

UNIVERSITY of CALIFORNIA
Santa Barbara

Electron-Hole Recollisions in Driven Quantum Wells

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

Doctor of Philosophy

in

Physics

by

Hunter Bennett Banks

Committee in charge:

Professor Mark S. Sherwin, Chair
Professor David M. Weld
Professor Cenke Xu

December 2016

The dissertation of Hunter Bennett Banks is approved:

Professor David M. Weld

Professor Cenke Xu

Professor Mark S. Sherwin, Chair

September 2016

Copyright © 2016
by Hunter Bennett Banks

To my family

Acknowledgements

First, I would like to thank my advisor Mark Sherwin. He let me into his group and put me on probably the coolest project a young graduate student could imagine, and for that I am truly thankful. His support, patience, and sage advice has kept me and this work focused, grounded, and moving forward.

The measurements in this dissertation certainly would not have been possible without the FEL, so I would next like to thank the people who kept the whole thing running for over thirty years. The machine and all of its parts that Dave Enyeart, Gerry Ramian, and Nick Agladze (and others) have designed, built and improved is a testament to their ingenuity and skill. Getting to work with and around Dave has certainly been something I'm better for, with his razor wit and uncanny ability to fix things.

I've been very fortunate to have worked with some incredibly talented students and postdocs. I'd like to thank Ben Zaks in particular—Ben took me under his wing when I joined the lab and taught me pretty much everything I know about optics. Thank you also to Dan Ouellette, Devin Edwards, Jessica Clayton, Andrea Hofmann, Jordan Grace, Andrew Pierce, Dominik Peller, Jonathan Essen, Nutan Gautam, Aaron Ma, Mengchen Huang, Darren Valovcin, Blake Wilson, and Chang Yun, to name just a few. The lab is in good hands with the new generation of students, Darren, Blake, and Chang.

The materials and the theoretical support provided by our collaborators have been wonderful. Thank you to Shawn Mack and Loren Pfeiffer for the amazing samples, they've been such source of interesting and fun physics. Thank you to Fan Yang, Qile Wu, and Ren-Bao Liu for the stimulating discussions and keen theoretical insight.

Thank you to John Leonard, Garrett Cole for sharing their hard-earned processing techniques and to Brian Thibeault, Aidan Hopkins and the rest of the cleanroom staff for the incredible facility they run. Thank you to Mike Deal, Jennifer Farrar, Guy Patterson, Rob Marquez, Rita Makogon, and the rest of the Physics Department staff who go above and beyond every day.

I wouldn't've made it this far in grad school without all of my friends and family. You've been roommates, labmates, teammates, trivia-mates, golf-mates, and kept me driven and happy and curious. Thank you all for the lunches, the late nights, and the good times. Thanks to my parents, Kathy and Ben, and my brother and sister, Brian and Katie, who have all been incredibly supportive and surprisingly willing to visit in August. And, of course, thanks to Miranda, who has been everything.

Curriculum Vitæ

Hunter Bennett Banks

Education

2016 Ph.D., Physics, University of California, Santa Barbara, California

2010 A.B., Physics, Washington University in St. Louis, St. Louis, Missouri

2006 St. Mark's School of Texas, Dallas, Texas

Professional Experience

2012–2016 Graduate research assistant, Physics Department, UCSB

2011–2012 Teaching assistant, Physics Department, UCSB

2010–2011 Graduate research assistant, Physics Department, UCSB

2008–2010 Undergrad research assistant, Physics Department, WUSTL

Publications

“Anomalous He-gas high-pressure studies on superconducting $\text{LaO}_{1-x}\text{F}_x\text{FeAs}$ ”, Bi, W., Banks, H. B., Schilling, J. S., Takahashi, H., Okada, H., Kamihara, Y., Hirano, M., and Hosono, H. *New Journal of Physics* Vol. 12, 023005 (2010)

“Dependence of the magnetic ordering temperature on hydrostatic pressure for the ternary intermetallic compounds GdAgMg , GdAuMg , EuAgMg , and EuAuMg ”, Banks, H., Hillier, N. J., Schilling, J. S., Rohrkamp, J., Lorenz, T., Mydosh, J. A., Fickenscher, T., and Pttgen, R. *Physical Review B* Vol. 81, 212403 (2010)

“Dimeric endophilin A2 stimulates assembly and GTPase activity of dynamin 2”, Ross, J. A., Chen, Y., Müller, Joachim, Barylko, B., Wang, L., Banks, H. B., Albanesi, J. P., and Jameson, D. M. *Biophysical Journal* Vol. 100, 729–737 (2011)

“Dependence of magnetic ordering temperature of doped and undoped EuFe_2As_2 on hydrostatic pressure to 0.8 GPa”, Banks, H. B., Bi, W., Sun, L., Chen, G. F., Chen, X. H., and Schilling, J. S. *Physica C: Superconductivity* Vol. 471, 476–479 (2011)

“High-order sideband generation in bulk GaAs”, Zaks, B., Banks, H., Sherwin, M. S. *Applied Physics Letters* Vol. 102, 012104 (2013)

“Terahertz Electron-Hole Recollisions in GaAs Quantum Wells: Robustness to Scattering by Optical Phonons and Thermal Fluctuations”, Banks, H., Zaks, B., Yang, F., Mack, S., Gossard, A. C., Liu, R., Sherwin, M. S. *Physical Review Letters* Vol. 111, 267402 (2013)

“Antenna-boosted mixing of terahertz and near-infrared radiation”, Banks, H. B., Hofmann, A., Mack, S., Gossard, A. C., Sherwin, M. S. *Applied Physics Letters* Vol. 105, 092102 (2014)

Abstract

Electron-Hole Recollisions in Driven Quantum Wells

by

Hunter Bennett Banks

Driving semiconductor quantum wells with terahertz electric fields strong enough to overcome the Coulomb attraction between bound electron-hole pairs leads to high-order sideband generation (HSG). In HSG, excitons are optically-injected into quantum wells by a weak near-infrared (NIR) laser while simultaneously being illuminated with a terahertz field from the UCSB Free Electron Laser. The phenomenon can be described by the so-called “three step model” developed in high-field atomic physics: (1) the electron and hole tunnel-ionize in the strong field, (2) the now-free particles accelerate in the field, and (3) they recollide, emitting a photon. The two lasers are continuous, so the emitted photons are sidebands on the NIR laser. Because of the large gain of kinetic energy before recollision, an HSG spectrum has a broad bandwidth with many more sidebands above the NIR frequency than below. The largest spectra span over one hundred nanometers, with over 100th order sidebands above and 20th order below.

The electron and hole must remain coherent throughout their trajectories, which can last hundreds of femtoseconds and extend for more than fifty nanometers, if they are to recollide. Sidebands have been observed that result from recollisions with kinetic energies

far above the threshold for optical phonon emission. These high orders persist up to room temperature. Not even quenched disorder in the quantum wells strongly attenuates the HSG signal.

Because of this coherence, the electron and hole are very sensitive to the complete band structure of the material. Excitation by linear NIR polarization creates both the electron and hole in a superposition of spin-up and spin-down states with complex coefficients given by the relative orientation of the NIR polarization and the THz polarization. Interference between these the spin-up and spin-down particles, particularly in the valence band and mediated by non-Abelian Berry curvature, has large effects on both the intensity and polarization state of the sidebands. The connection between HSG and complete band structure points to the possibility of directly measuring both the band dispersion relations as well as the non-Abelian Berry curvature of the host material.

Contents

1	Introduction	1
1.1	Perturbative nonlinear optics	3
1.2	A brief introduction to GaAs	5
1.3	The three step model	9
1.4	GaAs quantum wells	16
1.5	Photoluminescence and absorption	20
1.6	UCSB Free Electron Laser	23
2	The electron-lattice interaction in the three step model	26
2.1	Sample details	28
2.1.1	Sample design	28
2.1.2	Sample processing	28
2.1.3	Sample optical properties	29
2.2	Sideband measurements	32
2.2.1	Optical setup	32
2.2.2	Results	34
2.3	Theoretical model	38
2.3.1	Three step model: Classical interpretation	38
2.3.2	Full theoretical results	42
2.4	Conclusion and remaining questions	46
3	The Sideband Spectrometer	47
3.1	Optical setup	48
3.1.1	Linear measurements	48
3.1.2	Nonlinear measurements	53
3.2	FEL characterization	67
3.2.1	Power measurement	68
3.2.2	Frequency and linewidth	71
3.3	Outlook	76
4	Terahertz-induced birefringence and sideband polarimetry	77
4.1	Sample details	79
4.1.1	Sample growth	79

4.1.2	Sample processing	80
4.1.3	Sample optical properties	82
4.2	Sideband measurements	86
4.2.1	Experimental preparation	86
4.2.2	Results	93
4.3	Theoretical model	102
4.3.1	Separation into circular components	103
4.3.2	Semiclassical trajectories	106
4.4	Conclusion and remaining questions	111
A	Equipment details	115
A.1	Cryostat	116
A.2	Sample imaging camera	119
A.3	M ² SolsTiS Ti:sapphire laser	122
A.3.1	Polarizer attenuator	125
A.4	PhotonControl spectrometer	125
A.5	Acousto-optic modulator	126
A.6	Mechanical shutter	128
A.7	Absorption LED	128
A.8	SPEX monochromator	129
A.9	Hamamatsu photomultiplier tube	132
A.10	Acton spectrometer	132
A.11	Andor electron-multiplying CCD camera	134
A.12	Thomas-Keating energy meter	139
B	FEL transport details	140
C	Cleanroom processing	142
C.1	Quantum well membrane process	143
C.2	Epitaxial transfer process	146
C.3	Indium-tin oxide deposition	150
C.4	Dicing samples	151
	Bibliography	152

Chapter 1

Introduction

Almost thirty years ago, experimentalists studying the behavior of noble gases in strong laser pulses generated light at 32 nm when illuminating the atoms with excitation light at 1064 nm, the thirty-third harmonic [19]. They observed only odd order harmonics $\omega_n = n\omega_{\text{drive}}$, where n is the order and ω_{drive} is the excitation frequency. Most surprisingly, the harmonics from seventh to twenty-seventh order hardly decayed with increasing order, likely the first reported harmonic plateau. Within five years, intense near-infrared (NIR) laser pulses were generating harmonics of more than one hundredth order through a process called high-order harmonic generation (HHG) [36, 50, 43, 40]. Light with wavelengths around 1 μm could generate light with wavelengths less than 10 nm and the monochromator was the limiting factor! Despite the nonlinear nature of this process, the theoretical description for this phenomenon, the three step model, is actually quite simple and is based on the ionization, acceleration, and recollision of an electron with its

parent ion [14]. Many discoveries have been made using this phenomenon, including the creation of single 35 attosecond pulses using light with a period $T = 7$ fs ($\lambda = 2 \mu\text{m}$) and the seeding of x-ray lasers, both table-top and free electron lasers [15, 12, 75, 37].

Inspired by HHG from the atomic community, Liu, *et al.*, investigated the possibility of recollision physics in a semiconductor using excitons, atom-like bound electron-hole pairs, termed high-order sideband generation (HSG), which was later observed experimentally by Zaks, *et al.*, with the detection of eighteenth order sidebands [41, 66, 73]. Like HHG, HSG is a highly nonlinear mixing process. In HSG, however, two frequencies of light illuminate the material, one in the NIR wavelength range, ω_{NIR} , and the other in the terahertz (THz), ω_{THz} . Through the process of HSG, many THz photons combine with a single NIR photon to create a comb of sidebands, $\omega_{\text{SB}}(n)$ of integer order n with frequencies,

$$\omega_{\text{SB}}(n) = \omega_{\text{NIR}} + n\omega_{\text{THz}}.$$

The physics of atoms and excitons are very similar, despite existing at such distant scales. The dielectric constant and effective masses bring the binding energy from 20 eV for neon down to 10 meV for an exciton in a GaAs quantum well and the Bohr radius from less than 1 Å up to 10 nm. To stay in the limit of tunnel ionization, the driving frequency for HSG must be less than 1 THz (4 meV), but the required field strengths drop to 1 mV/nm (10 kV/cm).

This chapter will introduce the physics of HHG and HSG. First, it will put them into context by describing perturbative nonlinear optics. Second, it will explain the unique

properties of GaAs to understand how to extend the atomic models to explain HSG. Then, it will go into detail about the three step model and the modifications required to move from atoms to solids. Finally, it will give background on the generation of intense terahertz light by the UCSB Free Electron Lasers.

1.1 Perturbative nonlinear optics

To understand the significance of high-order nonlinear processes, it is important to first introduce perturbative nonlinear optics to demonstrate how unusual double-digit THz nonlinearities are. This subsection will follow the excellent resource *Nonlinear Optics* by Robert W. Boyd [5]. Electrons in free atoms or in crystals are not bound by exactly quadratic potentials, although the approximation is often made, and it is this anharmonicity that leads to nonlinear optics [20]. In perturbative nonlinear optics, the polarization $P(t)$ of the medium is expanded in a Taylor series in the applied electric field. Treating the electric field as a scalar for simplicity,

$$P(t) = \chi E(t) \tag{1.1}$$

$$= \chi^{(1)} E(t) + \chi^{(2)} E^2(t) + \chi^{(3)} E^3(t) + \dots \tag{1.2}$$

For conventional linear optics, the approximation $\chi \approx \chi^{(1)}$ is taken. For sufficiently intense electric fields, however, the second- and third-order terms can become significant.

For real, vector-natured electric fields, the susceptibility $\chi^{(n)}$ is a $(n+1)$ -rank tensor that

can couple electric fields polarized along arbitrary axes, but such a treatment is beyond the scope of this introduction. In general, only the lowest nonlinearity matters as the higher order $\chi^{(n)}$ coefficients are extremely small. For many interesting nonlinear optical phenomena, it is important to have the largest nonlinear coefficients possible. Usually, the lowest-order term, $\chi^{(2)}$, dominates. For a material to have a non-zero $\chi^{(2)}$, however, it must not be inversion-symmetric. Second-order mixing is a well-established technique for generating THz pulses with very strong fields [70, 58]. In inversion symmetric materials, such as (100)-cut GaAs, $\chi^{(3)}$ is the lowest-order nonzero nonlinearity. Other planes in GaAs, such as (110), do have non-zero $\chi^{(2)}$, however, but they will not be used in this dissertation.

The multiplication of sinusoidal fields leads to the generation of new frequencies based on the sum and difference of the initial frequencies. When the weak frequency is in the NIR, around 400 THz, and the strong frequency is in the THz, around 0.5 THz, then sum- and difference-frequency generation become sideband generation. Perturbative THz sideband generation, first discovered at UCSB using the Free Electron Laser, has been used to study semiconductor heterostructures since 1995 [8, 34, 6, 61, 60, 7]. Because sideband generation can be resonantly enhanced by nearby transitions, even electric-dipole-forbidden ones, many studies have probed the energy spectrum of excitons under various internal and external perturbations [34, 61].

To generate new frequencies efficiently, matching the phase of the participating laser fields and their products is important. The wavelength-dependent index of refraction

in a material means that the photons created throughout the propagation length of the material may not all add in phase. From the perspective of conservation laws, both energy (frequency ω) and momentum (wavevector k) must be conserved. The wavevector mismatch, $\Delta k = |k_f - \sum_i k_i|$ where k_f is the generated wavevector and k_i are the generating wavevectors, determines the lengthscale L_c over which there is efficient sideband generation;

$$L_c = \frac{2}{\Delta k}. \quad (1.3)$$

If there is perfect phase matching, $L_c \rightarrow \infty$, the entire length of the nonlinear crystal l contributes to generating new frequencies, assuming an undepleted pump. For finite L_c , destructive interference prevents all but the back portion of the nonlinear crystal, length $l \bmod L_c$, from contributing. It takes tremendous materials and optical engineering to efficiently generate THz sidebands in the perturbative regime [6]. Fortunately, sidebands in GaAs quantum wells are generated efficiently enough that it is straightforward to work in the limit where $l \ll L_c$ [73].

1.2 A brief introduction to GaAs

The electronic and optical properties of GaAs can be explained extremely well using band theory. In band theory, the available energy and momentum states are calculated, and then they are filled with electrons up to the Fermi level (i.e. until all the electrons have been accounted for). Because they form a degenerate Fermi gas, the electronic and optical properties are then determined by only the electrons near the chemical poten-

tial. The dynamics of these electrons are then treated as if the electrons and holes are quasiparticles.

From a more experimental standpoint, semiconductor growth and processing technologies for GaAs are second only to silicon. Material growth techniques like molecular beam epitaxy (MBE) are extremely mature, allowing for useful heterostructured materials, using other elements from groups III and V, to be grown of exceptional quality. By growing precision stacks of different materials, the electronic states of the semiconductor can be tailored to the wishes of the user. This section and later sections about semiconductor physics will follow derivations from the excellent *The Physics of Low-Dimensional Semiconductors* by J. H. Davies [17].

The optical properties of a semiconductor are determined by the band structure at the valence band maximum and the conduction band minimum. For GaAs, these two points are both at $k = 0$, the Γ point, and they are separated by about 1.5 eV, corresponding to a photon with a wavelength of 800 nm, see Fig. 1.1. The characteristic wavevector of light in this range (about $1/\mu\text{m}$) is much smaller than that of an electron (given by the lattice spacing, about $1/\text{nm}$). In order to conserve momentum, then, optical interband transitions are very nearly vertical on the scale of the Brillouin zone. A photon with wavelength $\lambda \sim 800$ nm can, therefore, excite an electron just barely across the gap, leaving an empty electron state behind in the valence band. This empty electron state, or hole, in an otherwise-full band can be thought of as a positively-charged quasiparticle.

These electronic states are more complex than plane-waves, they retain some symme-

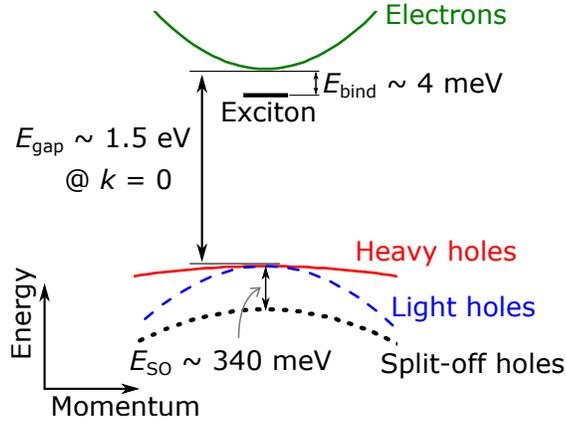


Figure 1.1: Band structure of GaAs around the Γ -point. The conduction band and valence band are separated by a minimum energy $E_{\text{gap}} \sim 1.5\text{eV}$ at the $k = 0$ in the Brillouin zone. The conduction band only contains one species of quasiparticle, electrons (green, solid line). The valence band contains three species: heavy holes, light holes, and split-off holes (solid red, dashed blue, and dotted black, respectively). The maximum of the split-off hole band is 340 meV below the valence band maximum. If electron-electron interactions are not neglected, the exciton appears, with a binding energy $E_{\text{bind}} = 4\text{ meV}$.

tries from the atomic orbitals they come from. The valence electrons in both gallium and arsenic are in p orbitals. The next shell of available orbitals are s orbitals. Through the mechanics of the tight-binding model, the p and s symmetry of the valence orbitals and the empty orbitals will be preserved in the valence and conduction bands. The top of the valence band is then six-fold degenerate, three for the p_x , p_y , and p_z orbitals times two for the electrons being spin $1/2$ (from now on all bands will be implicitly doubly degenerate from the electron spin). These angular momenta add by the Clebsch-Gordon rules to create three bands, two with total angular momentum $j = 3/2$ and one with $j = 1/2$. Spin-orbit coupling breaks the degeneracy between the two total angular momentum states, lowering the $j = 1/2$ band 340 meV relative to the $j = 3/2$ bands. That leaves

only the two hole bands with $j = 3/2$ at the very top of the valence band. These two bands, $|j, m_j\rangle = |3/2, \pm 3/2\rangle$ and $|3/2, \pm 1/2\rangle$, are degenerate at $k = 0$, but have different effective masses, and so are named the heavy hole and light hole bands, respectively. The bottom of the conduction band in GaAs is built from s orbital-like states, so there is only one band.

To fully explain the optical properties of a semiconductor near the band edge, the Coulomb interaction between charge carriers can no longer be ignored. An electron and a hole sitting on either end of the band edge have quadratic dispersion relations and a Coulombic attraction, creating a hydrogenic Hamiltonian,

$$H = \frac{\hbar^2 k_e^2}{2m_e^*} + \frac{\hbar^2 k_h^2}{2m_h^*} - \frac{e^2}{4\pi\epsilon |r_e - r_h|},$$

where k_i , m_i^* , and r_i are the momentum, effective mass and position coordinate of particle i , either an electron or a hole, and $\epsilon = \kappa\epsilon_0$ where κ is the dielectric constant of the material, $\kappa_{\text{GaAs}} \sim 13$, and ϵ_0 is the permittivity of free space. Because all of the masses here will be “effective,” the asterisk denoting such will be dropped. The only difference between this Hamiltonian and that of the hydrogen atom is that the effective masses have been rescaled by the bands and the permittivity of the material is rescaled by the dielectric constant due to screening. Transforming into the center of mass frame gives the familiar result;

$$H = \frac{\hbar^2 K^2}{2M} + \frac{\hbar^2 k^2}{2\mu} - \frac{e^2}{4\pi\epsilon r},$$

where K and M are the center of mass momentum and mass, respectively, k is the relative momentum of the electron and hole, μ is the reduced mass $m_e m_h / (m_e + m_h)$,

and ε is the dielectric constant of the material. This Hamiltonian allows for the formation of bound states called excitons, where the Bohr radius and binding energy are about 10 nm and 4 meV. If the electron and hole exist at the bottom and top of their respective bands, then the exciton will actually exist inside the previously-forbidden band gap. The exciton can even be excited directly with an optical photon with energy less than E_{gap} , see Fig. 1.1. Excitons are not long-lived, however, as the close proximity of the electron and hole encourages recombination on the order of nanoseconds after creation [18, 29].

1.3 The three step model

One of the early empirical discoveries in HHG was the existence of a harmonic plateau and cutoff. The harmonic spectrum would exhibit constant intensity for many dozens of orders, even into the hundreds, but then would disappear very quickly at the cutoff,

$$n_{\text{cutoff}} \sim \frac{3U_P + E_{\text{bind}}}{\hbar\omega} \quad (1.4)$$

$$U_P = \frac{e^2 F^2}{4m\omega^2}, \quad (1.5)$$

where n_{cutoff} is the cutoff harmonic order, U_P is the ponderomotive energy, E_{bind} is the ionization energy of the atom, and ω is the center frequency of the intense laser pulse [36]. The ponderomotive energy is the average kinetic energy of a charged particle in a sinusoidally oscillating electric field, where F is the electric field strength, and e and m the charge and mass of the electron. This cutoff law can actually be understood very

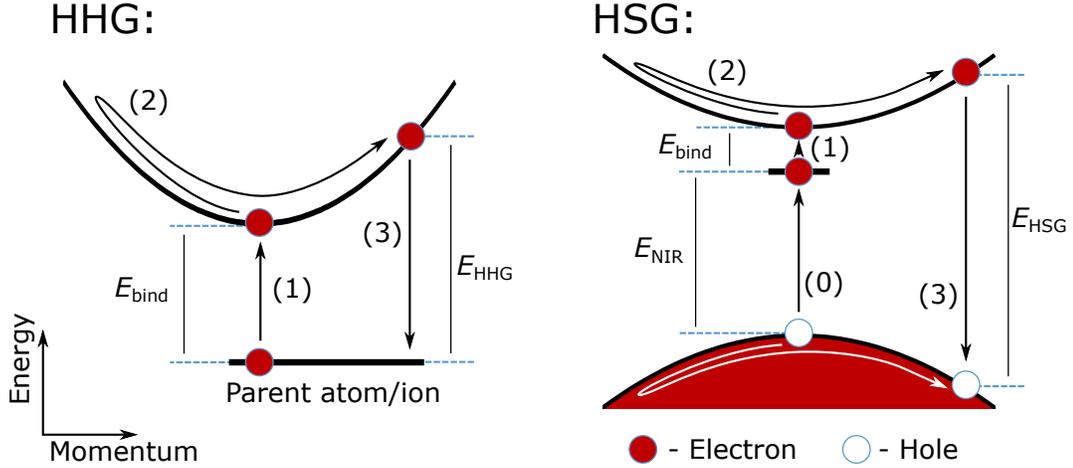


Figure 1.2: The three step model in HHG and HSG in momentum space. In HHG, the electron tunnel-ionizes from its parent atom/ion, (1). Next, the electron accelerates as a free particle in the driving electric field, (2). Initially, the electron moves away from the ion, but eventually the electric field turns around and so does the electron. The parent ion is assumed to be infinitely massive. Eventually, if the electron tunneled out at an allowed phase of the electric field, it recollides with the parent ion, (3). The electron and ion recollide at the origin in real space, but the electron has gained a significant amount of energy and momentum. In HSG, the exciton must first be created, (0). Steps (1), (2), and (3) are all much the same as in HHG, but now the hole has finite mass. This figure is not to scale.

simply using a semiclassical three step model proposed by Corkum [14]. This model was made fully quantum using the strong field approximation by Lewenstein, *et al.* [39]. This section follows *Atoms in Intense Laser Fields* by C. J. Joachain, N. J. Kylstra, and R. M. Potvliege for descriptions of HHG and other atomic phenomena [28].

The fundamental physical difference between HHG and HSG is that the exciton must be created before it can be ionized by the driving field, the “zeroth” step of HSG, see Fig. 1.2. This new laser field changes the frequency-order relationship in an HSG spectrum such that $\omega_{\text{SB}} = \omega_{\text{NIR}} + n\omega_{\text{THz}}$. Adding this excitation step means that the properties of the newly-created electron-hole pair can be controlled. The NIR frequency could be

tuned to create excitons, or a free electron-hole pair with non-zero individual momenta. The NIR polarization can be changed to excite electrons and holes with different spin states. Xie, *et al.* predict there will be a plateau at positive *and* negative orders when free electron-hole pairs are created [63].

In the first step, the exciton tunnel-ionizes in the strong driving field, see Fig. 1.2. Ionization via tunneling is not a resonant effect. The Coulomb potential can be warped by a strong enough electric field to actually cause the electron to tunnel through the Coulomb barrier from the hole. Tunnel ionization occurs in a regime defined by the Keldysh parameter,

$$\gamma = \frac{t_{\text{tunnel}}}{T_{\text{drive}}} \quad (1.6)$$

$$= \omega \frac{\sqrt{2\mu E_{\text{bind}}}}{eF} \quad (1.7)$$

where t_{tunnel} is the tunneling time and T_{drive} is the laser period [32]. Recall that in the center of mass frame, the hole is stationary and the electron has the reduced mass $\mu = m_e m_h / (m_e + m_h)$. If $\gamma < 0.5$, the exciton will ionize only via tunnel ionization, [26].

The tunnel-ionization step is largely similar in HHG and HSG as it was generalized by Keldysh. In HSG, however, the THz fields that can be easily achieved are much larger, relative to the Coulomb barrier in an exciton, than the pulsed laser fields are relative to the barrier in atoms. A sinusoidal field of 35 kV/cm in GaAs, easily generated by the UCSB Free Electron Laser, is strong enough to classically ionize, i.e. no tunneling is required, the most tightly bound excitons for the large majority of the field's period.

In the second step, the now-free electron accelerates in the field. As a free charged particle in an oscillating field, the ponderomotive energy becomes the dominant energy scale, defined in Eq. 1.5. Initially, the electron accelerates away from the hole, but eventually the driving field changes sign and drives the electron back towards the hole. Whether the two recollide depends on the phase of the driving field when tunnel ionization occurred, based on solving Newton's equation for the electron:

$$\mu\ddot{x} = eF \cos(\omega t + \varphi) \quad (1.8)$$

$$\dot{x}(t = 0) = 0 \quad (1.9)$$

$$x(t = 0) = 0 \quad (1.10)$$

where $x(t)$ is the position of the electron as a function of time and the initial position and velocity of the electron is zero at $t = 0$, and φ is the phase of the driving electric field at $t = 0$. If φ is in the interval when the electric field magnitude is increasing, $[\pi/2, \pi]$ or $[3\pi/2, 2\pi]$, then the electron gains too much kinetic energy before the driving field turns around, and the electron never recollides with the hole. If tunnel ionization occurs during the other two quarters of the cycle, then the electron can recollide with the hole, possibly several times depending on φ .

The acceleration step is far more complicated in HSG, and a large portion of this dissertation will address what happens during this step. The parent ion in HHG, being thousands of times heavier than the electron, is assumed to be stationary. The hole, however, has a finite mass, so its contribution to the final sideband energy must be

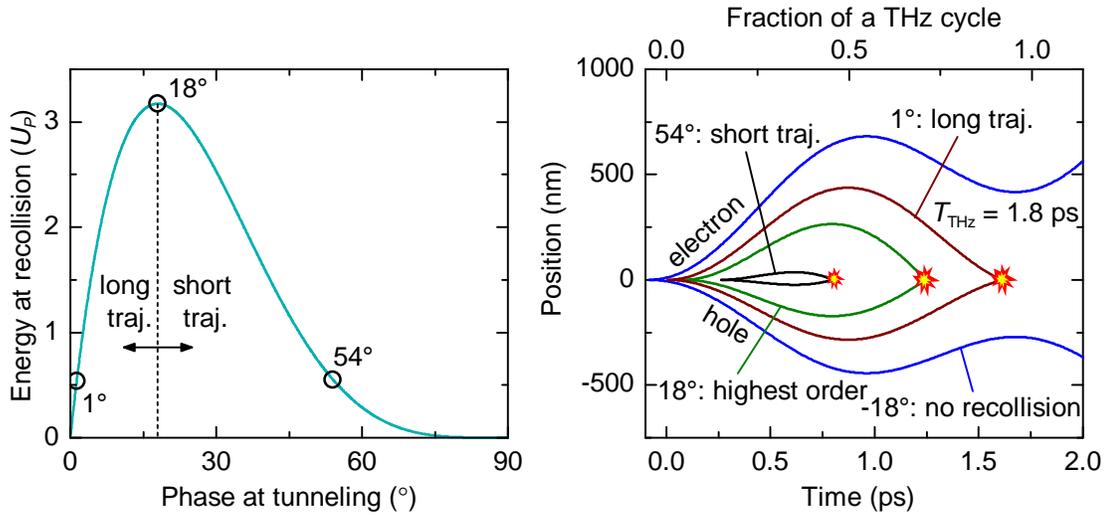


Figure 1.3: Solutions of the classical three step model in GaAs. (left) The kinetic energy at recollision as a function of the phase φ of the driving electric field at tunneling, $F_{\text{driving}} \propto \cos(\omega t + \varphi)$, in units of the ponderomotive energy U_P . The electron and hole have several opportunities to collide, this derivation assumes that they do not survive past the first and that the bands are parabolic. Every solution here has two possible initial phases, differentiated by their trajectories, long and short. (right) Different possible trajectories based on tunneling time for $F_{\text{THz}} = 11$ kV/cm, $f_{\text{THz}} = 570$ GHz, $m_e = 0.063$ and $m_h = 0.1$. For $\varphi = 18^\circ$, the electron and hole never recollide. For $\varphi = 1^\circ$ or $\varphi = 54^\circ$, the recollision has the same total kinetic energy, but the trajectories are very different. The long trajectory explores almost $1 \mu\text{m}$ of space and lasts over 1 ps. For $\varphi = 18^\circ$, the maximum energy trajectory, the whole duration is relatively short compared to the $\varphi = 1^\circ$ trajectory.

considered at recollision. The hole contribution can be accounted for by transforming into the center of mass frame and using the reduced mass, a fairly trivial though important adjustment. As the electron and hole accelerate in the field, they can interact with the host lattice, utterly unlike in HHG. During their trajectories, they gain a substantial amount of kinetic energy, so there are many mechanisms for them to relax back across the band non-radiatively. Accounting for scattering will be discussed in Chap. 2. The large kinetic energies also mean that complicated subband mixing influences recollisions. A substantial amount of theoretical analysis has gone into this step as well to understand the possible observation of Berry phase physics and other quantum coherence effects [68, 16, 69, 67]. Accounting for these effects will be discussed in Chap. 4.

