization time quoted by Klaassen, et al. [1] as a typical time for decay to the conduction band. The longer time was chosen since the $2p^-$ state is below the conduction band and the ionization may be slower than for excited states above the conduction band. The results are shown in Figure 7.7, where Rabi oscillations remain visible for values of γ_{2c} as low as $6.00 \times 10^9\ s^{-1}$.

Varying γ_{21}

Next I will vary the value of γ_{21}, the dephasing rate between states 1 and 2. I fix the following parameters:

$$\begin{aligned} \gamma_{2c} &= 1.00 \times 10^{11}\ s^{-1} \\ CR &= 10^6 \end{aligned} \tag{7.44}$$

$$E_{max} \;=\; 1.98 \;\times\; 10^4 \; V/m.$$

γ_{21} is varied from $1.98 \times 10^{11}\ s^{-1}$ to $1.00 \times 10^{10}\ s^{-1}$. The upper value corresponds to the fitted value for observed Rabi oscillations at 84 cm^{-1}; the lower value corresponds to a dephasing time of 100 ps, a plausible estimate of the dephasing time that could be expected for the 1s to $2p^-$ transition. The results are shown in Figure 7.8. As the dephasing time increases (rate decreases) more cycles of Rabi oscillations become visible.

Modelling the conditions of the 34-cm^{-1} experiment

A major goal of experimentation on the hydrogenic donor system over the course of this thesis was the observation of Rabi oscillations between the 1s and $2p^-$ states at a magnetic field of approximately 4.5 T. At this magnetic field, the $2p^-$ state lies well below the conduction band and the resonant frequency is around 34 cm^{-1}. The output power of the UCSB FEL at this frequency is approximately one tenth of the power at 84 cm^{-1}. A decrease in intensity of this magnitude corresponds to a reduction in the electric field amplitude of approximately 3. In Figure 7.6 I show that a reduction of the electric field amplitude to a magnitude of $\sim 6 \;\times\; 10^3\ V/m$ would make observing Rabi oscillations impossible. However, since the excited ($2p^-$) state is well separated from the conduction band, the ionization rate to the conduction band should decrease and it is possible that the dephasing rate would also decrease.

I modelled the behavior of the system under conditions similar to those I

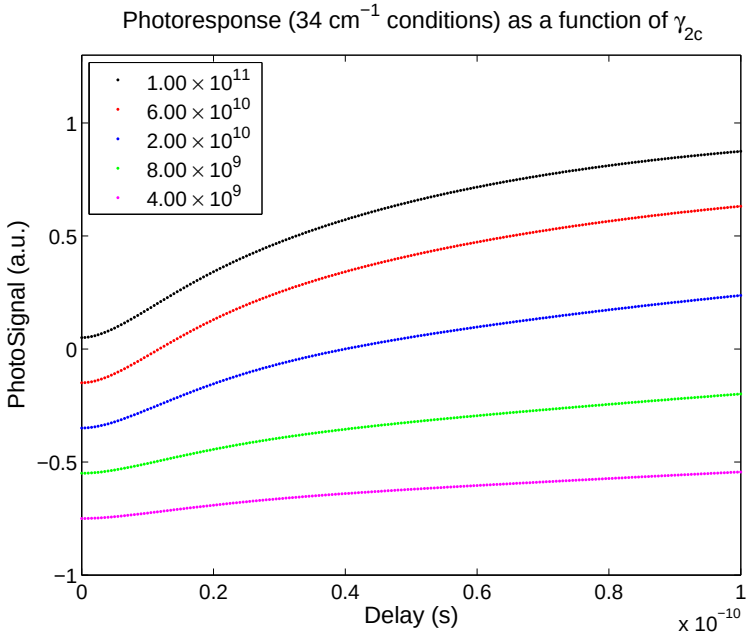

Figure 7.9: Calculations of the theoretical photocurrent as a function of varying γ_{2c} with fixed parameters that emulate the conditions of the 34-cm^{-1} experiment.

believe I obtained for the 34 cm^{-1} experiment by fixing

$$CR = 7 \times 10^4 \tag{7.45}$$

$$E_{max} = 6.2 \times 10^3 \; V/m.$$

This value of CR is chosen as the largest CR measured for the pulse slicer system and the value of E_{max} is chosen to represent an intensity one tenth of that fitted to the system behavior at 84 cm^{-1}. I then varied the two remaining parameters separately.

Fixing γ_{21} at 1.98×10^{11} s^{-1}, I varied γ_{2c} from 1.00×10^{11} s^{-1} to 4.00×10^9 s^{-1}. The results are shown in Figure 7.9. It is clear that a significant decrease in the ionization rate does not make the observation of Rabi oscillations possible. In this case, the primary limitation is that γ_{21} is large, while E_{max} is

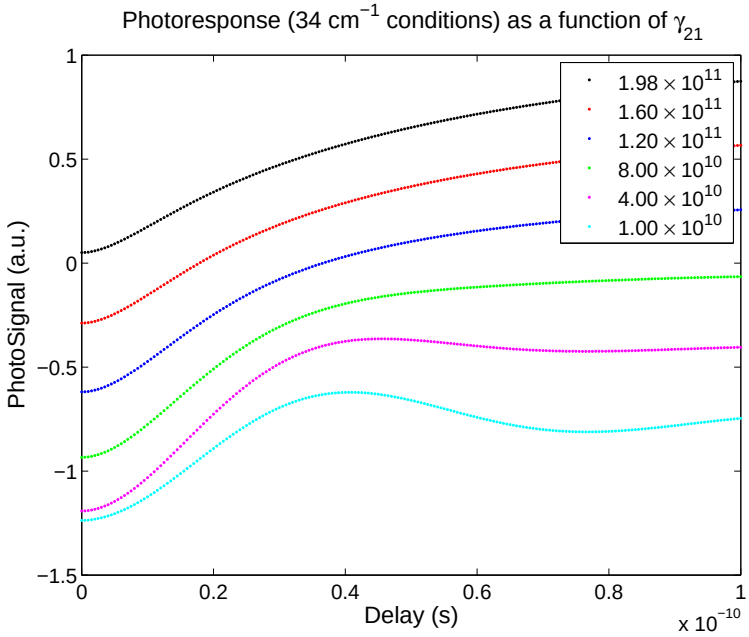

Figure 7.10: Calculations of the theoretical photocurrent as a function of varying γ_{21} with fixed parameters that simulate the conditions of the 34-cm^{-1} experiment.

relatively small; electrons dephase more quickly than they are driven through a Rabi oscillation.

Fixing γ_{2c} at $1.00 \times 10^{11}\ s^{-1}$, I varied γ_{21} from $1.98 \times 10^{11}\ s^{-1}$ to $1.00 \times 10^{10}\ s^{-1}$. The results are shown in Figure 7.10. Rabi oscillations become visible for a dephasing rate of about $4.00 \times 10^{10}\ s^{-1}$, which corresponds to a dephasing time of 25 ps. This modelling could be interpreted to suggest that Rabi oscillations between the 1s and $2p^-$ states must have a dephasing time of less than 25 ps since they have not been seen by experiments to date. However, this conclusion is not justified without a well established value of the ionization rate, γ_{2c}.

As an example, in Figure 7.11 I plot the populations of the conduction band and $2p^-$ states using the E_{max} and CR values that simulate the experimental

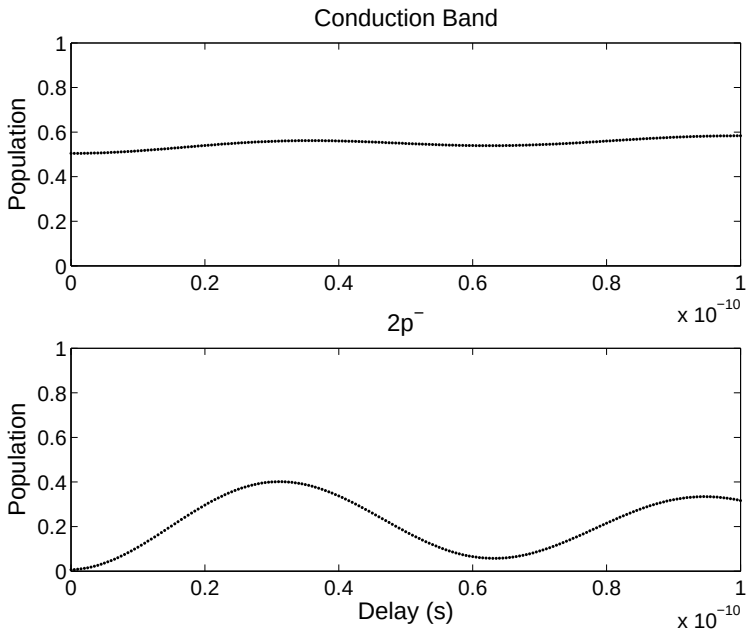

Figure 7.11: Calculations of the populations of the conduction band and $2p^-$ states with $\gamma_{21} = 1.00 \times 10^{10}\ s^{-1}$ and $\gamma_{2c} = 2.00 \times 10^{9}\ s^{-1}$.

conditions. I take $\gamma_{2c} = 2.00 \times 10^{9}\ s^{-1}$, which corresponds to an ionization time of 500 ps, and $\gamma_{21} = 1.00 \times 10^{10}\ s^{-1}$, which corresponds to a dephasing time of 100 ps. The $2p^-$ state population clearly oscillates, while the conduction band population is nearly constant. The constant conduction band population (and thus photoconductivity signal) exemplifies how the absence of observed oscillations can not be used to draw conclusions about the limits of γ_{21} without a reliable value for γ_{2c}.

7.2.4 Discussion of the three-level model

The three-level density matrix equations of motion used here include an imperfect contrast ratio, lifetimes of all states, a dephasing rate and an ionization rate.

However, the model still neglects some effects that may be present in the actual experiment. For instance, the model makes no allowances for an imperfect turn-off of the sliced pulse. Although no conclusive evidence for an imperfect turn-off has been found, it is suggested by certain system behavior and is discussed in more detail in Appendix C.3.

In Section 5.2 I described the method for mounting a silicon solid-immersion lens to a sample. This was intended to increase the intensity of the THz on the sample, providing higher electric-field amplitudes for the experiment at 34 cm^{-1}. The increase in intensity from the solid-immersion lens should be approximately a factor of ten, offsetting the decrease in intensity from the FEL at this wavelength. According to the three-level model, resonant illumination of the 1s to $2p^-$ transition with electric-field amplitudes comparable to those used in the 1s to $2p^+$ experiment should have resulted in the observation of Rabi oscillations. However, as discussed in Chapter 6, the behavior of the lens-mounted samples appears to be limited by effects that I have not considered in the model, for example carrier heating. Particulary after electrons have been ionized into the conduction band, illumination by THz radiation can heat the electrons, and therefore change the sample conductivity. This model does not contain any terms that could generate the exponential-like increase in photoconductivity that was observed using the lens-mounted sample (Figure 6.26).

If we neglect the possibility of carrier heating, it is still difficult to conclude whether Rabi oscillations could have been observed with the lens-mounted sample.

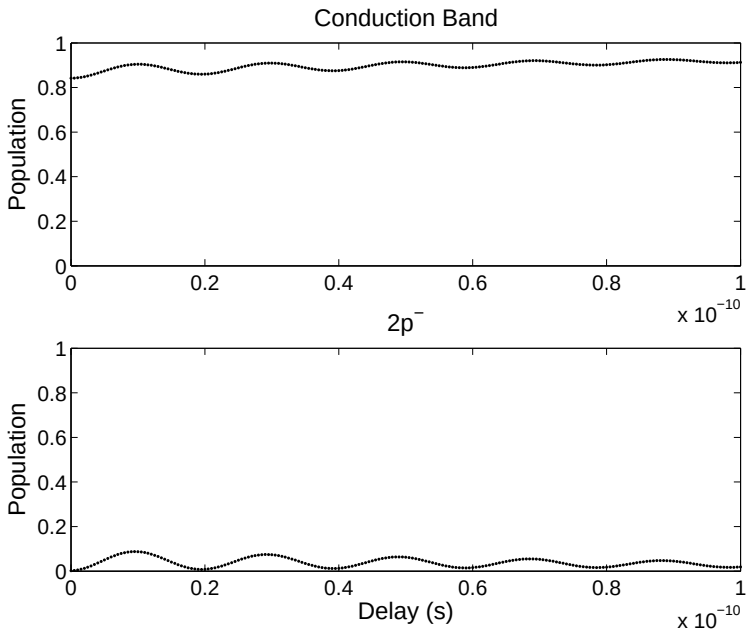

Figure 7.12: Calculations of the populations of the conduction band and $2p^-$ states with $E_{max} = 2 \times 10^4\ V/m$, CR $= 7 \times 10^4$, $\gamma_{21} = 2.00 \times 10^{10}\ s^{-1}$ and $\gamma_{2c} = 1.00 \times 10^{10}\ s^{-1}$.