In the third and final step, the electron and hole recollide. Being particle and antiparticle, the electron and hole can radiatively recombine, emitting a photon. Part of that emitted photon's energy comes from reverse ionization, which explains why E_{bind} appears in the expression for the cutoff order, although in HSG reverse ionization is simply energy detuning of the NIR excitation laser from the band gap. The other part comes from the kinetic energy gained in the driving field. If the kinetic energy at recollision is plotted as a function of φ and the bands are parabolic, there is a peak in recollision kinetic energy at approximately $\varphi = 18^\circ$, where the kinetic energy $E_K = 3.2U_P$, see Fig. 1.3(left). This result is in very close agreement with the empirical cutoff law. Because of the complex valence band of GaAs and again utterly unlike HHG, it is possible for the hole to change state during its trajectory such that it can no longer radiatively recombine with the

electron.

This model also explains why the harmonic spectrum has a plateau in energy. During one full cycle of the driving field, there are four instants when an electron can tunnel out to create a harmonic of a given order. All of the highest orders then have the same number of opportunities to be created. One factor of two comes from the inversion symmetry in the system, that the positive and negative halves of the optical cycle contribute the same. The other factor of two comes from what are termed the “short” and “long” trajectories, see Fig. 1.3(left). The short trajectory occurs when the electron tunnels out between 18° and 90° after the peak of the driving field. The electron gains a small amount of kinetic energy before turning around and recolliding, and the whole trajectory occurs in less than half a cycle, see Fig. 1.3(right). The electron has the most kinetic energy at recollision. The long trajectory occurs when the electron tunnels out between 0° and 18° after the peak of the driving field. These trajectories can take more than a half-cycle to complete, and the electron can gain substantially more kinetic energy during the trajectory than it has at recollision.

Recent pulsed THz work

The observation of recollision physics in solids is not limited to GaAs and free electron lasers. Strong, subresonant fields have created HHG in bulk ZnO, GaSe, and WSe₂ [22, 51, 38]. In these works, coherent populations of electrons tunnel across the band gap when driven with very strong pulsed mid-infrared or THz pulses. Harmonic photons are then created via two pathways, non-parabolic band structure, including Bloch oscillations,

and electron-hole recollisions, depending on the frequency and field strengths involved. Langer, *et al.*, even observed HSG in WSe₂ using a NIR excitation pulse and a THz driving pulse [38]. By using a very short NIR pulse time-delayed through the THz pulse, a slight delay between NIR excitation and sideband emission was observed, consistent with the delay predicted by the acceleration step in the three step model.

1.4 GaAs quantum wells

To study excitons in GaAs, it is useful to build quantum wells (QWs) into the structure of the material. Growing QWs restricts the effective dimensionality of the electrons and holes, affecting both the band and the exciton structure in nontrivial ways. The binding energy of excitons, for example, can increase by more than a factor of two based on confinement alone, $E_{\text{bind}} \sim 10$ meV for 5–10 nm QWs. Development in growth technologies like molecular-beam epitaxy allows for controlled deposition of single atomic layers in a large plane. For GaAs, the very highest quality quantum wells are grown alternating several nanometers of GaAs and Al_{0.3}Ga_{0.7}As. Both materials have a direct band gap and very similar lattice constants, but Al_{0.3}Ga_{0.7}As has a larger band gap by about 400 meV. About 65% of the increase is in the conduction band, and 35% is in the valence band. Electrons at the bottom of the conduction band in the GaAs well region are not allowed to propagate into the AlGaAs barriers, the z -direction, and so electrons are restricted to motion in the x - y plane.

Breaking the semi-continuous lattice symmetry in the growth direction forces electrons

and holes into a discrete energy spectrum. The simplest model for this situation is the QW with infinite barriers,

$$V(z) = \begin{cases} +\infty & z \leq -L/2 \\ 0 & -L/2 < z < L/2 \\ +\infty & z \geq L/2 \end{cases}$$

where $V(z)$ is the potential and the QW has a width L . This potential requires any wavefunctions to be free particles in the well and zero outside of it. The full set of orthonormal wavefunctions $\psi_n(z)$ and their energies E_n ,

$$\psi_n(z) = \begin{cases} \sqrt{\frac{2}{L}} \sin\left(\frac{n\pi}{L}z\right) & n \in 1, 3, 5, \dots \\ \sqrt{\frac{2}{L}} \cos\left(\frac{n\pi}{L}z\right) & n \in 2, 4, 6, \dots \end{cases}$$

$$E_n = \frac{n^2\pi^2\hbar^2}{2mL^2}.$$

There are still two other dimensions, x and y , to add to the complete wavefunction. Because the potential does not depend on x or y , it is separable and the the wavefunctions are plane waves in the plane of the quantum well. The total energy of the particles then is the sum of the energies in all three dimensions,

$$E_n(k_x, k_y) = \frac{n^2\pi^2\hbar^2}{2mL^2} + \frac{\hbar^2}{2m} (k_x^2 + k_y^2).$$

The total energy of the particle is actually continuous, because k_x and k_y are continuous, even though the z subbands are discrete. When the barriers are finite, the wavefunctions and energies require numerical calculations. Lowering the barriers has important effects on the properties of the wavefunctions. The energy of a subband still scales approximately as n^2 , but the number of bound states is finite.

Calculating the electronic states in a GaAs QW requires careful consideration of the band physics and the QW physics. Well widths in GaAs are generally tens of nanometers, but they can be as small as a few nm. Even at their smallest, the wells are several atomic layers thick, so the QWs perturb electron states at a much smaller energy scale than the band gap. For GaAs/Al_{0.3}Ga_{0.7}As QWs, confinement modifies the conduction band in a straightforward manner. The two electron states only differ by spin direction, which does not couple to the confinement potential. The energy at $k = 0$ is raised above the GaAs conduction band due to the confinement energy, but the effective mass remains approximately the same. Electron states in different subbands simply stack upwards in energy. The valence band, however, changes dramatically as a result of confinement. First, the two hole states will shift down in energy (“up” from a hole’s perspective). Taken with the increase in the conduction band energy, the band gap will increase relative to bulk material. The two hole states couple differently to the confinement potential, however, because their masses are different. The QW then breaks the heavy hole-light hole degeneracy at $k = 0$. The holes are p -like, so the effective masses are actually anisotropic. A hole excited in a p_x orbital has good overlap with nearby p_x orbitals in

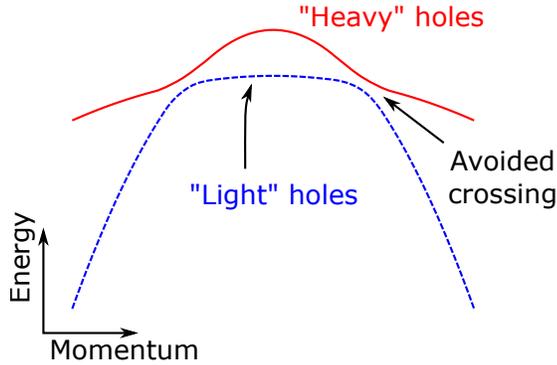


Figure 1.4: Schematic valence band structure in the plane of the QW. The confinement energy at $k = 0$, the peak of the subbands, is determined by the effective mass in the growth direction, so heavy holes do not shift downward as much compared to light holes. The effective mass normal to the growth direction, in the plane of the QW, changes from heavy to light, and vice versa. Because the hole states are coupled, the bands do not cross but avoid each other. The dynamics of holes accelerating through the avoided crossing must be calculated carefully.

the x -direction, and so it is classified as “light.” That same p_x hole, however, has poor overlap with the neighboring orbitals in the y - and z -directions, and so it is “heavy.” The energy level for the QW is based on the z -direction mass. In the plane, the light and heavy holes “switch” in a sense. The dispersion relations then actually will cross at some finite momentum near $k = 0$, but because the two holes are coupled, there will be an avoided crossing, see Fig. 1.4.

Optical interband transitions in QWs follow three selection rules. First, the transitions must be vertical on the scale of the Brillouin zone because the photon momentum is so small, which is a general rule of optical transitions in semiconductors. Next, also similar to bulk GaAs transitions, the the spin states must follow the standard electric dipole transition rule where the angular momentum of the states must change by one. The hole states $|3/2, +3/2\rangle$ and $|3/2, -1/2\rangle$ therefore couple to the electron state $|S, \uparrow\rangle$.

Finally, unique to interband transitions in QWs, the subband quantum number n must not change to ensure good wavefunction overlap. This rule is not strictly true, however. Because the valence band and conduction band QW Hamiltonians have different masses and barrier heights, different subband envelope functions are not strictly orthogonal, although requiring $\Delta n = 0$ is typically a good approximation.

1.5 Photoluminescence and absorption

With the knowledge of the types of states in bulk or heterostructured GaAs on paper, it is important to be able to count the number of those states at a given energy E per unit volume. This number is called the density of states $n(E)$. Measuring the density of states provides information about the quality of a sample. Dirty or imperfectly-grown material will manifest as unexpected features. In general, the density of states for an arbitrary system with available energies ε_i is,

$$n(E) = \frac{1}{\Omega} \sum_i \delta(E - \varepsilon_i),$$

where $1/\Omega$ is the volume normalization factor. For a free fermion gas, which is a good approximation for the electrons and holes near the band edge, the density of states

depends strongly on the dimensionality of the system,

$$n_{3\text{D}} = \frac{m\sqrt{2mE}}{\pi^2\hbar^3}$$

$$n_{2\text{D}} = \frac{m}{\pi\hbar^2}.$$

The density of states is particularly useful because when several bands or transitions contribute, their respective densities of states simply superpose. In undoped bulk GaAs, physical density of states will have a small peak just below the E_{gap} on top of the $n_{3\text{D}}$.

Optical measurements are only sensitive to transitions between these states, which requires a slight modification of $n(E)$. The optical joint density of states $n_{\text{opt}}(\hbar\omega)$ is then,

$$n_{\text{opt}}(\hbar\omega) = \frac{2}{\Omega} \sum_{i,j,\mathbf{k}} \delta(E_{\text{gap}} + \varepsilon_{v,i}(\mathbf{k}) + \varepsilon_{c,j}(\mathbf{k}) - \hbar\omega),$$

where $\varepsilon_{v,i}(\mathbf{k})$ and $\varepsilon_{c,j}(\mathbf{k})$ are the energies of the hole in state i and electron in state j , respectively. They are both at the same momentum because the photon has negligibly small momentum.

The optical properties of GaAs QWs are traditionally studied using two measurements, photoluminescence (PL) and optical absorbance. In a PL measurement, electron-hole pairs are injected with an energy much larger than the band gap of the material, generally using a laser. If the lifetime of the two particles is longer than the scattering time, then they will fall down to a state near their respective band edges. The ensemble

of electron and hole pairs will come to thermal equilibrium with the lattice before somehow recombining. The particles must recombine radiatively if they are to be observed. PL does not, therefore, directly measure the density of states of the lowest radiative transition(s). Instead it measures the optical density of states modified by a Boltzmann factor from the electron temperature,

$$S_{\text{PL}}(\omega) \sim n_{\text{opt}}(\omega) \cdot e^{-\omega/kT}, \quad (1.11)$$

where $S_{\text{PL}}(\omega)$ is the PL signal as a function of emission photon energy ω . The equation also includes factors related to the strength of the transition at photon energy ω , but those have been ignored as the transitions in focus in this dissertation have similar dipole moments. The Boltzmann weighting is asymmetric, it favors the lower energy side of otherwise symmetric peaks. At very low temperatures in undoped GaAs-based quantum wells ($T < 20$ K), the heavy hole exciton will be the primary feature in PL [29, 1]. As the temperature increases, the high-energy tail of the Boltzmann factor will get longer, making higher-energy states such as the light hole exciton or free electron states apparent [9, 54]. Eventually, as the temperature increases, the PL signal will decrease substantially because injected electrons and holes will be much more likely to recombine nonradiatively.

Another useful spectrum is the optical absorbance of a sample. Here, a broadband light source, such as an incandescent lamp or pulsed laser perhaps, illuminates a sample, usually around the band gap. Light that gets absorbed, to create an exciton or a free electron-hole pair, will either get re-radiated out of the beam path or the quasiparticles

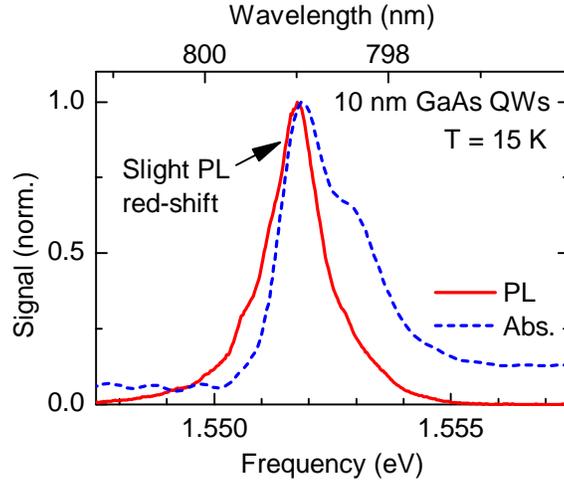


Figure 1.5: Comparison of PL and absorbance spectra in QWs. The PL and absorbance data for 10 nm GaAs QWs are plotted, normalized to the peak of respective signal. The PL signal is slightly red-shifted from the absorbance peak due to weighting by the Boltzmann factor.

will relax non-radiatively. Because the light that is detected was not absorbed and re-radiated, the absorbance signal S_{abs} is directly proportional to the optical joint density of states without thermal weighting factors, see Fig. 1.5. Relatedly, the absorbance signal is only sensitive to the absorbed light; the re-radiated light is not measured. These measurements are then very applicable to resonant excitation experiments, the same process that absorbance measures, such as HSG.

1.6 UCSB Free Electron Laser

Entering the regime of such nonlinear THz physics requires the use of the UCSB Free Electron Laser (FEL). The FEL generates THz light by steering a 1 A beam of electrons with up to 5 MeV through an permanent magnet undulator inside a high-Q laser cavity

[48]. The FEL cavity is approximately 6 m long, leading to a longitudinal mode spacing of 25 MHz. Before the electron beam turns on, those modes are populated with thermal radiation. The early electrons in the 5 μ s-long electron beam emit synchrotron radiation as they pass through the undulator and into the laser cavity at a frequency based on their kinetic energy and the magnetic period of the undulator in a 1 GHz band. This radiation interacts with the thermally populated modes, amplifying them. As the fields in the modes get stronger, later electrons become more likely to radiate in phase with the radiation in the cavity. This coherent amplification increases the power in the cavity as in a more traditional laser, and, if the electron beam is long enough, the FEL lases. The initial thermal modes and the quality of the electron beam determine the final frequency and phase properties of the laser emission. There are currently two undulator/laser cavities in the FEL electron beamline, the Far Infrared laser, tunable from 1–4 THz, and the Millimeter (MM) laser, tunable from 200–800 GHz. The MM laser was used for all experiments in this thesis. The laser light is directed from the FEL experiments using a periodic focusing transport system that is pumped to about 200 mtorr to minimize atmospheric absorption. The lenses and windows of the transport system are made of a transparent plastic.

There are two mechanisms for coupling the THz power out of the FEL, [30, 48]. The most straightforward is the hole coupler. Because of the relatively long wavelengths of the THz radiation emitted by the FEL, the traditional “partially-silvered end mirror” of the cavity is actually a curved metal mirror with a small hole bored through the center, the

“hole coupler.” The THz pulse duration from the hole coupler is up to about $1.5 \mu\text{s}$ long with a peak power of about 5 kW. Because the hole coupler is a circle, the beam profile in the transport system is nearly Gaussian. The beam profile also does not change over the course of the pulse. The other output coupler is the silicon coupler “cavity dump” (CD) [53]. A thin silicon wafer is placed at a beam waist in the FEL cavity. The wafer is oriented at nearly Brewster’s angle to minimize the reflection out of the cavity. Once the FEL is lasing, a Q-switched 532 nm pulsed laser fires at the silicon wafer, exciting an electron-hole plasma in the silicon with a plasma frequency above the lasing frequency, creating a temporary mirror. This mirror couples all of the THz light out of the cavity. The THz pulse duration from the cavity dump is 40 ns long, twice the length of the cavity, with a peak power of about 60 kW. A pulse duration of 40 ns is relatively short given the 1 Hz repetition rate of the FEL, but it is about 10,000 periods of the THz light, making it a quasi-continuous wave source. Because the silicon wafer changes from insulating to pseudo-metallic, its reflection properties change dramatically. Recall, the silicon wafer is mounted at approximately Brewster’s angle, but at a beam waist. The incident light hits it at a range of angles, sweeping across Brewster’s angle. The electron-hole plasma, however, reflects all angles approximately equally, so the beam profile is much more Gaussian during the cavity dump than before it. This spatial profile change has important consequences when measuring the instantaneous power in the cavity dump FEL pulse. Failure to account for this change can lead to power measurements that are off by as much as 50%.

Chapter 2

The electron-lattice interaction in the three step model

Unlike in high-order harmonic generation (HHG), where the electron and ion accelerate through free space, the electron and hole in high-order sideband generation (HSG) must accelerate through a complex and dynamic lattice. Phonon interactions provide a particularly strong mechanism for the quasiparticles to lose coherence and energy to the lattice. There are two distinct phonon classes in III-V semiconductors like GaAs, acoustic phonons and optical phonons. For the energy scales of HSG, fortunately, acoustic phonons actually should not have a large effect, with scattering rates on the order of 1 THz due to their linear dispersion relation making momentum conservation in electron-acoustic phonons scattering difficult to achieve [33, 49, 17]. Optical phonons, however, interact very strongly with electrons and holes, with scattering rates on the order of 10

THz, much faster than the 0.5 THz driving frequency, because of their relatively large energy (36 meV for longitudinal optical phonons) and flat dispersion relation [4]. The thermal population of optical phonons is only significant above liquid nitrogen temperatures, but very hot charge carriers can emit such phonons very efficiently. It is, therefore, reasonable to think that the full plateau of the three step model will not be observed because of optical phonon scattering. Previous experiments only observed sidebands up to almost exactly this 36 meV cutoff, but the host InGaAs quantum wells were highly strained [73].

Here, we investigate the consequences of a dynamic lattice in HSG using clean, strain-free GaAs/AlGaAs quantum wells. By minimizing other scattering pathways, we are able to focus on the the electron-optical phonon interaction. Furthermore, the exciton absorption wavelength of these new quantum wells fit in the tuning range of the Ti:sapphire laser from low temperature (15 K) up to room temperature. This temperature range allows for control of the thermal population of optical phonons, meaning control of the scattering rate that the electrons and holes experience during the acceleration step of HSG.

2.1 Sample details

2.1.1 Sample design

Two quantum well samples with twenty QW repetitions were grown for this experiment. Both samples were grown on semi-insulating GaAs substrates by molecular beam epitaxy with no intentional doping. The 5 nm GaAs QW sample consists of 5 nm GaAs QWs with 20 nm $\text{Al}_{0.3}\text{Ga}_{0.7}\text{As}$ barriers and the 10 nm QW sample consists of 10 nm $\text{Al}_{0.05}\text{Ga}_{0.95}\text{As}$ QWs with 20 nm $\text{Al}_{0.3}\text{Ga}_{0.7}\text{As}$ barriers. For the experimental convenience of matching excitonic absorption peaks in the different samples, the band gap of the wider well was raised by adding 5% aluminum to match the 2D band gap of the narrower well. The growth machine was not calibrated to grow 5% AlGaAs directly. Instead, it discretely deposited five parts GaAs then one part 30% AlGaAs. Growing twenty QWs increases the absorption of NIR photons, increasing the number of excitons created and, presumably, the number of sideband photons created. Because quantum confinement raises the exciton absorption lines above the band gap of bulk GaAs, optical transmission experiments like sideband generation and absorption experiments will be entirely blocked by the substrate unless it is removed.

2.1.2 Sample processing

The substrate was removed by a series of etches. Etching through 500 microns of substrate requires fast etch but highly selective etch rates. A series of three wet etches were

	50 nm GaAs	50 nm GaAs	
	40 nm Al _{0.3} Ga _{0.7} As	40 nm Al _{0.3} Ga _{0.7} As	
	20 nm Al _{0.3} Ga _{0.7} As	20 nm Al _{0.3} Ga _{0.7} As	} 20×
QW region	5 nm GaAs	10 nm Al _{0.05} Ga _{0.95} As	
	10 nm Al _{0.3} Ga _{0.7} As	10 nm Al _{0.3} Ga _{0.7} As	
	50 nm Al _{0.3} Ga _{0.7} As	50 nm Al _{0.3} Ga _{0.7} As	
	50 nm GaAs	50 nm GaAs	
Etch stop region	300 nm Al _{0.7} Ga _{0.3} As	300 nm Al _{0.7} Ga _{0.3} As	
	50 nm GaAs	50 nm GaAs	
Substrate	500 μm GaAs	500 μm GaAs	

Table 2.1: The growth structure of the two samples. The only difference is the material in the quantum well region. Both samples are grown analog except for the Al_{0.05}Ga_{0.95}As layers in the 10 nm QW sample, which were grown digitally. The 5% aluminum alloying raises the band gap of the quantum well to cancel out the smaller confinement energy of the wide quantum well.

used. First, diluted H₂SO₄ and H₂O₂ etches Al_xGa_{1-x}As quickly regardless of aluminum content. Then, a mixture of citric acid and H₂O₂ etches until an etch stop layer of Al_{0.7}Ga_{0.3}As [45, 10]. Finally, the etch stop layer is removed with buffered HF. The final result of the process was a sample approximately 5 × 7 mm in size with a thin membrane “window” approximately 1 × 2 mm in size less than 1 μm thin, supported by the surrounding substrate frame. Full processing details can be found in Appendix C. This membrane is very fragile, but was able to withstand many cycles of vacuum pumping and venting, thermal cycling from room temperature to 12 K, and sideband experiments in air without damage.

2.1.3 Sample optical properties

Both photoluminescence and absorption spectra were measured to characterize the samples, see Fig. 2.1 for the optical schematics. Excitation for PL was a 3 mW HeNe laser

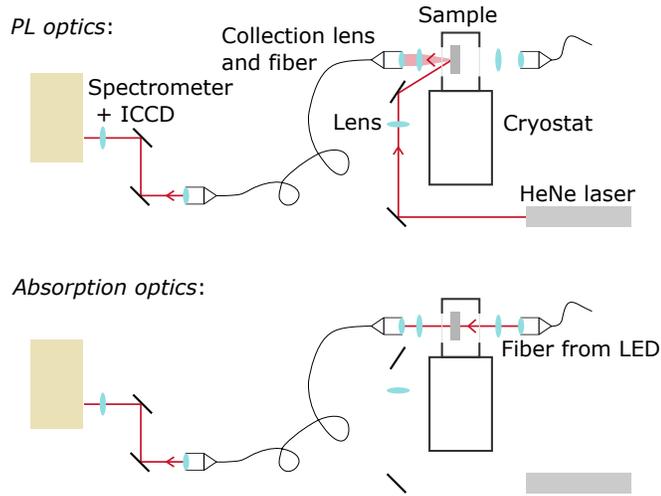


Figure 2.1: Experimental schematic of PL and absorption measurements. *PL optics*: to measure PL, a lens focuses a HeNe laser onto the sample at an angle, the thin red line. A lens oriented normal to the sample collects the incoherent PL signal, the thick, faint red line. The collected light is sent into a fiber that is coupled to an imaging spectrometer to be measured by an intensified CCD. *Absorption optics*: For the absorption measurement, an LED spot centered at 810 nm is focused on the sample through it and into a collecting lens.

focused on the sample at an angle. An objective lens normal to the sample collimated the light into a fiber coupler that brought the light to the intensified CCD. The light for absorption spectra was generated by an 810 nm LED focused on the sample and then collected by an objective lens.

Photoluminescence and absorption spectra of the samples reveal the states near the band edge, see Fig. 2.2. In the PL spectrum at low temperature, the lowest-energy radiative state is the heavy hole exciton, which is the same in both samples, as designed. Excitation for PL was a 3 mW HeNe laser focused on the sample at an angle. An objective lens normal to the sample collimated the light into a fiber coupler that brought the light to the intensified CCD. The two samples can, however, be differentiated by their

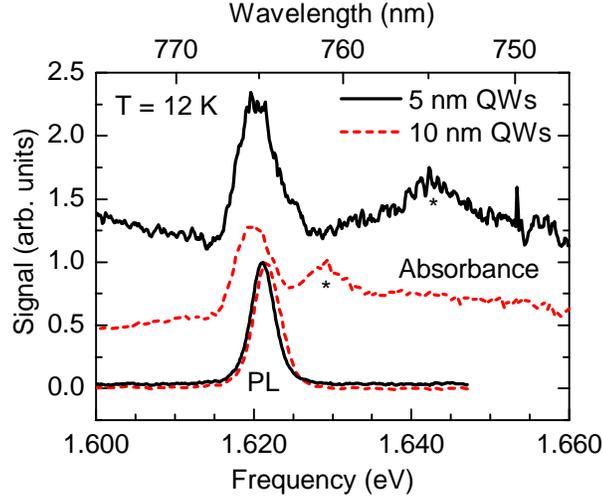


Figure 2.2: The PL and absorption spectrum at low temperature. Both samples have nearly identical PL spectra, as only the heavy hole exciton peak is observed. The absorption spectra of both samples, shifted vertically for clarity, are much different, owing to the observation of the light hole exciton (the higher-energy, smaller peaks, denoted by *). The half linewidth of the heavy hole exciton in both samples is approximately 2.8 meV. The data was published in [4]. Copyright 2013 by the American Physical Society.

absorption spectra. The light for absorption spectra was generated by an LED, centered at 810 nm, focused on the sample and measured similar to PL. The heavy hole-light hole (HH-LH) splitting in the valence band scales as the ratio of the effective mass to the square of the well width, so the 5 nm GaAs QWs have a 25 meV splitting while the 10 nm AlGaAs QWs have only 10 meV splitting.

The PL and absorption spectra can be followed up to higher temperatures, see Fig. 2.3. As temperature increases, band gap of GaAs redshifts approximately quadratically, and the $\text{Al}_{0.3}\text{Ga}_{0.7}\text{As}$ shifts similarly. The exciton lines from both experiments follow that redshift very closely. Photoluminescence peaks become too weak to be detected by about 200 K, because the photoinjected electrons and holes relax nonradiatively much more efficiently at high temperatures. Excitons in absorption measurements are still detectable

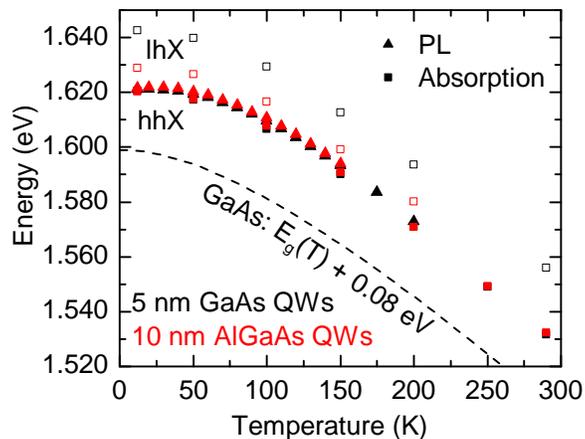


Figure 2.3: Temperature dependence of the exciton peaks measured by PL and absorption. The heavy hole exciton (HHX, closed points) peaks and light hole exciton (LHX, open points) redshift with increasing temperature as measured both by PL and absorption. Exciton peaks in PL are no longer observed above approximately 200 K, but absorption remains observable up to room temperature. The trend almost exactly following the bulk GaAs band gap dependence (plotted shifted up 0.08 eV for clarity). Because the hole masses do not depend on temperature, the HH-LH splitting remains constant through the temperature range, although the light hole excitons become difficult to resolve at the higher temperatures. The data was published in the Supplementary Materials of [4]. Copyright 2013 by the American Physical Society.

up to room temperature, see Fig. 2.4. The heavy hole-light hole splitting remains constant with temperature, at least as far as light hole excitons are observable.

2.2 Sideband measurements

2.2.1 Optical setup

Sidebands were generated by collinear illumination of the sample with a NIR and THz beam. The NIR beam was generated a continuous-wave Ti:sapphire laser (Coherent 890). The THz beam was generated by cavity dumping the millimeter beam line of the FEL

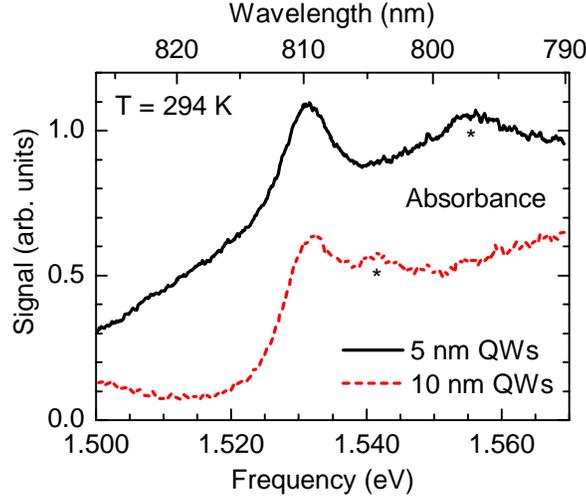


Figure 2.4: Room temperature absorption spectra. The heavy hole exciton lines still dominate the absorption spectra of both samples, but the light hole excitons (denoted by the *), are still observable. The half linewidth of the heavy hole exciton line is 4 meV in the 5 nm GaAs QWs and 5 meV in the 10 nm AlGaAs QWs. The data was published in [4]. Copyright 2013 by the American Physical Society.

tuned to 570 GHz (2.35 meV). This frequency was chosen because of its high power and the accelerator's stability. The laser field from both sources was linearly copolarized and focused onto the same spot. Radii were approximately $100 \mu\text{m}$ and 0.5 mm for the NIR and THz spots, respectively. The strength of the terahertz electric field was 11.5 kV/cm inside the quantum wells. The power of the NIR laser was 39 mW (60 mW) at $T = 12$ K (294 K). The transmitted NIR laser power and generated sidebands were sent via free space to a monochromator and detected by a photomultiplier tube (PMT). Sidebands with order $n \leq 16$ were measured by the average current output from the PMT during the 40 ns FEL cavity dump pulse. Sidebands of order $n > 16$ were too weak to detect by averaging the current, so the PMT signal from sideband orders $n \geq 14$ was read out with a high-speed photon counter. The overlap between analog readout and photon

counting scaled the two together. Finally, the sample was removed from the NIR beam path and the laser line was heavily attenuated and measured with the PMT to determine the sideband creation efficiency for the full sideband spectrum measurements.

2.2.2 Results

Sideband spectra spanning two dozen orders and tens of meV were observed from both samples. In the 5 nm GaAs QW sample at 12 K, we observed down to -8th order below and up to 28th order above the NIR laser, see Fig. 2.6(top left). The intensity of the sidebands with negative order decreases exponentially with n . The intensity of the sidebands with positive order decreases much more gently with no clear systematic dependence on n . We observed a similar spectrum of sidebands in the 10 nm AlGaAs QW sample, 2.6(top right). However, the intensity of each sideband is approximately 3 times smaller than the corresponding sideband in the 5 nm GaAs QW sample, and the highest sideband is 22nd order rather than 28th.

Remarkably, HSG persists to room temperature in both samples. Sidebands with orders as high as 22nd and 18th for the 5 nm GaAs and 10 nm AlGaAs QW samples, respectively, were observed 2.6(bottom row). The sideband conversion efficiencies are approximately one order of magnitude smaller at 294 K compared to 12 K, but otherwise the HSG spectra have similar shapes.