The difficulty arises from the delicate balance between coupling in sufficient energy (to drive oscillations faster than the dephasing time) and increasing the pre-pulse (to the level where the ground state has been significantly depopulated before the short time-resolved pulse arrives). As an example, in Figure 7.12 I show a simulation of the population of the conduction band and $2p^-$ states using the best CR I observed (7×10^4) and assuming that the lens increased the intensity to that obtained at 84 cm^{-1} ($E_{max} = 2 \times 10^4\ V/m$). I chose reasonable estimates for γ_{21} ($2.00 \times 10^{10}\ s^{-1}$, corresponding to 50 ps) and γ_{2c} ($1.00 \times 10^{10}\ s^{-1}$, corresponding to 100 ps). With these parameters, the pre-pulse is sufficient to excite nearly all of the electrons into the conduction band: the conduction band population before the arrival of the short, intense pulse is around 85%. When

the intense pulse arrives, it drives oscillations of the small population of electrons remaining in the 1s and $2p^-$ states. The photoconductivity signal (conduction band population) does oscillate, but the oscillations are quite small and would be difficult to distinguish from noise. The key parameter that could be improved is the contrast ratio. An increase of the contrast ratio to 10^6, for example, would make Rabi oscillations clearly visible in the example of Figure 7.12. In the absence of the exponential-like growth in photoconductivity that I suspect arises from carrier heating, it is likely that the lens-mounted sample would have allowed me to observe Rabi oscillations using the $2p^-$ state.

Chapter 8

Conclusions

I am interested in hydrogenic states of electrons bound to shallow donor impurities in GaAs because of their potential applications in quantum information processing. As described in Chapter 1, the relatively simple quantum mechanical nature of the states, the ease of incorporating them into semiconductor devices and the advantages of THz cavity quantum electro-dynamics are the primary motivation for my interest. In Section 2.3 I describe variational calculations of the energy levels of the hydrogenic electrons in strong magnetic fields performed by Makado and McGill [13]. The spectroscopy described in Section 6.2, as well as the work of many other researchers over the years, confirms that the observed state-transition energies closely match the calculated values, though I do observe a small offset in energy of the 1s to $2p^-$ transition.

In Section 6.3 I described experiments on the selection rules for the hydrogenic transitions. The experiments verify that the selection rules for the bare hydrogen

atom do hold for hydrogenic electrons. However, the photoconductive response to polarized radiation can be distorted by the proximity of the cyclotron transition or a parametric two-photon absorption. This experiment provided a first step toward the use of selection rules to drive selective coherent manipulations of the states and may aid in eventually designing a THz cavity that will couple to the state transitions.

In Section 6.4 I describe the observation of Rabi oscillations between the 1s and $2p^+$ states of an ensemble of electrons in unintentionally-doped GaAs. By fitting to the results using a two-level density-matrix model, described in Section 7.1, I can extract information on the ionization rate and dephasing rate of the ensemble. I find that the dephasing time of the ensemble has a maximum value of about 80 ps and that electrons remain in the excited state for as long as 25 ps before ionization to the conduction band. The experimental data and modelling strongly suggest that the dephasing time is limited by the ionization rate and/or the presence of the excited state in the conduction band.

In Section 6.4 I also describe the use of polarization selection rules to make a selective manipulation of the states. The selective manipulations are a first step toward the implementation of more complicated quantum logic operations. For example, it may be possible to implement a quantum logic operation similar to that implemented by Li, et al. [27] by using the 1s, 2p and 3s states of a hydrogenic donor. This would be a challenging, but interesting, experiment.

Since the observed dephasing times of Rabi oscillations seem to be limited by

the proximity of the conduction band to the excited state, I devoted a great deal of effort in this thesis to driving and observing Rabi oscillations between the 1s ground state and an excited state below the conduction band. I was not successful in observing any coherent transitions between the 1s and $2p^+$ states at 52.3 cm^{-1} or between the 1s and $2p^-$ states at 35.0 cm^{-1}. The numerical modelling described in Section 7.4 suggests that the maximum electric field that could be coupled in to the sample at 35.0 cm^{-1} is not sufficient for observing coherent behavior.

Data taken using a lens-mounted sample, which should significantly increase the intensity at the sample, is presented in Section 6.5. An exponential-like rise in the photoconductivity is observed and no coherent state manipulations are evident. It appears that other effects, possibly carrier heating, are dominating the sample response.

In Section 7.2.3 I presented a simulation of the three-level model where the excited state is well separated from the conduction band and the ionization rate is slow. In the example, coherent oscillations between the ground and excited states exist without evidence of such oscillations appearing in the conduction band population. Dan Allen has been working to develop a method of determining the electron orbital state by donor-bound exciton spectroscopy. When this is operational, it may be possible to undertake experiments that are similar to those I have done here but use a direct probe of the excited state population rather than a photoconductive measurement. This would make it possible to resolve coherent transitions occurring with a well-isolated excited state.

Another possibility for measuring the dephasing time of the hydrogenic electrons, and possibly even resolving coherent manipulations, is to develop a time-resolved measurement of the sample absorbtion. Sangwoo Kim will be working on this project. If he is successful, it will be possible to monitor the state transitions without making the ionization to the conduction band part of the signal.

Bibliography

[1] Tjeerd O. Klaassen, Janette L. Dunn, and Colin A. Bates. Shallow donor states in a magnetic field. In Schmelcher and Schweizer, editors, *Atoms and Molecules in Strong External Fields*, pages 291–300. Plenum Press, 1998.

[2] G.E. Stillman, C.M. Wolfe, and J.O. Dimmock. Magnetospectroscopy of shallow donors in gaas. *Solid State Communications*, 7:921–925, 1969.

[3] J. Burghoorn, J.P. Kaminski, R.C. Strijbos, T.O. Klaassen, and W.Th. Wenckebach. Generation of subnanosecond high-power far-infrared pulses by using a passive resonator pumped by a free-electron laser. *J. Opt Soc. Am. B*, 9(10):1988, October 1992.

[4] Frank A. Hegmann and Mark S. Sherwin. Generation of picosecond far-infrared pulses using laser-activated semiconductor reflection switches. In *International conference on millimeter and submillimeter waves and applications III.* SPIE, August 1996.

[5] Walter Kohn. Shallow impurity states in silicon and germanium. *Solid State Physics: Advances in Research and Applications*, 5:257–320, 1957.

[6] P. W. Shor. In S. Goldwasser, editor, *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, pages 124–133, Los Alamitos, CA, 1994. IEEE Press.

[7] D. D. Awschalom, D. Loss, and N. Samarth, editors. *Semiconductor spintronics and quantum computation.* Nanoscience and technology. Springer, 2002.

[8] C. H. Bennett and D. P. DiVincenzo. Quantum information and computation. *Nature*, 404(6775):247, March 2000.

[9] A. Imamoglu, D. D. Awschalom, G. Burkard, D. P. DiVincenzo, D. Loss, M. Sherwin, and A. Small. Quantum information processing using quantum dot spins and cavity qed. *Phys. Rev. Lett.*, 83(20):4204, 1999.

[10] M.S. Sherwin, A. Imamoglu, and T. Montroy. Quantum computation with quantum dots and terahertz cavity quantum electrodynamics. *Phys. Rev. A*, 60(5):3508–3514, November 1999.

[11] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics*, volume 1. Wiley and Sons, 1977.

[12] J. H. Davies. *The physics of low-dimensional semiconductors.* Cambridge University Press, 1998.

[13] P.C. Makado and N.C. McGill. Energy levels of a neutral hydrogen-like system in a constant magnetic field of arbitrary strength. *J. Phys. C: Solid State Phys.*, 19:873–885, 1986.

[14] R.J. Heron, R.A. Lewis, P.E. Simmonds, R.P. Starrett, A.V. Skougarevsky, R.G. Clark, and C.R. Stanley. Central-cell corrections for si and s in gaas in a strong magnetic field. *Journal of Applied Physics*, 85(2):893, 1999.

[15] D. Cabib, E. Fabri, and G. Fiorio. The ground state of the exciton in a magnetic field. *Solid State Communications*, 9:1517, 1971.

[16] D. M. Larsen. Inhomogeneous broadening of the lyman-series absorbtion of simple hydrogenic donors. *Phys. Rev. B*, 13(4):1681, 1976.

[17] A. J. Kalkman, H. P. M. Pellemans, T. O. Klaassen, and W. Th. Wenckebach. Far-infrared pump-probe measurement of the lifetime of the $2p^-$ shallow donor level in n-gaas. *Int. Journ. of Infrared and Millimeter Waves*, 17(3):569, 1996.

[18] P. C. M. Planken, P. C. van Son, J. V. Hovenier, T. O. Klaassen, W. Th. Wenckebach, B. N. Murdin, and G. M. H. Knippels. Far-infrared picosecond time-resolved measurement of the free-induction decay in gaas:si. *Phs. Rev. B*, 51(15):9643, 1995.

[19] Tune in to terahertz. *Economist*, page 67, August 8 2002.

[20] G. Ramian. The new ucsb free-electron lasers. *Nuclear Instruments and Methods in Physics Research A*, 318:225, 1992.

[21] Born. *Principles of Optics*. Pergamon Press, Inc., 6 edition, 1980.

[22] A. J. Alcock, P. B. Corkum, and D. J. James. A fast scalable switching technique for high-power co_2 laser radiation. *Appl. Phys. Lett.*, 27:680, 1975.

[23] Sorab K. Gandhi. *VLSI Fabrication Principles, Silicon and Gallium Arsenide*. John Wiley and Sons, Inc., 2nd edition, 1994.

[24] H. J. A. Bluyssen, J. C. Maan, L. J. van Ruyven, F. Williams, and P. Wyder. Mechanism of cyclotron resonance induced conductivity in n-gaas. *Solid State Comm.*, 25:895, 1978.

[25] T. H. Stievater, X. Li, T .Cubel, D. G. Steel, D. Gammon, D. S. Katzer, and D. Park. Measurement of relaxation between polarization eigenstates in single quantum dots. *App. Phys. Lett.*, 81(22):4251, 2002.

[26] B.E. Cole, J.B. Williams, B.T. King, M.S. Sherwin, and C.R. Stanley. Coherent manipulations of semiconductor quantum bits with terahertz radiation. *Nature*, 410:60, March 2001.

[27] X. Li, Y. Wu, D. Steel, D. Gammon, T. H. Stievater, D. S. Katzer, D. Park, C. Piermarocchi, and L. J. Sham. An all-optical quantum gate in a semiconductor quantum dot. *Science*, 301:809, 2003.

[28] Robert W. Boyd. *Nonlinear Optics*. Academic Press, 1992.

[29] H.S. Brandi, A. Latge, and L.E. Oliveira. Rabi oscillations, coherent properties, and model qubits in two-level donor systems under terahertz radiation. *Phys. Rev. B*, 68:233206, 2003.

Appendix A

Processing Recipes

A.1 Overall Process Order

For the lens-mounted sample used in this thesis, I did the mesa etch (step 11-13) before the anneal (step 4). I had difficulty removing the metal after the contact pad metal deposition because the gold was on top of the mesa and the photoresist was in the trench below the mesa. Therefore the gold overlapped the mesa edge and was hard to remove. The order presented here is better.

	Action
1	Two-layer photolithography of finger contacts
2	Finger contact metal deposition
3	Lift off
4	Anneal
5	4110 lithography of capping layer mask
6	Capping Layer etch
7	Lift off
8	Two-layer photolithography of contact pads
9	Contact Pad metal deposition
10	Lift off
11	4330 lithography of mesa mask
12	Mesa etch
13	Lift off

A.2 Clean

	Action	Time	Comments
1	Acetone	3 m	
2	Methanol	3 m	
3	Isopropanol	3 m	
4	DI water	3 m	
5	Blow dry with Nitrogen		

A.3 Two-Layer Photolithography

Note: Edge Bead removal (steps 8-10) is needed only for small finger features, not contact pads. The Oxygen Plasma Descum is particularly important for removing residual photoresist between small features like those of the finger contacts.