The NIR frequency, f_{NIR} dependence was also measured to both peak up the sideband signal, and as an initial experiment to understand the “zeroth” step of HSG, exciton

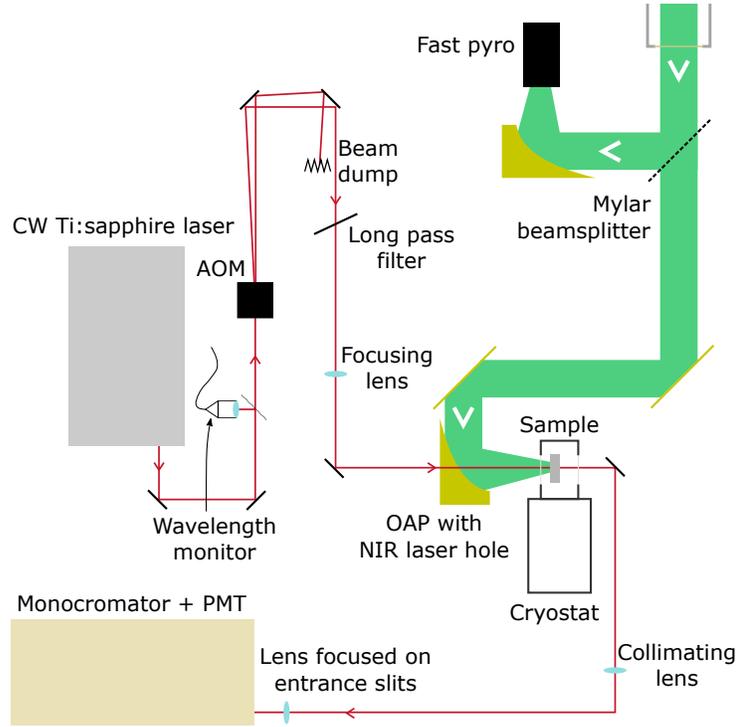


Figure 2.5: Experimental schematic of the sideband measurement. First, we follow the path of the excitation NIR light, the thin red line. The Ti:sapphire laser emits a continuous NIR beam. A pellicle splits off a few percent of the light and into a fiber-coupled spectrometer to measure the NIR wavelength. To keep the laser from overheating the sample, the laser is electronically shuttered by an AOM, triggered only for the duration of the FEL pulse. The zeroth order beam is sent into a beam dump. The NIR laser is then focused by a lens before traveling through a small hole bored in the off-axis parabolic mirror (OAP) to become collinear with the THz beam. The cryostat window is made of a NIR- and THz-transparent plastic, TPX. The laser spot on the sample is about $200\ \mu\text{m}$ in diameter. The THz path, the thick green line with white arrows, leaves the FEL transport system exit port. Immediately about 5% of the beam is picked off by a thin mylar beamsplitter and focused with an OAP onto a fast pyroelectric detector to monitor the THz pulse shape. The majority of the beam gets focused onto the sample by a 5" OAP onto the same location as the NIR beam. Sidebands are then sent through a collimating lens before being focused onto a monochromator and detected a PMT.

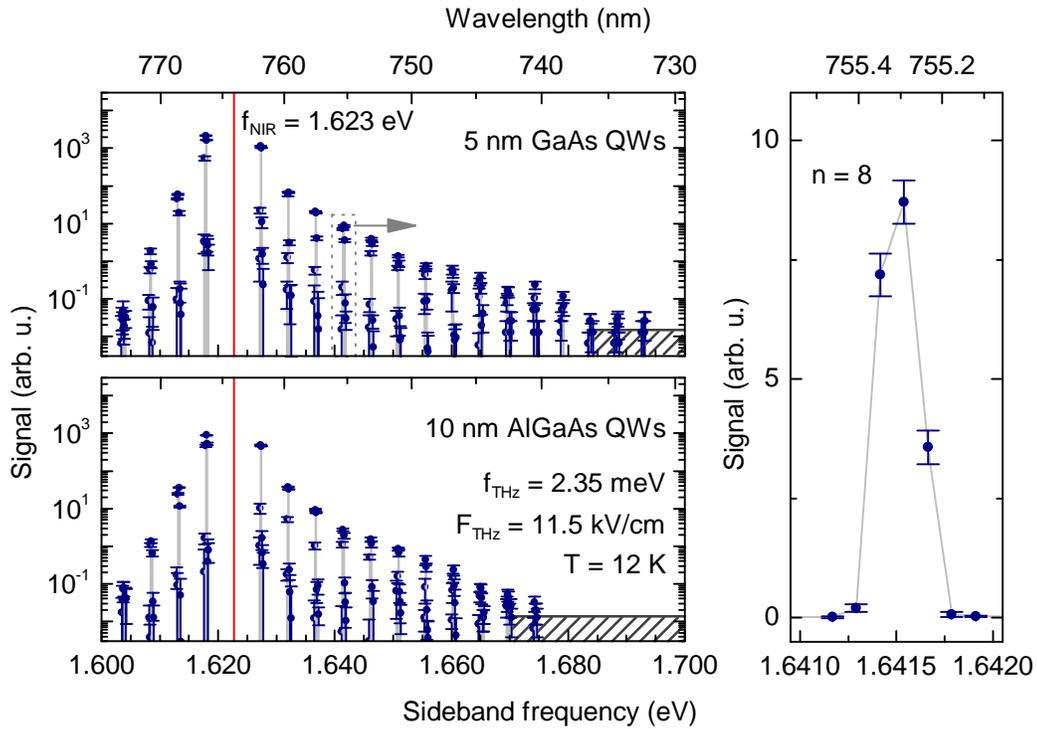


Figure 2.6: Measured sideband spectra at $T = 12$ K. The sideband sideband strength read out by the PMT versus frequency for the 5 nm QW sample at $T = 12$ K (top left) and for the 10 nm QW sample at $T = 12$ K (bottom left) are plotted as blue circles. The hatched region represents the noise floor. Sideband order $n = 8$ is plotted on the left to show the narrow width of the individual sidebands. This data was published in [4]. Copyright 2013 by the American Physical Society.

creation, see Fig. 2.7. The NIR laser frequency was scanned across as many as 60 meV (30 nm) in approximately 1.5 meV steps. At every f_{NIR} step, sideband orders $n = 2, 4,$ and 6 were measured. At room temperature, the NIR laser could scan over the entire exciton absorption peaks region. At low temperature, however, the laser was unable to tune over the 5 nm GaAs QW light hole exciton. In all samples at all temperatures, there is no evidence of sideband generation when f_{NIR} is deep in the band gap. Without a finite density of states, no excitons can be created for HSG. Even $\chi^{(3)}$ processes will be weak because there is no resonant enhancement even from exciting the excitons. The high-frequency “in-the-band tail” is much longer than the low-frequency “in-the-gap” tail because electron and holes are still generated, though not bound as excitons but as free pairs. They can even still be driven to recollide, but with a lower overall probability [63].

The f_{NIR} dependence near the exciton lines has many peaks depending on the sideband order for both samples and temperatures. The $n = 2$ sideband at low temperature in both samples is peaked approximately 5 meV below the heavy hole exciton (HHX) absorption peak, see 2.7(a, c). The HHX peak will redshift slightly in the strong THz field due to AC Stark effect [46], but not enough to explain the full 5 meV. The THz frequency is 2.35 meV, so the 5 meV offset is likely due to resonant enhancement of sideband generation by the HHX absorption peak and has been seen before in HSG from bulk GaAs and perturbative sideband generation in QWs [34, 61, 71] As the sideband order increases, though, a secondary peak grows midway between the two exciton absorption lines in both

samples, eventually becoming the dominant peak, likely because resonant enhancement is not as large an effect at higher orders. At room temperature, the two peaks are smeared out. In the 5 nm GaAs QWs, all three orders are peaked just below the heavy hole exciton, and the signal falls off slowly as f_{NIR} is tuned across the light hole exciton. In the 10 nm AlGaAs QWs, the high-frequency tail falls off much more quickly because the hole splitting is much smaller.

2.3 Theoretical model

2.3.1 Three step model: Classical interpretation

Many features of the data can be explained qualitatively by adapting the three-step model for HHG from atoms in intense laser fields to the phenomenon of HSG. Step zero, which is not present in the case of HHG for atoms, is the creation of the electron-hole pairs by the NIR laser at a particular phase of the THz field. For these experiments, the NIR laser is tuned to the exciton resonance stemming from the highest heavy hole subband and the lowest electron subband, thus creating carriers with in-plane masses of $m_e^* = 0.067m_e$ and $m_h = 0.103m_e$ for the electron and hole, respectively (The hole mass is appropriate for a 5 nm GaAs QW within the Kohn-Luttinger model [42]). Figure 2.8 illustrates the next three steps that lead to the emission of the 28th order sideband. In the first step, the exciton tunnel ionizes in the strong THz field at a phase 54° after the maximum of the terahertz field. At this point, the carriers are assumed to be point particles with zero

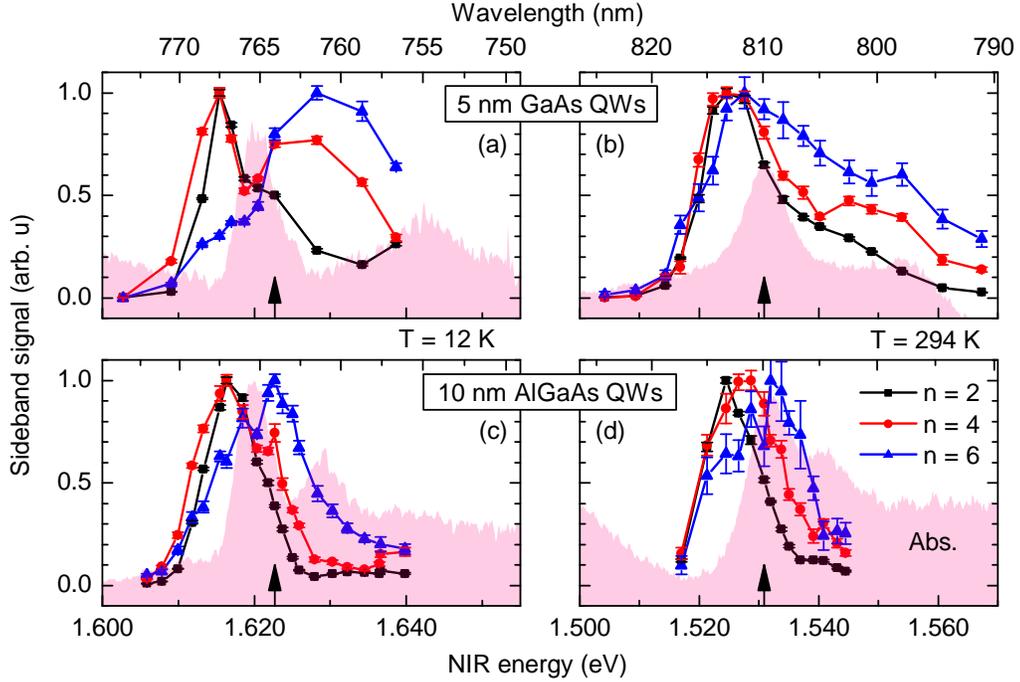


Figure 2.7: Low-order sidebands vs. NIR frequency at low temperature and room temperature. For each plot, the low order sideband spectra were taken at several NIR frequencies around the two exciton lines of the given sample. The zero-field absorption spectrum is the filled region. The sideband data for each order are normalized to strongest of that order. (a) The sideband signal is peaked near the heavy hole exciton line of the 5 nm GaAs QWs at low temperature. The $n = 2$ sideband is most strongly peaked when f_{NIR} is about 5 meV below the heavy hole exciton, but $n = 4$ and 6 are peaked very broadly 10 meV above the heavy hole exciton. Sideband creation drops off quickly as f_{NIR} is tuned into the band gap. (b) The same sample at room temperature has a much broader peak around the heavy hole exciton that is similar in each sideband. (c, d) The sideband signal in the 10 nm AlGaAs QWs shows a very similar f_{NIR} dependence to the narrow quantum wells, but the peaked region is much smaller because the heavy hole-light hole splitting is much smaller. The small black arrow denotes the f_{NIR} used for full sideband spectrum measurements, Figs. 2.6 and 2.10.

initial momentum. The Keldysh parameter, $\gamma = \omega_{\text{THz}}\sqrt{2\mu^*E_{\text{bind}}}/(eF_{\text{THz}})$, the ratio of the average tunneling time to the period of the driving field, is 0.2, well within the regime of tunnel ionization [32, 26]. Here, $\mu^* = 0.041m_e$ is the reduced mass of the electron-hole pair, E_{bind} is the exciton binding energy, e is the electric charge, and ω_{THz} and F_{THz} are the angular frequency and electric field strength of the terahertz radiation, respectively. In the second step, the electron and hole are assumed to be independent and to accelerate in the THz field classically. In the third step, the electron and hole recollide to emit an NIR photon with larger energy than the photon that created them.

The high-order sidebands with positive n can be analyzed to extract the kinetic energy of the electron-hole system at recollision. The sideband energy offset is given by $E_{\text{SB}} = E_{K,e} + E_{K,h} + \Delta$, where $E_{K,i}$ is the kinetic energy for the i th particle and Δ is the detuning of the NIR laser from the 2D band gap. Because this is an excitonic effect, Δ should be close to the binding energy of the exciton, approximately 9 meV for the 5 nm GaAs QWs and 8 meV for the 10 nm AlGaAs QWs [24]. The 28th order sideband is 66 meV above the NIR laser energy. Following this detuning assumption, the total kinetic energy of the electron and hole is ~ 57 meV. This energy is significantly higher than the threshold for a charge carrier to emit a longitudinal optical (LO) phonon. Figure 2.8 illustrates the classical trajectories (top) and kinetic energy (bottom) for the electron and hole associated with the 28th order sideband between ionization at $t = 0$ and recollision. The kinetic energy of the electron-hole system is above the LO phonon energy only for the final 50 fs before reollision. Since this is below the typical electron-LO phonon scattering time

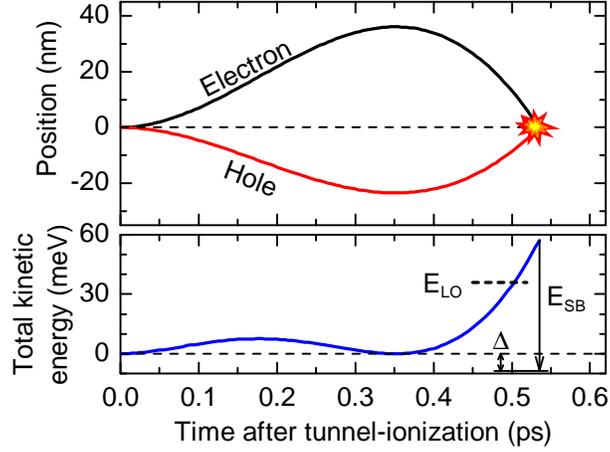


Figure 2.8: Classical calculation of electron-hole dynamics associated with 28th order sideband. Top: trajectories of the electron (black) and hole (red) show electron-hole separation reaches a maximum of 60 nm between tunnel ionization (0 ps) and recollision (flash). Bottom: total kinetic energy of the electron and hole. The total kinetic energy exceeds the threshold for LO phonon emission E_{LO} (short dashed line) for only 50 fs before emitting a sideband with sideband offset energy E_{SB} (long vertical arrow) = Total kinetic energy + detuning Δ (double arrow). The long dashed horizontal lines denote 0 nm (top) and 0 meV (bottom). This calculation was published in [4]. Copyright 2013 by the American Physical Society.

of 100 fs, a large fraction of electron-hole pairs are likely to recollide before they emit a LO phonon—even though the electron-LO phonon scattering time is much shorter than the period of the terahertz radiation. This analysis suggests that the fundamentally ballistic process of HSG persists to room temperature because thermal fluctuations simply do not have enough time to disrupt electron-hole recollisions.

Although there is no abrupt cutoff due to LO phonons, we observe far fewer sidebands than are predicted by the ballistic three-step model of HSG. This model predicts a sharp cutoff for the sideband offset energy: $\hbar\omega_{\text{cutoff}} = 3.17U_P + E_{\text{bind}}$, where U_P is the ponderomotive energy [= $e^2F_{\text{THz}}^2/(4\mu^*\omega_{\text{THz}}^2)$]. Such a cutoff is observed for the harmon-

ics in high-order harmonic generation experiments on atoms. Inserting the parameters appropriate for our experiments, the ballistic three-step model predicts a cutoff 130th order, approximately 4 times larger than the highest-order sideband reported here.

2.3.2 Full theoretical results

To investigate this discrepancy, our collaborators Fan Yang and Renbao Liu have also performed the simplest calculation that captures the relevant quantum mechanical nature of the recollision process together with the most important scattering mechanisms, exciton dephasing and electron-LO phonon scattering. In this calculation, the intensity of the n th order sideband is proportional to the integral in Eq. 2.2, which can be integrated numerically and is plotted in Figs. 2.9 and 2.10.

$$P_n \propto \int e^{iS(\mathbf{k},t,\tau)} e^{-\gamma\tau} e^{-W(d,T,\mathbf{k})\tau} d\mathbf{k} dt d\tau \quad (2.1)$$

$$S(\mathbf{k}, t, \tau) = - \int_{t-\tau}^t \frac{1}{2\mu^*} [\mathbf{k} + e\mathbf{A}(t'')]^2 dt'' - \Delta\tau + n\omega_{\text{THz}}t. \quad (2.2)$$

Here, $S(\mathbf{k}, t, \tau)$ is the action, where μ^* , e and ω_{THz} were defined previously, \mathbf{k} is the momentum (working in units where $\hbar = 1$), and \mathbf{A} is the vector potential of the oscillating terahertz field. The exciton is created at $t-\tau$, the electron and hole evolve through τ , and eventually annihilate at t , emitting sidebands with order n . The detuning Δ is assumed to be E_{bind} . Eq. 1 represents an extension of the model used in Ref. [41] to include the exciton dephasing rate γ , which was estimated from the half-lidewidth of the heavy-hole

absorption peak (2.8 meV for both samples at 12 K; 4 meV and 5 meV for the 5 nm GaAs and 10 nm AlGaAs QWs at 294 K, respectively), and the scattering rate due to emission or absorption of LO phonons $W(d, t, \mathbf{k})$. W is calculated from Ref. [13] and depends on the QW width d , the temperature T and the kinetic energy $|\mathbf{k}|^2/(2\mu^*)$ of the carrier, which is assumed to be an electron with mass μ^* . In the calculation, $\omega_{\text{THz}} = 2\pi \times 0.58$ THz (photon energy $E_{\text{THz}} = 2.4$ meV), $F_{\text{THz}} = 11$ kV/cm, and the photon energy of the NIR laser $E_{\text{NIR}} = 1.623$ eV. The LO phonon energy is taken as 36.6 meV. In order to isolate the effects of scattering on high-order sideband generation, many interactions, such as the Coulomb interaction and heavy hole-light hole mixing, have been neglected.

These calculations refine the simple idea that HSG works “too fast” to be affected by LO phonons. At low temperature, sideband intensities calculated with LO phonon scattering begin to fall below those calculated without such scattering very close to the threshold for LO phonon emission (see Fig. 2.9a). Above this threshold, sideband intensities are suppressed but not eliminated. At room temperature (Fig. 2.9b), all positive sidebands are suppressed because there is now a finite chance of absorbing an LO phonon from the thermal bath. As at low temperatures, sidebands persist above the threshold for LO phonon emission.

Calculations are compared with experimental data in Fig. 2.10. At low temperature, the calculations overestimate the intensities of high-order sidebands for positive n relative to the low-order positive sidebands. At both temperatures, the calculations approximate the observed exponential decay with n of the sidebands below the NIR laser

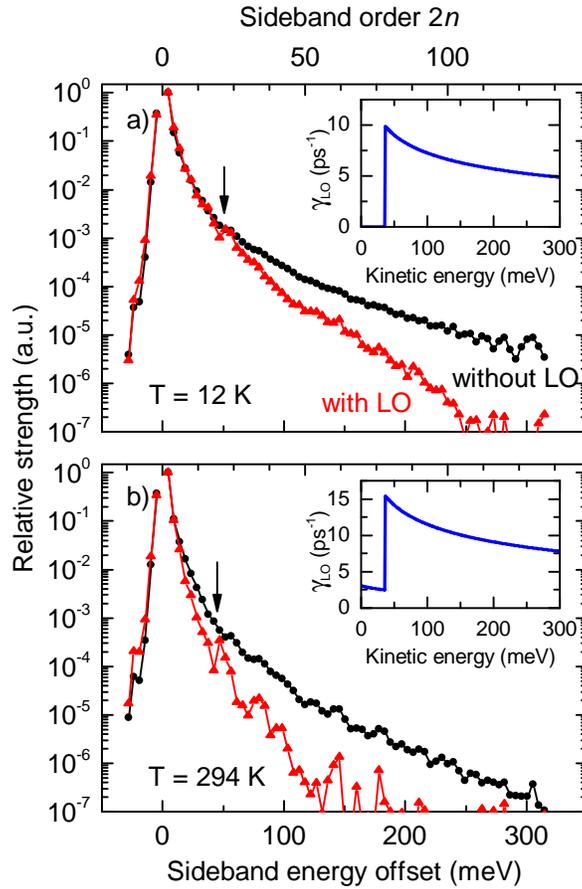


Figure 2.9: Calculations of the effects of longitudinal optical phonon scattering. HSG calculations at (a) $T = 12$ K and (b) $T = 294$ K with (black circles) and without (red triangles) LO phonon scattering. An arrow denotes the threshold for emitting a LO phonon in the 5 nm QW sample. Excitonic dephasing is included in all calculations. The inset is the calculated total energy-dependent scattering rate given by [13]. This calculation was published in [4]. Copyright 2013 by the American Physical Society.

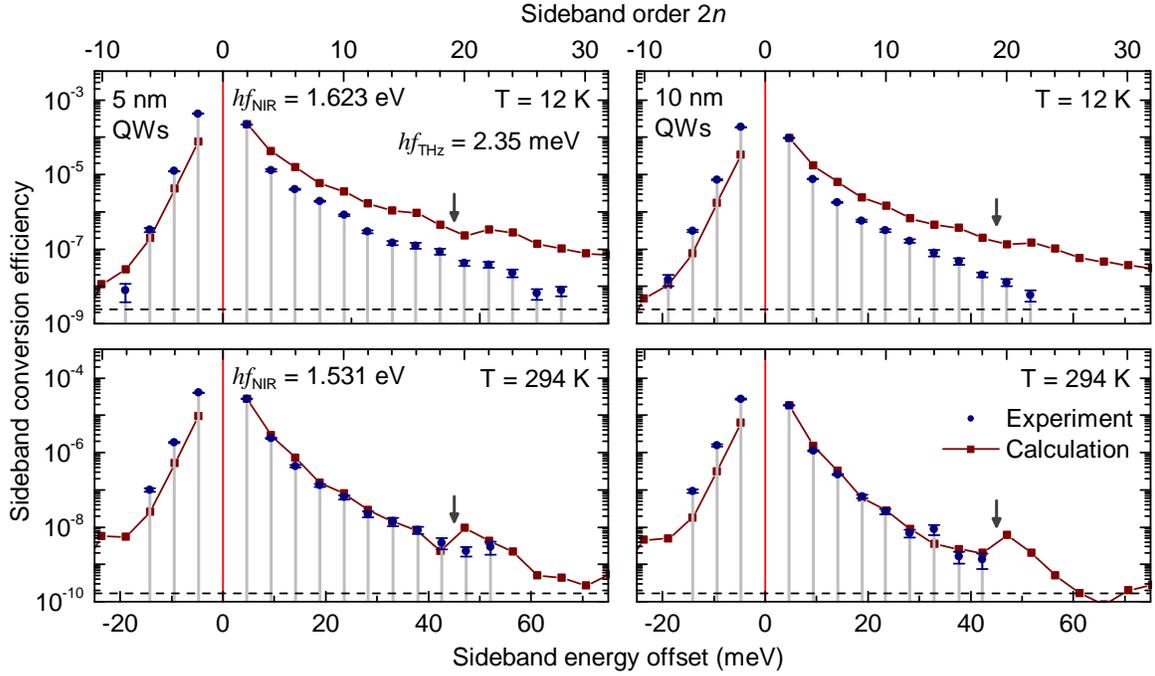


Figure 2.10: Theory vs. experiment. The sideband sideband strength read out by the PMT versus sideband energy offset ($hf_{\text{SB}}hf_{\text{NIR}}$) for the 5 nm QW sample at $T = 12$ K (top left) and $T = 294$ K (bottom left) and for the 10 nm QW sample at $T = 12$ K (top right) and $T = 294$ K (bottom right) are plotted as blue circles. The hatched region represents the noise floor. The vertical arrows represent the energy offset when the recolliding electron and hole have enough kinetic energy to emit a longitudinal optical phonon. This data was published in [4]. Copyright 2013 by the American Physical Society.

line. However, at room temperature, the calculated intensities of the positive sidebands agree remarkably well with experiment. The agreement of experiment and theory at room temperature suggests that the decoherence at room temperature is so strong that it dominates the physics missing from the model.

2.4 Conclusion and remaining questions

We have studied HSG in two QW samples demonstrating that electron-hole recollisions in semiconductors are surprisingly robust, though still affected, to electron-lattice interactions. In both samples, sidebands were generated from electron-hole pairs that had more than enough kinetic energy to emit an LO phonon. Large sideband spectra were observed in both samples at even room temperature, where the thermal population of LO phonons is high enough that phonon absorption by electrons becomes probable. In the 5 nm GaAs QWs at room temperature, we observed sidebands from electron-hole pairs capable of emitting LO phonons. The classical three step model demonstrates that recollisions happen too quickly to be affected by phonon scattering. A more complete quantum saddle point calculation that includes both phonon scattering and excitonic dephasing agrees very nicely with the data. These results imply that further reductions in scattering and dephasing could increase the number of observed sidebands.

Chapter 3

The Sideband Spectrometer

High-order sideband generation (HSG) experiments demand a great deal of a spectrometry set up. Terahertz (THz) and near-infrared (NIR) wavelengths differ by three orders of magnitude, so few optics work in both wavelength ranges. Because the THz wavelengths are so long, and the focal spot must be as small as possible to maximize the electric field, very large optics and lenses must be used which get in the way of NIR optics. On the detection side, NIR sensitivities must be very high, as the highest order sidebands can have creation efficiencies less than 10^{-10} and the cavity-dumped free electron laser has a duty cycle of less than 10^{-7} . Furthermore, HSG depends a lot on the material to be studied, so basic optical characterization measurements must be able to be performed in the same space to avoid unnecessary duplication of expensive equipment and software.

3.1 Optical setup

A complete optical measurement apparatus has been developed for use between 700 and 1000 nm. Linear experiments such as photoluminescence (PL) and differential transmission (DT) to measure absorption can characterize potential samples. Then nonlinear THz experiments such as sideband generation and driven absorption will probe the light-matter interaction at THz scales. Both excitation and detection optical paths have been streamlined to make switching between measurement techniques fast. The THz beam path has also been simplified from previous experiments to make initial alignment easy and very stable. Samples are mounted on a three-axis translation stage in a closed-cycle cryostat with a temperature range of 15–295 K. The sample translation stage can be manipulated even at low temperature using specialized feedthroughs. The entire cryostat itself can be translated linearly by as much as a foot to switch between samples, measurements, or alignment. The front window of the cryostat is made of NIR- and THz-transparent suprasil quartz, although TPX plastic or c-cut sapphire are also useful windows. The back window is glass that has been anti-reflection coated for the NIR.

3.1.1 Linear measurements

The two linear measurements, PL and DT, are very important for understanding the properties of a candidate material or processing technique for sideband experiments. It is extremely important that these can be performed rapidly without misaligning the more challenging nonlinear THz measurements. The optical table has been set up to

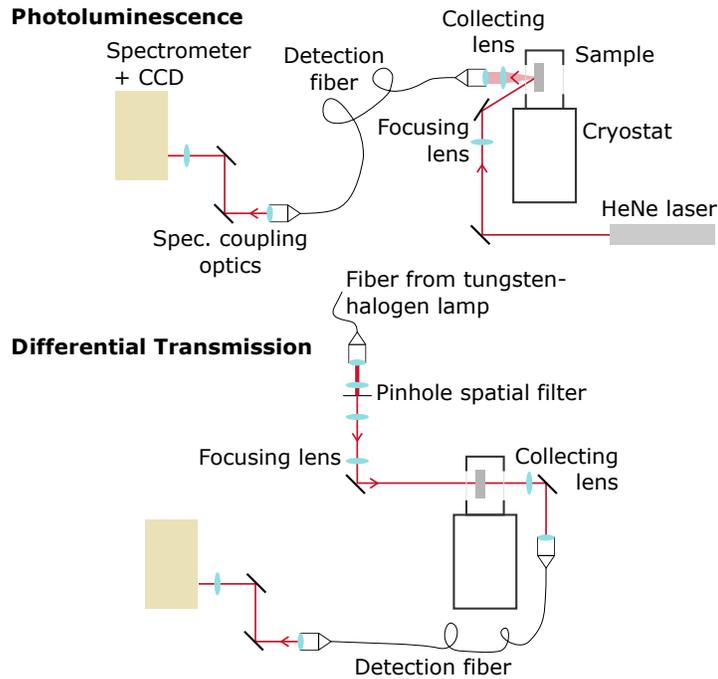


Figure 3.1: The optical schematic for photoluminescence (PL) and differential transmission (DT) measurements. **Photoluminescence:** A HeNe laser is focused on the sample in a cryostat at an angle of about 30° . A collecting lens is placed as close to the sample as possible to maximize signal. The PL emission is then coupled into a detection fiber which brings the light to an imaging spectrometer to be measured by a CCD. **Differential Transmission:** Light from a tungsten-halogen lamp is brought by a fiber to a pinhole spatial filter. The collimated light from the spatial filter is focused on the sample. Transmitted light is collected by lens and coupled to a fiber to be brought to the spectrometer.

incorporate these spectroscopic tools alongside the nonlinear measurements.

Photoluminescence

As described in chapters 1 and 2, PL is a particularly attractive optical measurement because of its simplicity. The improved apparatus is shown in Fig. 3.1(top). The whole cryostat can be translated linearly between optical setups so that the optics to measure PL can remain set up while other experiments are being performed in the THz beam

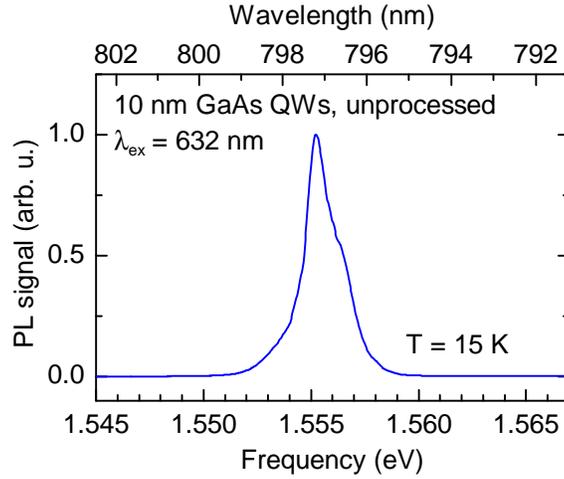


Figure 3.2: A PL spectrum from undoped quantum wells (QWs) as-grown. A HeNe laser excites the PL, $\lambda_{\text{ex}} = 632 \text{ nm}$. The single peak results from radiative recombination of heavy hole excitons. The small shoulders in the peak result from well-width fluctuations in the QWs distributed across the ten-well stack and the $200 \mu\text{m}$ laser spot. At 15 K , there is no significant population of light hole excitons about 10 meV above the heavy hole exciton peak.

path. New samples can be quickly tested to determine if they are good enough for further experiments. A 5 mW HeNe laser ($\lambda_{\text{HeNe}} = 632 \text{ nm}$) is focused on the sample at an angle of about 30° to the sample normal with a spot size of about $200 \mu\text{m}$. The excited electron-hole pairs luminesce in all 4π steradians, but because of Snell's law most of the light that escapes does so normal to the sample surface. The luminescence is collected by a lens with the shortest focus allowable, usually 60 mm , and fiber-coupled to an imaging spectrometer to be measured by a Peltier-cooled, electron-multiplying CCD array. An example PL spectrum can be seen in Fig. 3.2.