	Action	Time	Comments
1	Clean		
2	Dehydration Bake	5 m	120^O C in oven
3	OCG-825	30 s	4 krpm
4	Bake	1 m	95^O C
5	Flood Expose	6 s	
6	AZ 4110	30 s	5.5 krpm
7	Bake	1 m	95^O C
8	Edge bead expose	2.5 m	use silicon chips, well centered on sample
9	Develop AZ 400K:DI 1:4	30s	light agitation by hand
10	DI rinse	30 s	light agitation by hand
11	Pattern Expose	15 s	
12	Develop AZ 400K:DI 1:4	25 s	light agitation by hand
13	DI rinse	30 s	light agitation by hand
14	Blow dry with nitrogen		
15	Oxygen Plasma Descum		removes residual resist

A.4 Oxygen Plasma Descum

1	Open vent, load sample when at atmospheric pressure
2	Close vent
3	Turn on solenoid, wait until pressure drops to 0.030 Torr
4	Turn on O_2 Gas, pressure rises to 0.300 Torr
5	Turn on power (100W) for 30 s, then turn off
6	Turn off gas, wait until pressure returns to 0.030 Torr
7	Turn off solenoid
8	Open vent, remove sample when at atmospheric pressure
9	Close vent
10	Turn on solenoid, wait until pressure drops to 0.030 Torr
11	Turn off solenoid

A.5 AZ 4110 Photolithography

Used for capping layer etch.

	Action	Time	Comments
1	Clean		
2	Dehydration Bake	5 m	120^O C in oven
3	AZ 4110	30 s	5.5 krpm
4	Bake	1 m	95^O C
5	Pattern Expose	15 s	
6	Develop AZ 400K:DI 1:4	25 s	light agitation by hand
7	DI rinse	30 s	light agitation by hand
8	Blow dry with nitrogen		

A.6 Capping Layer Etch

The etch rate is about 600 Åper minute.

	Action	Time	Comments
1	Combine H2O2:H3PO4:DI 3:3:200 (mL)		
2	Put on stir plate to mix	5 m	210 rpm
3	Remove from stir plate		
4	Dip sample	1 min	light agitation by hand
5	DI rinse	30 s	light agitation by hand
6	Blow dry with nitrogen		

A.7 AZ 4330 Photolithography

Used for mesa etch.

	Action	Time	Comments
1	Clean		
2	Dehydration Bake	5 m	120^O C in oven
3	AZ 4330	30 s	4 krpm
4	Bake	1 m	95^O C
5	Pattern Expose	30 s	Use pattern closest to “3” of “Bryan3”
6	Develop AZ 400K:DI 1:4	1 m	light agitation by hand
7	DI rinse	30 s	light agitation by hand
8	Blow dry with nitrogen		

A.8 Mesa Etch

The etch rate is about 9156 Åper minute. To avoid damaging the samples with the stir bar, I recommend putting the beaker on the stir plate so that the stir bar is all the way to one side of the beaker. Then put the samples on the opposite side of the beaker. Because the flow of etchant is not uniform to all areas, rotate the samples by 180^o halfway through the etch to end up with a uniform etch.

	Action	Time	Comments
1	Mix H2O2:H3PO4:DI 8:40:112 (mL)		
2	Put etchant on stir plate to mix	5 m	
3	Etch mesa	16 m 23 s	rotate sample at ∼8 m
5	DI rinse	30 s	in DI beaker under slow tap
6	Blow dry with nitrogen		

A.9 Anneal

Use rapid thermal annealer. My recipe is named 510.rcp and anneals for 2 minutes at 510^O C in forming gas.

A.10 Finger Contact Metal Deposition

I used E-Beam 3 to do this.

	Thickness	Metal
1	150 Å	Ge
2	350 Å	Au
3	100 Å	Ni
4	3000 Å	Au

A.11 Contact Pad Metal Deposition

I used E-Beam 3 to do this.

	Thickness	Metal
1	100 Å	Ti
2	3500 Å	Au

A.12 Lift Off

	Action	Time
1	Acetone soak	10-15 m
2	Spray acetone over features until metal comes loose	
3	Isopropanol	3 m
4	DI Water	3 m
5	Blow dry with nitrogen	

Appendix B

Regenerative Amplifier Alignment

As described in Section 3.4, the NIR regenerative amplifier ("regen") system consists of four separate lasers. This appendix is not intended to be a substitute for the Spectra-Physics manuals for the three lasers that pump and seed the regen cavity itself. For those lasers, I intend only to provide supplementary information on the operational settings and any additional tricks or procedures not in the manual that I have learned from service engineers. For the regen itself, there is no actual manual. There are notes originally written by Spectra-Physics and supplementary notes written by Jon Williams. These notes do not, in my opinion, do a sufficient job of detailing how to align the regen, which is necessary on a regular basis. My notes on how to do this alignment therefore form the bulk of this appendix.

All four of the lasers in this system are Class IV lasers, which means that direct radiation can damage your skin and both direct **and scattered** radiation can damage your eyes. It is therefore extremely important to follow good laser safety procedures. For example, always remove all watches or jewelry that could possibly reflect before putting your hands anywhere near the beam. We have laser-safety glasses that block the 532-nm light from the Millenia and Quanta-Ray. These probably will **NOT** block a direct illumination, but they are a good safety precaution for scattered light.

B.1 Millenia

The Millenia laser is a CW diode-pumped Nd:YAG laser that produces up to 5 W of output power. To operate the laser, simply turn on the cooling water, hit the power button on the remote to warm up the diode bars and then hit the power button again to turn on the power once the diodes are warm. From that point on, only the total power output can be adjusted by varying the current going to the diodes. Normally we run this laser in a power-control mode where the desired output power is set and the laser automatically adjusts the current to maintain that output power.

B.1.1 Maintenance

There is very little maintenance that can be done on the Millenia. Most repairs have to be done by a service engineer. Here are just a few things to keep an eye on:

Cooling water The cooling water for the Millenia should be set for 18^o C. You should periodically check to make sure that the chiller is filled to the appropriate level with DI water. The cooling water loop actually runs through the Millenia, Tsunami and regen cavities, so periodically checking this is important to all three lasers. I've installed a little flow meter under the regen table so you can check if the water is actually flowing.

Diode bars There are two diode bars in the Millenia and they are extremely expensive. Fortunately, they are expected to last approximately 8000 hours. The laser electronics automatically adjust the temperature and operating current over the life of the diodes. When the diodes die you have to call Spectra-Physics to replace them. Spectra-Physics did produce a bad production run of diodes that lasted only about 500 hours, but they replaced those diodes for free and I have had no problems with the current diodes.

SHG temperature The SHG crystal is used to double the 1064-nm light produced by the Nd:YAG crystal to the desired output of 532 nm. The oper-

ating temperature of this crystal is very important to stable power output. The instructions for changing this temperature are in the manual. In early 2004 we started to have a problem with the Millenia requiring a very long time to start up to high power. If you turned the laser on and requested 4.5 W of output power, the laser would not be able to ramp the power up to that level in a sufficiently short time and therefore flagged error 146: power adjust timeout. This problem was completely solved by adjusting the SHG crystal temperature.

B.2 Tsunami

The Tsunami laser is a mode-locked NIR laser that operates at 76 MHz with an output center frequency and bandwidth that can be adjusted. The temporal width of the output pulses is about 150 fs; the temporal and spectral characteristics of the Tsunami pulse determine the temporal width of the output regen pulse. For the regen, you want to set the Tsunami so that the output pulses have a center frequency of 795 nm with a bandwidth from shoulder-to-shoulder of about 10 nm. This can be checked by fiber-coupling the output over to a spectrometer in the electro-optics setup.

The lasing action in the Tsunami is passive, so it will start as soon as the

Millenia pump beam is turned on. The mode-locking of the Tsunami will not start spontaneously; it can be turned on with the control box for the Tsunami, which sends a signal to the AOM in the Tsunami cavity. 4.5 - 4.65 W of Millenia pump power is necessary to achieve mode-locking. When operating the regen, we want to run the Tsunami in passive mode-locking mode. To do this, simply turn on the mode-locking from the control box and then turn it back off. The Tsunami will remain passively mode-locked for at least 8 hours.

B.2.1 Aligning the Cavity

The manual for the Tsunami does not contain good instructions for realigning the Tsunami cavity if it does not lase properly. This does not happen often. I got these instructions from a service engineer during a visit to work on the Tsunami. All references to components (M5, P3, etc.) refer to the cavity diagram in the Tsunami manual, on p. 4-5 and 4-6. Also see p. 7-6.

Note 1 Virtually all problems with the Tsunami (mode-hopping, bad mode-locking) arise from dirt on the optics. There is a nitrogen purge on the cavity, which helps to keep dirt out. The optics most likely to get dirty are M7, M8 and the prisms (since they face upwards). To clean them, use a methanol-soaked lens tissue to wipe across the surface.

Note 2 Most of the optics in the cavity are extremely stable because everything is mounted directly to an Invar steel bar that has zero coefficient of expansion and is therefore extremely stable against temperature fluctuations. After cleaning, only M1, M10, the slit and the prisms are likely to need adjustment (and possibly any component that was cleaned). The alignment procedure below should only be used in full if the laser is very badly misaligned (for example, it doesn't lase).

Note 3 M6, M7, M8, M9 are very stable and their alignment is quite critical (since it affects the passage through the prisms). Leave these components alone. You probably only need to adjust M1, M5, M10 and the slit and prisms.

1. Verify that the Millenia is operating properly with an output power of $\geq$ 4.5 W and is getting into the Tsunami box.

2. Align the Millenia injection:

   (a) Center the beam on P1 using the position of the Millenia.

   (b) Center the beam on P2 using P1.

   (c) Center the beam on M3 using P2

   (d) Focus the beam on the center of the crystal by walking P1 and P2.

3. Use the Millenia light (532 nm) that passes through the crystal to align M1 and M2. The beam should be centered on these optics.

4. Clean (with Methanol) the faces of the rod, any dirty mirrors and the prism faces.

5. Using the IR viewer, check that the beam is centered on M4 and M5.

6. Place an index card after M10 and look at the "glow" of M10. Center the beam in this glow by adjusting M5. Note that the beam passes off M6-9, but they are not adjusted.

7. Place an index card after the AOM and before M10. Looking at the card, you may be able see a faint cross-hair that indicates the center of the AOM. (I have never been able to see this). Center the beam in the glow of the AOM (or on the cross-hair) by adjusting M1.

8. Cut a small slit ($\sim$ 2 mm wide) into the middle of a business card, with the slit running along the long axis. Place this card into the space between the mounting blocks so that it stands up in the path of the light between M4 and M5. Position the card so that the light coming from M4 is centered on the slit. Then look at the opposite side of the card and adjust M10 so that the light coming back off of M5 also hits the center of the slit. You may

need to use the IR viewer to see the light. This step overlaps the forward and return beams so that the cavity forms a closed path.

9. At this point, you should have lasing. If not, go back and check all the previous steps.

10. Monitor the output power with a power meter and peak it up by adjusting M1 and M10. First adjust each single horizontal or vertical component individually and then walk the horizontal and vertical pairs.

11. Peak up the signal to the mode-lock monitoring photodiode by first walking the beam with M5 and M10 to maximize the bars of photodiode signal shown on the controller box without losing the lasing action. Then unscrew the photodiode card and adjust its position slightly to get the maximum number of bars.

12. Check for good mode-locking. It should be possible to move the slit position around without losing the mode-lock. (Note that the phase knob on the control box does nothing when the system is passively mode-locked. When actively locked, adjusting the phase can cause the laser to lose mode-locking.)

13. Verify that the slit width is about 3 mm. Otherwise don't adjust this width.

Note that adjusting the slit width can shift the center frequency.

14. Using a spectrometer, adjust the tuning slit position to get the desired center frequency.

15. Use the prism adjust to control the bandwidth of the pulse. This may also shift the center frequency, but you can use the slit position to get it back where you want it.

B.2.2 Maintenance and Tricks

The least stable component in the Tsunami is M10, the cavity end mirror and output coupler. This should be your first adjustment if anything happens.

Cooling water See Millenia notes.

Power output When pumped by 5 W of green (532 nm) from the Millenia, the Tsunami normally puts out between 730 and 760 mW. Although the total power output is not really critical to the regen, the mode-locking tends to drop out if the total power drops too far. Adjust M10 and possibly M1 to regain power.

Bad Mode-locking Problems with mode-locking are almost always a result of low power in the cavity. If fixing that doesn't work, make sure that no re-

flections from the Tsunami beam are getting back into the cavity as these back reflections can cause the laser to frequency lock onto the spectral component that is getting back into the cavity. The Faraday Isolator in the regen should prevent back reflections, but sometimes it isn't perfect if you intentionally retro-reflect the beam.

B.3 Quanta-Ray Nd:YAG

The Quanta-Ray is a Nd:YAG laser. It is just like the Millenia, but operates at 10 Hz using flashlamps instead of diode bars. The build up in the lasing cavity is dumped using a Q-switch to get high power short pulses. These are used to pump the gain crystals in the regen.

There is essentially no alignment on the Quanta-Ray cavity that you can do yourself. There are, however, some basic maintenance and operation tips. When operating properly, the Quanta-Ray should put about 2.6 W into the regen cavity.

We run the Quanta-Ray, and thus the regen, at 10 Hz because we want high power in each pulse. There are kHz regen systems that are more stable, but don't produce as much power per pulse. See the description of synchronization in the chapter on the pulse slicer for more details about the 10-Hz signal that is sent to the flashlamps and the Q-switch.