Differential Transmission

Undriven DT is slightly more complicated than PL, but the resulting absorption spectra are more directly relevant to sideband experiments. To truly measure absorption requires measuring both the transmitted and reflected components of the incident light. In practice measuring absorption only with differential transmission is sufficient as the reflected component is generally small and independent of wavelength. The absorbance of a sample α_{samp} is the negative log of the light transmitted through the optics with and without the sample:

$$\alpha_{\text{samp}}(\lambda) = -10 \log \frac{T_{\text{samp}}}{T_{\text{ref}}} \quad (3.1)$$

where $T_{\text{samp}}(\lambda)$ and T_{ref} are the amount of light transmitted by the sample and by the reference (no sample), respectively, and α_{sample} is the absorbance of the sample, reported in dB. Differential transmission can be difficult if the sample has other NIR-absorbing layers or if the light reflected off the front of the sample is larger than the light absorbed. Sideband measurements already require sample substrates to be transparent to NIR light, so that requirement must be met anyway.

The optical schematic for measuring DT can be seen in Fig. 3.1(bottom). The light source is a tungsten-halogen lamp coupled to a multimode optical fiber. The lamp provides ample light so that the spot can be spatially filtered using a pinhole to minimize the spot on the sample. The image of the pinhole on the sample is four times larger than the pinhole itself. Pinholes as small as $10 \mu\text{m}$ in diameter are trivial to align, and enough light is focused through the sample to get plenty of signal. The beam is then collected

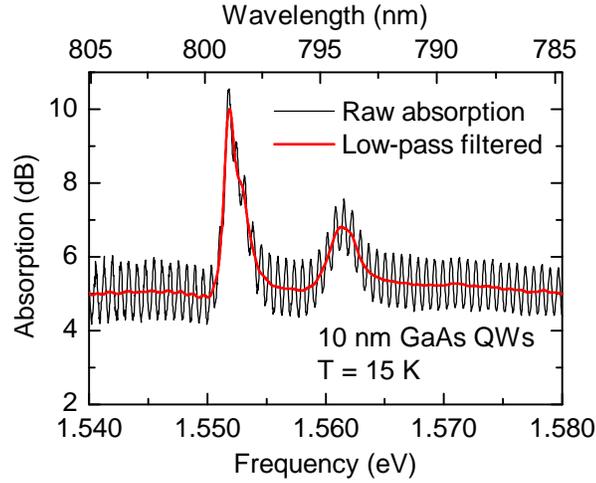


Figure 3.3: A raw and processed absorption spectrum of QWs epitaxially transferred to a sapphire substrate. The raw spectrum has a fast Fabry-Perot oscillation on top of the desired absorption spectrum. Applying a low-pass filter in software removes the Fabry-Perot oscillation with no significant artifacts. The QWs are the same as in Fig. 3.2, and the same shoulder in the heavy hole exciton peak can be resolved post-filtering. The light hole exciton can also be clearly seen at about 1.565 eV. (This an example Abs spectrum from 1/12/16)

after the sample and coupled to a fiber sent to the imaging spectrometer to be measured by the CCD.

An absorbance spectrum measured by DT can be found in Fig. 3.3. The primary artifact in the spectrum, the high-frequency sinusoid on top of the whole spectrum, is directly related to the fact that R_{sample} was not measured. Samples with extremely smooth, transparent substrates exhibit Fabry-Perot interference that modulates the transmission and reflection coefficients as a function of wavelength. The highly polished, 500 μm -thick, c-cut sapphire substrates have a Fabry-Perot fringe period of about 0.7 meV with an amplitude of approximately 2 dB overlaid on the actual transmitted spectrum. Because this period is much smaller than even the narrowest excitonic absorption lines measured,

the fringes can be removed in software by short-pass Fourier filtering. With the fringes removed, the heavy hole and light hole excitons become readily apparent. To measure narrow resonances or where the signal-to-noise is not high enough to accurately measure the fringes, depositing an anti-reflection coating on the backside of samples is necessary.

3.1.2 Nonlinear measurements

Nonlinear THz-optical measurements allow for the THz-material interaction to be measured using far more mature optical technology. Sideband experiments, such as HSG, can probe quasiparticle dynamics at picosecond-times at the single-quasiparticle level. Measuring differential NIR transmission in a strong driving THz field (driven differential transmission, or DDT) can probe THz-exciton interactions.

Sideband measurement

First, the NIR and THz light must be aligned collinearly on the sample, see Fig. 3.4. Starting with the NIR beam path, a continuous-wave Ti:sapphire laser outputs approximately 1 W of laser power across its entire tuning range, 700–1000 nm. To safely attenuate the beam without introducing thermal lensing, a rotating half-wave plate (HWP) is placed before a Glan-laser polarizer. The HWP controls how much of the NIR light transmits through the polarizer. The reflected portion of the beam is dumped in a beam dump (BD). A pellicle beam splitter (BS) then picks off approximately 8% of the laser towards a multimode fiber coupler connected to a small spectrometer to monitor the wavelength of the NIR laser. The NIR laser then passes through an acousto-optic modulator (AOM)

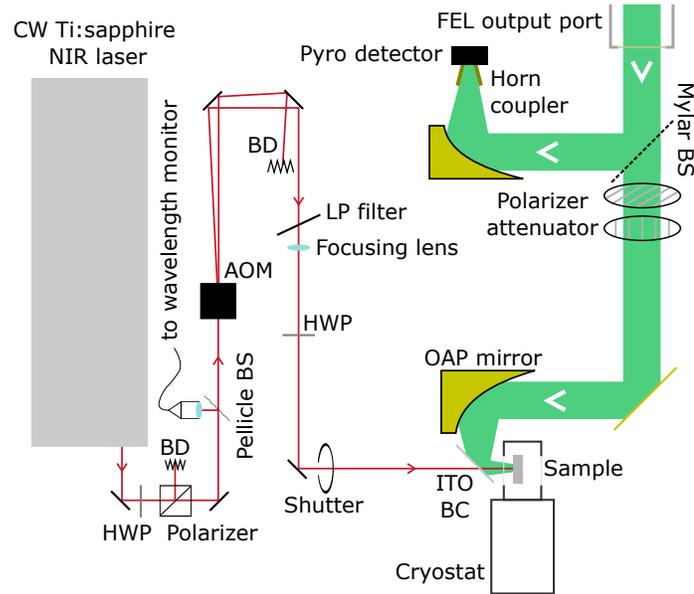


Figure 3.4: The optical schematic for the NIR and THz excitation paths. The continuous wave (CW) Ti:sapphire laser is tuned to a particular NIR wavelength. A rotatable half-wave plate (HWP) followed by a linear polarizer provides consistent attenuation. Power reflected by the polarizer is sent into a beam dump (BD). A pellicle beamsplitter (BS) picks off a few percent of the laser power that is sent to a wavelength-monitoring spectrometer. The laser then goes through an acousto-optics modulator (AOM). The zeroth order beam is sent to a BD, the first order beam then goes through a longpass (LP) filter and an achromatic focusing lens. A second HWP controls the direction of linear polarization at the sample. A mechanical shutter blocks the beam when the FEL is not lasing. Finally, the laser transmits through the indium-tin oxide (ITO) beam combiner (BC) to be made collinear with the FEL. The THz beam path starts at the FEL output port. A mylar BS picks off some of the power and an off-axis parabolic (OAP) mirror focuses the beam on a horn-coupled pyroelectric (pyro) detector. The majority of the beam passes through a wire-grid polarizer attenuator. Finally, the THz gets focused down with a second OAP mirror, and the ITO beam combiner reflects the THz beam into the cryostat, collinear with the NIR laser. Both lasers are focused on the sample.

that picks off as much as 60% of the beam into the first-order diffracted beam. The first-order beam has an on-off contrast ratio of 3000:1 and a rise and fall time of approximately 200 ns to quickly shutter the NIR laser so that it is only on during the FEL pulse. The zeroth-order beam is sent to a BD. The first-order beam is then sent through a long-pass filter. The sidebands are weak enough that amplified spontaneous emission from the NIR laser can overwhelm the sideband signal, so it must be filtered out. The NIR laser is then gently focused with a 500 mm achromatic lens. A rotatable achromatic HWP controls the linear polarization of the beam, then the beam passes through a mechanical shutter. The NIR laser then passes through a glass slide coated with indium-tin oxide (ITO) that serves as a beam combiner (BC) and the suprasil cryostat window before being focused on the sample with a 200 μm -diameter spot.

The THz beam path is far simpler. The THz beam exits the FEL transport system output port before being split by a mylar BS, see App. B. About 5% of the THz power is picked off and sent to an integrating pyroelectric detector to monitor and count FEL pulses. The majority of the horizontally-polarized beam then passes through two wire-grid polarizers. The first rotates to tunably attenuate the THz beam, and the second is fixed to pass the remaining horizontally-polarized component. The THz laser is then focused by an off-axis parabolic (OAP) mirror to a spot of about 1 mm in diameter. Before the focus, the THz beam reflects off the ITO BC to become collinear with the NIR beam. The ITO BC is mounted on a tip-tilt stage so that THz alignment can be peaked up without disturbing the sensitive focusing optics.

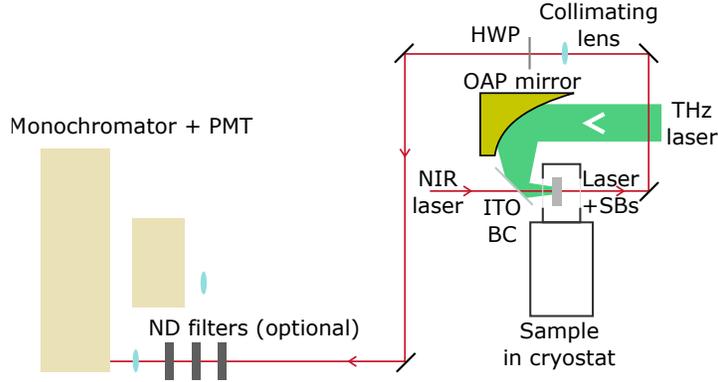


Figure 3.5: The optical schematic for sideband PMT detection path. Sidebands are created with a similar beam profile as the NIR laser. They are collimated by a lens after the sample, and the polarization is rotated using the HWP for maximum grating efficiency in the monochromator. Because PMT is easily saturated, calibrated neutral density (ND) filters allow for measurements of the NIR laser line and the strong low-order sidebands.

If NIR-THz alignment is good and the experiment is well-designed, sidebands around the NIR wavelength will be generated. There are two detectors used to measure the sidebands, depending on their proximity to the NIR laser wavelength, a photomultiplier tube (PMT) and a CCD. For low-order sidebands, $n \leq 10$, the PMT, mounted on a double-grating monochromator, is used. The monochromator has a very high rejection of light more than 0.1 nm from the selected wavelength (about 100 GHz at 750 nm), allowing for sensitive detection of low order sidebands generated from intense THz light with frequency as low as 240 GHz. It is also a very sensitive, very fast detector, with a single photon rise and fall time of about 10 ns. In addition, it allows for rapid THz-NIR alignment by maximizing the signal from low-order sidebands because of its fast readout ($n = 8$ is approximately the highest order that can be used). The primary drawback of the PMT is that it is a single channel detector, so integrating enough to get good

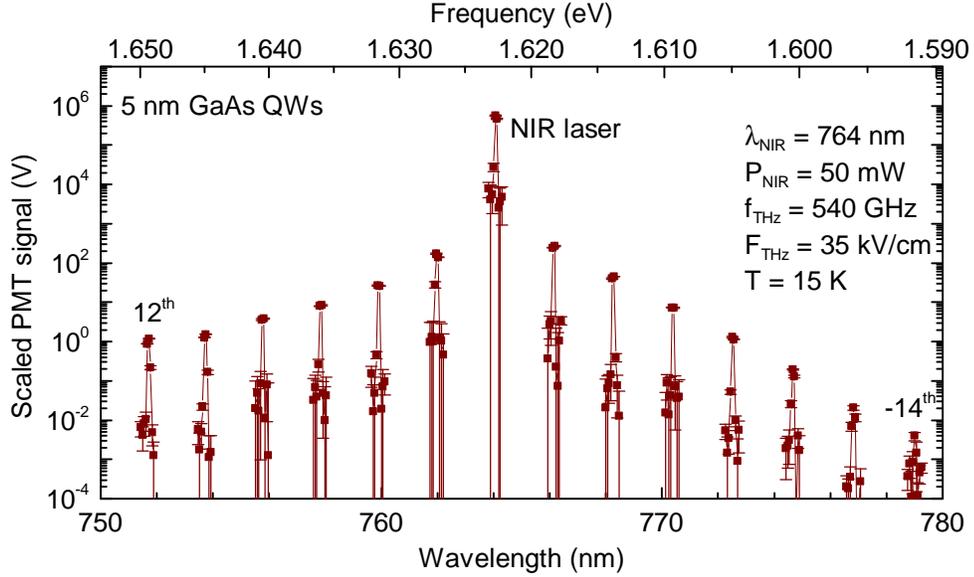


Figure 3.6: Example sideband spectrum measured by the PMT. Sidebands spanning from orders $-14 \leq n \leq 12$ are measured. The sideband spectrum does not end at +12, but the CCD is a better detector for higher-order sidebands. The noise floor can be seen in the -12th and -14th sidebands, about 5 mV. Sidebands closer to the NIR laser, and the NIR laser itself, are much stronger and must be attenuated, raising the local noise floor substantially. The strength of an individual sideband is measured by integrating the signal around a peak.

statistics is difficult.

To get to the PMT, the NIR laser and sidebands are collimated by another 500 mm achromatic lens and the polarization is rotated to horizontal by another rotating achromatic HWP, see Fig. 3.5. Before being focused on the monochromator slits with a 25 mm lens, the laser and sidebands may have to be attenuated. The PMT signal is amplified with a 300 MHz bandwidth amplifier and read out by an oscilloscope with two readout “modes,” integrating mode and photon-counting mode. In integrating mode, the time trace is sliced to sum the PMT signal just when the FEL is active. The noise floor of integrating mode is a few nanowatts and has a dynamic range of only about 15. To

measure strong signals such as the NIR laser (> 10 mW) or low order sidebands (10–100 μ W), the PMT gain is turned down and the light is attenuated with calibrated neutral density (ND) filters placed just before the final lens. By measuring the NIR laser with a strong enough attenuator, the PMT response can be calibrated so that the sideband efficiency, the number of n th order sideband photons created per NIR laser photon, can be measured, see Fig. 3.6 for an example sideband spectrum. The monochromator then scans over a sideband, one wavenumber ($1/\text{cm}$) at a time. It takes approximately 70 FEL pulses to measure a single sideband with decent statistics, but single FEL pulses can be used to peak up the alignment by tweaking the ITO BC. When PMT signal is very small, it can be used in photon counting mode, where a fast peak detector (in software or hardware) is used to count the 50 mV spikes characteristic of a single photon from the PMT output. Weaker sidebands can be photon counted, but a better detector for these sidebands is the CCD.

The optical schematic for measuring a sideband spectrum with the CCD shares many features with the PMT setup, see Fig. 3.7. The initial optics for both the NIR laser and the THz laser are identical, but now the shutter is engaged. Even with the AOM contrast ratio maximized, there is about a factor of 10^3 more NIR laser light between FEL pulses than during. By using a mechanical shutter with an infinite contrast ratio and a open-close time of 10 ms, the excess light can be brought down to an acceptable factor of ten larger.

The CCD detection path changes as well. The achromatic HWP after the collimating

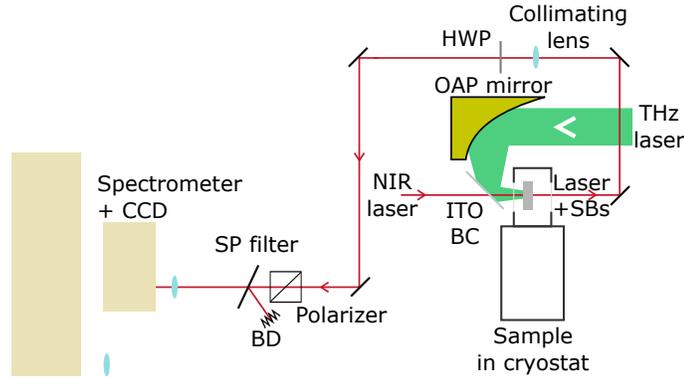


Figure 3.7: The optical schematic for the sideband CCD detection path. Similar to the PMT detection path, the sidebands are collimated by the same lens and the polarization is rotated using the HWP for maximum efficiency through the spectrometer. To filter out the NIR laser line so that it doesn't overwhelm the sideband signal, a polarizer cleans up the sideband polarization and a short pass (SP) filter attenuates the laser line by approximately 70 dB without affecting sidebands above approximately sixth order.

lens rotates the NIR laser and sideband polarizations to vertical. Then, to ensure maximum NIR laser rejection by the angle-tuned short pass (SP) filter, the NIR laser and sidebands are sent through a polarizer. Angle-tuning optical edge filters does not move the edge for both s- and p-polarizations at the same rate, so a final polarizer is needed to achieve the highest extinction coefficients. For high-quality measurements and the safety of the CCD chip, more than 50 dB attenuation is required (the SP filter is theoretically capable of 70 dB). A 60 mm achromatic lens focuses the transmitted sidebands onto the entrance slits of the imaging spectrometer to be measured by the CCD.

The CCD allows for very rapid HSG spectrum acquisition rates. A sideband spectrum spanning 6th to 48th order and and 40 dB of dynamic range can be taken in just five FEL pulses, about six seconds, see Fig. 3.8. Because the detector is cooled to -90° C, it can also integrate for as long as five minutes without significant dark count signal,

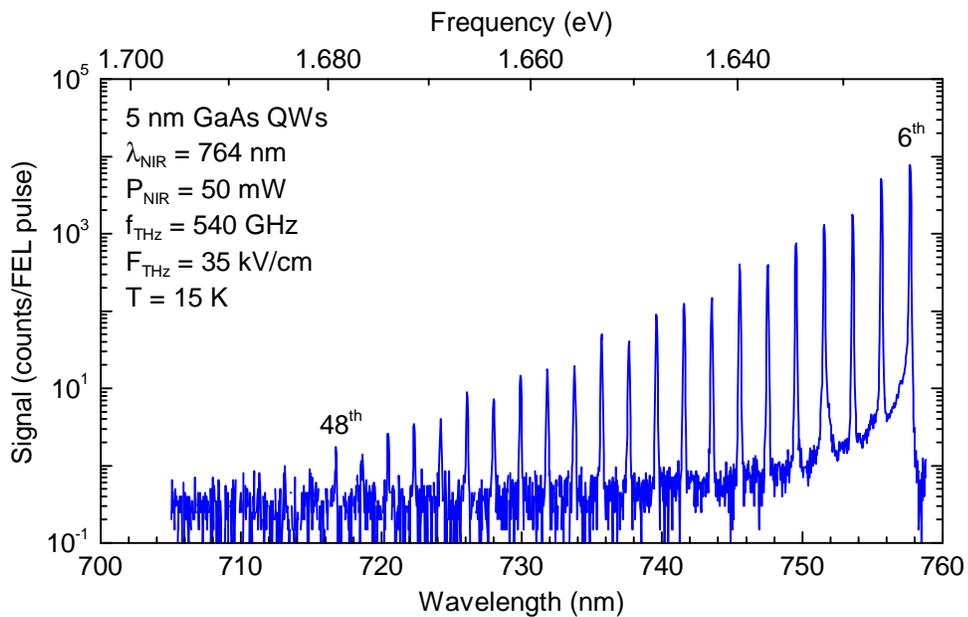


Figure 3.8: Example sideband spectrum measured by the CCD. Sidebands of order $6 \leq n \leq 48$, spanning 40 nm, can be observed in four images covering 20 total FEL pulses, or 800 ns of THz power. The strongest four sidebands spill a little bit out of their pixels during readout, but 30 dB down from the peak. Overlap between this spectrum and the PMT spectrum in Fig. 3.6 to calculate the efficiency of these sidebands as well.

dramatically decreasing the readout-noise-dominated noise floor. The dynamic range of the entire experiment can thus be extended to over 100 dB. The effect of the short pass filter cutoff is apparent in this spectrum in the fast drop off above the 6th order sideband at about 759 nm. Sideband strengths and linewidths are calculated from the area and standard deviation of a Gaussian peak fit, respectively. Sidebands measured by both the PMT and the CCD can be used to calibrate the CCD to measure sideband efficiency over the full spectrum.

Taking the image of a sideband spectrum requires several steps. The CCD chip is oriented such that the horizontal axis (1600 pixels) maps to wavelength, and the vertical axis can be integrated over (400 pixels). Before any image is taken, the camera is set to a temperature of -90° C, to minimize dark current, and a electron-multiplying gain of 110, which corresponds to a signal gain of about 7 (these are simply the standard settings). An un-binned image of a sideband spectrum contains a number of two dimensional peaks. The peaks are very narrow in the horizontal (wavelength) direction, spanning no more than ten pixels, but very broad in the vertical direction, spanning as many as one hundred pixels. The size of the vertical hardware bins are then set so that the entire vertical extent of a binned pixel can contain the peak of the strongest sideband in the image. Vertical binning in hardware minimizes the effective readout noise by allowing a slower, more accurate readout amplifier to be used. If the bins are made too large, however, charge noise integrated over the many physical pixels can add significantly to the total noise. It is also possible to bin in software, but the noise characteristics were found to be worse.

Two vertical bins are recorded with every image (so that the array output by the CCD is 1600×2). The “top” row contains the sideband or background measurement. The bottom row is never illuminated, even by background sources, and so it provides a measurement of the total noise per binned-pixel for that individual image.

With the camera and image settings determined, the next step is to take a background. The FEL is turned off and the NIR laser is blocked so that background images can be taken. Because images can be very long (often 40 s, but as long as five minutes), cosmic rays can interfere with the signal (they look a lot like a sideband when vertically binned). To remove cosmic ray artifacts, four identical images are taken, although three is the minimum. These images are compared and any large signal spikes that are only found in one image are removed and the remaining pixels are processed to return a mean and standard error for every binned pixel. Taking more images allows for effective cosmic ray removal (and better error statistics) at the expense of FEL time.

With the background saved, data is ready to be taken. An identical set of images are taken with the NIR and FEL back on, and the same cosmic ray removal process is performed. The background array is then subtracted from the background array. The background-corrected signal data is then divided by the number of times the FEL fired to correct for data with different integration times or FEL repetition rates. Finally, the sideband spectrum data is saved with three columns of data: wavelength, sideband signal, and error. The error is the quadrature sum of the dark pixel noise and the standard error of that binned pixel.

Sideband polarization measurement

The large optical bandwidth and dynamic range of a single CCD spectrum allow for measurements of the polarization state of many sidebands in parallel. The complete polarization state of light is measured by a Stokes polarimeter, a rotating achromatic quarter-wave plate (QWP) in front of a linear polarizer [23]. Both optics are placed immediately after the cryostat, as the high-reflectivity dielectric mirrors in the beam path scramble polarization as a function of wavelength because the dielectric stacks induce propagation delays that are a very complicated function of wavelength, see Fig. 3.9. The strength of the sidebands measured in parallel on the CCD as a function of QWP rotation angle can be fit using the Mueller calculus to obtain all four Stokes parameters that describe the complete polarization state of the sidebands [23]. To measure the NIR laser polarization state, the intensity transmitted through the polarizer is measured using a standard laser power meter instead of the CCD.

Calculating the polarization state from the data requires fitting to the power measured after the polarimeter, see Fig. 3.9. The four Stokes parameters are defined as follows:

$$S_0 = I$$

$$S_1 = pI \cos 2\alpha \cos 2\gamma$$

$$S_2 = pI \sin 2\alpha \cos 2\gamma$$

$$S_3 = pI \sin 2\gamma,$$

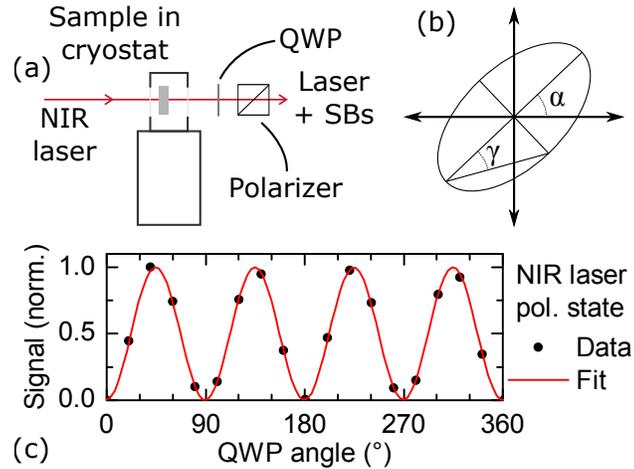


Figure 3.9: The Stokes polarimeter schematic and an example measurement. (a) The QWP and polarizer are placed immediately after the cryostat so that all polarization manipulations can be done with as few intervening optics as possible. (b) Two angles α and γ define the complete polarization ellipse. (c) An example Stokes measurement, the power of the NIR laser transmitted through the polarizer as a function of the QWP rotation angle. For this measurement $\alpha = 90.8 \pm 1^\circ$ and $\gamma = 0.4 \pm 1^\circ$.

where I is the intensity of the light, p ($\in [0, 1]$) is the degree of polarization, α ($\in [0, \pi]$) is the angle of the major axis of the polarization ellipse, and γ ($\in [-\pi/4, \pi/4]$) is the “roundness” of the polarization ellipse, see Fig. S2(a). For example, (α, γ) of horizontal light would be $(0, 0)$, vertical $(\pi/2, 0)$, and right circular light $(0, \pi/4)$. Let us define the QWP angle θ as counter-clockwise rotation starting at the positive x-axis and orient the polarizer so that it transmits horizontally polarized light. By measuring the power transmitted through the polarizer as a function of QWP angle θ , the four Stokes parameters can be determined by fitting the transmitted power data $I(\theta)$ with the following

equation:

$$I = \frac{1}{4}(2S_0 + S_1 + S_1 \cos \alpha) + \frac{1}{4}(S_1 + S_1 \cos \alpha) \cos(4\theta) - \frac{1}{2}S_3 \sin \alpha \sin(2\theta) + \frac{1}{4}(S_2 - S_2 \cos \alpha) \sin(4\theta).$$

The different phases and frequency in θ allow for one measurement to fit all four Stokes parameters simultaneously. For a complete derivation and more details, consult Goldstein [23].

The NIR laser polarization state was measured by placing the QWP at the beam waist and using a standard optical power meter to measure the intensity as a function of QWP rotation angle, see Fig. S2(c). The sideband polarization states were measured in parallel using the CCD spectrometer to monitor the intensity of individual sidebands. The accuracy of this broadband polarimeter was tested by placing a rotatable polarizer at the sample and shining collimated light from a tungsten-halogen lamp through the beam path. The polarization state as a function wavelength was able to predict the α of the test polarizer to within 1° .

Driven absorption measurement

Sideband photons in HSG can only be created if excitons or even unbound electron-hole pairs are created first. The exciton absorption line, in general, shifts non-monotonically when illuminated by a strong terahertz field, so ensuring that the NIR laser wavelength is appropriately tuned is not trivial [46, 6, 60, 74]. The difficulty of measuring driven

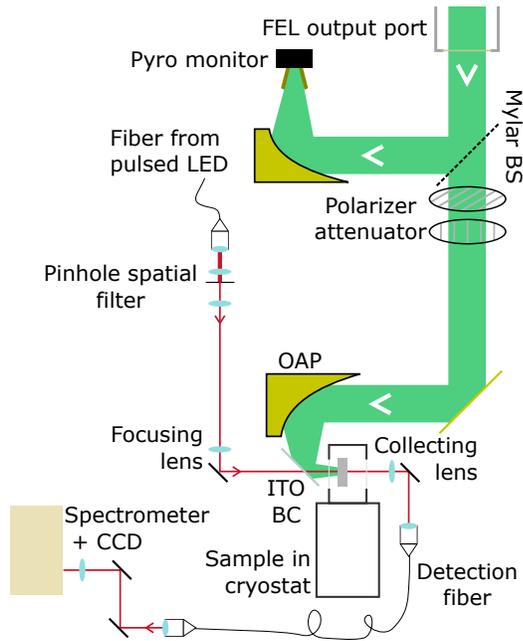


Figure 3.10: The optical schematic of a driven absorption experiment. This measurement combines the DT measurement with THz illumination. Instead of the tungsten-halogen lamp, however, an electronically-switched LED provides the NIR power only while the FEL is lasing.

absorption lies entirely in getting enough optical power through the sample while also ensuring spatial and temporal overlap between the NIR and THz light.

A light-emitting diode (LED) can be pulsed very quickly at high currents, see Fig. 3.10. The pulse-generating circuit is capable of driving about 200 mA in pulses with 10 ns rise and fall times. The pinhole spatial filter ensures that the LED spot has a top hat profile with 200 μm diameter, well within the 1 mm THz spot. The spatial filter and FEL duty cycle severely limit the amount of light at the sample, so integrating enough signal is an issue. The FEL has to be operated in hole coupler mode, capable of μs -long pulses, to get enough signal, so the THz power is limited to a few kilowatts, or less than 5 kV/cm in GaAs. Even that low a field is enough to substantially red-shift the exciton absorption

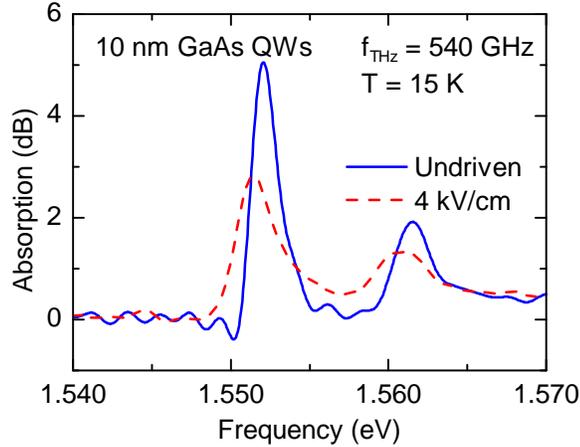


Figure 3.11: Example driven absorption measurement. Driving the QWs with a 4 kV/cm THz field in the plane of the QWs redshifts the undriven QW absorption by about 2 meV and lowers the absorption coefficient. The FEL is lasing using the hole coupler, so the THz pulse length is approximately 1 μ s long. Even the large duty cycle of 10^{-6} does not provide high-enough signal-to-noise to perfectly filter out the Fabry-Perot fringes as in Fig. 3.3, but the peak is still well resolved.

lines, see Fig. 3.11.

3.2 FEL characterization

Rapid developments in the generation of strong THz light have motivated new understanding of the FEL to discover its comparative advantage. On the optical side, THz time-domain setups have opened coherent THz spectroscopy [27]. New techniques in optical rectification have increased peak fields in single-cycle picosecond pulses above MV/cm [70, 51]. On the electronics side, new developments have pushed commercial sources and detectors above even 1 THz. The advantages of the FEL lie in its high power, narrow bandwidth, and long pulses.

3.2.1 Power measurement

Measuring the instantaneous power of the FEL at arbitrary frequencies is important for nonlinear effects where the peak power dominates the response. A pyroelectric (pyro) detector can be a very simple option. These use a ferroelectric material which has a large temperature-dependent polarization response. By heating them, for instance by painting them black and illuminating them, they output a voltage approximately proportional to the intensity of the light on the element. They can be made very fast, with bandwidths of approximately 100 MHz, but must then be very small. That speed is just fast enough to measure the instantaneous FEL power, but the size means that they are relatively sensitive to FEL alignment and spatial profiles that change in time, see Sect. 1.6. Large area pyro detectors which are more robust to alignment changes are necessarily much slower. The sensitivity of these detectors is very dependent on ambient temperature as well, and so make poor long-term absolute power meters. A more robust large-area detector is the Thomas-Keating (TK) detector. This photo-acoustic detector contains a thin gas cell with one side coated with a thin film of metal with known THz absorption. When the metal absorbs the THz power, it creates an acoustic wave in the cell. A pressure transducer registers the pressure change, which is directly proportional to the energy in the energy in the THz beam. The large size of the mylar sheet means that the bandwidth is very small, less than a kHz. Furthermore, this detector is easily calibrated by driving a DC voltage across the metal film and measuring the response.