B.3.1 Maintenance

Cooling water The cooling water for the Quanta-Ray is self-contained. You can see the water level on the back side of the main electronics box for the laser. If it gets low, fill it with DI water.

Flashlamps The Flashlamps are supposed to be good for somewhere upwards of 40 million flashes. Tom King installed a counter so we can keep track of the actual number of flashes. I purchased a replacement set of flashlamps since the ones we had were getting old, but as of this writing they were still working and I have not installed the replacements. See the manual for information on how to replace the lamps and be careful with the high voltage.

Flash cavities One thing that can cause the power output of the laser to go down even if the lamps are fine is dirt on the flash cavities. The flashlamps are at one focus of an elliptical reflective cavity and the Nd:YAG crystals are at the other focus. If the reflective cavity gets dirty, not all the light makes it to the crystals. To clean these cavities, follow the instructions in the manual to disconnect the high voltage, drain the cooling water and open the cavity. The cavities can then be cleaned with regular white vinegar and cotton swabs.

Peak and Hold The total power output of the Quanta-Ray is partially determined by the voltage going to the flashlamps. If the total power output of the laser is low, you can turn up this voltage by adjusting the PFN cal (R175 on the main circuit board) while monitoring the power. The flashlamp charging voltage can be monitored between TP3 and TP5 on the board. The voltage across these pins should be roughly 7.45 V and should never exceed 8 V. I once had a problem where this voltage seemed to jump back and forth between two values. The problem went away when I took out and then reinserted the main circuit board, so I presume it was a loose connection.

A Warning I once turned up the PFN cal to adjust for a sudden drop in the charging voltage and thus the power output of the laser. When the charging voltage suddenly jumped back up, it exceeded its normal level. The resulting laser output was too high and damaged the face of one of the gain crystals in the regen. If you see a sudden change like that, try to diagnose the problem before adjusting the power.

B.4 Regen

The regen takes a single output pulse from the Tsunami, chirps it so that the pulse width is roughly 200 ps and uses that pulse to seed a lasing cavity that is

pumped by the Quanta-Ray. The buildup of this lasing cavity is dumped, passed through a double-pass amplifier and then re-compressed to get the high energy, short-pulse output. The alignment of the regen is very critical and I find I have to make adjustments every few months. I got the full alignment procedure from Joe Jueneman, the Spectra-Physics service engineer who works on the regen systems. He's extremely good at working with the systems and also a very nice guy. If you have a problem, do not hesitate to call or email him.

B.4.1 Aligning the Cavity

Note that there are several potentially dangerous steps in this alignment procedure. Take proper precautions to protect your eyes and those of anyone else in the room. All of the part references in this section refer to my diagram of the regen optics, which is shown in Figure B.1. This alignment procedure could be used to align the regen totally from scratch. Usually things aren't that bad, but if you're going to do any major adjustments to the system, it is good to start at the beginning and check all these things.

1. Make sure that the Tsunami (red) is passing through P1 and P2. If necessary, adjust the mirror right outside the cavity and M1 to get the beam centered.

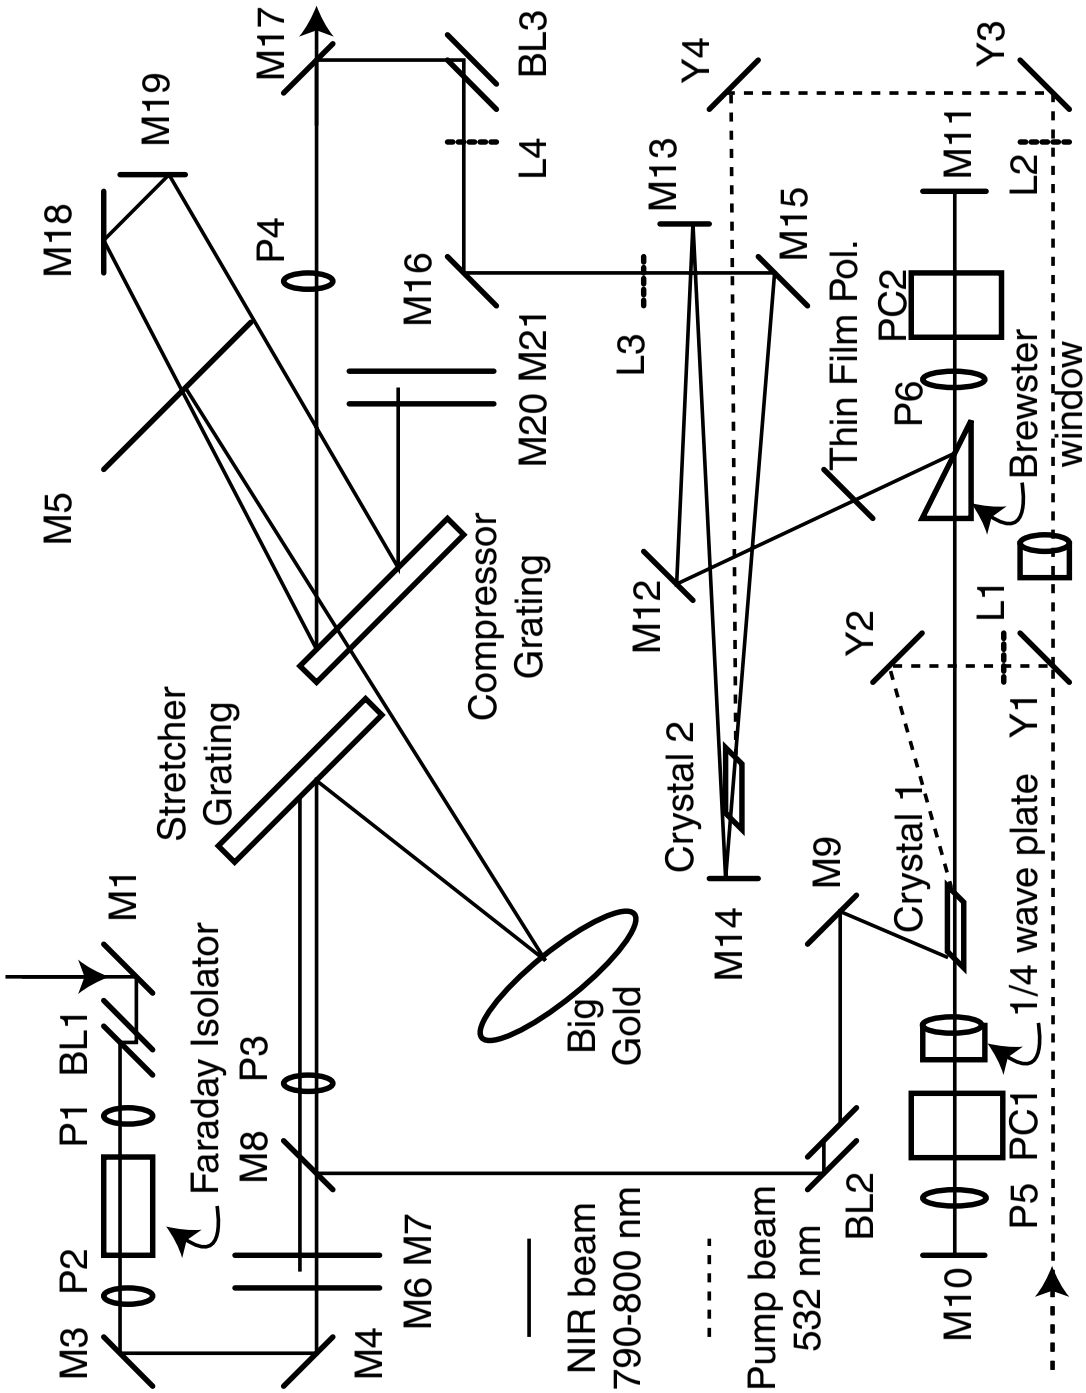

Figure B.1: A diagram of the optics in the regen.

2. Put the red through P3 and P4 by adjusting M3 and M4.

3. Remove the grating.

4. Check that the M6-M7 assembly and M20-M21 are properly aligned. Note that this should not have to be done more often than every few years. Don't do this step unless a total realignment is necessary.

   (a) Remove the assemblies from where they are screwed down and clamp them down in the path of the P3-P4 beam.

   (b) Check each mirror individually by screwing the assembly into the wrong hole of the mount and checking for reflection in the vertical plane. See Figure B.2.

   (c) Reattach the assembly into the proper holes of the mount and check that the reflection off both mirrors is a beam on axis at 3.5" by moving the alignment jig back and forth with the incoming red on the center hole. The beam off the mirror assembly should stay on the top hole.

5. Check that the gratings are parallel, not rotated and 180^o opposed.

   (a) Turn down the Millenia so the the Tsunami is NOT mode-locked.

   (b) Using the Oth order reflection, check that the grating is reflecting the beam in the same horizontal plane as the incident beam. See Figure

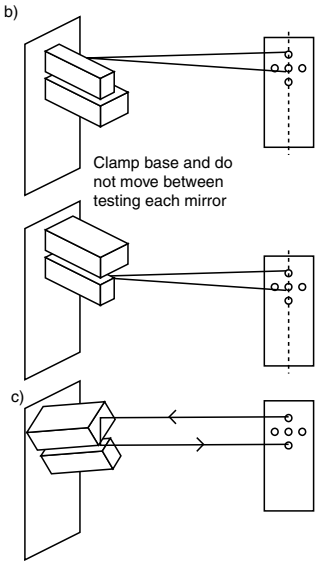

Figure B.2: How to check the M6-M7 or M20-M21 assemblies.

B.3. Be careful not to put the beam straight back through the pinhole, which can cause the Tsunami to frequency lock. This is annoying but not dangerous. If there is a vertical offset between the incident and reflected beams, adjust the grating's vertical tilt to correct it.

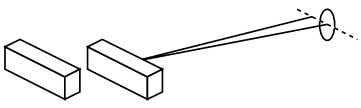

Figure B.3: How to check alignment of the gratings.

(c) Rotate the grating stage so the first-order reflection goes back to the

pinhole (again watch the frequency locking). Verify that this is also in the same horizontal plane as the pinhole. Correct any vertical offset by loosening the grating housing and rotating the grating.

(d) Repeat the last two steps for the other grating.

(e) Rotate the stage so that one grating is reflecting directly back to the pinhole (or a fixed point near the pinhole). Read the stage angle using the Vernier scale. Rotate so the other grating is now reflecting to the same point and read the stage angle again. If these two numbers are not separated by 180^o, adjust one of the gratings horizontally until they are separated by 180^o.

6. Loosen the mount holding the grating rotation stage to the base of the regen. Rotate the gratings to roughly the correct angle and then position the entire assembly so that the incident beam hits the left-right center of the grating and the reflected beam is centered on the Big Gold mirror. Clamp the base.

7. Adjust the micrometer on the Big Gold mirror so that the reflection comes back onto the grating. Adjust the angle of the Big Gold mirror so that the reflection is aligned in the same vertical plane as the first spot where the incident beam hits the grating.

8. Adjust the micrometer on the Big Gold mirror so that the beam once again

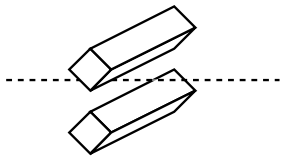

Figure B.4: The position of the beam through M6-M7.

goes up to M5 and back. Look at P3. The beam returning from M5, off the Big Gold mirror and off the grating should line up in the same vertical plane as the center of P3. If it is off, adjust the angle of M5 until it is right.

9. Reinsert the M6-M7 assembly and mount. Center the assembly around the beam so that the beam passes through the left-right center of the assembly and does not clip on the edge of either mirror. Adjust the angle of the the assembly so that all four spots on the grating are in a vertical line. See Figure B.4

10. Mode-lock the Tsunami and fine tune the angle of the grating to center the stripes. This is necessary because the non-mode-locked beam is not necessarily at the center of the mode-locked frequency spectrum.

11. Start up the Quanta-Ray and allow about 20 minutes for warm up.

12. Check the unseeded build-up time (block the red before Crystal 1) of the the lasing cavity using the photodiode after M11.

(a) Adjust Y2 to shorten the unseeded build-up time.

(b) Adjust the cavity end mirrors (M10 and M11) to shorten the unseeded build-up time.

(c) If (and only if) there is no unseeded build up, center the green pump beam from Y2 on the face of Crystal 1. Look at the spontaneously emitted NIR with the viewer and adjust M10 to get the NIR onto P6. Then adjust M11 to get the return beam on P5. Scan carefully until you get lasing.

(d) Once you have lasing, check to make sure that the beam is on P5 and P6. If not, walk M10 and M11 to center the beam on the pinholes.

13. Let the Tsunami seed beam into the cavity, Adjust M9 until the seed beam is on P6. Look at the seeded build up time and walk M8 and M9 to optimize this time. Be careful not to interrupt the stretcher by blocking a part of the path of the wide beam when you try to adjust M8. You will probably need to use an Allen wrench. It is okay to interrupt the path around P2; the danger comes if you block part of the wide beam, thus distorting the stretcher. This can result in sending an unstretched pulse into the crystal and possible damage to the crystal.

14. Adjust Pockel's Cell 2 timing to dump a pulse out of the lasing cavity.

15. Adjust M12 so that the output beam is hitting the edge of M13.

16. Remove the cover from the second gain crystal. Hold a card between the crystal and M14 and adjust M13 so that the beam is centered on the crystal. You can identify the edges of the crystal because the beam will scan off the crystal and disappear.

17. Adjust M14 so that the return beam is hitting the edge of M15. Use a card to check that the beam leaving M13 for the crystal and the beam coming back from the crystal are at the same height. Adjust M14 if they are not. This is done so that the solid angle between the beams going to and coming from M14 is as small as possible. This way there can be maximum overlap between the crystal area actually pumped by the green and the two paths through the crystal.

18. Block the beam and remove the expansion telescope (L3 and L4)

19. Block the Tsunami seed beam. With the Quanta-Ray firing, the cavity will lase. Dump from the end of the unseeded gain curve to get a weak pulse to use for alignment of the compressor:

    (a) Adjust M15 so that the beam hits the middle of M16 at 2.5" (use the little pinhole on a stand).

(b) Adjust M16 so that the beam is level at 2.5" and hits near the center of the bottom mirror of BL3.

(c) The beam should now reflect off BL3, hit M17 and reflect toward the compressor. Remove the gratings so that you can see P3 and P4.

(d) Walk M16 and M17 so that the beam is on the axis defined by P3 and P4.

20. Replace L4 and adjust its position so that the beam is still on the P3-P4 axis. You will need to adjust the angle, horizontal position and vertical position of the lens.

21. Replace L3, again adjusting angle and position so that the beam remains on P3-P4.

22. Replace the gratings.

23. Take out the M18-M19 assembly.

24. Place a card where the M18-M19 assembly was. Using the IR viewer to look at the wide beam on the card, wiggle M9 vertically. If the pulse is unevenly chirped (some spectral components more intense than others) you will see a brighter portion of the beam that moves left and right as you wiggle M9 vertically. If the chirp is even, the beam will get brighter and

dimmer uniformly as you wiggle M9 vertically (the uniform change will also happen if you wiggle M9 horizontally). If the pulse is unevenly chirped:

(a) Note if the bright spot moves left or right with a CCW rotation of M9 vertical.

(b) Unclamp the M5 assembly and slide it either forward or backward. Adjust the angle of the assembly so that the spots and stripes on the grating are always properly aligned (vertically). You should probably block the seed beam into the cavity while moving M5 to prevent unstretched pulses from getting in.

(c) Continue to adjust the M5 position until the chirp is even. If M5 is moved to the other side of the ideal position, the direction of the motion you identified in step (a) will be reversed.

25. Replace the M18-M19 assembly.

26. Remove the M20-M21 assembly.

27. Verify the alignment of the M18-M19 assembly. Note that this should not have to be done more often than every few years. It is a dangerous step because it involves intentionally directing beams out of the regen cavity.

(a) Unscrew the rail on which M18-M19 are mounted and prop up the front

end on a thin block so that the reflected beams off M18-M19 go out of the cavity and onto some kind of screen (tape up a piece of paper).

(b) Position the rail and assembly so that the stripe coming to the assembly from the grating is centered on the apex of the assembly. See Figure B.5a.

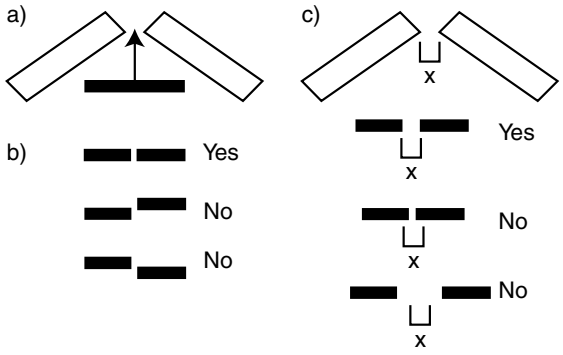

Figure B.5: a) The position of the stripe at the apex of M18-M19. b) The two halves of the stripe should be in plane. c) The gap should be equal to the apex hole.

(c) Look at the reflections on the screen with the IR viewer. Adjust the vertical of either M18 or M19 so the two halves of the stripe are in plane. See Figure B.5b.

(d) Adjust the M18-M19 horizontal spacing so that the gap between the two halves of the stripe has the same width as the apex hole. If the apex hole is negligible, adjust so that the two halves are just touching. See Figure B.5c.

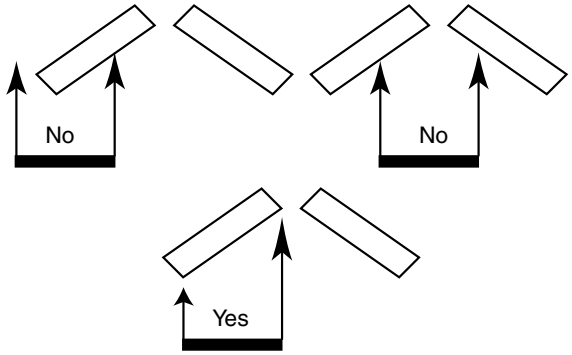

Figure B.6: Stripe should hit M18 right next to the apex.

(e) Reattach the rail and replace the assembly.

28. Adjust the position of the M18-M19 assembly (i.e. the position of the mirror assembly relative to the mount that rides the rail) so that the stripe from the grating hits M18 right next to the apex. See Figure B.6.

29. There should now be a spot on the grating from the beam coming from M17 and a stripe from the wide beam returning from M19. If they are at slightly different heights, adjust the compressor grating vertical to put them at the same height. See Figure B.7. Although we checked earlier (Step 5) that the gratings were both reflecting in-plane, the rotation stage is not necessarily completely flat and so a fine adjustment here may be necessary to compensate.

30. Replace M20-M21 assembly and position it at a proper height to capture

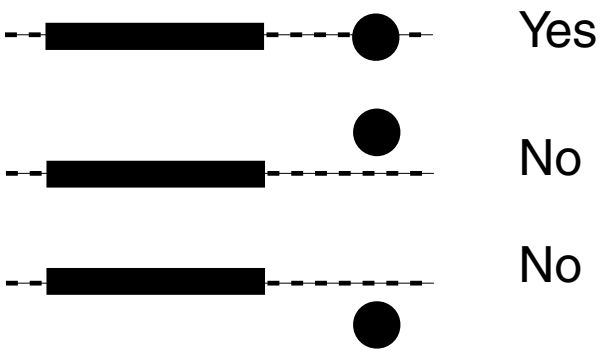

Figure B.7: Spot and stripe should be even.

the stripe from the grating on the lower mirror. Position the assembly horizontally as close as possible to the beam from M17 to the grating without clipping. See Figure B.8a.

31. Adjust the angle of the M20-M21 mount so that the compressed beam that comes off the grating is in the same vertical plane as the beam coming to the grating from M17. See Figure B.8b.

32. Monitor the power of the compressed beam with a properly seeded and dumped cavity. Adjust Y4 to optimize the output power.

33. Use a prism to look at the output compressed pulse. When the compressor is set properly, the temporally-short, spectrally-broad pulse will generate a white light flash on passing through the prism. Adjust the position of M18-M19 along the rail until you see this white light generation. You could use

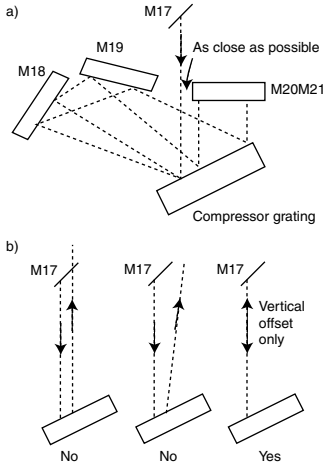

Figure B.8: The position of M20-M21.

an autocorrelator to do this, if you have one. The correct spacing between the compression grating and M18-M19 is (Distance between Big Gold and M5) minus (Distance between Big Gold and stretcher grating). You can use this as a rough positioning measurement to get close and then fine tune with the white light generation.

34. Measure the output power and check the mode quality by looking at the spot shape. A funny shaped mode (particularly sharp edges) is often the result of a beam clipping somewhere. Likely candidates are M17, M18-M19 and M20-M21. Check these and adjust their position if necessary to get the full beam on each of the mirrors.

35. At best, the compressor is about 70% efficient. This is measured by comparing the power before M16 to the total output power of the regen.

36. You have now finished aligning the regen. I recommend a celebratory frothy beverage.

B.4.2 Maintenance and Tricks

Cooling water See Millenia notes.

Cavity stability The build-up time in the oscillator cavity depends strongly on any flexure in the regen cavity. The weight of the metal lids does cause some flexure and the cavity is optimized to run when the lids are on. If you have taken off all the lids, you may find that the build up is very long, or unstable. I have had a special frame made that weighs the same amount as the right most lid so that you can put this on to keep the cavity flat while adjusting components inside.

Adjusting PC 1 The timing of Pockel's Cell 1 is set so that one and only one pulse from the Tsunami is in the cavity providing a seed for the oscillator. The timing of this should not need to be varied unless you make a drastic change to the system, like moving the Tsunami to a new location. It is currently set at 12 ns. If you do need to adjust this time, look at the seeded

build up. As you scan the PC 1 time, you will sometimes see two separate pulse trains in the cavity. Set PC1 for a time that is in the middle of the range where there is only one pulse.

Adjusting PC 2 Pockel's Cell 2 is adjusted routinely throughout any given run day to dump the maximum output of the oscillator cavity. In fact, for most stable output, you usually want to dump one pulse AFTER the maximum. You can normally see that this has the most stable amplitude on the scope. As you decrease the PC2 delay, you will see that the pulse you are trying to dump drops down, rises again slightly, and then drops away for good. To get the best pre-pulse suppression (i.e. maximum ratio between the output pulse and leak-through pulses that come out before PC 2 fires), set PC2 to fire in the the first dip (i.e. the dip corresponding to a slightly longer PC2 delay).

Thin-Film Beam Splitter Pre-pulses from the regen cause problems with the silicon switches because they can generate an electron hole plasma that, while not dense enough to reflect, can absorb and thus distort the wafer interference effect that we use to get good contrast ratios. To suppress these pre-pulses I added a thin-film beam splitter after the Brewster window and before M12. This passes only vertically-polarized light and so helps

to suppress any pre-pulses leaking out of the oscillator cavity, which are horizontally polarized.

Appendix C

Pulse Slicer Operation

In Section 3.3 I described the operation of the UCSB Pulse Slicer, which involves temporally and spatially overlapping the THz beam arriving from the FEL with the NIR pulse arriving from the regenerative amplifier. It should come as no surprise that this is a tricky thing to do. In Section 3.3 I tried to focus on the overall theory of the operation without including too many details that might clutter the description. However, many of these details are very important to actually operating the system, so in this Appendix I will attempt to provide a more comprehensive description of the alignment and operation of the system.

As I said in Section 3.3, the pulse slicer was originally built by Frank Hegmann. Many of the things I will describe here were originally figured out or set up by Frank. Bryan Cole was a postdoc who overlapped and then succeeded Frank.

He also gets credit for devising many of the tricks I use when working with the system. Neither Tom King nor I overlapped with Bryan and so what we know about the system comes from word of mouth and Bryan's and Frank's notebooks. I hope that this thesis will bring together everything I have learned and discovered about the system so that it can serve as a fairly comprehensive description.

I will break the appendix into two separate pieces for the sake of clarity. First there is a section about the "distribution optics" that take the NIR pulse from the regen, split it into three beams (one for each switch) and establish the correct time delay between each beam. Then comes the pulse slicer itself, which consists of the THz optics and the switching wafers. For each of these pieces there will be both a description and an alignment procedure.