Instantaneous power measurements can be made by combining the speed of a pyro

detector with the absolute precision of the TK detector, see Fig. 3.12. First, about 5% of the THz beam is picked off by a mylar beamsplitter and focused on a horn-coupled fast pyro. The horn makes the detector more insensitive to alignment or spatial profile changes. The pyro detector is connected directly to an oscilloscope with a 1 M Ω input impedance so that it operates as a THz power integrator. The height of the discontinuous jump relative to the total height (not including the ringing) is the proportion of energy of the FEL pulse in the 40 ns cavity dump refChapter1. The TK detector is placed in the beam path behind a rotatable wire-grid-polarizer attenuator. With the polarizer attenuator set to full transmission, the power in the FEL is measured simultaneously by the pyro and the TK detectors, thereby calibrating the pyro detector for the day's measurements. By rotating the front polarizer, the THz power can be attenuated easily and repeatably while the rear polarizer maintains the polarization at the sample. The TK can then calibrate the attenuator as well.

After recent improvements inside the accelerator, the FEL is generating record cavity dump powers. On an average day at 540 GHz, an FEL pulse contains about 2.9 mJ, and 90% of that is in the cavity dump. For those 40 ns, the THz power is 65 kW. The proportion of energy in the cavity dump depends on the quality of the electron beam tuning. Poor tuning leads to slow lasing times, and the proportion of energy in the cavity dump can be as low as 60%. Also, by calibrating the pyro at the beginning of the day, any long-timescale drifts can be easily monitored and corrected for using the polarizing attenuator.

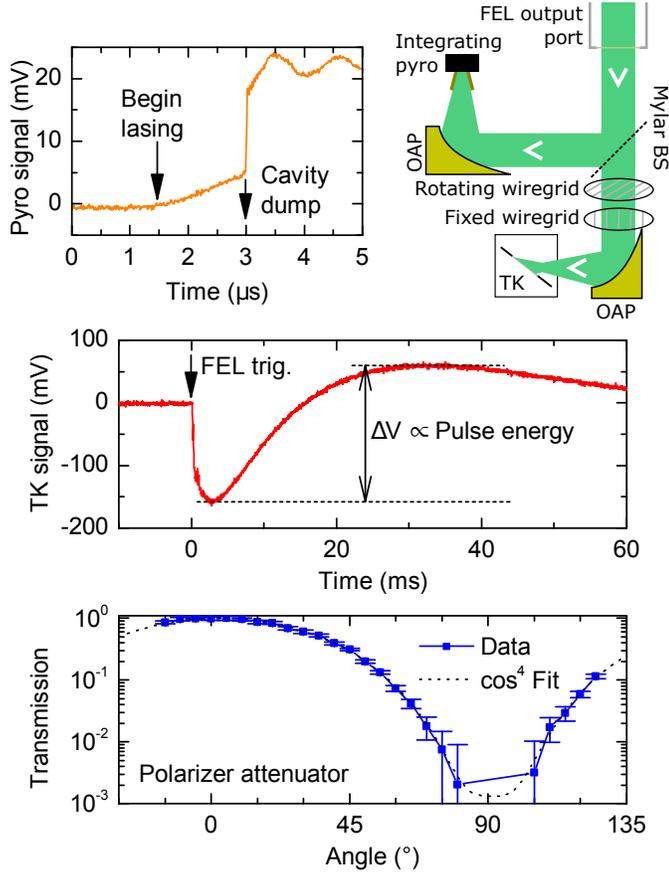


Figure 3.12: Measuring FEL power. (a) Schematic for measuring FEL power. The THz beam exits the port and travels along the optical table. A mylar BS picks off about 5% of the power and focuses it onto a horn-coupled fast integrating pyro detector. The rest of the beam passes through the polarizer attenuator consisting of a rotating wire-grid polarizer and a fixed wire-grid polarizer. The THz beam is then focused on the TK detector. (b) Time trace of a cavity dumped (CD) FEL pulse measured by the integrating pyro. Initially, the silicon coupler reflects a few percent of the THz light out of the cavity, and the integrated signal rises approximately linearly. At 3 μs , an intense green pulsed laser creates an electron-hole plasma in the silicon, dumping all of the THz power out of the FEL cavity in 40 ns. The pyro signal then jumps nearly discontinuously. After 3 μs , the oscillations are due to impedance mismatch. (c) Time trace of the TK detector after absorbing an FEL pulse. The peak-to-peak amplitude of the signal is directly proportional to the total power in the FEL pulse. Traces (b) and (c) combine to measure the power only in the cavity dump. (d) Results of calibrating the polarizer attenuator using the TK. The FEL output horizontally polarized and the fixed polarizer transmits horizontally polarized light, so the transmission through the two polarizers is proportional to $\cos^4 \theta$ where θ is the angle between the two wire grids. Because of imperfections in the polarizers, the minimum attenuation is about 1%.

3.2.2 Frequency and linewidth

Because of the low duty cycle and low repetition rate of the FEL, measuring the frequency is not trivial. The FEL electron paramagnetic resonance lab measures the FEL frequency using superheterodyne detection [52]. This method can measure the frequency and phase of the FEL very accurately, but the bandwidth is limited to a few GHz around a single frequency that is not tunable. Another method is to mix the FEL output with a weak NIR beam in a nonlinear optical crystal [72]. Using a material with a large $\chi^{(2)}$ for efficient sum-frequency generation, the output NIR light and first order sideband frequencies can be measured, $f_{\text{FEL}} = f_{\text{SB}} - f_{\text{NIR}}$. This technique has a large bandwidth, capable of measuring from 0.4–2.5 THz with large signal-to-noise ratios. With only first order mixing, the precision for measuring the FEL frequency is limited to the resolution of the spectrometer, $\Delta_{\text{FEL}} = \Delta_{\text{spec}} \sim 3 \text{ GHz}$ [72]. Increasing the order of the mixing makes this measurement much more accurate, so that even lower resolution spectrometers can leverage this effect to be extremely precise.

With dozens of orders of sidebands, HSG is an ideal mechanism to measure the FEL frequency very accurately. Measuring the location of the n th order sideband increases the resolution by a factor of n :

$$f_{\text{FEL}} = \frac{f_{\text{SB}n} - f_{\text{NIR}}}{n} \quad (3.2)$$

$$\Delta_{\text{FEL}} = \frac{\Delta_{\text{spec}}}{n} \quad (3.3)$$

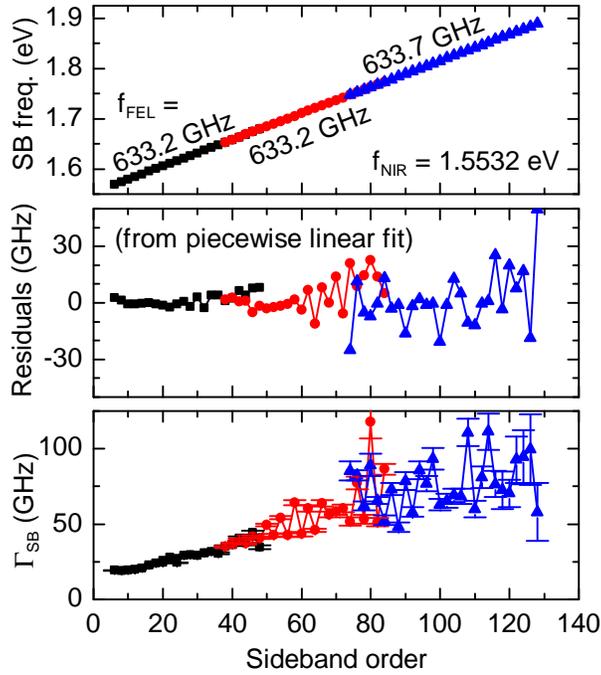


Figure 3.13: Three spectrometer steps measuring the frequency characteristics of a generic sideband spectrum. The FEL was set to $f_{\text{FEL}} = 640 \text{ GHz}$, and the THz field strength in the QWs was $F_{\text{THz}} = 57 \text{ kV/cm}$. (Top) The sideband frequencies increase almost perfectly linearly, yielding the actual THz frequency, $f_{\text{THz}} = 633 \text{ GHz}$. (Middle) Residuals from the linear fits to the sideband frequencies. The residuals increase with increasing order, likely because the highest order sidebands are quite weak. Even the largest residuals are less than 30 GHz, which means that the frequencies of the sidebands deviate from linearity by less than one part in 10^4 . (Bottom) The linewidths of the sidebands from their Gaussian fits. These also increase approximately linearly with order, but they are several times larger than the residuals.

Because there are sidebands at every order, the frequency can be known very well, see Fig. 3.13(top). With the FEL tuned near 640 GHz, sidebands up to order $n = 126$ are observed, requiring three spectrometer steps to measure them all. The image at each spectrometer step has about 25 individual sidebands (because n must be even), and so a single spectrometer step can measure the FEL frequency quite well. For context, the slight change in the FEL frequency between the three steps results from a voltage drift of 4 kV on the 5 MV accelerator terminal voltage over about twenty minutes (assuming $f_{\text{FEL}} \propto V_{\text{terminal}}$). Plotting the residuals shows that the sideband frequencies measured by the spectrometer, which was home-calibrated using neon and argon lamps, Fig. 3.13(middle), deviates by less than 30 GHz from linear until the last measured sideband. These residuals are far smaller than the linewidths of the sidebands measured by the spectrometer, Fig. 3.13(bottom).

These sideband linewidths, Γ_{SB} , provide a measurement of the FEL linewidth. They increase approximately linearly with order n , although noisily. The FEL is broadened nonhomogeneously by terminal voltage fluctuations, and so the sideband linewidths add incoherently, $\Gamma_{\text{SB}} = \sqrt{\Delta_{\text{spec}}^2 + n^2\Gamma_{\text{THz}}^2}$ (The NIR laser has a linewidth of less than 5 MHz, so it is ignored). For small n , $\Gamma_{\text{SB}} \sim \Delta_{\text{spec}}$, and for large n , $\Gamma_{\text{SB}} \sim n\Gamma_{\text{THz}}$. The slope from a linear fit of $\Gamma_{\text{SB}}(n)$ represents Γ_{THz} and the y-intercept represents Δ_{spec} . Linewidth results from several FEL frequencies are shown in Fig. 3.14. All but the 247 GHz sideband spectra required multiple spectrometer steps. The linewidth of all sidebands was calculated and were averaged where sidebands were measured twice (in

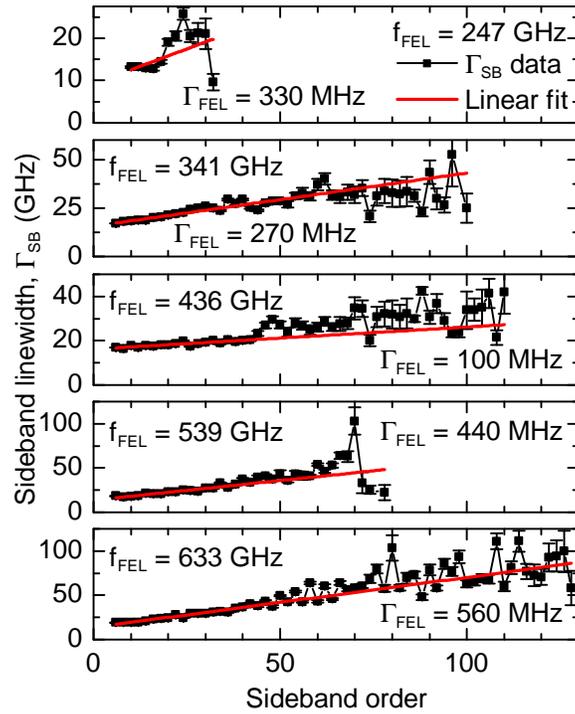


Figure 3.14: The fits of the linewidths for various FEL frequencies. The black squares are the linewidths from fitting sideband peaks with Gaussian functions, and the reported f_{FEL} is measured from sideband spacings as discussed in Fig. 3.13. The red line is a linear fit to linewidths as a function of order. The y-intercept for all but the 247 GHz spectrum is consistently 15 GHz. The slope is the linewidth of the FEL, Γ_{FEL} , reported in each subfigure.

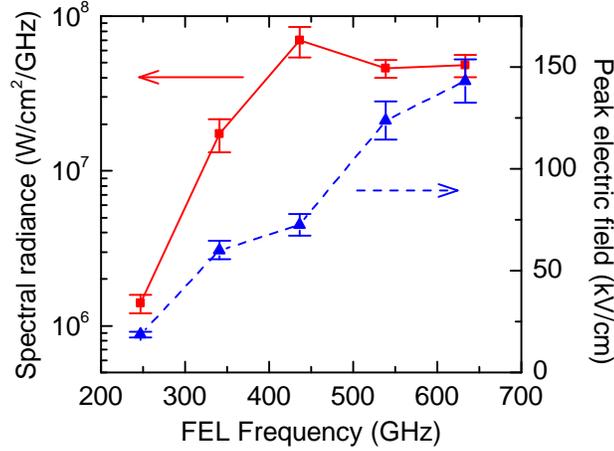


Figure 3.15: Spectral radiance and peak electric field in vacuum. The spectral radiance increases with frequency until it saturates in the mid- 10^7 range, solid line, left axis. The peak electric field (in air) increases linearly, dashed line, right axis. Diffraction at low frequencies (long wavelengths) causes substantial losses in the FEL. At higher frequencies, however, the FEL becomes less stable and the linewidth gets broader. The peak electric fields keep increasing but the spectral radiance plateaus.

overlapping regions).

These results allow the spectral radiance of the FEL to be measured accurately. The spectral radiance (sometimes called spectral “brightness”) is the power flux of an optical beam per solid angle per area per bandwidth. It is invariant along the beam path (assuming propagation through a uniform medium) and is best understood as the highest intensity possible at the best possible focus of that beam. For a moderately-focused Gaussian beam (so that the paraxial approximation is valid),

$$\mathcal{R} = \frac{P}{\pi w_0^2 \pi \theta^2} = \frac{P}{\lambda^2}, \quad (3.4)$$

where \mathcal{R} is the radiance, P is the power in the beam, w_0 is the beam radius at the focus, θ is the divergence angle ($\theta = \lambda/(\pi w_0)$), and λ is the wavelength of the light. The

spectral radiance is therefore the radiance per bandwidth. The FEL, being a high-power, quasi-continuous wave source, has very large spectral radiances. Large spectral radiances are useful for sideband experiments where the highest orders might be washed out without it.

3.3 Outlook

The sideband spectrometer can be expanded and extended to encompass other important measurements to study nonlinear THz-NIR interactions. A complete set of optical edge filters are required to fully make use of the Ti:sapphire laser tuning range for sideband measurements. There are many materials that support excitons in this range, but currently the filters only allow for excitons between about 750–850 nm. The pulsed LED also does not span the full 700–1000 nm range, and not all materials will have such well-understood excitonic responses to strong THz fields. A super-continuum laser source for driven absorption measurements would be a big upgrade. Because the NIR laser is highly tunable, photoluminescence excitation can be added to the repertoire to characterize samples, and it could also be performed under FEL excitation. Also, for more efficient polarization state measurements, including finely controlling the NIR laser polarization, a commercial polarimeter is necessary.

Chapter 4

Terahertz-induced birefringence and sideband polarimetry

A recent observation of high-order harmonic generation (HHG) in ZnO has set off a flurry of new research on high-order nonlinear processes in solids [22, 51, 55, 56]. Using extremely strong pulsed mid-infrared (MIR) or terahertz (THz) laser pulses, these studies create interband polarization by causing electrons to tunnel across the band gap in semiconductors like ZnO and GaSe. The nonlinear response comes from two different mechanisms, one interband and one intraband. Harmonics created by the interband process result from electron-hole recollisions, much like like HHG in atoms. Harmonics created by the intraband process result from the non-parabolic dispersion relation of the bands. In the works by Ghimire, *et al.*, Schubert, *et al.*, and Vampa, *et al.*, the electrons (holes) accelerate up to nearly the top (bottom) of the conduction (valence) band, where

the bands are extremely non-parabolic. These papers suggest that the driving frequency determines what mechanism dominates, MIR for interband and THz for intraband. These mechanisms are important to differentiate, as the recollision physics could be a useful tool for measuring band structure from electron and hole trajectories [41, 56].

High-order sideband generation (HSG) is only sensitive to recollision physics. In HSG, the generation of interband polarization results from weak, resonant NIR excitation instead of electric field-induced tunneling. Resonant excitation restricts the creation of electron-hole pairs to a small region of momentum space with well-defined spin properties, whereas tunneling generates electrons and holes throughout the Brillouin zone with little regard to spin [73, 38]. The quasiparticles created by the NIR laser then evolve, driven by the strong THz electric field. Because the state of these quasiparticles can be controlled independently by the NIR excitation, it is possible to separate the recollision physics from the non-parabolic band physics. Langer, *et al.*, showed that the behavior of the high-order nonlinearities could be controlled by switching between HHG and HSG experiments in the same sample. Rotating the bulk WSe₂ crystal under only strong THz illumination (keeping the THz polarization direction constant) leads to HHG spectra that reflect in a six-fold rotational symmetry of the crystal. Under NIR and THz illumination, the dependence on lattice angle was largely smeared out. In the HHG case, the number of excited carriers depends on how steep the bands are around the band edge. Steeper bands (lower masses) in certain k -space directions leads to more difficult tunneling and so lower interband polarization. In HSG, however, the number of electron-hole pairs is roughly

constant as a function of lattice angle, since the bottom of the bands is approximately spherically symmetric [38].

This chapter will show how changing the polarization of the NIR laser and the angle of the [100] direction of the lattice with respect to the polarization of the THz laser affects HSG. Linear NIR polarization along the THz polarization creates a superposition of spin-up and spin-down electrons (and holes) that are in phase, but when the NIR polarization is perpendicular to the THz polarization, the superposed spins are 180° out of phase. This different initialization will lead to substantial changes in sideband intensity and polarization state.

4.1 Sample details

The large increase in sideband order observed by moving to unstrained, clean GaAs/AlGaAs quantum wells (QWs) and the much stronger THz field produced by the FEL inspired the request for growth of a new sample and the development of a new, more robust processing technique.

4.1.1 Sample growth

Three quantum well (QW) samples with varying degrees of quenched disorder were grown by molecular beam epitaxy. The 5 nm GaAs sample consists of twenty 5 nm GaAs QWs with 20 nm $\text{Al}_{0.3}\text{Ga}_{0.7}\text{As}$ barriers. The 10 nm AlGaAs sample consists of twenty 10 nm $\text{Al}_{0.05}\text{Ga}_{0.95}\text{As}$ QWs with 20 nm $\text{Al}_{0.3}\text{Ga}_{0.7}\text{As}$ barriers. Those two samples are the same

	5 nm GaAs QWs	10 nm AlGaAs QWs	10 nm GaAs QWs	
	50 nm GaAs	50 nm GaAs	150 nm GaAs	
QW	40 nm Al _{0.3} Ga _{0.7} As	40 nm Al _{0.3} Ga _{0.7} As	50 nm Al _{0.3} Ga _{0.7} As	} N
	20 nm Al _{0.3} Ga _{0.7} As	20 nm Al _{0.3} Ga _{0.7} As	15 nm Al _{0.3} Ga _{0.7} As	
	5 nm GaAs	10 nm Al _{0.05} Ga _{0.95} As	10 nm GaAs	
	20 nm Al _{0.3} Ga _{0.7} As	20 nm Al _{0.3} Ga _{0.7} As	15 nm Al _{0.3} Ga _{0.7} As	
	40 nm Al _{0.3} Ga _{0.7} As	40 nm Al _{0.3} Ga _{0.7} As	50 nm Al _{0.3} Ga _{0.7} As	
ES	50 nm GaAs	50 nm GaAs	150 nm GaAs	
	300 nm Al _{0.7} Ga _{0.3} As	300 nm Al _{0.7} Ga _{0.3} As	300 nm Al _{0.735} Ga _{0.265} As	
	50 nm GaAs	50 nm GaAs	100 nm GaAs	
Subst.	500 μ m GaAs	500 μ m GaAs	500 μ m GaAs	

Table 4.1: The growth structure of all three samples. The left two samples were used in Chap. 2, and the right sample is new. The QW region, etch stop (ES), and substrate (Subst.) layers are labeled. The layers labeled N are repeated, $N = 20$ for the 5 nm GaAs and the 10 nm AlGaAs samples, and $N = 10$ for the 10 nm GaAs sample. The only difference is the material in the quantum well region. Both samples are grown analog except for the Al_{0.05}Ga_{0.95}As layers in the 10 nm QW sample, which were grown digitally. The 5% aluminum alloying raises the band gap of the quantum well to cancel out the smaller confinement energy of the wide quantum well.

used for the experiments in Chap. 2. The 10 nm GaAs sample is new and consists of ten 10 nm GaAs QWs with 15 nm Al_{0.3}Ga_{0.7}As barriers, grown with interrupts at AlGaAs-on-GaAs interfaces for very smooth walls [21, 31, 35, 25]. All samples were grown on (100) semi-insulating GaAs substrates.

4.1.2 Sample processing

The old membrane sample geometry, discussed in Ch. 2, could not withstand the intense fields of the rebuilt FEL, so an epitaxial transfer process was developed so that the epitaxial layer could be supported by a NIR- and THz-transparent sapphire substrate¹ [65, 64]. This process involves bringing two extremely smooth surfaces (\AA -scale roughness) into

¹With significant instruction from Garrett Cole from Crystalline Mirror Solutions

contact, here a very clean c-plane, (0001), sapphire wafer and an initially-internal surface of an epitaxial layer exposed by etching away the substrate and sacrificial layers, see App. C. Because both surfaces are so clean and so smooth, Van der Waals forces spread across millimeters are strong enough to bond the surfaces. No extra materials are needed to create the bond, so there are no confounding problems like strong terahertz absorption or stress-induced birefringence. Compared to the membrane geometry, the sapphire substrate is much more robust and easier to handle. GaAs and sapphire also have similar thermal expansion coefficients, so thermal cycling between 15 K and 294 K does not harm the sample.

With a total sample thickness on the order of the THz wavelength, interference between the incident field and reflections off the back surface affect the THz field strength at the front surface, in the QWs. This interference leads to a periodic dependence of the THz field strength in the QWs on THz frequency which is usually undesirable. This interference could be reduced by depositing a THz anti-reflection coating to the back surface of the sapphire [44]. Instead, we deposited a 300 nm, $12.6 \Omega/\square$ thin film of indium-tin oxide (ITO). The large THz reflection off the film exaggerates the interference at the front surface, enhancing the electric field at certain frequencies inside the QWs by more than 50%, see Fig. 4.1, though other frequencies suffer almost complete destructive interference. The THz source, the UCSB Free Electron Laser (FEL), is broadly tunable in frequency with a very narrow linewidth, and HSG does not depend on resonant effects, so the THz frequency can be chosen solely based on the results of processing. The ITO

film is very transparent at NIR wavelengths, so it does not attenuate the laser or the sidebands.

The field enhancement as a function of frequency, $\kappa(f_{\text{THz}})$, was measured using a frequency-extended vector network analyzer (VNA) as a quasi-optical vector reflectometer [59, 3]. Schottky diode multipliers and dividers in the transmitter and receiver extend the 0–40 GHz frequency domain of the VNA up to 70–700 GHz, and precision horns couple the light to and from free space. The reflectometer is set up so that a transmitter illuminates a planar sample. The reflected light is then split off by a beam splitter and directed towards a receiver. Because the reflectometer operates in the frequency domain, Fabry-Perot interference develops that can confound the signal. By running a calibration routine through the VNA on a metal mirror, the phase changes from the beam path can be subtracted out. The reflection spectrum of the sapphire-ITO cavity then represents the electric field at the sapphire surface, where the QWs will be [47]. This measurement has been repeated both with and without QWs on the front surface, and there is little difference. The THz field in the QWs is then $F_{\text{in QWs}} = \frac{\kappa}{\sqrt{\epsilon_{\text{GaAs}}}} F_{\text{applied}}$ where ϵ_{GaAs} is the dielectric constant of GaAs.

4.1.3 Sample optical properties

The thickness and composition of the QWs determine the locations of the exciton absorption peaks. Absorbance spectra were taken using the experimental setup described in Sect. 3.1.1. The sample surfaces are so smooth that the sapphire substrate still func-

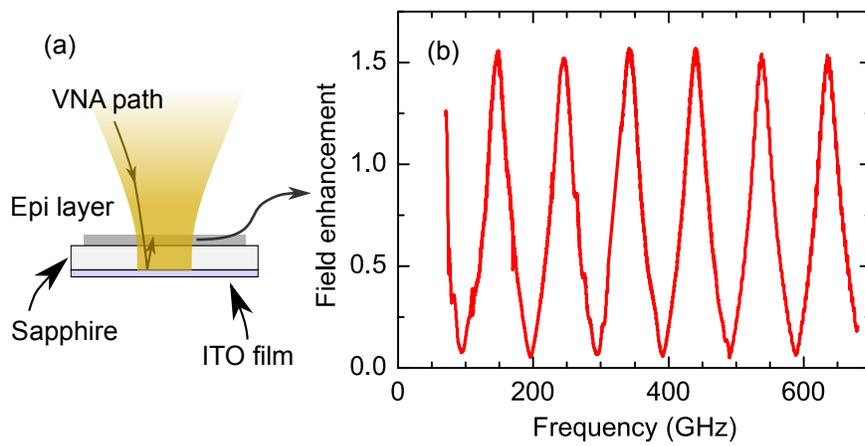


Figure 4.1: The THz field enhancement measurement geometry and results. (a) The focused quasi-optical light from the VNA illuminates the sample. The light transmits through the epitaxially-transferred QWs (epi layer), through the sapphire, and reflects off of the ITO surface back at the QWs. Interference between the incident and ITO-reflected THz field creates the field enhancement $\kappa(f_{\text{THz}})$. (b) The field enhancement as a function of frequency. The enhancement is peaked at about 1.5 with a period of approximately 100 GHz in the measurement range, but there are also frequencies with almost total destructive interference in the QWs.

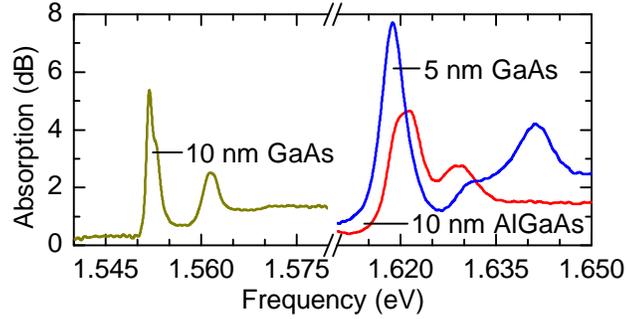


Figure 4.2: Optical absorption spectra of the three samples measured by differential transmission. In all three samples, the heavy hole exciton line (largest peak) and the light hole exciton line (slightly weaker, blue-shifted by 10 or 25 meV depending on well width) are both clearly resolved. Absorption and all other experiments are performed at 15 K.

tions as an etalon at optical wavelengths. As previously mentioned in Sect. 3.1.1, recall that the short-period sinusoid with an amplitude approximately 30% of the transmission signal can be removed in software using a short pass Fourier filter. The samples were held at 15 K for all experiments. Absorbance spectra for all three samples are found in Fig. 4.2. In each spectrum, two peaks are clearly visible, the heavy hole exciton (HHX) and the light hole exciton (LHX) absorption lines. The 5 nm and 10 nm samples can be differentiated by the size of the HHX-LHX splitting. Quantum confinement breaks the HH-LH degeneracy at the top of the valence band such that narrower wells lead to larger splittings. The two 10 nm samples can be differentiated by the aluminum content in the well region. The 5% Al content increases the 2D band gap so that the HHX absorption line is blue-shifted up to the 5 nm GaAs sample HHX absorption line. The 5 nm GaAs and 10 nm AlGaAs samples are from the same wafers used in [4].

The linewidths of the exciton absorption lines reflect the quenched disorder of the

samples and so depend on the growth materials and conditions. In the two GaAs samples, the linewidths of the HHX peaks are largely determined by well-width fluctuations [21, 31, 35, 25]. Because it is half as wide, the 5 nm GaAs QWs should exhibit four times the nonhomogeneous broadening of the 10 nm GaAs QWs for the same size well-width fluctuations. The 10 nm GaAs sample was grown, however, to have particularly smooth well-barrier interfaces, so it should be even more narrow than the 5 nm GaAs sample from well-width fluctuations alone. Indeed, the 5 nm GaAs sample has an HHX full width at half max (FWHM) of 6.3 meV where the 10 nm GaAs sample has an HHX FWHM of 2.0 meV. From well-width fluctuations alone, the 10 nm AlGaAs sample should have a narrower HHX peak than the 5 nm GaAs sample, but it also has a FWHM of 6.3 meV. The additional nonhomogeneous broadening in the 10 nm AlGaAs sample is due to alloy disorder in the well region. The effect of alloy disorder is modeled by Poisson-distributed aluminum content fluctuations over the 10-nm-diameter exciton wavefunction leading to local fluctuations in the band edge. At large length scales, these fluctuations lead to slight blurring of the band separation. We assume that the electron and hole are spheres 10 nm in diameter and that the aluminum atoms in the well region are randomly distributed.

The lattice constant of GaAs is 0.57 nm.

$$V_e = \frac{4\pi}{3} r_e^3 \sim 520 \text{ nm}^3$$

$$\langle N_{\text{Al}} \rangle \sim 140 \text{ atoms}$$

$$\sigma_{\text{Al}} = \sqrt{\langle N_{\text{Al}} \rangle} \sim 12 \text{ atoms}$$

$$n_{\text{Al}} = 5.0 \pm 0.4\%$$

$$E_{\Gamma} = 1.5777 \pm 0.0048 \text{ eV},$$

where V_e is the volume of the quasiparticle wavefunction, $\langle N_{\text{Al}} \rangle$ is the average number of aluminum atoms in that volume, σ_{Al} is the standard deviation of the aluminum atom number and n_{Al} is the number density of the aluminum atoms. From this derivation, we should expect alloy disorder to contribute about 5 meV of nonhomogeneous excitonic linewidth broadening from 5% aluminum concentration.

4.2 Sideband measurements

4.2.1 Experimental preparation

Sidebands are generated by focusing the NIR and THz beams at the same spot on the sample with collinear propagation normal to the sample, see Sect. 3.1.2. Recall that the radii of the NIR and THz focal spots are approximately 100 μm and 0.5 mm, respectively. The NIR polarization is switched between horizontal and vertical by a half-wave plate

(HWP), and the THz polarization is kept horizontal. First, the ideal NIR and THz frequencies and powers were chosen to maximize signal and experimental convenience. Then, the THz polarization state was measured to ensure it was highly linearly polarized.