C.1 Distribution Optics

A diagram of the distribution optics is given in Figure C.1. It is difficult to draw an accurate diagram since the distribution optics reflect the beam in three dimensions. There are two elements of Figure C.1 that change the vertical height of the beam: beam ladders (BL) and beam towers (BT). A beam ladder accepts a beam at a low height and sends out a beam at a larger height. A beam tower accepts a beam at a high height, passes it through a lens and sends out a beam

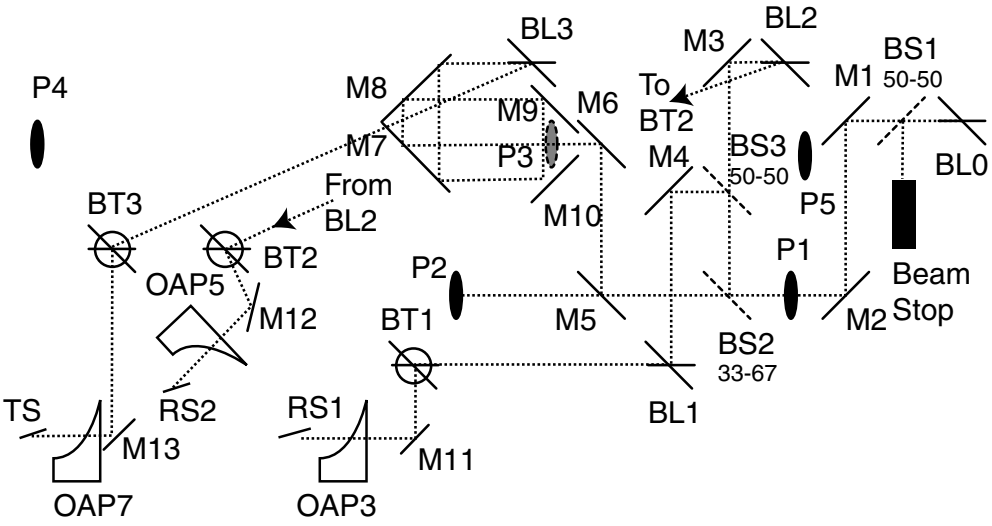

Figure C.1: A diagram of the distribution optics.

at a low height.

At each of the switches of the pulse slicer, the NIR radiation activates the wafer, making it reflective. Since the FIR is incident at a grazing angle (Brewster's Angle), different parts of the wave front hit the wafer at different times, as shown in Figure C.2. If the NIR were to be incident normally, it would activate all of the wafer simultaneously, and therefore activate different parts of the wave front of the FIR. To eliminate this effect, the NIR is brought onto the wafer collinearly with the FIR; the wave front of the NIR activates the wafer in synch with the wave front of the FIR. To accomplish this collinear NIR excitation, a tapered hole

is drilled in each OAP before a switch. The hole is large at the backside of the OAP and tapers to a small pinhole on the front side. The lens in each beam tower is positioned so that the NIR focus is at the pinhole in the OAP. The NIR then expands slightly before hitting the wafer. Since the FIR beam diameter is quite a bit wider than the hole, there is minimal loss of THz intensity when it reflects from the OAP.

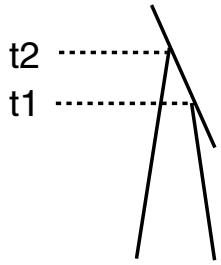

Figure C.2: One side of the FIR wavefront hits the wafer at time t1, the other at t2. To reflect a flat wavefront, the NIR should activate the different sides of the wafer with this same time difference.

C.1.1 Distribution Optics Elements

All labels refer to Figure C.1.

BS1 This beam splitter transmits 50% of the incident power, which goes into the rest of the distribution optics. The reflected 50% goes into a beam dump. The beam splitter is used because the pulse slicer normally needs only about 5 mJ, rather than the full power of the regen (10 mJ), to operate well. At

high powers, the optics can be damaged, particularly those at the end of long optical paths (like M13). This is especially likely if the output of the regen is not well collimated, as self-focusing can reduce the beam diameter and thus increase the power density. However, running the regen below maximum power often causes the total power output to be unstable. Thus, BS1 allows us to run the regen at full power while still allowing only the necessary power into the slicer. If the regen power drops over the course of a day, this beam splitter can be removed to regain sufficient power to the pulse slicer. I have found that the beam splitter introduces negligible shifts to the transmitted beam, so it can be removed without realigning any subsequent optics.

BS2 This beam splitter transmits 33% of the incident power, which eventually goes to the Transmission switch. The reflected 67% goes to BS3.

BS3 This beam splitter transmits 50% of the incident light, which goes to reflection switch 2. The reflected 50% goes to reflection switch 1.

BS3 - M4 These two optics are mounted on a translation stage. Adjusting the position of this stage determines the relative arrival time of the beams at reflection switch 1 and reflection switch 2. I have found that this is extremely stable and will only need adjustment in the case of drastic changes to the

slicer.

M5 This mirror is on a flipper mount that allows you to tip it out of the way to allow the beam to pass over the mirror and go to P2. The flipper mount is kinematic and returns to the same position extremely reliably as long as you are gentle.

M7-M8 These mirrors are mounted on a computer-controlled delay stage. The position of this stage determines the relative timing between the beams to the reflection switch and the beam to the transmission switch. The position of this stage therefore determines the length of a sliced pulse.

BL# At each of the beam ladders, the NIR is reflected off the bottom mirror and goes up to the top mirror. If the top mirror were damaged or if the beam missed the mirror, this would result in NIR light shining straight up toward the metal frame over the table. To avoid any possibility of an uncontrolled beam like this, you will see that I have placed either a metal plug or a piece of index card on the top side of the upper mirror.

C.1.2 Distribution Optics Alignment

The following procedure is used to align the distribution optics. You should remember that the NIR beam coming out of the regen is extremely intense and

quite capable of damaging your eyes. You should therefore take all the usual precautions for working with a Class IV laser whenever working with the distribution optics. A full realignment like this is rarely necessary, and you will probably only have to periodically perform a few steps to realign a particular component. However, I will present the full procedure assuming that the system has been totally misaligned. On a given day of experiments, you will only want to do a quick check of the alignment, so I will follow this full procedure with a "quick check" version. All names refer to components labelled in Figure C.1.

1. The output of the regen should hit the bottom mirror of BL0. If necessary, adjust this bottom mirror so that the beam is reflected vertically to the top mirror of BL0.

2. Take out BS1 and adjust the top mirror of BL0 so that the beam hits the center of M1.

3. Insert BS1, being careful not to reflect an uncontrolled beam. Adjust the angle of BS1 so that the reflected beam is being safely dumped in the beam stop.

4. Adjust M1 so that the beam hits the center of M2.

5. Tip M5 down so that it is out of the optical path and place a card over BS3

so that no light gets through it.

6. P1 and P2 define the desired axis of injection into the rest of the distribution optics. Adjust M1 and M2 so that the beam is centered on P1 and that when P1 is closed, the spot is centered on P2.

7. Adjust BS2 so that the light reflected from it hits the center of M3. The beam will pass through BS3 to get to M3.

8. Adjust M3 so that the beam hits the center of the lower mirror of BL2.

9. Adjust the lower mirror of BL2 so that the beam goes vertically up to the upper mirror.

10. Adjust the upper mirror of BL2 so that the beam hits the center of the upper mirror of BT2.

11. Close the pinhole in BT2, which identifies the center of the lens. The beam reflecting from the top mirror of BT2 should be centered on the pinhole. If necessary, adjust the top mirrors of BL2 and BT2 so that the beam is centered without clipping on the top mirror of BT2.

12. The beam should now be hitting the center of the bottom mirror of BT2. Adjust this mirror so that the beam is centered on M12.

13. Adjust M12 so that the beam is going into the hole in the back of OAP5. See the last step for instructions to align the NIR through the pinhole in the front of OAP5.

14. Adjust the angle of BS3 so that the beam is centered on M4. Adjust M4 so that the beam is centered on the bottom mirror of BL1.

15. Adjust the lower mirror of BL1 so that the beam goes vertically up to the upper mirror.

16. Adjust the upper mirror of BL1 so that the beam hits the center of the upper mirror of BT1.

17. Close the pinhole in B12, which identifies the center of the lens. The beam reflecting from the top mirror of BT1 should be centered on the pinhole. If necessary, adjust the top mirrors of BL1 and BT1 so that the beam is centered without clipping on the top mirror of BT1.

18. The beam should now be hitting the center of the bottom mirror of BT1. Adjust this mirror so that the beam is centered on M11.

19. Adjust M11 so that the beam goes into the hole in the back of OAP3. See the last step for instructions to align the NIR through the pinhole in the front of OAP3.

20. Adjust M5 so that the beam transmitted through the back of BS2 hits the center of M6.

21. Remove the kinematic base holding M9 and M10 in position and replace them with P3. Remove the assembly holding M7 and M8 on the translation stage.

22. Use M5 and M6 to adjust the beam so that it is centered on P3 and P4. This ensures that the beam is travelling parallel to the direction of motion of the translation stage.

23. Replace the M7 and M8 assembly and remove P3. The beam now passes by the edge of M6 and hits P5. P5 is used to check if the M7-M8 assembly is retro-reflecting the beam. Move the delay stage until it hits the X-axis limit closest to P5. Adjust M7 to center the beam on P5. Move the delay stage to the other end of its range (farthest from P5). Adjust M8 to re-center the beam on P5. Repeat this process until the beam remains centered on P5 at all delay stage positions.

24. Place the M9-M10 assembly on the kinematic base in front of P2. Lower M5. Adjust M9 and M10 so that the reflected beam is parallel to the incident beam (same height and constant horizontal spacing).

25. Replace the M9-M10 assembly near the delay stage and raise M5. The beam now makes two round trip passes through M7-M8 and finally comes off to the lower mirror of BL3.

26. Adjust the lower mirror of BL3 so that the beam goes vertically up to the upper mirror.

27. Adjust the upper mirror of BL3 so that the beam hits the center of the upper mirror of BT3.

28. Close the pinhole in BT3, which identifies the center of the lens. The beam reflecting from the top mirror of BT3 should be centered on the pinhole. If necessary, adjust the top mirrors of BL3 and BT3 so that the beam is centered without clipping on the top mirror of BT3.

29. Move the delay stage to the end of its travel closest to P5. Place an index card with a target over the lens of BT3 and position the card so that the beam hits the center of the target. Move the delay stage to the other end of its travel. If the beam has translated off the center of the target, you will need to fine tune the alignment of M9 and M10. Adjust M9 when the delay stage is closest to P5 and M10 when the delay stage is farthest from P5. Repeat until the beam does not translate off the target when the delay stage moves.

30. Remove the target from the lens of BT3. The beam should now be hitting the center of the bottom mirror of BT2. Adjust this mirror so that the beam is centered on M13.

31. Adjust M13 so that the beam goes into the hole in the back of OAP9. See below for instructions to align the NIR through the pinhole in the front of OAP9.

32. Finally, you need to align the beam so that it passes through the pinhole in the front face of each OAP. The procedure is the same for all three cases, but I will use OAP3 as an example.

    (a) Place a dentist's mirror behind M11 (on a line drawn from OAP3 to M11) and remove the card from the back of M11 so that you can see through M11 and into the hole in the back of OAP3. DO NOT place the mirror on a line drawn from the bottom mirror of BT1 to M11, as this is the path that would be taken by any beams that leak through M11.

    (b) When the NIR hits the metal inside the hole in OAP3, it will spark, giving off white light. This light illuminates the hole and allows you to see the shape of the hole. You can often also see a faint white pinpoint at the center of the hole where light is coming in from the front face of

the OAP.

(c) Place a ceramic doubler at RS1.

(d) Adjust M11 to position the white spark at the center of the hole. When it hits the center, the white spark will get significantly fainter.

(e) Look at the doubler crystal. The NIR should be passing through the pinhole, so you will see it appear on the doubler. Adjust M11 to get the full NIR spot through the pinhole.

C.1.3 Distribution Optics Quick Check

I use this procedure every day I run to check the alignment of the distribution optics.

1. Check that the reflected beam from BS1 is being safely dumped in the beam stop.

2. Tip M5 down so that it is out of the optical path and place a card over BS2 so that no light gets through it.

3. P1 and P2 define the desired axis of injection into the rest of the distribution optics. Adjust M1 and M2 so that the beam is centered on P1 and that when P1 is closed, the spot is centered on P2.

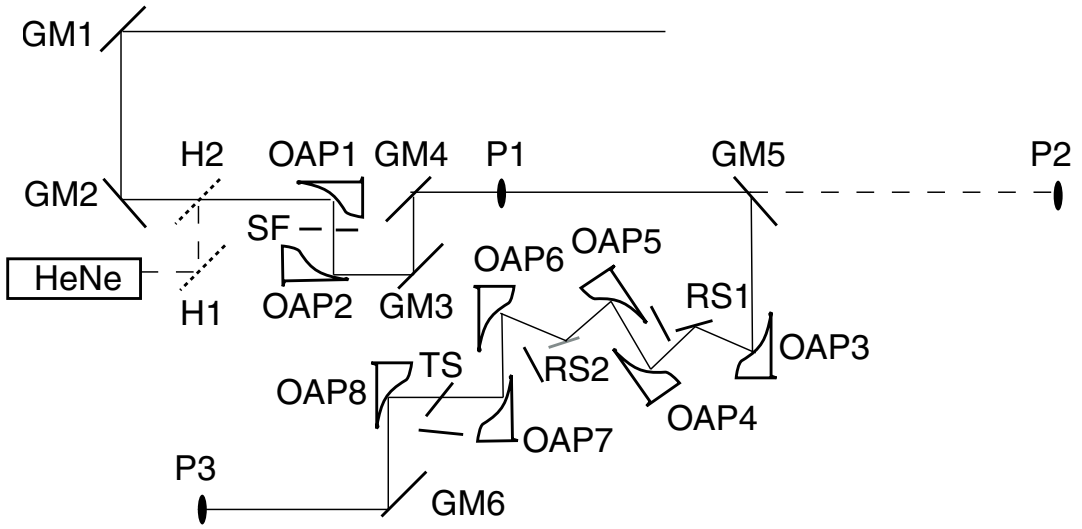

Figure C.3: A reference diagram of the pulse slicer for describing individual components and overall alignment.

C.2 The pulse slicer layout

The pulse slicer has many components. There is usually a reason that each component exists or is laid out in a particular way. Rather than attempting to write one very long description of the slicer that covers everything, in this section I will individually describe each of the elements that has some explanation behind it. I will refer throughout to Figure C.3.

Spatial Filter Denoted by SF on the diagram. The output beam of the FEL is nominally Gaussian, but not perfect. It can also be distorted somewhat as it goes through the transport system. When a non-Gaussian beam is

brought to a focus, the higher-order components are spatially separated from the focus. If a beam like this were to hit one of the switches, the spatially-separated components would not overlap with the NIR and would also contribute to the pre-pulse reflectivity. To avoid this, we spatially filter the beam, which simply consists of putting a pinhole at the focus. The lowest-order spot goes right through, but the spatially-separated higher orders are clipped off. The focusing for the spatial filter is provided by OAP1. The size of the spatial filter that should be used depends on the wavelength (because the diffraction-limited spot size depends on wavelength). I find that a diameter of around 1300 μm works well.

Beam Reducer/Expander Denoted by OAP 1 and OAP 2 on the diagram. OAP 1 focuses the light down onto the spatial filter. OAP 2 then re-collimates the light. By choosing different focal lengths for these optics, you can reduce or expand the beam. A standard OAP has a focal length of 4.8", but we also have one that is 3.2". By placing the 3.2" (4.8") at OAP 1 (OAP 2), you end up with an expanded beam, which brings the FIR to a tighter focus and can give you somewhat higher intensities at the focus on the sample. By reversing the order so that the 4.8" (3.2") is at OAP 1 (OAP 2), you end up with a reduced beam diameter. A smaller diameter

leads to a smaller range of incident angles when focused onto a reflection switch and thus reduces the pre-pulse reflectivity, improving the contrast ratio. The reduced beam is preferable for getting good performance out of the slicer if you can tolerate the reduced intensity at the sample.

GM 5 GM 5 is on a translation stage because it is important that the beam coming off of GM 5 comes in to OAP 3 at the correct angle. Using a translation stage gives you full control of both the position and angle of the beam to make sure it is correctly aligned with OAP 3.

Switch angles In order to get maximum suppression of the pre-pulse, the FIR is incident on the switches at Brewster's angle. The OAPs surrounding each reflection switch are set up to have the FIR incident at that angle and then to recapture the beam that is reflected off the switches. Since both GaAs (Brewster angle of 74.2^o) and silicon (73.7^o) are possible switch materials, Frank set up the slicer to be in between these two: $\sim 74.1^o$. It may be possible to improve the slicer by getting closer to Brewster's angle for silicon, but that would require removing and replacing all the components from OAP 4 onward. The steep incident angle also results in a FIR spot that is significantly elongated on the wafer.

Reflection switch mounts Located at RS1 and RS2, these mounts are designed

to allow you to scan the part of the reflection switch that is actually illuminated without changing the angle of reflection. This is important because the thickness of a switch varies across the surface of the wafer. This scanning allows you to maximize the pre-pulse suppression by finding a spot on the wafer with optimal thickness.

Reflection switch frames The frames on which the reflection switches are mounted are shaped as a "C" and placed so that the open side of the "C" is on the downstream end of the wafer. I originally used a complete square, not a "C" shape, but because the FIR hits the wafer at such a steep angle, the FIR that passed through the wafer would scatter off the edge of the frame, increasing the pre-pulse. Cutting out the downstream wall reduced this scatter significantly.

Polarizers There are two polarizers (not shown in Figure C.3) that can be inserted between OAP3 and RS1, or OAP 5 and RS2. Even if the initial beam coming into the slicer is well polarized (horizontally), reflection off the OAPs causes the polarization to be somewhat distorted. Adding these polarizers after the OAP improves the contrast ratio by making sure the beam actually hitting the switches is properly polarized to take advantage of suppression at Brewster's angle. To allow the NIR beam to get through the polarizers,

they each have some of the wire strands cut out of the center.

Transmission switch orientation There is a proper orientation for the TS, which is given in Frank's paper. I show a diagram of the correct orientation in Figure C.4. This orientation is necessary because the OAP is not symmetric and so the angle subtended by one half of the beam is larger than that subtended by the other. You want the smaller subtended angle to be the one that is closest to the grazing angle of the wafer in order to minimize any chance of having some of the FIR bypass the wafer.

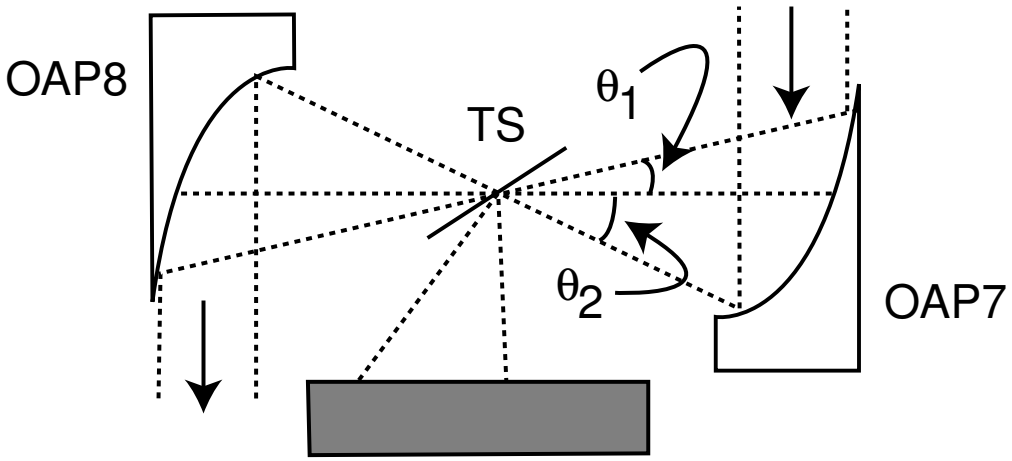

Figure C.4: The correct orientation of the transmission switch. Because of the asymmetry of the OAP, Θ_1 is smaller than Θ_2.

C.3 Aligning the pulse slicer

As described in Section 3.1, aligning the pulse slicer is difficult because you can't actually see the FIR beam. I first describe how to use a beam profiler to verify the overlap between a visible beam and the FIR beam. Then, for the alignment procedure, most of the steps will be done with this collinear visible beam.

For the purposes of this description, I will assume that the pulse slicer has somehow been horribly misaligned and you have to redo everything. In everyday usage this is not necessary, so at the end I will give an abbreviated procedure that I use for everyday checks of the alignment.

C.3.1 Beam profiling

The beam profiler is nothing more than a pyroelectric detector that has been mounted on a frame with stepper motors to control its horizontal and vertical position. Borys Kolasa built this stepper motor and frame system. The output of the pyro is sent to a Stanford pre-amplifier and then to an oscilloscope. I wrote software that moves the pyro through a specified pattern, reads off the pyro value and builds up a display of the beam intensity. Since the FEL only runs at roughly 1 Hz, it can take an extremely long time to build up a full profile of the beam,

especially if you set the step size to be small for high resolution. However, for the purposes of verifying alignment, it suffices to take a vertical and horizontal slice across the beam and see that the peak of the Gaussian lies at the desired position. The beam profiler could be set up anywhere and I have used it at many positions in the pulse slicer. At the places where I use it regularly (for example at P1, P2 and P3 on Figure C.3). I have installed kinematic stops that position the beam profiler so that its center position is right on the intended axis of the pulse slicer. That way zero on the data from the beam profiler corresponds to the desired beam axis.

C.3.2 Alignment Procedure

1. Make sure that the transport system is set up so that the THz comes to Port 1 in Room 1380D.

2. Remove OAP1, GM4 and GM5 from their kinematic bases.

3. The desired injection axis for the FIR runs through P1 and P2. Place the beam profiler at P1 and then at P2. Check that the FIR is centered on the desired axis at 5" above the table. If it is off, walk GM1 and GM2 until the beam is on the desired axis.

4. Put in the green alignment beam for the FEL and check that it is also on

axis at P1 and P2. If it is off, work with David Enyeart to adjust the mirrors that inject the green into the FEL cavity so that the beam is centered on P1 and P2. This makes sure that the green alignment laser is collinear with the THz.

5. Turn on the red HeNe laser on the magnet table. This serves as a local alignment beam. Insert the silver mirror (H2) on the kinematic base plate that directs the red beam toward P1 and P2. Adjust H1 and H2 until this beam is on axis through P1 and P2.

6. Replace OAP1 and GM4 in their positions. Check that the red beam (which now passes off OAP1, through the spatial filter and then off OAP2, GM3 and GM4) is passing through P1 and P2. Adjust GM3 and GM4 if necessary to get the beam on this axis.

7. Replace GM5. Using a tall card with an indicator line, make sure that the beam from GM5 to OAP3 is on axis (it should remain 5" above the table and over the axis line drawn on the the table. If the beam is off axis, adjust both the position and angle of GM5 to get it on axis.

8. Remove the wafer mounting plate from the stages at RS1. There is a spring-loaded clip mount in a large screw sitting near RS1 with a spacer ring around the screw. Screw this into the rotation stage at RS1, using the spacer ring

to raise the level of the clip (i.e. don't screw it all the way down). The center of the rotation stage is indicated by the milling mark on the center of the clip. Using this mark, position the center of the stage at the focus of the FIR, which can be seen by placing a liquid crystal at the focus and looking for the spot that changes color in sych with the FEL pulse. This spot will only be visible near the focal point. This step is rarely necessary and should only be performed if something drastic has happened.

9. Replace the wafer mount on the rotation stage, place a gold mirror on the mount and rotate the mount so that the mirror is at approximately the correct angle. Adjust the stage position slightly so that the focus is on the face of the gold mirrors.

10. Fine tune the angle of the mirror (at RS1) so that the red alignment beam hits the center of OAP4 and OAP5

11. Repeat the last three steps for RS2, setting the angle for good reflection on OAP6 and OAP7.

12. Finally, repeat these steps to align the center of the transmission switch (TS) stage under the focus of the FIR beam.

13. The red beam should now be coming off of OAP8. Be sure that the beam

is coming out perpendicular to OAP8, above the axis mark drawn on the table. If it is not, go back and check the angles of RS1 and RS2. These are the only movable parts in the alignment of the actual slicer.

14. Check the position of the HeNe beam on GM6, or the variable polarizer if it is used in place of GM6. The beam should come off GM6 and hit the center of P3.

15. Now, with a gold mirror at both RS1 and RS2, you can see the focus of the FIR at the location of the TS. Place a liquid crystal on top of the TS stage so that you can see this focus. Allow the NIR to get to the TS only. You will be able to see a large spot where the NIR heats the crystal. Center the NIR on the FIR with the following:

    (a) The NIR could be steered using only M13 of Figure C.1. However, the NIR must BOTH pass through the pinhole in OAP7 AND hit the desired spot on the liquid crystal, so it is necessary to use both M13 and the mirror at the base of BT3 to completely control the path of the NIR.

    (b) From the liquid crystal, determine which direction the NIR spot needs to move. Using a ceramic doubling crystal to look at the NIR, adjust M13 in the appropriate direction. You will see the intensity of the beam

on the doubling crystal reduce as the beam is clipped at the pinhole on the surface of OAP7.

(c) Adjust the mirror at the base of BT3 to bring get the full beam through the pinhole of OAP7.

(d) Repeat this procedure, walking the beam with M13 and the mirror at the base of BT3, until the NIR beam comes cleanly through the pinhole and overlaps the area of the FIR focus on the liquid crystal.

16. Remove the liquid crystal and insert the silicon transmission switch.

17. Remove the gold mirror from RS2 and place a liquid crystal into the mount so that you can again see the FIR focus. Repeat the procedure described above to overlap the NIR and FIR, now using M12 and the mirror at the base of BT2.