NIR frequency dependence

With the new sample design with dramatically increased THz field in the quantum wells, it was important to remeasure the HSG dependence on NIR frequency. In the 5 nm GaAs sample, Fig. 4.3(left), the excitonic absorbance spectrum spans 40 meV (20 nm), from the low energy wing of the HHX absorption line to the onset of the E1-LH1 continuum. The sideband spectrum was measured in a 50 meV (25 nm) band encompassing these absorbance features, using a THz frequency, $f_{\text{THz}} = 540$ GHz, total field in the QWs, $F_{\text{THz}} > 15$ kV/cm, and horizontally polarized NIR and THz lasers². Because it was an early measurement in the development of the sideband spectrometer, only three orders were measured, $n = 10, 12$ and 14 . They are reported at the NIR excitation frequency and wavelength, not those of the sideband. The sideband strengths of all three samples are peaked at the HHX line, similar to the results in the membrane samples, see Chap. 2. At lower frequencies, as f_{NIR} shifts into the gap, the sideband signal dies quickly. The sideband signal remains strong up to the high-frequency edge of the measurement, 25

²The rigorous THz power measurement routine from Sect. 3.2.1 had not been developed yet. This and other early experiments measured the energy in the FEL cavity dump by attempting to subtract away the power output before the cavity dump by triggering the 532 nm green laser pulse exactly at the end of the electron beam pulse or long after the electron beam pulse. The difference between an FEL pulse with and without a green pulse would be the energy in the cavity dump, if the ringdown of the FEL cavity was not a substantial amount of energy and it was easy to align the end of the electron beam with the green pulse trigger. While this approach seems elegant, neither of these assumptions are valid.

meV higher in energy than the nearest exciton absorption peak, suggesting that even free electron-hole pairs can generate HSG, although excitons seem to be more efficient. Unfortunately, the sideband signal also seems to jump up and down by approximately 30-50% as f_{NIR} is tuned through the experiment. This effect results from the sapphire substrate forming an etalon for the NIR laser *and* the sideband, which has been observed in absorbance spectra, discussed in Sect. 3.1.1. For continuous sources like white light, the effect can be filtered out using software, but to do the same by stepping the NIR laser more finely than the etalon fringes would require unreasonable stability of both lasers, not to mention extreme patience. This effect only limits the precision of measurement of changes smaller than 30% in sideband signal, not the measurement of a whole sideband spectrum which can span ten orders of magnitude. Because this sample and the 10 nm AlGaAs sample were both discussed in Chap. 2 and are very similar, similar data for the 10 nm AlGaAs sample was not taken.

A NIR wavelength scan of the 10 nm GaAs sample reveals similar qualitative physics, see Fig. 4.3(right). The f_{NIR} scan is much narrower here, 20 meV (10 nm), because the excitonic features are also much narrower and the HHX-LHX splitting is much smaller, but the THz frequency and field are the same. For this measurement, however, the NIR laser is vertically polarized, perpendicular to the THz polarization, but the polarization orientation has little effect on this measurement. The sideband signal is again peaked near the HHX absorption line, but the peak is much broader than the line itself. The exciton line is likely observably broadened by the strong THz field because the zero-field

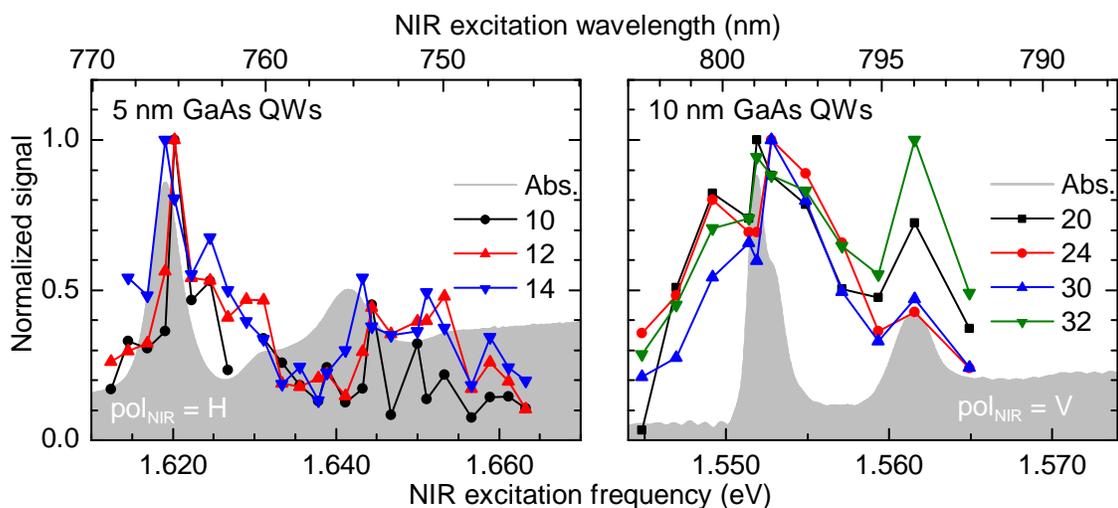


Figure 4.3: Measuring the NIR frequency dependence of HSG. (left) The signal of the $n = 10, 12, 14$ sidebands plotted as a function of f_{NIR} and normalized to the frequency with the highest signal in the 5 nm GaAs QWs. The zero-field absorption spectrum is plotted as a reference. The NIR laser is horizontally polarized, parallel to the THz polarization. The sideband signal peaks sharply at the zero-field HHX absorption line. Sideband signal remains strong over a very broad frequency band, however, suggesting that HSG does not require excitons. (right) The signal of several sidebands between $20 \geq n \geq 32$ plotted as a function f_{NIR} and normalized to the frequency with the highest signal in the 10 nm GaAs QWs. The zero-field absorption spectrum is plotted as a reference. The NIR laser is vertically polarized, perpendicular to the THz polarization. The sideband signal is not peaked as strongly around the HHX absorption line, and there is a clear secondary peak near the LHX absorption line. For both plots, $F_{\text{THz}} > 15$ kV/cm and $f_{\text{THz}} = 540$ GHz.

line is so much narrower in this sample. There is also a second peak in the sideband signal around the LHX absorption line, which fits into the hypothesis that HSG signal is more dependent on the optical density of states than on electron-hole pair character in such strong THz fields.

NIR power dependence

It is also important to determine the best NIR power, P_{NIR} , to use for measurements. For all of these experiments, the NIR spot is not very focused; the diameter is about $220 \mu\text{m}$. In the 5 nm GaAs sample, Fig. 4.4, $F_{\text{THz}} > 45 \text{ kV/cm}$ (still using the difference technique for measuring FEL power). Sideband signal for several different orders up to twentieth is linear in P_{NIR} up to 150 mW. In the 10 nm GaAs sample, Fig. 4.4(center), the NIR power dependence is very complicated. Here, $F_{\text{THz}} > 30 \text{ kV/cm}$. The sideband signal actually peaks at about 125 mW. There are two possible hypotheses to explain this kind of power dependence: exciton-exciton interactions or sample heating. Assuming that the absorption per QW is about 10% and that the excitons live for approximately 1 ns [18], which is most certainly a large overestimate for excitons in THz fields even above 10 kV/cm, the peak density of excitons is 10^{11} cm^{-2} . The average distance between excitons at that density is 26 nm, about two and a half times the diameter of the excitons (given by the well width). At these (very underestimated) separations, there should not be any exciton-exciton interactions. It is more likely that the interaction of the very strong FEL pulse, featuring average powers of several kW, and the NIR laser, which creates carriers that can absorb the THz radiation efficiently, heats up the sample and decreases sideband

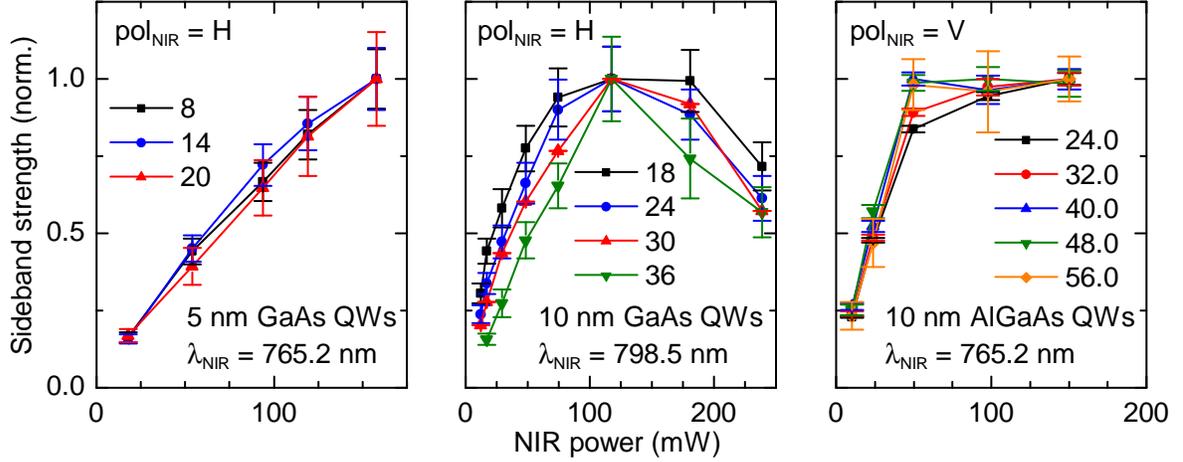


Figure 4.4: The NIR power dependence of HSG in the three samples. (left) The sideband signal in the 5 nm GaAs sample is almost exactly proportional to the NIR power for sideband orders between 8 and 14. (center) The sideband signal in the 10 nm GaAs sample, however, has a strong peak at 125 mW for orders between 18 and 36, but it departs from linear at around 50 mW. (right) The sideband signal in the 10 nm AlGaAs sample actually plateaus at 50 mW.

intensity. A similar phenomenon is seen in the 10 nm AlGaAs sample, Fig. 4.4(right), here with perpendicular NIR and THz polarizations and $F_{\text{THz}} > 45$ kV/cm. Above 50 mW of NIR power, the sideband signal flattens out as the P_{NIR} is increased.

The NIR power and frequency are chosen for the polarization experiments from the results of the preceding subsections. For all following experiments in this chapter, the NIR laser will be tuned to the HHX absorption line, 1.6227 eV (764.0 nm) for the 5 nm GaAs and the 10 nm AlGaAs samples, and 1.5530 eV (798.4 nm) for the 10 nm GaAs sample. The NIR power will be 50 mW for all experiments so that the sideband signal is strong but still approximately linear in P_{NIR} . For ease of FEL electron beam tuning, the FEL frequency will be set to 540 GHz. Also, to ensure that all experiments are directly comparable, the THz field will be attenuated down to a total field of 35 kV/cm inside

the quantum wells, accounting for losses from the ITO beam combiner and the cryostat window. The highest-order sidebands are measurable by averaging over one thousand FEL pulses.

Terahertz polarization

Before final measurements of the THz-induced birefringence and sideband polarimetry, we measured linearity of the polarization of the THz beam to ensure that there was no significant THz electric field component in the vertical direction which could feasibly confound the HSG results and interpretation. To do this, we inserted a rotatable wire grid polarizer at the focus of the beam and a horn-coupled pyroelectric detector immediately behind the polarizer, see Fig. 4.5(a). The angle is defined such that $\phi = 0^\circ$ when the wires are vertically oriented (they transmit horizontally polarized light). It should be noted that this setup is only a linear polarimeter, as a wire grid polarizer cannot distinguish between circularly polarized light and unpolarized light. The transmitted power as a function of wire grid angle, Fig. 4.5(a), shows that the THz field is linearly polarized to a high degree. The contrast ratio between the minima and maxima of the sinusoidal fit indicate that the THz field is $> 99\%$ polarized. The locations of the extrema also show that the polarization is horizontal. Deviations from perfect horizontal linear THz polarization are likely due to wire grid quality or alignment in either the linear polarimeter or the wiregrid polarizers in the THz attenuator discussed in Sect. 3.2.1

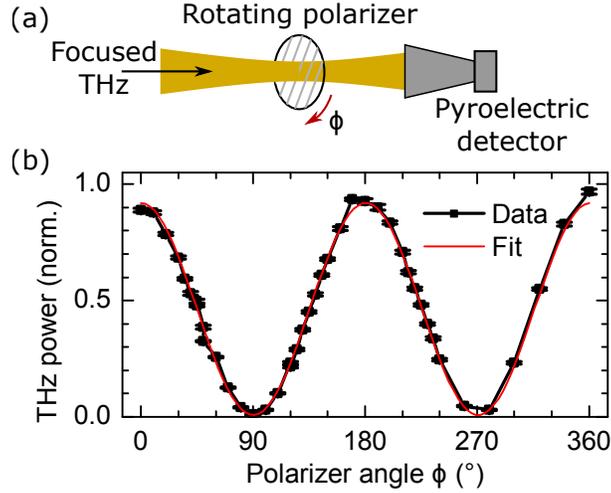


Figure 4.5: Measuring THz polarization. (a) A wire grid polarizer is mounted to a rotatable stage and placed at the focus of the THz beam. A fast, horn-coupled pyroelectric detector is placed immediately behind the wire grid. (b) The THz power measured behind the polarizer as it was rotated. The signal was normalized the the signal when the polarizer was removed. The peaks at 0 and 180° and minima at 80 and 270° reveal that THz beam is $> 99\%$ horizontally polarized.

4.2.2 Results

NIR polarization dependence

The strong terahertz fields lead to the observation of sidebands of more than sixtieth order, but the shape of the sideband spectra are very different between samples, see Fig. 4.6. Because these samples are grown on (100) GaAs, a plane with inversion symmetry, no odd-order sidebands are observed. For the rest of the chapter, the sideband spectra will be referred to by the relative polarization of the NIR and THz light. All three samples show many more sidebands of positive order than negative orders, and each spectrum spans more than 150 meV. In the classical three step model, the electron and hole tunnel ionize from the exciton and begin to accelerate in their respective bands

with zero momentum and kinetic energy. As they accelerate in the THz field, therefore, they can only gain energy, so that photons from recollisions must contribute to positive sidebands. In the 5 nm GaAs sample, Fig. 4.6(a), the negative order sidebands fall off quickly and purely exponentially. The positive order sidebands also fall off exponentially, quickly at first but then much more slowly, such that the $n = 66$ order sideband is the highest order clearly observed. In the 10 nm AlGaAs sample, Fig. 4.6(b), the sideband spectrum is almost identical to spectrum in the 5 nm GaAs sample. The negative orders again fall off quickly and purely exponentially, and the positive orders also fall off quickly at first but then much more slowly. Despite the sideband signal being a little bit weaker than the spectrum in 5 nm GaAs sample overall, the $n = 66$ order sideband is clearly observed. In the 10 nm GaAs sample, Fig. 4.6(c), the sideband spectrum has a much different shape. The negative order sidebands fall off exponentially at first, but between $-8 \leq n \leq -16$, the sideband signal flattens out before falling away again. This negative order sideband “shoulder” is likely related to interference between QW HSG signal and weak HSG signal generated by the relatively thick GaAs end caps, see Table 4.1 for growth details. The positive order sidebands also shows three distinct regions. Up to about the $n = 12$ order sideband, the sidebands fall off quickly and exponentially. At this order, they flatten out significantly until about the $n = 30$ when they fall off at a more moderate exponential rate. This very slow fall off looks suspiciously like the classic three step model plateau, but it is not nearly large enough. The ponderomotive energy for the THz field strength and frequency is 1.2 eV, meaning the plateau should extend

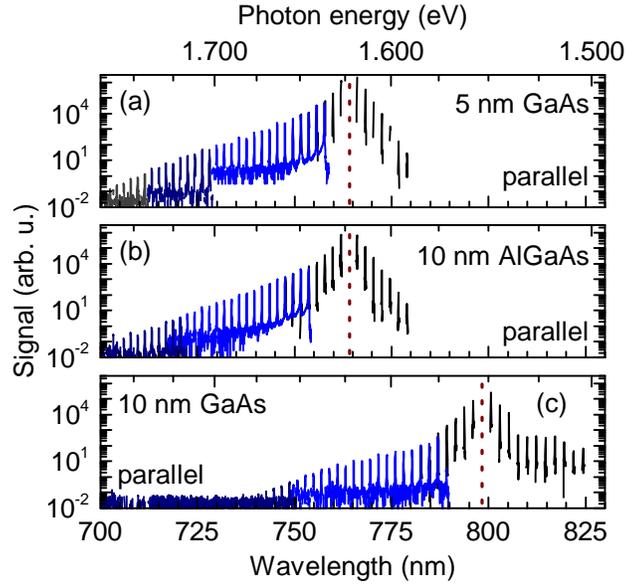


Figure 4.6: The sideband spectra created by parallel NIR and THz polarizations. The NIR laser is represented by a dotted vertical line. The colors represent different detectors. Black discontinuous peaks are very close to the NIR laser and must be measured with a PMT in a high-rejection monochromator. Blue continuous spectra are measured with a CCD in an imaging spectrometer, darker having longer exposure times. Sidebands measured by the CCD are scaled to match those measured by the PMT by ensuring the detectors measure several sidebands of overlap. To match CCD images of the same sideband spectrum with different integration times, the signal is normalized to the number of FEL “pulses.” The readout noise, which is the dominant noise source, is then diminished by a factor of the number of pulses. The noise floor of the PMT is higher, but no extra negative sidebands are observed using the CCD to detect those sidebands.

for *twelve hundred* orders with a cutoff wavelength of 230 nm.

Changing the relative polarization of the NIR and THz lasers has a large effect on the sideband spectra, leading to the observation of up to 86th order sidebands with total spectrum bandwidth of more than 200 meV, see Fig. 4.7. The sideband conversion efficiency is proportional to the area of a Gaussian function fit to a sideband peak. More details are shown in Fig. 2.6 and discussed in Sect. 3.1.2. In all three samples, the negative order sidebands and the positive low-order sidebands do not change with

relative polarization. These sidebands result from both perturbative sideband generation and very low energy electron-hole recollisions which do not extend far into the bands [57]. Above about tenth order, however, sidebands from the perpendicular geometry fall off much more slowly than sidebands from the parallel geometry for all samples. That different NIR polarizations exhibit very different HSG response means the strong THz field induces birefringence in the sample. In the 5 nm GaAs sample, Fig. 4.7(a), the two polarization geometries separate in HSG intensity between twentieth and thirtieth order. In the 10 nm AlGaAs sample, Fig. 4.7(b), the two spectra separate at much lower order relative to the 5 nm GaAs sample, between tenth and twentieth order. In the 10 GaAs sample, Fig. 4.7(c), the perpendicular HSG spectrum looks like the parallel spectrum has been stretched horizontally. The two spectra separate at a similar order to the other 10 nm sample, suggesting that this phenomenon is connected to the well width.

Lattice angle dependence

The observation that the well width determines at what order the HSG spectra from the two excitation geometries diverge suggests that subband separation could play a role in this polarization effect. Although the subband energy differences at $k = 0$ cannot be changed without growing new samples, the energies at finite k depend strongly on lattice direction. According to the Luttinger model for bulk GaAs, the $\langle 100 \rangle$ family of directions has a different effective mass than the $\langle 110 \rangle$ family. This non-spherical bulk effective mass reduces to non-cylindrical effective masses in a quantum well with complicated effects on hole subband avoided crossings. By rotating the sample relative to the THz driving field,

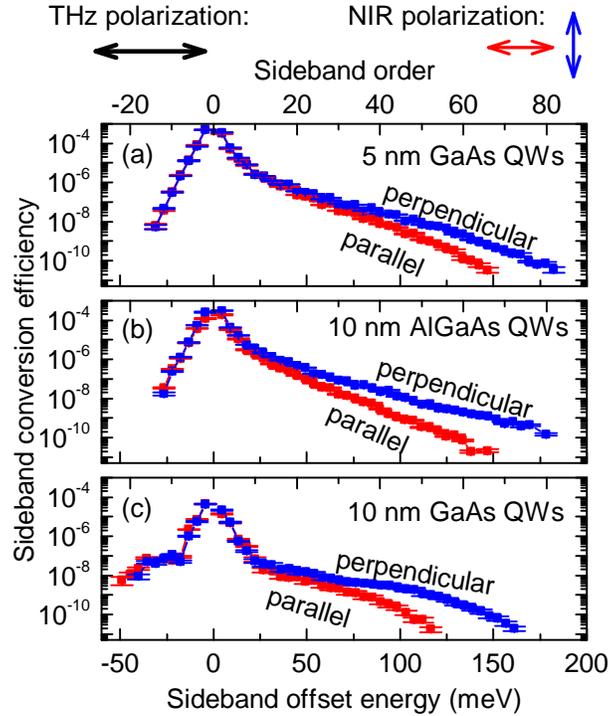


Figure 4.7: Sideband conversion efficiencies for two THz-NIR relative polarizations. Polarization definitions are drawn on the top. The perpendicular-excited sidebands are much stronger than the parallel-excited sidebands above a certain order threshold for all three samples, $n_t \approx 30$ for the 5 nm GaAs sample and $n_t \approx 10$ -20 for the two 10 nm samples. This threshold is directly related to the subband spacing of the QWs. Above this threshold, the perpendicularly-excited sidebands fall off much slower. In both the 5 nm GaAs and 10 nm AlGaAs samples, the sidebands above n_t fall off almost exactly exponentially with order. High order sidebands from the 10 nm GaAs sample, however, show a much more complicated relationship with increasing order.

it is possible to isolate how HSG spectra and the THz-induced birefringence depend on the subbands at finite k .

Rotating the sample relative to the terahertz polarization reveals the effects of the lattice on HSG, a strong sensitivity to the band structure and a robustness to quenched disorder, Fig. 4.8. In the upper left of Fig. 4.8, the rotation of the lattice is defined as the angle between the [100] axis of the GaAs lattice and the THz polarization by the angle θ looking at the sample along the propagation direction of the excitation light. In the 5 nm GaAs sample, the upper right of Fig. 4.8, two angles are measured and their dispersion relations are calculated, 85° and 55° . At 85° , as has been shown in Fig. 4.7, sideband spectra from the two excitation geometries diverge at a relatively high sideband order, about $n = 30$. At 55° , the two excitation geometries actually do not diverge significantly at all, up to the highest order observed sidebands. The small THz-induced birefringence in this sample can be traced to the large avoided crossing between the HH1 and HH2 subbands, almost 10 meV wide at relatively high momentum, $k \sim 0.1$ 1/a. In the 10 nm GaAs sample, bottom right of Fig. 4.8, two angles are measured and their dispersion relations are calculated, 91° and 63° . At 91° , as has been shown in Fig. 4.7, sideband spectra from both excitation geometries diverge at a relatively low sideband order, $n < 20$. At 63° , sideband spectra from both excitation geometries change significantly, the sideband conversion efficiencies become stronger and the shapes change. Compared to the 5 nm QW dispersion relations, the HH1-HH2 avoided crossing is much smaller, only 5 meV, and at lower momentum, $k \sim 0.08$ 1/a. In the 10 nm AlGaAs

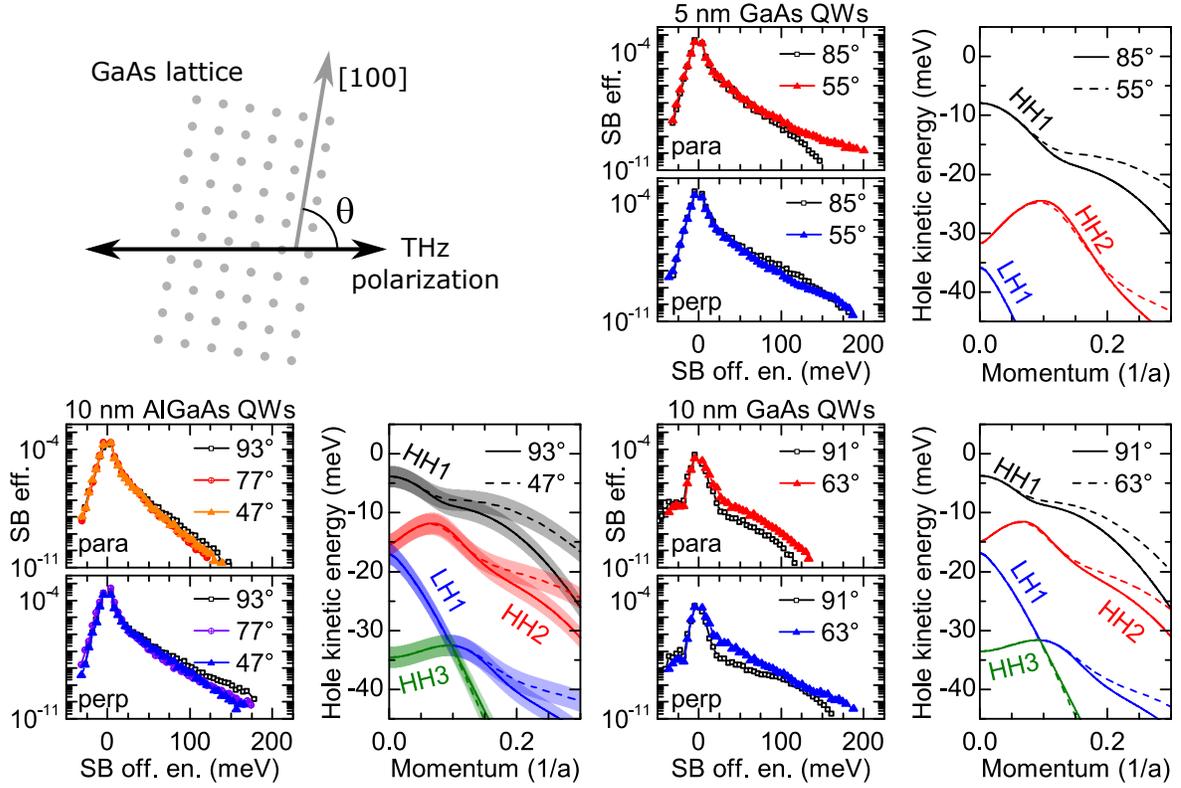


Figure 4.8: The ratio I_{\perp}/I_{\parallel} in all three samples for different lattice orientations, and calculated valence subband dispersion relations. The THz polarization was kept horizontal and the sample was rotated clockwise, shown in the upper left. The empty black square scatter points are replotted from Fig. 4.7. The valence subband dispersion relations are plotted for two lattice directions to demonstrate the relationship between hole dispersion and sideband spectrum. For comparison, the hole kinetic energy contributes approximately 30-40% of the sideband offset energy (SB off. en.). **5 nm GaAs:** The large order threshold n_t for parallel and perpendicular sideband spectra to diverge is due to the HH1-HH2 avoided crossing being at much larger momentum with a much larger splitting than in the 10 nm QWs. At 55° orientation, so the THz is nearly along the [110] direction, the splitting is even larger, and so there very little difference between the parallel and perpendicular HSG spectra. **10 nm AlGaAs:** The SB efficiencies show little dependence on lattice orientation. Alloy disorder, which blurs the dispersion relation at moderate length scales, smooths out the cubic symmetry of GaAs and is represented in the dispersion relation as a broadening of the curves. The ensemble of recollisions therefore average out to have the same response at different angles despite the local lack of cylindrical symmetry. **10 nm GaAs:** The sideband spectra depends substantially on angle because the avoided crossings in the dispersion relation are at relatively low energy, and there is little macroscopic disorder to average out hole state interference.

sample, bottom left of Fig. 4.8, three angles are measured and two of their dispersion relations are calculated, 93° , 77° , and 47° . Unlike the other two samples, the sideband spectra from this sample are almost independent of orientation. Related to how alloy disorder broadens the exciton absorption lines, it also blurs the hole subbands at large length scales. At the length scale of a recollision, a few tens of nanometers, alloy disorder does not mask the local avoided crossings, so the THz-induced birefringence is similar to that in the 10 nm GaAs sample. At the length scale of the $200\ \mu\text{m}$ NIR spot, however, alloy disorder smears out the hole subbands. The ensemble of recollisions averages out the four-fold symmetry to nearly cylindrical symmetry.

Sideband polarimetry

So far, these experiments have been performed assuming the NIR laser is perfectly linearly polarized and that the sidebands have identical polarization states to the laser. Using the polarimeter and techniques discussed in Sect. 3.1.2, the actual polarization states of the NIR laser and sidebands were measured, see Fig. 4.9. In all of the measurements, the NIR laser and all of the sidebands were measured to be perfectly polarized. They can then be described by two angles, α , the angle that the major axis of the ellipse makes with the THz polarization, and γ , the “degree” of ellipticity, see the center of Fig. 4.9. In the 10 nm AlGaAs sample, the left half of Fig. 4.9, the polarization states of the sidebands were measured for three NIR polarization states, one in the parallel excitation geometry and two in the perpendicular excitation geometry. In the 10 nm GaAs sample, the right half of Fig. 4.9, the polarization states of the sidebands were measured for four

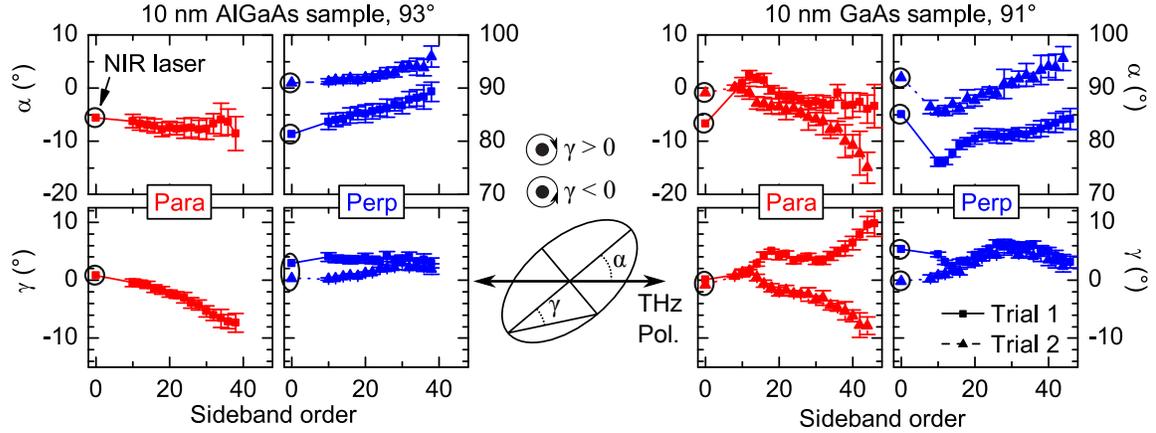


Figure 4.9: The measured polarization states of sidebands for various NIR laser polarization states. The angles of the NIR polarization ellipse, α_{NIR} and γ_{NIR} , are circled for clarity. (center) The definition of the angles in the polarization ellipse, α is the angle the major axis makes with THz polarization, and γ is the arctangent of the ratio of the semiminor and semimajor axes. The top row of the graphs plot α and the bottom row plots the associated γ for the particular measurement. (left) Sideband polarization states for three different NIR polarization states, one with the NIR laser approximately parallel to the THz polarization and two with the NIR laser approximately perpendicular. The sideband polarization states are systematically different from the NIR laser. (right) Sideband polarization states for four different NIR polarization states, two with the NIR laser approximately parallel to the THz polarization and two with the NIR laser approximately perpendicular. The sideband polarization states again are systematically different from the NIR laser.

NIR polarization states, two for each polarization excitation geometry. For all data sets, the sideband polarization states are measurably different, sometimes substantially so, from the NIR polarization state. Furthermore, the sideband polarization states depend sensitively on the NIR polarization state. Small changes in either angle of the polarization ellipse of the NIR laser lead to large changes in the polarization ellipses of the sidebands. It is difficult to construct a simple model to explain these observations, as these results show that two linearly polarized lasers are combining to make elliptically polarized light with seemingly random angles.

4.3 Theoretical model

In collaboration with theorists Qile Wu and Renbao Liu, full quantum mechanical and semiclassical simulations of the HSG process were developed³. These simulations show that, under experimental conditions, electron and hole trajectories explore roughly 10% of the Brillouin zone around the Γ point. The trajectories of the spin 1/2 electrons are entirely contained in the lowest electron subband, which is nearly parabolic. However, the spin 3/2 holes travel through highly non-parabolic subbands that are strongly coupled away from the Γ point. We are unable to account for the experimental results within the simplest semiclassical model of electron and hole motion, in which $\hbar\dot{\mathbf{k}} = e\mathbf{E}(t)$ and the phase velocity at any instant is determined solely by the dispersion relation $E(\mathbf{k})$, where \hbar is the reduced Planck constant, \mathbf{k} is the crystal momentum, e is the electric charge, \mathbf{E} is the electric field, and E is the energy of the state [2]. Within the last 20 years, it has become generally appreciated that such a description is incomplete and requires the consideration of Berry physics [11, 62]. The Berry curvature $\mathbf{F}_{mn}(\mathbf{q})$ is defined by calculating the Berry connection $\mathbf{R}_{mn}(\mathbf{q})$, which is analogous to the Berry phase in one-dimension:

$$\mathbf{F}_{mn}(\mathbf{k}) = \nabla_{\mathbf{k}} \times \mathbf{R}_{mn} - i\mathbf{R}_{mn} \times \mathbf{R}_{mn} \quad (4.1)$$

$$\mathbf{R}_{mn}(\mathbf{k}) = i \langle u_{m\mathbf{k}} | \nabla_{\mathbf{k}} | u_{n\mathbf{k}} \rangle, \quad (4.2)$$

³As the theory was developed by Qile Wu and Renbao Liu, and the former performed all of calculations, the following paragraphs will be an overview to explain the relevant physics rather than a full derivation of the theory.

where $|u_{m\mathbf{k}}\rangle$ is the non-plane wave or cellular part of the Bloch wavefunction of an electron ($e^{i\mathbf{k}\cdot\mathbf{x}}|u_{\mathbf{k}}\rangle$), m and n are the band indices spanning the bands included in the calculation, and \mathbf{k} is the crystal momentum [11]. The Berry curvature is zero in the conduction subband but nonzero in the hole subbands. For this calculation, the Berry connection is a 2×2 matrix, and therefore non-Abelian. The Berry connection rotates the spinor that combines the HH1 and HH2 subbands, $\vec{\eta} = [|\text{HH1}\rangle, |\text{HH2}\rangle]^T$, and describes the hole state. Inclusion of the non-Abelian Berry curvature and connection in the semiclassical equations of motion captures the crucial effects of spin-orbit coupling and how the hole state mixes with other valence subbands during its trajectory. When the electron and hole recombine to emit a sideband, its polarization state is determined by the angular momentum of the electron and hole just before recombination. This calculation reproduces all the main features observed in experiments, supporting the role of Berry curvature in the high-order sideband generation.

4.3.1 Separation into circular components

To better consider how non-Abelian Berry connection in the valence subbands affects HSG, it is useful to think of the electron-hole pair excitation process in the circular polarization basis to isolate individual electron and hole species, see Fig. 4.10. An incoming NIR photon is separated into left and right circularly polarized components, σ_{NIR}^- and σ_{NIR}^+ , with complex coefficients given by the initial NIR photon. This conceit allows the use of Berry physics. In “traditional” manifestations of Berry phase, only closed paths

create detectable a Berry phase because the dynamical phase of two interfering particles otherwise is random. The HSG process is at its core not a closed path. The non-parabolic hole bands mean the recollision will not necessarily be at the same real-space location as creation, and the electron and hole will be at much different momenta. By initializing every HSG event with the same relative phase, however, the Berry connection becomes measurable. Once the electron and hole recollide, the electron state remains the same, but hole state has changed dramatically.

Consider an exciton excited by a σ_{NIR}^- photon resonant with the HHX state, the top of Fig. 4.10. Recall from Ch. 1 that subbands are defined by an envelope function f_n of the n th QW bound state and that, at $k = 0$, the electrons have spin-1/2, the heavy holes have $j = 3/2$, $m_j = \pm 3/2$, and the light holes have $j = 3/2$, $m_j \pm 1/2$. The electron state is in the first electron subband with spin-up, $|E_i\rangle = f_1 |S, \uparrow\rangle$, and the hole state begins as a heavy hole in the first subband $|\text{HH}_i\rangle = f_1 |3/2, +3/2\rangle$. The electron remains in the same subband with the same spin throughout its trajectory, $|E_f\rangle = |E_i\rangle$. As the hole accelerates, however, it mixes with nearby states in the valence subband manifold to become a superposition $|\text{HH}_f\rangle = \sum_{n,j} \eta_{n,j} f_n |u_j\rangle$ where $\eta_{n,j}$ are coefficients calculated from the non-Abelian Berry connection and $|u_j\rangle$ is the spin state label. Not all of the terms of this sum can recombine with the electron state, however. For a dipole-allowed transition, the two states must be in the same subband and have an angular momentum difference of ± 1 . In $|\text{HH}_f\rangle$ there are components with $|3/2, -1/2\rangle$ spin, meaning that there is nonzero probability for a σ_{NIR}^- to create a σ_{HSG}^+ photon. The bottom of Fig. 4.10

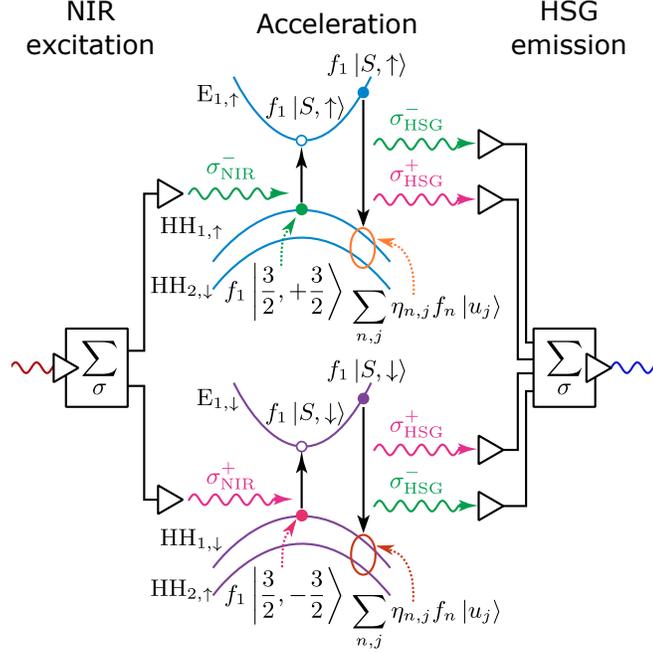


Figure 4.10: Mechanism for the creation of a single HSG photon including the Berry connection of the valence band. There are three steps to this mechanism: NIR excitation, tunneling and acceleration, and HSG emission. First, the NIR excitation light is decomposed into left and right circularly polarized light to isolate its coupling to a single electron species species, each with complex coefficients determined by the NIR polarization. A σ_{NIR}^- (σ_{NIR}^+) photon resonant with HHX creates a $f_1 |S, \uparrow\rangle$ ($f_1 |S, \downarrow\rangle$) electron and a $f_1 |3/2, +3/2\rangle$ ($f_1 |3/2, -3/2\rangle$) heavy hole. These particles then accelerate, driven by the THz. During acceleration, the electron wavefunction does not change character. The heavy hole wavefunction, however, mixes with the nearby hole subbands. Upon recollision, the hole state is no longer simply $f_1 |3/2, +3/2\rangle$ ($f_1 |3/2, -3/2\rangle$), but a superposition of several different hole subbands $\sum_{n,j} \eta_{n,j} f_n |u_j\rangle$ with $\eta_{n,j}$ coefficients determined by the non-Abelian Berry curvature and the initial phase of the photon that created the hole, f_n the subband label, and $|u_j\rangle$ the spin state label. Some of these states have allowed electric dipole moments with the electron, and so create photons σ_{HSG}^\pm that then interfere to create the HSG emission. Controlling the relative phase of the σ_{NIR}^+ and σ_{NIR}^- light influences the output by changing how the emitted photon states interfere.

demonstrates the same physics but starting with excitation by a σ_{NIR}^+ photon.

4.3.2 Semiclassical trajectories

To demonstrate how the non-Abelian Berry connection relates to HSG trajectories, two semiclassical trajectories that are influenced very differently by relative NIR and THz polarization are shown in Fig. 4.11. As can be seen in Fig. 4.7(c), the relative strength of the parallel- and perpendicular-excited sidebands of twentieth and sixtieth orders are very different. The sixtieth order is not even observed in the parallel-excited HSG spectrum, for example. The phase of the THz field for ionization and recollision for the two orders is labeled in the upper right of Fig. 4.11. The trajectory resulting in the sixtieth order sideband is over 100 fs longer (more than 30% longer) than the twentieth order trajectory, so it has more time to accelerate in the THz field.

The details of the twentieth order trajectory are plotted in the lower left two graphs in Fig. 4.11. Over the course of the acceleration step, the electron and hole paths separate by almost 30 nm at their farthest. The spinor state of the two particles is plotted at seven different times (the spinor direction is chosen for excitation by a σ_{NIR}^+ photon). The electron spinor direction does not change, but the hole spinor rotates by a large amount in the last 50 fs. The same seven instants for the hole state are shown on the dispersion relation below, where the area of the circles represents the norm of the spinor coefficient. At instant #1, the hole is entirely in the HH1 subband. As the hole accelerates to the left, instant #2, it only slightly mixes with the HH2 subband. As the hole accelerates

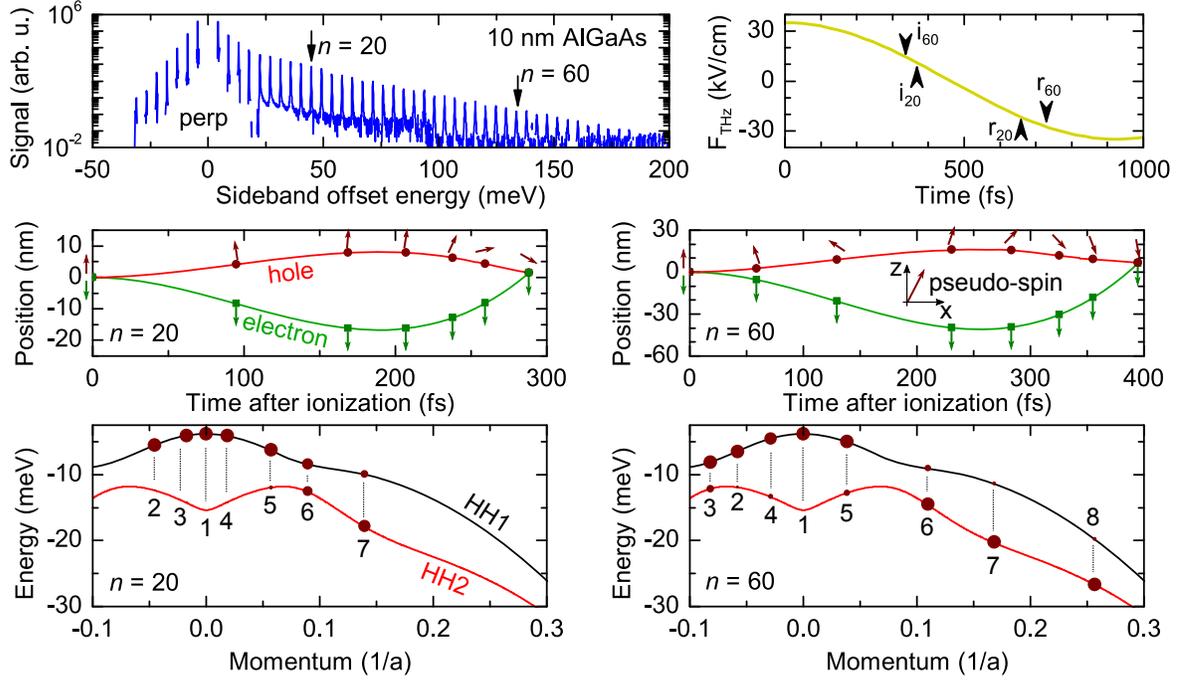


Figure 4.11: Semiclassical trajectories for the $n = 20$ and $n = 60$ sidebands in the 10 nm AlGaAs sample. (upper left) Full perpendicular HSG spectrum from the 10 nm AlGaAs sample oriented at 93° . (upper right) Time trace of the THz electric field. The arrows point to the time of ionization and recollision for the two sidebands modeled here, i_n and r_n for the time of ionization and the time of recollision. (middle left) The real-space trajectories of the electron and hole for the 20th order sideband. The arrow represent the spinor direction in x - z plane for seven instants throughout the trajectory. The hole spinor state is an approximately equal superposition of HH1 and HH2 at recollision. (lower left) The location of the hole in the valence subbands at those seven instants with the area of the maroon circle representing the relative weight in either subband. (middle right) The real-space trajectories of the electron and hole for the 20th order sideband, with the spinor direction drawn for eight instants. The hole spinor almost entirely flips to the HH2 state. (bottom right) The location of the hole in the valence subbands at those eight instants.

to the right, instants #3–7, the hole spinor shifts significantly as it accelerates through the avoided crossing. The majority of the spinor weight is then in the HH2 subband where it cannot recombine with the electron via an electric dipole transition. Also, in the position plot again, notice that the position at recollision is not at exactly zero, it is slightly positive. This translation results from the hole becoming “more” massive as it enters the region of the HH1 subband with less curvature. If the electron and hole remained the same mass the entire time, the recollision would occur at exactly zero.

The details of the sixtieth order trajectory are plotted in the lower right two graphs in Fig. 4.11, and they tell a far more exaggerated story. The electron and hole travel much farther apart for this trajectory, but, more significantly, the hole spinor state rotates substantially twice, shown now at eight instants in the trajectory. At $t = 120$ fs, instant #3, the hole sits right in the avoided crossing, and some weight is transferred to the HH2 subband. As it accelerates to the right, instant #4, some of the weight remains in HH2. As it continues to accelerate to the right in the last 100 fs, the spinor weight almost completely shifts to the HH2 subband. Also, notice that the location of recollision has shifted almost 10 nm.

Understanding how the two polarization geometries can result in such different sideband strengths requires looking into how the trajectories excited by σ_{NIR}^+ and σ_{NIR}^- photons interfere. The semiclassical results above only demonstrate the large scale of the mixing calculated by the non-Abelian Berry curvature. The entirety of the HSG process including non-Abelian Berry curvature can be written in one simple matrix for every

sideband order n , $T_{ij,n}$ in the equation:

$$\begin{pmatrix} \sigma_{\text{HSG}}^+ \\ \sigma_{\text{HSG}}^- \end{pmatrix}_n = \begin{pmatrix} T_{++} & T_{+-} \\ T_{-+} & T_{--} \end{pmatrix}_n \begin{pmatrix} \sigma_{\text{NIR}}^+ \\ \sigma_{\text{NIR}}^- \end{pmatrix}. \quad (4.3)$$

The elements of $T_{ij,n}$ contain all of the terms that create the $i = \pm$ -polarized HSG photon from the $j = \pm$ -polarized NIR photon for a given THz polarization and lattice orientation. The NIR polarization inputs different complex coefficients on the σ_{NIR}^\pm components and causes different terms in the matrix product to constructively or destructively interfere.

The full calculation of $T_{ij,n}$ is very complicated. In addition to including the HSG physics, it includes a full calculation of the dispersion relations and non-Abelian Berry curvature for the six QW subbands for all four hole species and scattering as a simple dephasing rate that does not couple subbands. It is an open question as to what kind of scattering to include and how to actually include it. By taking the ratio of sidebands in the two polarization geometries, however, some of the possible scattering mechanisms can be divided out. Theoretical ratios and experimental ratios are compared in Fig. 4.12. The top row of Fig. 4.12 compares theory and experiment for the 5 nm GaAs sample at both lattice angles. The calculated and measured ratios do not agree very well, particularly at 55° , likely due to the hard wall approximation which breaks down for narrow QWs and leads to an overestimate of the Berry connection. The middle row of Fig. 4.12 compares theory and experiment for the 10 nm GaAs sample. Theory and experiment agree much better here. The calculated ratio predicts a large peak at a sideband offset energy of about 125 meV, which is partially observed in the measurement. The source of this peak, however, is not understood. The bottom row of Fig. 4.12 compares theory and

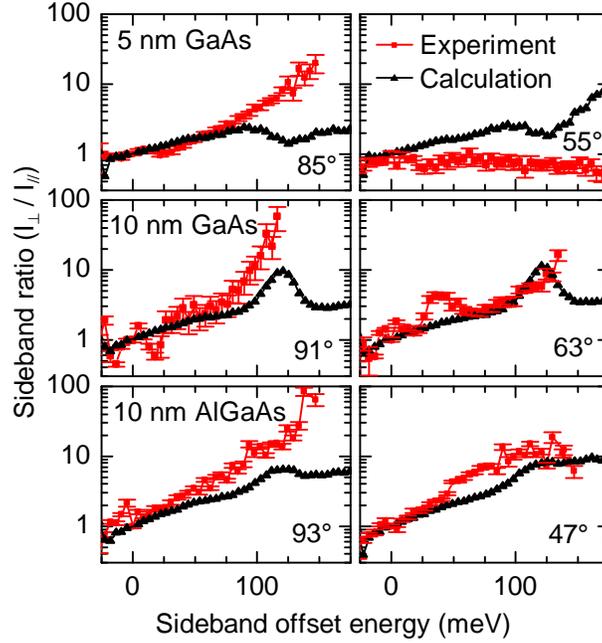


Figure 4.12: Comparison of experiment and theory. The ratio I_{\perp}/I_{\parallel} is compared in all three quantum wells at different lattice angles. The non-Abelian Berry connection is responsible for the qualitative agreement. The quantum theory of sidebands only includes macroscopic dephasing, no k - or t -dependent scattering. Comparing intensity ratios instead of absolute intensities cancels out some of the complicated scattering behavior.

experiment for the 10 nm AlGaAs sample. Theory and experiment agree here as well. The calculated ratio has a much smaller peak than in the other 10 nm sample, and the measured ratio follows this trend. Overall, this model matches the measured data quite well.

A dramatic, and surprising, prediction of this non-Abelian Berry connection model is the rotation of sideband polarization states relative to the incident NIR polarization state. What initially began as a simple (if technically difficult) exploration of the source of the stark difference between HSG spectra for different polarization excitation geometries became the strongest evidence in support of this model, see Fig. 4.13 for a comparison of

experiment and theory (the experimental data is replotted from Fig. 4.9). The detailed interference of different electron and hole trajectories results in a sensitive dependence of the final sideband polarization state on the incident NIR polarization state. Calculations in the $T_{ij,n}$ formalism, when the NIR polarization state as the input variable and no free parameters, show surprising agreement with experiment in the 10 nm AlGaAs sample at 93° , Fig. 4.13(left) It should be noted that the models contained in $T_{ij,n}$ are incomplete and that quantitative agreement should not be expected, recall that scattering is not perfectly accounted for and the QW properties are calculated assuming hard walls. The qualitative agreement that individual sideband polarization states change with order and depend sensitively on NIR polarization state is significant, though, and shows the importance of considering Berry physics when thinking about HSG.

4.4 Conclusion and remaining questions

In conclusion, we have studied how the interplay between the relative orientations of the NIR polarization, THz polarization, and lattice affect HSG. We have measured HSG spectra of up to 86th order and spanning over 200 meV, a bandwidth of over 12% of the NIR laser frequency, by manipulating those relative orientations. We have shown conclusively that electrons and hole accelerate coherently through the the lattice before recolliding. This coherence allows for interference between different electron and hole states, initialized by the NIR polarization and caused by non-Abelian Berry curvature in the hole subbands, and leads to large changes in sideband strength and sideband

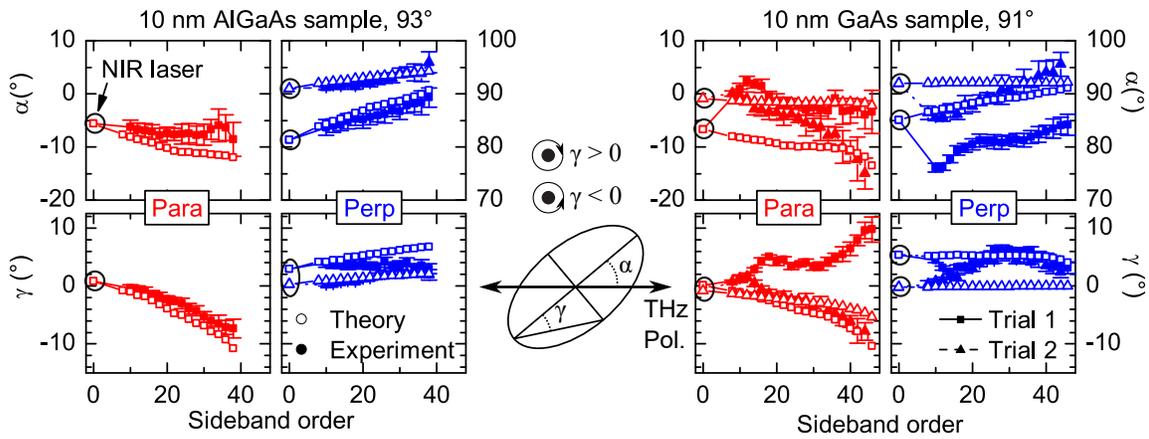


Figure 4.13: Polarization state of the sidebands. Experimental measurements are filled scatter points (replotted from Fig. 4.9, theory are empty scatter points). Polarization state measurements were performed on both 10 nm samples for both parallel and perpendicular geometries, and all the measurements were duplicated with approximately the same NIR laser polarization except for one. The NIR laser polarization state is plotted as the order zero sideband and circled in black for each measurement. Overall, both theory and experiment agree that the polarization state of a given sideband is extremely sensitive to the NIR laser polarization state. In the 10 nm AlGaAs sample, theory and experiment agree almost quantitatively.

polarization state.

The results presented here point to several interesting future experiments. The relationships between HSG and either the NIR power or frequency are poorly understood. The nonlinear dependence of HSG on the NIR power presented in Fig. 4.4 is likely due to heating. If heating could be removed, however, then there are potentially interesting effects that occur at high exciton densities. Furthermore, the injection of non-excitonic electron-hole pairs by tuning the NIR frequency could also result in interesting phenomena such as the observation of a large number of negative sidebands predicted in [63].

In the next experiments, the observations discussed here should lead to a new generation of complete band structure measurement because HSG is inherently sensitive to both elements of the Bloch wavefunction, the plane-wave dispersion relation, $e^{i\mathbf{k}\cdot\mathbf{x}}$, and the cellular function, $|u_{\mathbf{k}}\rangle$. By clever control of the NIR polarization, the THz frequency and field, and lattice orientation, a self-consistent algorithm can be developed for the direct measurement of the electron and hole dispersion relations, non-Abelian Berry curvature, and even k - and t -dependent scattering rates in any material. For example, exciting with purely circularly polarized light will isolate any polarization changes to just one electron and hole species, measuring the non-Abelian Berry curvature more directly by eliminating the complex interference generated by linearly polarized light. Then, by tuning the THz frequency and field strength, the timescales for sidebands of the same order or offset energy can be changed to measure different scattering rates. Because the THz polarization must be linear for high-probability recollisions, all the proposed experiments would

have to be repeated for several different THz polarization-lattice orientations.

Appendix A

Equipment details

This appendix includes descriptions of most of the equipment used in the experiments for this dissertation. Many descriptions of the equipment used in this appendix comes from the theses of Ben Zaks and Sam Carter.

A note on software In redesigning the experimental setup, a decision was made to move from using LabVIEW to Python and PyQT. Although development time in LabVIEW from new instrument to functional graphical user interface (GUI) can be quite fast, development of more advanced interfaces and even modest number crunching routines is much simpler in Python. Furthermore, Python works on all major operating systems and is free as in speech and as in beer, so installing on new and home computers is trivial. Programs are kept up to date and consistent using Github, a cloud-based git repository service, under a collaboration named SherwinGroup. It is a little bit like Dropbox but aimed at collaborative programming. It allows for editing the same file

by multiple people at the same time and handles merging all of the changes. The old LabVIEW code is still on the computer in 1380D, called 1380D-PC.

A.1 Cryostat

The cryostat is a CTI Cryogenics Model 22 closed-cycle helium refrigerator that can cool to a minimum temperature of 10 K in about an hour. The cryostat is connected to a CTI Cryogenics Model 8200 helium compressor through two high pressure helium lines. The cold head is mounted on a linear bearing to allow tight optical access as well as simple sample access. Bungie cords and a micrometer provide fine left-right control of the sample in the optical path. The cryostat currently has two electrical connections, one to the compressor and one to a Lakeshore 330 temperature controller. The samples are glued to a copper bar with rubber cement. For quick sample changes, the sample bar is bolted onto a copper mount that is semi-permanently attached to the cryostat. A Lakeshore temperature diode is attached to this mount to monitor the temperature of the samples. The copper mount is attached to a three-axis dove-tail translation stage, an Elliot Scientific ultra-small micropositioner stage MDE269, with 3 mm of travel in x, y, and z. This stage can be translated when the sample is cold using Allen wrenches fed through the vacuum shroud. The temperature of the sample will spike to 30 K or so because the wrenches are room temperature, but it should settle back down. A molybdenum-based spray dry lubricant prevents sticking and eases sliding at all temperatures. The stage is a poor conductor of heat, so the outer conductor of

a BNC cable provides thermal contact between the copper sample mount to the copper cold finger mount. The cold finger mount has a $22\ \Omega$ nichrome wire heater attached at the base with stycast. The copper cold finger mount bolts to the cold finger of the cryostat. All interfaces have Apiezon N grease for improved thermal contact. The cold finger itself has wires for electrical contact, but these are not currently set up.

Several pieces make up the thermal radiation shield and the vacuum shroud. Both the shield and the vacuum shroud are made up of two main pieces, a cylindrical part and a head part. The cylindrical parts are semi-permanent and attach to the actual cryostat. The head parts attach to the cylindrical parts and allow for easy sample access. The vacuum shroud head has three rotary feed-throughs for cryogenic control of the three-axis translation stage, as mentioned earlier. There are also two large windows for optical access. The current front window is most likely made of suprasil quartz, which has decent THz transmission and is not birefringent nor optically active at either THz or NIR frequencies. The transmission has been measured by the VNA, see Fig. A.1. It is only attached using vacuum, and so it must be held in place during pump down. The back window is glass with a NIR anti-reflection coating. It is clamped in place using plastic straps. With all of these o-ring seals and large optical windows in the thermal radiation shield, the temperature can only reach about 14.5 K. The o-rings of the feed-throughs and the windows need to be checked and re-greased every few months.

The current vacuum shroud and heat shield are designed to allow for THz-NIR beam alignment using an indium-tin oxide (ITO)-coated glass slide as a dichroic beam com-

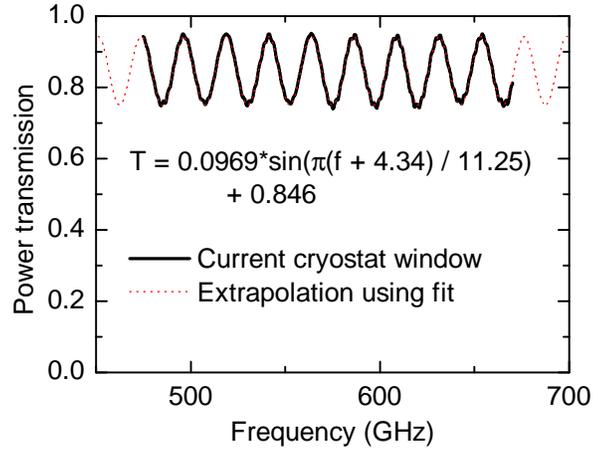


Figure A.1: The power transmission of the cryostat window as measured by the vector reflectometer. The fit only includes one band from the VNA, a more careful measurement may be required for frequencies too far away.

biner. The ITO slide is 50 mm by 75 mm with a sheet resistance of about $10 \Omega/\square$ and is mounted to a tip-tilt stage for independent NIR and THz alignment.

Cool down procedure

1. While holding the suprasil on the cryostat, rough pump the cryostat. Once the window is fixed, pump the cryostat to 10^{-6} torr or so for at least 2 hours, overnight is best.
2. Turn on building cooling water to the compressor.
3. Turn on compressor and wait until the the sample is cold, approximately 1.5 hours. Close pump valve when thermometer goes below 77 K. This protects turbo pump from being contaminated with rough pump oil, at least theoretically.
4. Once stabilized at low temperature, the temperature controller can hold the tem-

perature anywhere up to 150 K or so.

5. When finished or maintaining temperatures above 60 K, it is important open the pump valve to a rough pump because vacuum leaks will bring the pressure up to about 10^{-2} torr when the air boils off.

The cryostat will generally stay below 17 K for at least a day. It is possible with more vacuum-tight shrouds for the cryostat to remain cold for two or three days. Manipulating the translation stage makes the leaks more severe, but the cryostat will hold below 20 K.

Maintenance

If the temperature seems to be creeping up more quickly than usual, regreasing or replacing the o-rings on the vacuum shroud head should make it more stable, though the feed-throughs will be more slippery and prone to being sucked into unintentional contact with the cold parts inside. Every six months, the helium pressure in the compressor needs to be checked, and the manual has instructions for filling it and the parts for doing this should be in the cabinet nearby. If the cryostat needs to be moved, there are special wrenches to remove the high-pressure lines.

A.2 Sample imaging camera

A major quality-of-life improvement the experiment is the sample imaging camera, introduced to us by Kyle Seyler of the X. Xu Lab at the University of Washington, Seattle.

This camera allows for simple sample-NIR alignment as well as convenient sample monitoring during experiments. For excitation sources with long focal lengths, the basic camera schematic is shown in Fig. A.2. For short focal length experiments, the excitation lens can be used as Lens B, but this arrangement is less common for large-area quantum well samples. The optics are tucked away in between and behind the important THz and NIR optics, so that the camera can be used without needing realignment or disrupting a measurement. Lens A, the “eyepiece” lens, is mounted on a lens tube to the camera to block out unwanted light and so the camera is fixed to be at the focus of Lens A, f_A . Lens B, the “objective” lens, is mounted so the sample is at its focus, f_B , when reflected off a 50:50 pellicle beam splitter (BS) (1” diameter). There is deliberately nothing at the focus of the transmission path to avoid double-imaging. A scale bar of known size could feasibly be mounted at the focus, if more quantitative imaging was desired. The ratio of the focal lengths f_B/f_A is the magnification. The actual sizes of f_A and f_B determine the amount of light collected. For broad illumination, an bright lamp (likely halogen) is coupled into the beam path between Lenses A and B using a 2” diameter 50:50 pellicle BS. The indium-tin oxide beam combiner (ITO BC) overlays some small interference fringes, but they generally aren’t a problem.

Alignment for long excitation focal lengths starts by fixing the location of the camera, then Lens A, then the lamp and illumination pellicle, then Lens B and the imaging pellicle. To start, mount Lens A on a lens tube to the camera that is approximately the focal length of the lens. An adjustable lens tube is recommended for peaking up

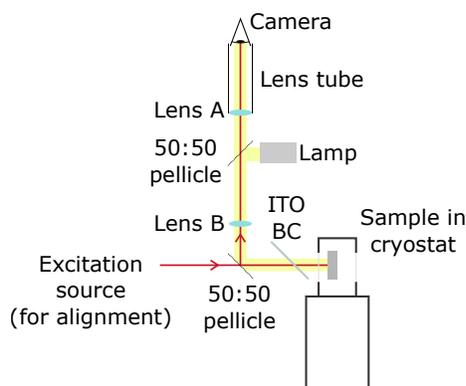


Figure A.2: The optical setup for the sample imaging camera.

the image focusing later. Then, align Lens B to be collinear with the lens tube. Try to mount Lens B close to where it will be in the final setup, but this is unnecessary. Turn the camera on. Using a beam block or a magnetic ruler with some lines or words written on it, find the location in front of Lens B that has a vaguely sharp image on the camera (do not use a pellicle BS at this point, just keep the optical path straight). A business card taped to a beam block works really well for this. Peak up the locations of both lenses to get the sharpest image on the camera. Use the lamp to illuminate the beam block from an angle if more light is needed (it usually is). Once the image is sharp, lock down the position of the camera and Lens A as much as possible, place posts to defend from accidental bumps, and measure the distance between Lens B and the beam block. Now, mount the illumination pellicle on a tip-tilt stage between Lenses A and B, and couple the lamp into the beam path. Peak up the brightness of the image on the camera by tweaking the pellicle and translating the lamp (and maybe the pellicle) around. Once

the image is good, lock both down. At this point, the camera is at the focus of Lens A, and the axis of the lens tube and the lamp are collinear. Now, using the distance measured between Lens B and the beam block, mount Lens B and the imaging pellicle so that the sample is that distance from Lens B. Be careful to ensure that the excitation source (the Ti:sapphire laser or the differential transmission source) are not blocked by the pellicle, and put the imaging pellicle on a magnetic clip or flip mount so it can be removed reliably. With a sample in place, peak up the location and angle of Lens B and the imaging pellicle to get the sharpest image on the camera. Only then should the alignment of the lamp or the position of Lens A be tweaked to ensure the highest-quality image.