18. Remove the liquid crystal from RS2 and insert the appropriate reflection switch for the wavelength you are using.

19. Remove the gold mirror from RS1 and repeat the overlap procedure using M11 and the mirror at the base of BT1.

20. Remove the liquid crystal and insert the appropriate reflection switch.

21. The slicer should now be aligned for proper operation.

C.3.3 Pulse Slicer alignment check

This is the procedure that I use on a daily basis to make sure the slicer alignment is okay before I start to do experiments.

1. Make sure that the transport system is set up so that the THz is coming to Port 1 in Room 1380D.

2. Remove OAP1, GM4 and GM5 from their kinematic bases.

3. Put in the green alignment beam for the FEL and adjust GM1 and GM2 so that it is on axis at P1 and P2.

4. Turn on the red HeNe laser on the magnet table. Insert H2 and adjust H1 and H2 until this beam is on axis through P1 and P2.

5. Replace OAP1 and GM4 in their positions. Check that the red beam is still passing through P1 and P2. Adjust GM3 and GM4 if necessary.

6. Replace GM5. Using a tall card with an indicator line, make sure that the beam from GM5 to OAP3 is on axis. Adjust both the position and angle of GM5 to get it on axis, if necessary.

7. Insert gold mirrors at RS1 and RS2.

8. The red beam should now be coming off of OAP8. Check the position of the HeNe beam on GM6, or the variable polarizer if it is used in place of GM6.

9. Now, with a gold mirror at both RS1 and RS2, you can see the focus of the FIR at the location of the TS. Place a liquid crystal on top of the TS stage so that you can see this focus. Allow the NIR to get to the TS only. You will be able to see a large spot where the NIR heats the crystal. Center the NIR on the FIR with the following:

   (a) The NIR could be steered using only M13 of Figure C.1. However, the NIR must BOTH pass through the pinhole in OAP7 AND hit the desired spot on the liquid crystal, so it is necessary to use both M13 and the mirror at the base of BT3 to completely control the path of the NIR.

   (b) From the liquid crystal, determine which direction the NIR spot needs to move. Using a ceramic doubling crystal to look at the NIR, adjust M13 in the appropriate direction. You will see the intensity of the beam on the doubling crystal reduce as the beam is clipped at the pinhole on the surface of OAP7.

   (c) Adjust the mirror at the base of BT3 to bring get the full beam through the pinhole of OAP7.

(d) Repeat this procedure, walking the beam with M13 and the mirror at the base of BT3, until the NIR beam comes cleanly through the pinhole and overlaps the area of the FIR focus on the liquid crystal.

10. Remove the liquid crystal and insert the silicon transmission switch.

11. Remove the gold mirror from RS2 and place a liquid crystal into the mount so that you can again see the FIR focus. Repeat the procedure described above to overlap the NIR and FIR, now using M12 and the mirror at the base of BT2.

12. Remove the liquid crystal from RS2 and insert the appropriate reflection switch for the wavelength you are using.

13. Remove the gold mirror from RS1 and repeat the overlap procedure using M11 and the mirror at the base of BT1.

14. Remove the liquid crystal and insert the appropriate reflection switch.

15. The slicer should now be aligned for proper operation.

C.4 Optimizing Contrast Ratio

The contrast ratio of the pulse slicer is critical to performing a variety of experiments. As described is Section 3.3, the contrast ratio is severely limited by the range of incident angles when FIR radiation is focused onto a switch. As mentioned in Section 3.3, you can use an aperture to reduce the spread of incident angles. This limits the total power throughput, but reduces the pre-pulse (which arises primarily from the outer portion of the beam) faster than the total power. Thus you can achieve a net improvement in contrast ratio.

Bryan Cole and I both used an aperture to achieve Rabi Oscillations at 84 cm^{-1}. We both used a 1"-diameter aperture placed on the axis of the FIR between OAP4 and OAP5. I have also investigated the effect of an aperture on contrast ratio at both 34 and 51 cm^{-1}. I have put my observations on this into this Appendix rather than Section 3.3 because I still have a lot of uncertainties about what actually happens when I aperture the beam.

Figure C.5 shows a characteristic photoconductive response of the B434 unintentionally-doped GaAs sample to a sliced pulse of duration ∼600 ps using different aperture values. The earliest rise in the signal is a response to the pre-pulse, coincident with the beginning of the FIR pulse. When the pulse slicer turns on, there is a rapid rise in the signal. This signal decays to zero in approximately 30 ns: long in

comparison to the actual duration of the FIR pulse, but short on the time scale of the ~ 5 μs FIR pulse. It is clear that for smaller aperture values, the pre-pulse response is significantly reduced. I have sometimes observed that the response to the sliced pulse **increases** with decreasing aperture diameter. I attribute this to a larger number of electrons available for photo-excitation, since the pre-pulse is reduced.

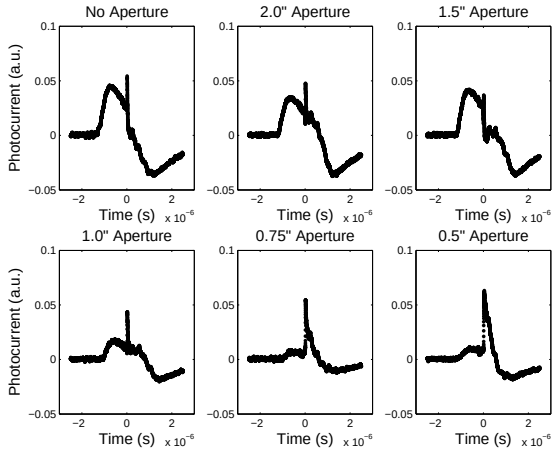

Figure C.5: Characteristic photoconductive response to a sliced pulse of duration 600 ps using different aperture diameters.

Figure C.6 shows a similar set of photoconductive traces at different aperture values, this time for a sliced pulse of duration ~ 3 ns. As in Figure C.5, it is clear that the pre-pulse reduces in amplitude as the aperture diameter is decreased.

In both Figure C.5 and Figure C.6 it is clear that for small aperture diameter, the sliced pulse response decays less rapidly than for large diameters. This is particularly troubling, as it implies that the transmission switch is not fully cutting

off the trailing edge of the FIR pulse. I have seen behavior similar to this when measuring sliced pulses with the bolometer. In that case, I observe a large photo-signal upon activation of the reflection switches. This photo-signal is still present even when the transmission switch is activated **before** the reflection switches. As the activation time of the transmission switch is increased, there is an increase in the magnitude of the photo-signal. However, the change with sliced pulse duration has a magnitude of $\sim 10\%$ of the total observed photo-signal. Since the time constant of the bolometer is 200 ns, far longer than the actual sliced pulse, I have not been able to resolve the origin of this signal.

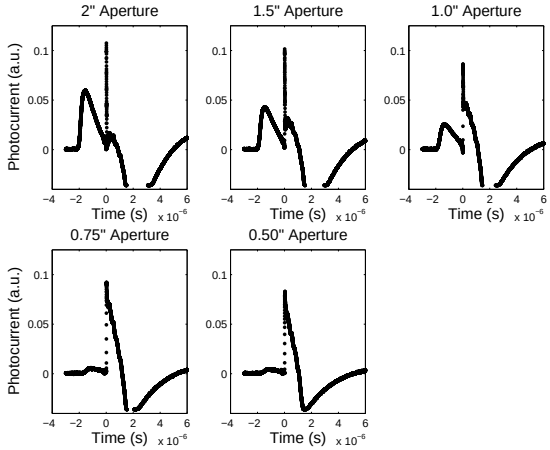

Figure C.6: Characteristic photoconductive responses to a sliced pulse of duration 3 ns using different apertures.

Finally, Figure C.7 shows the photoconductive signal as a function of the position of a 0.75" aperture across the beam center. This data was obtained by placing the aperture after GM6. The displacement value indicates the position

of the center of the aperture relative to the nominal center of the FIR beam. The vertical center of the aperture remained at the vertical center of the FIR for all positions. A strong dependence of the fall time of the sliced pulse on the position of the aperture is evident. This implies a spatial dependence of the excess FIR appearing at the end of the sliced pulse. I have not been able to conclusively determine whether this arises from diffraction around the aperture, changing the coupling into the sample, or if it is a true representation of the beam characteristics. There has been discussion of developing a time-resolved method of observing the sliced FIR pulse using electro-optic detection. I believe this detector would help to determine the origin of these unusual signals.

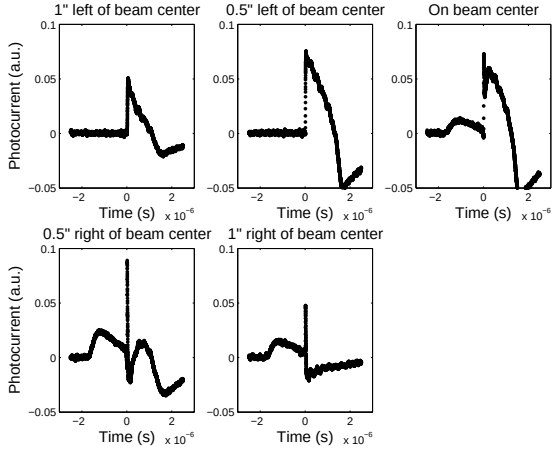

Figure C.7: Characteristic photoconductive responses to a sliced pulse with a 0.75" aperture at various positions across the FIR beam.

Appendix D

Start-up and Shut-down procedures

The Rabi Oscillation experiment described in this thesis requires an extremely large number of components, each with their own idiosyncracies. As an aid to anyone continuing with this sort of experiment, this appendix provides an ordered list of steps to take when starting up or shutting down the experiment. I assume for this list that everything is operating normally. If it is not, you should look at the other appendices for information on the alignment or maintenance of specific components

D.1 Start-up Procedure

1. Refill the nitrogen jacket of the cryostat.

2. Refill the helium space of the cryostat.

3. Start the large vacuum pump and open a small valve on the cryostat to begin pumping down slowly. Continue to open this valve a little bit at a time over the next hour.

4. Start the chillers for the Millenia and YAG lasers.

5. Warm up the diodes on the Millenia laser.

6. Start up the FEL GVM, oplink and shaft and allow the system to warm up for about 20 minutes. (If you're nice, Dave will do this for you.)

7. Start up the Millenia laser and make sure that the Tsunami is properly modelocked.

8. Start the electronics boxes that provide the triggering signals to the FEL and the YAG laser.

9. Start the YAG laser and allow the regen to warm up for 20 minutes.

10. Make sure that the proper output port of the transport system has been

selected. Using the FEL green alignment beam, check that the FEL is on the incident axis of the pulse slicer. (Removing OAP 1, GM 4 and GM 5.)

11. Verify that the pulse slicer red alignment laser is also on the proper axis.

12. Replace OAP 1, GM 4 and GM 5 and verify that the red laser is still on the correct axis.

13. Verify that the injection into OAP 3 is on axis.

14. Start up the FEL (again, Dave will do this if you're nice).

15. Start up the regen and pick an appropriate pulse.

16. Using the liquid crystal, overlap the FIR and NIR at each switch.

17. If you have been opening the small valve, the cryostat should be down to about 100 Torr by now. Once it's below 100 Torr, close the small valve and open the big valve a crack. Once the pressure is below 37 Torr (the superfluid transition point) open the big valve all the way.

18. Connect the magnet leads and ramp the magnetic field up to the desired value. Make sure to turn on the persistent switch heater.

19. If necessary, engage the persistent switch and ramp the field back down to zero in order to detach the magnet leads and put on the RF shielding.

20. Set the bias voltage on the sample to 50 mV. If you attempt to do this before the field is on, you will not be able to reach 50 mV. I believe this is related to some magneto-resistance of the sample.

21. Turn on the oscilloscope and the fast amplifier power supply. Note: do not power up the amplifiers unless they are loaded, i.e. the output is connected to the scope.

22. Switch the FEL to external triggering to synchronize the FIR and NIR.

23. Make sure that the scope is in smart triggering mode so that the trigger is armed by the FEL trigger pulse and fires when the photodiode observes a regen pulse.

24. Put in the proper switches at the reflection switch positions.

25. You should now be ready to start doing actual experiments.

D.2 Shut-down Procedure

1. Block the YAG to crystal 2 of the regen and block the seed beam from getting to crystal 1. Turn off the Millenia and Yag, but allow the cooling water to continue to flow.

2. Make sure that the magnet is fully discharged (zero field).

3. Turn the sample bias voltage down to zero.

4. Turn off fast amplifier power supply and then scope.

5. Close all valves that pump on the cryostat.

6. Remove the switches and place them in their protective cases.

7. Turn off the FEL.

8. Turn off the vacuum pump that pumps on the cryostat.

9. Replace the windows around the pulse slicer table.

10. Make sure the Helium gas is off.

11. Turn off alignment HeNe and power supply to any pyroelectric detectors.

12. Turn off the cooling water to the Millenia and Yag after about 20 minutes of cooling.