This setup makes sample-NIR alignment extremely easy. The Ti:sapphire laser at low power and 50% duty cycle should be readily apparent on the sample now. The FEL alignment laser can technically be seen, but its spot is usually less than perfect. Both the tungsten-halogen lamp and the LED (from the differential transmission experiments) can be seen with this camera, even with a 30 μm spot size.

A.3 M² SolsTiS Ti:sapphire laser

The M² SolsTiS Ti:sapphire laser is a continuous-wave NIR laser pumped by a 7 W, 532 nm Sprout laser from Lighthouse Photonics. It outputs more than 1 W of power across its prodigious tuning range of 700–1000 nm, and it is a single mode laser with a linewidth of less than 5 MHz. The nicest feature, however, is that it cannot be opened by

the user. Occasionally the pump alignment needs to be peaked up, but otherwise there is no maintenance that even can be done to the laser itself beyond cleaning the windows (which hopefully never has to happen).

The laser is controlled via ethernet and a local network IP address. The laser is tuned using software. The birefringent filter provide the coarsest tuning. It works using a look up table to associate wavelength and filter angle, and can get to within 0.5 nm of the desired wavelength. The etalon yields slightly finer tuning. It can move the laser about 0.5 nm total in steps of about 0.05 nm. The laser can technically lock itself using the etalon, but this mode is prone to mode hopping. Finally, one of the resonator mirrors can be tuned with two settings, coarse and fine. These can theoretically move the laser wavelength about 0.5 pm and 0.005 pm, respectively, but we don't have a detector precise enough to actually measure that kind of tuning. This level of tuning should give mode-hop free tuning.

The software is also capable of diagnosing certain electrical problems within the laser, but only with M² support close at hand.

For safety and the elimination of detectable stray light, both laser heads and much of the optics are kept in a laser housing made of anodized aluminum sheets and 80-20 posts. There is a hole cut into the top sheets covered by square of laser curtain material to allow easy but safe access to the half-wave plate attenuator.

Alignment

About once a month, the pump alignment should be peaked up to ensure that it does not drift too far. There is one mirror accessible externally with three knobs. The diagonal knobs are tip-tilt, and only they should be touched. The middle knob, the point of the “L,” should never be touched. The alignment is extremely sensitive, so be careful when peaking it up. To measure the laser power easily, there is a flip mirror about two inches in front of the laser head. It redirects the laser into a power meter head with a 1000:1 attenuator on it, and the power meter readout is kept on top of the SPEX. It is highly recommended never to move this power meter as alignment of the attenuator is not trivial.

Maintenance

The only other maintenance that needs to be done on this laser and its pump laser is to change the fluid in the chillers every six months. The SolsTiS chiller is a dedicated unit. The Sprout chiller is contained in the power supply. The plumbing needed to drain and refill the chillers are found on top of the Sprout power supply. Both chillers use distilled water with 10–20% isopropanol? Distilled (*not* de-ionized) water can be bought at the grocery store if it’s not available in the chemistry storeroom.

A.3.1 Polarizer attenuator

The Ti:sapphire laser outputs as much as 1.8 W continuous-wave power. That much power is certainly too much for most experiments, but it is also too much for some optics. Any optic with even a small amount of absorption will heat up wherever the beam contacts it and expand. This expansion is called thermal lensing, and it causes Gaussian beams to become donut-shaped. One might think that expanding the beam diameter would limit thermal lensing, but the increased surface area of the thermal lens roughly cancels out the decrease in thermal expansion. The easiest way to reduce thermal lensing is to attenuate the beam using optics designed for high powers. In order to keep the power controllable, the polarized Ti:sapphire light is sent through a half-wave plate and then through a high-power polarizer. The component of polarization that passes through the polarizer can be controlled by rotating the half-wave plate. Be aware that rotating polarizers causes slight lateral beam shifts based on the prism spacing, but rotating a half-wave plate does not. The unwanted power is dumped into a razor blade stack.

A.4 PhotonControl spectrometer

The PhotonControl SPM-001-N spectrometer is a compact, fiber-coupled spectrometer capable sensing light from 680-920 nm. It is used to quickly measure the wavelength of the Ti:sapphire laser. The manual claims the resolution is 0.24 nm, but experience reveals that the Gaussian fit algorithm can tell when the laser moves by 1 1/cm (about

0.06 nm at 800 nm). The M^2 Ti:sapphire laser can theoretically be sent to any wavelength between 700 and 1000 nm, but it only works using by rotating a birefringent filter based on a lookup table. The actual wavelength can differ by as much as 0.5 nm based on etalon and resonator settings. A very small amount of laser power is picked off and sent to the PhotonControl for precision laser tuning. By “very small amount of laser power,” a pellicle picks off part of the beam and sends it at the side of a fiber coupler. The diffuse reflection off the pellicle is enough to measure the laser wavelength.

Software

The PhotonControl spectrometer is connected to the computer via USB, and it is controlled with a proprietary LabVIEW-esque interface that we cannot seem to talk to with external software.

A.5 Acousto-optic modulator

An IntraAction AOM-602N acousto-optic modulator (AOM) is used to as a fast electronic shutter for the NIR beam. The IntraAction Model ME signal processor drives a piezo in the AOM with an RF signal, creating a density Bragg grating that is capable of diffracting approximately 65% of the input power into a first-order diffracted beam, also sometimes called a sideband. The first-order beam has a contrast ratio ($P_{\text{on}}/P_{\text{off}}$) of at most 3000, which is enough for the SPEX but often not enough for imaging spectrometers like the CCD and the ICCD. The driver is controlled by a TTL pulse. The switching

time depends on the beam size, but it is approximately 100 ns.

Alignment

1. Ensure that the AOM is level and the aperture is 5" above the table. For simpler alignment peaking up in the future, it is useful to mount the AOM on a transverse horizontal translation stage and a z-axis rotation stage.
2. Without driving the AOM, send the laser straight through the AOM and peak up the power. Ensure the laser is 5" above the table before and after the AOM.
3. Apply a 50% duty cycle TTL signal to the AOM and turn the carrier level to 10.
4. Rotate and translate the AOM around until a sideband appears. The AOM deflects the sideband by about 2° . Depending on the direction of rotation, this could be to the left or the right of the zeroth-order beam.
5. Block all sidebands (I've seen as many as three total sidebands at any given time) except the first-order one. Measure the power and peak up the laser power going into the sideband the rotation and translation stages. It should be possible to get about 65% total power into the first order beam (legends claim some AOMs are capable of $> 85\%$ efficiency).
6. Be sure to block the zeroth-order beam using a razor stack or other absorber. The higher-order beams can be blocked with an iris if necessary.

A.6 Mechanical shutter

A Uniblitz shutter driven by a Uniblitz model SD-10 shutter drive timer controls the output of the laser out of the laser enclosure. The contrast ratio of the AOM alone is not high enough to block the laser given the duty cycle of the FEL. With a duty cycle of 10^{-6} and a contrast ratio of 10^3 , approximately one thousand times more laser light hits the sample, possibly causing photoluminescence, and enters the CCD, possibly damaging the CCD or at least raising the noise floor substantially, when the AOM is “off” than when it is “on.” Although the shutter has an opening time of about 8 ms, it completely blocks the laser when closed, bringing that factor of a thousand to less than ten. The shutter driver is controlled with a TTL pulse.

A.7 Absorption LED

Driven absorption measurements require fast pulses of broadband light. A Thorlabs LED on a metal-core printed circuit board can be driven with the HP 214B pulse generator to make pulses as fast as 25 ns with rise and fall times of about 10 ns. The pulse generator can drive a $50\ \Omega$ load up to 10 V. Above 10 V, it can only drive “high Z ” loads, but up to 200 V. The LED impedance drops very quickly at about 1.5 V, so the pulse generator acts more as a current source above this voltage.

A.8 SPEX monochromator

Any light within about 10 nm of the laser has to be detected using the SPEX 1403 0.85 m double monochromator. The SPEX has two 1800 line/mm gratings, making its efficiency very low, about 10%, but its rejection extremely high. The detector won't see any laser power at all when the SPEX is about 10 wavenumbers away (this number has not been tested). The entrance slits are almost always set to 400 μm , giving a resolution of 1 cm^{-1} . Incident light must be horizontally polarized or the efficiency will be much lower but not zero, so polarization-sensitive measurements need external polarizers.

Besides measuring the low-order sidebands and the laser itself, the SPEX is extremely useful for ensuring good alignment between the NIR and THz spots for sideband measurements. The ITO beam combiner can move the THz spot around easily without affecting the NIR spot or alignment. First, peak up the power in the second order sideband, then move to a higher order one, the sixth or the eighth, and peak up again. This alignment process needs to be done first thing in the morning before any experiment, and may need to be done several times depending on the day's experiments. If more careful alignment is necessary, check with the CCD as well. The feedback with the CCD is much slower, though.

Alignment

1. Find the white mirror cover, a 6" diameter metal plate with five small holes, which should be under the SPEX whenever not in use. This plate isn't necessary, but it

will make alignment much easier.

2. Ensure the laser beam is collimated going towards the SPEX, but block it from entering the SPEX. Put a lens on a three-axis stage with a focal length to match the $f/7.8$ of the SPEX.
3. Turn the detector off, and set the SPEX wavenumber to far, far away from the alignment laser line.
4. Open the hatches on top of the SPEX and put the mirror cover over the first mirror.
5. Unblock the laser and send through the 400 μm slit, don't try too hard.
6. Adjust the beam to be centered horizontally and vertically on the slits and on the mirror cover.
7. Place the focusing lens at the appropriate distance from the slits, and repeat the alignment, centering the beam on the slits and the mirror cover.
8. For best alignment, which may not be necessary, choke the slits down to the calculated spot size and repeat the alignment.
9. With the detector blocked, or even removed, set the SPEX to pass the laser and ensure that the light is hitting the detector slits. Hope that it is.
10. Remove the mirror cover and put back under the SPEX, and put everything back together.

11. Using either the laser (attenuate by $\lesssim 60$ dB) or even photoluminescence/sidebands, peak up the signal on the detector by translating the last lens or the last mirrors.
12. Set up irises so that you don't have to do this process again any time soon.

Software

The SPEX gratings can only be controlled with a computer. Being so old that GPIB was new, the SPEX is controlled by a SPEX MiniStepDriver which itself plugs into another box which can then connect to the computer via GPIB. These are left on all the time. Occasionally they must be power-cycled, or the power goes off. They must be turned on in the correct order, so if they don't work the first time, turn them off and turn them on in the other order.

To minimize backlash when rotating the gratings, the SPEX is designed to always step down in wavenumber when moving. If a higher wavenumber is needed, the SPEX will go to twenty cm^{-1} above the desired location and then rotate back. The physical readout on the side of the SPEX is higher than the actual wavenumber setting of the optics path by about 4 cm^{-1} , so be aware if things look wrong.

There are two sets of software to control the SPEX. Ben Zaks modified the Optilab code, Vinh Nguyen's LabVIEW code for the VNA, to operate it and the oscilloscope. This software has been replaced with Python code in SherwinGroup/SPEXSidebands-fuzzy. They both work in the same way, but the Python code is much more advanced. This software also controls reading from the oscilloscope.

A.9 Hamamatsu photomultiplier tube

A Hamamatsu R7400U-20 photomultiplier tube (PMT) detector was mounted on the SPEX. The response time is very fast, approximately 1 ns, and the noise-free gain is large, allowing for fast single photon detection. The efficiency of the PMT is low, around 10%, and the response dies off dramatically above 830 nm. The full response curve is taped to the detector. A voltage of -700 to -1000 V is applied using an SRS high voltage source. The gain ratio between 1 kV and 700 V is almost exactly 10. The output current is sent to an SRS 445 preamp set to 50Ω sent through two channels of 5 gain. Single photon peaks are approximately 50 mV tall.

A.10 Acton spectrometer

For broadband optical measurements, an Acton SpectraPro 2750 spectrometer is used. The total focal length is 275 mm, and there are two input ports and two exit ports. Currently, only the side entrance slit and the end exit port are used. The mirror that flips in for the side exit port does not block the end port entirely, so detectors mounted there could be damaged by errant bright light. There are currently three gratings mounted in the spectrometer, the bandwidth and resolution are approximations:

Grating no.	Groove density (g/mm)	Blaze (nm)	Bandwidth (nm)	Resolution (GHz)
1.	1800	500	13	8
2.	1200	1000	24	16
3.	600	1000	53	30

Table A.1: Gratings currently installed in the Acton SP-2750. Resolution is approximated by the spacing of three pixels in the data array, the approximate linewidth of a low-order sideband (which itself should have a linewidth of approximately 1 GHz, far smaller than the resolution of the spectrometer).

All of the gratings generally prefer vertically polarized light, the efficiency is higher and flatter throughout the center wavelength range. Acton/Princeton Instruments have a website with this information for lots of gratings.

The grating turret inside the spectrometer has two “stable” positions, one that has been calibrated, and one that is approximately 1 nm off. When measuring sidebands, it is generally easy to tell when the spectrometer has “jumped” as long as the laser wavelength is well known externally. It is less obvious when measuring PL or absorption. This jumping seems to be entirely random, but possibly more common at certain center wavelengths. Getting the spectrometer to jump back to the calibrated position requires shifting it around a half a nanometer or so at a time until it is where it says it is based on the measurement of a known wavelength. This bug is very, very frustrating.

Software

Software to control the Acton spectrometer in SherwinGroup/InstrumentLibrary-yolo.
The CCD camera software more directly controls the spectrometer.

A.11 Andor electron-multiplying CCD camera

The Andor NewtonEM electron-multiplying (EM) CCD camera is the primary HSG detector because it is very low-noise and can measure many sidebands in parallel. The primary advantages of the camera are the Peltier cooling down to -90° C requiring only minor water cooling, an EM amplifier, and hardware binning. Cooling to -90° C allows for 5 minute integration times with minimal dark current noise. The EM amplifier allows for low light signal, signal on the order of single photons, to be amplified “noise-free” above the readout noise. For sideband measurements, readout noise is the limiting factor of the detector. Hardware binning allows for binning before amplification or readout, so small signals can be summed before the multiplication and readout. Current limits of detection are about 0.5 sideband photons per cavity dump pulse, or about 0.5 photons per second when integrating for twenty minutes (four 5 minute images). Achieving this sensitivity requires 5 minute integration times, the chip to be at -90° C, and a gain setting of 110, which corresponds to real gain of about 6-7 (it is extremely sensitive to temperature). In addition, the 1380D ceiling lights, the hallway lights, and the FEL vault lights must be off, and the transport system windows in room 1380C must be covered in aluminum foil. It is likely that a shutter before the spectrometer blackout box triggered to open

only during the FEL pulse would relax these requirements.

The background of the CCD is complicated. Increasing the EM gain It is not truly noise free, as increasing the gain increases the amplitude of the background noise. The cause of this background is not well understood, as it is most apparent when binning a large number of pixels and integrating for a substantial time. It is likely from individual electrons of dark current summed over tens of pixels. Unfortunately, this background noise is substantially larger, though usually still small, on the left side of the chip, which is the shorter-wavelength side. It is as if the bottom left side, which does not seem to be the readout corner, is warmer than the right three quarters of the chip. It is unlikely that it is actually warmer, but the noise behaves similar to dark current. It is currently unknown what the longest integration time is to maximize the signal-to-noise ratio when the camera is at -90° C.

Optical edge filters

Safe and effective use of the CCD for HSG measurements requires decent long pass filters and very good short pass filters. To measure positive sidebands, the long pass filter blocks amplified spontaneous emission from the laser and the short pass filter blocks the laser from the CCD. It is difficult to say with certainty, but the short pass filter must be able to attenuate the laser by at least 60 dB, likely more 70 dB to suppress laser scatter in the spectrometer. In general, Semrock makes the best filters. Rotating filters away from normal decreases the transition edge, and good filters can be shifted almost 100 nm, although at that point the filter mount physically blocks the light more than the

filter fails. A more reasonable tuning range is 40 nm. There is currently one very good long pass filter with an edge of 808 nm. It can be angle-tuned down to 765 nm easily. There are two very good short pass filters, 847 nm and 785 nm. These both have very sharp transitions from transmission to reflection. The 847 nm short pass filter is good enough to allow CCD measurement of sidebands just 6 nm from the laser line. The 785 nm short pass filter is sharp enough to allow measurement of sidebands 4 nm from the laser line. When changing filters, never pull by the optic holder, only by the post. The optic slips easily and spins if pulled by the frame, often resulting in huge fingerprints on either face of the filter. Cleaning Semrock filters is annoying because of their shape.

Alignment

1. Block the spectrometer.
2. Send the laser with 50% duty cycle through the laser polarizer and through the short pass filter. Place a power meter behind the filter and note the power, usually tens of milliwatts.
3. With the filter reflecting light towards the edge of the table (not at the laser box wall), rotate the filter so that only some of the light is transmitting through.
4. Turn off the room lights and have a flash light (or cell phone) ready.
5. Rotate the filter with the rotation stage until about 100 μW are transmitting through the filter.

6. Using the fine-tuning screw of the polarizer, minimize the signal transmitting through the attenuator. The filter edge for horizontal (p-polarized) light moves much faster as a function of rotation than vertical (s-polarized) light. Any horizontal component of the laser light will then not be removed by the filter, so the polarizer must be tuned along with the filter.
7. Iterate between rotating the filter and the polarizer, minimize the light transmitting through the filter. Be careful not to rotate the filter farther than necessary, as too much rotation will start to block sidebands you may want to measure.

At this point, it may be useful to ensure quality alignment to the spectrometer and CCD. If the spectrometer lens never moves, the two irises in the spectrometer path should get the alignment very close. The best data, however, comes when the laser spot is aligned to the CCD that day. Be very careful doing this, it is better to assume the CCD is very fragile. I don't know how much laser power it can handle, so err on the safe side and block the laser when in doubt.

1. Put one of the ND20A filters, called "blue" or "white," in the laser path. In front of the polarizer is a good place.
2. Set the center wavelength of the spectrometer to the laser wavelength, and open the spectrometer slits a lot, as large as 1 mm.
3. With the laser blocked and AOM at zero, take a quick (say 0.2 s, no gain, no binning, fastest readout) picture with the CCD.

4. Unblock the laser, keep the AOM at zero, and take another picture. No laser spot should be visible on the screen. If there is light, check to make sure everything is in place and that the slits are open enough.
5. Set the CCD shutter to “always open” and change CCD to alignment mode. Slowly increase the AOM setting until a laser spot is visible on the CCD.
6. Using the last two mirrors before the spectrometer, put the laser at the right place vertically. Traditionally, it has been around 210–250, but that may change.
7. For horizontal alignment, close the slits until the laser is just barely visible. Scan the laser left or right to find the other edge of the slits. Move the laser to approximately the center and repeat this process iteratively until the slits are about 50–65 μm wide and the laser is dead center.

The NIR laser alignment should now be rock solid. The usual culprits responsible for misaligning the laser to the spectrometer are changing/rotating the long or short pass filters. Rotating the half-wave plates shouldn't affect the alignment at all, and tilting the ITO beam combiner doesn't seem to, either.

Software

There are two methods to use this camera, a LabVIEW software development kit (SDK) and a C++ SDK. A set of Python wrappers using the ctypes library were written for the C++ SDK for easy inclusion into a PyQT GUI in SherwinGroup/EMCCD-psyhic. This

GUI handles all CCD-based experiments, including photoluminescence, absorbance, and sideband generation. It is extremely simple to use and is capable of running complex experiments rather simply, but the code itself is very complex. It is not recommended to edit the code without a very complete understanding of the software and PyQT in general.

A.12 Thomas-Keating energy meter

The Thomas-Keating (TK) energy meter is used to measure the absolute energy of an FEL pulse. A more detailed description can be found in Sect. 3.2.1. The main points to worry about when using the TK are not to focus the FEL directly onto it and to use AC coupling on the oscilloscope for accurate measurements. There is a lot of high-frequency noise on the TK output, averaging a few FEL shots is usually necessary. This thesis used the peak-to-peak of the output, but other sources have just used the positive or negative amplitude.

Software

There is software for automatically measuring and calibrating the integrating pyroelectric detector using the TK is found in `SherwinGroup/InstrumentLibrary-yolo`.

Appendix B

FEL transport details

Bringing the intense light from the FEL to the sample requires transmitting long-wavelength, therefore fast-diffracting, light from the vault to the optical table of the experiment, tens of meters and several walls away. The FEL transport system was developed to accomplish this task without significant attenuation and while maintaining a high quality spatial profile. The FEL transport system is a series of 6"-diameter tubes pumped down to about 200 mtorr to minimize absorption by atmospheric water. It features several output ports and computer-controlled retractable mirrors and lenses. The lenses are placed with a regular spacing throughout the transport system with a focal length of 3 m to maintain a collimated beam for tens of meters for wavelengths between about $60 \mu\text{m} - 1 \text{ mm}$. Depending on the frequency, some mirrors are curved and others are flat. The lenses are made of the THz-transparent plastic TPX and the mirrors are coated with a thick layer of gold to minimize losses due to optics. The particular transport system setup for the

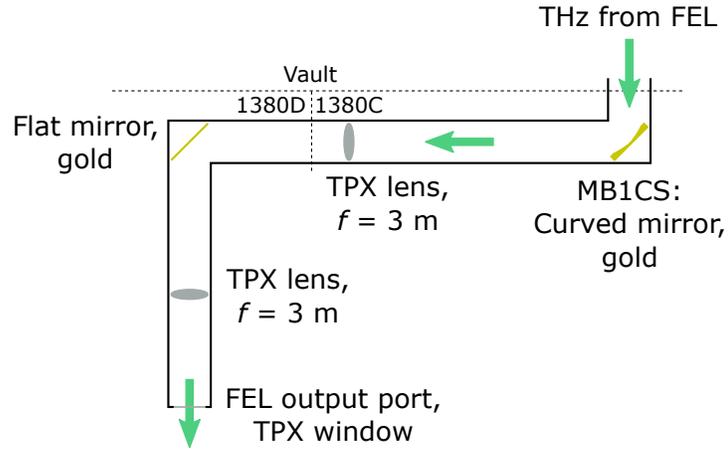


Figure B.1: The setup of the transport system for bringing the FEL to the optical table. The THz light is generated in the vault behind a very thick concrete wall. That light is brought through a transport system tube into room 1380C. The light reflects off of a curved mirror, named MB1 Curved South (MB1CS). The light is then focused/collimated by a TPX lens, brought through the wall into room 1380D, reflected off a flat mirror and focused/collimated again before exiting the transport system through an FEL output port, made of TPX plastic.

work in this dissertation is shown in Fig. B.1.

Appendix C

Cleanroom processing

The cleanroom on campus has a large collection of tools and well-maintained wet benches available for sample preparation and making/modifying the occasional optical component. This appendix contains several recipes and procedures for common processes and tools. The sample processing recipes are for making epitaxial growth membranes and performing epitaxial transfer. It is fairly straightforward to modify these recipes to fit the needs of a given process. For newer or more difficult processes, however, the cleanroom staff are extremely helpful and can often point you to another user who is experienced in the process. There are also procedures for depositing high-quality indium-tin oxide and for using the dicing saw.

C.1 Quantum well membrane process

For experiments where the highest terahertz fields are not necessary and where an optical etalon would be a problem, it is useful to etch a small window in the GaAs substrate and expose a thin membrane of epitaxially-grown quantum wells. The membrane is not very robust to high terahertz fields, but, with a thickness on the order of $1\ \mu\text{m}$, it does not form a Fabry-Perot cavity for NIR wavelengths.

Instructions

The etching steps of this process can take a long time. If making multiple samples, put them on different slides because some may etch faster than others. The Physics Department storeroom sells $1\times 3''$ glass slides, and the cleanroom has $2\times 3''$ glass slides. Glass slides are robust to all of the etches except HF. The glass will etch noticeably during the HF dip, but dip is not long enough to substantially weaken the slide.

- Spin acid-resistant photoresist (PR) like AZ4210 on growth side of the sample for protection, 3 krpm for 30 seconds, followed by a 60 second, 100° post exposure bake (even though no exposure). If worried, test that it's hard enough that it won't scratch off and bake longer if necessary.
- Spin the same acid-resistant PR on substrate side with the same recipe.
- Use a real mask, or tape four pieces of foil together to make a 1.5–2 mm square hole, a make-shift mask. Taping pieces is better than cutting a hole in aluminum

foil because the corners will be sharper. Time spent on this step saves much more time later. Smaller is better, but 1 mm is too small because of sloping etch walls.

- Expose the substrate side of the sample with the preferred mask. The rectangle needs to be at least 1.5 mm from the edge of the sample, and it's best if everything makes right angles.
- Do the post exposure bake and develop according to the PR instructions.
- Put one drop of PR on a glass slide, and lay the sample growth-side down on the drop. Bake for a while so that the sample is firmly attached to the glass slide.
- Drip some PR around the now-stuck on sample so that the PR covers everything but the hole in the middle that will be etched. Double check under microscope for pinholes in the PR, especially the corners and edges. Bake slowly for several minutes, until PR is dry and rigid. Bake too quickly and the PR will bubble and expose pinholes. Check under microscope again.
- This is a stopping point. The sample can be exposed to ambient lights and stored.
- Prepare to etch the sample. There are two options here. Use either the dilute piranha etch, 1:8:8 $\text{H}_2\text{SO}_4:\text{H}_2\text{O}_2:\text{H}_2\text{O}$, or Garrett Cole's epi transfer etch, 30:1 $\text{H}_2\text{O}_2:\text{NH}_4\text{OH}$. Dilute piranha is very fast, about 4 $\mu\text{m}/\text{min}$. It is not at all selective, though. As for any wet etching procedure, always have sample facing down while etching, this geometry keeps etch by-products from interfering with the etching process.

- Dektak at least twice while etching, especially if using a non-selective etch. The dilute piranha etch can undercut the PR layer substantially, so only drag the Dektak needle from high to low to prevent it catching on an overhang. Any post-Dektak measurement is a stopping point, so long as no AlGaAs with over 80% Al is exposed.
- If using dilute piranha to etch the substrate, then need to switch to a selective etch with about 50 μm remaining. Garrett Cole's etch, 30:1 $\text{H}_2\text{O}_2:\text{NH}_4\text{OH}$, works, as does citric acid and H_2O_2 . Check progress with a microscope.
- Do quick buffered HF dip to remove the high-aluminum-content etch stop layer. Again, check progress with a microscope. 70% AlGaAs can take about two minutes to etch through 200 nm, but 73.5% AlGaAs etches much faster in buffered HF. AlGaAs with less than 60% Al does not etch at all.
- Submerge entire slide in acetone, sample side up. Cover to prevent evaporation and wait several hours. Do not try to rush things by agitating the solution or pulling on the sample. The membrane is very fragile and will tear extremely easily

It may be safer for particularly thin epi-layers, such as the 10 nm GaAs QWs grown by Loren Pfeiffer which are only about 500 nm thick, to have the window region overhang the glass slide. The sample will be a little more fragile during etching, but removing the PR should be much easier. When removing the PR, it will dissolve quickly and evenly, hopefully eliminating any shearing that occurs when the PR is dissolved slowly between the sample and the glass slide.

C.2 Epitaxial transfer process

For experiments where very strong or even enhanced terahertz fields are useful and Fabry-Perot interference is acceptable, the epitaxial transfer process is extremely useful. This process was introduced to us by Garrett Cole of Crystalline Mirror Solutions.

Instructions

Before going into the clean room:

- Sketch the wafer geometry, including which direction the flat was. Lapping will remove all reference points, and the sample may break in two.
- Lap down all samples to 150 microns thick.
- Dice sapphire to appropriate size(s).

Once in the clean room:

- Clean 1165-Acetone-Isopropanol-DI water-dry N₂.
- Do an O₂ descum, 100 W, 1 minute to finish cleaning the sample.
- Take a full set of pictures of the growth side of the sample and record the size if any pieces broke off during lapping.
- Heat up on hot plate to 130° deg C to finish drying. Silicon makes a good wafer carrier for this process because thermal contact is occasionally important.
- Remove the sample and let cool to room temperature.

- Put large wax piece on the growth side of the sample, then put the whole wafer carrier on the hot plate, still at 130° C. The wax takes a few minutes to melt. If more is needed, remove the wafer to cool so that if the wax falls off the tweezers at the wrong time, it won't mess up the sample. Take this step slowly as removing black wax is difficult and very dirty.
- After melting fully, drop the temperature to 100° C in 5° C increments. Once the temperature gets to 100° C, cover with an aluminum boat and let anneal for an hour.
- Decrease temperature to 80° C in 5° C increments.
- Lay/balance the silicon carrier wafer or glass slide over the wax. If it is off to one side, do your best to right it while the wax is still soft. Needs to be > 1¹/₄" to use with vacuum evaporator.
- Drop the temperature to 50-60° C in 5° C increments, then remove from the hot plate and let cool completely. The samples should be glued to the wafer carrier with black wax with the lapped sides exposed
- Prepare 30-100:1 H₂O:NH₄OH and 30:1 H₂O₂:NH₄OH solutions (200:5 and 300:10 mL, respectively, works nicely). We call the latter Garrett Cole's etch.
- Etch in water-ammonium hydroxide solution for 30 seconds, then immediately place in peroxide-ammonium hydroxide solution.

- Etch until the substrate is removed. It will be obvious as the entire sample will be very smooth. This takes about 90 minutes for 150 microns. Occasionally a white scum builds up on the substrate, just swab it off in situ.
- While etching, clean a sapphire substrate on a solvent spinner if possible (500-1000 rpm, but substrate may be too small) following this order: (tergitol+swab)-DI-IPA-DI, Nomarsky microscopy inspect and pictures, repeat until very clean, no dust at all. Some dots may remain that (often) get annihilated in the descum. Then 200 W 300 mtorr O₂ descum the sapphire wafer, 2 minutes. Double check with Nomarsky microscopy to ensure cleanliness.
- Once peroxide-ammonium hydroxide etch is complete, take pictures, etch in buffered HF until the color is very uniform, no interference fringes, 3 minutes. Every 30-45 seconds, rinse in DI water. Then take pictures of the entire sample.
- After HF, check if any weird dust/filaments are there. If so, wash in AZ400K in a clean beaker, could possibly take 5 minutes. Be patient here, this dust prevents good Van der Waals contact.
- Add a very small drop of water to the exposed epi, then place the sapphire on top. It works best if the sapphire is appropriately sized. Be very careful that it doesn't slide off at first. Wick away any drops of water that are squeezed out with a clean wipe. Press firmly with the wooden end of a qtip, and you should see Fabry-Perots disappear. Wick away as much moisture as possible without touching the sample

with the wipe.

- IF IF IF there is an obvious bubble (of water, not air), it is maybe possible to squeeze the bubble out by covering with a Berkshire wipe and using a qtip like a rolling pin.
- Dont touch them after that, let the water evaporate. Takes about 2 hours when pulling rough vacuum, or a day or two in air.
- In aluminum boat, dissolve black wax in chloroform. May take a few fresh chloroform baths to get it all. Make sure to dispose of the used chloroform solution safely.
- Once no more black wax is visible, IPA-DI-chloroform-IPA-Ace-IPA-DI- gentle blow dry to clean

There are several improvements that could be made to this process. Using etching to define the transferred region would be very useful. Dicing introduces far too many microcracks at the edges to be useful, and cleaving is imprecise and can induce microcracks at the initiating corner. By defining a sample size with a photolithography mask, samples will be much more uniform and less likely to be severely damaged at the edges. Also, it may be possible to skip lapping entirely by simply etching the back surface faster. Also, to ensure an extremely smooth epitaxial bonding interface, more etch stop layers could be added in growth. A triple layered structure, 73.5% AlGaAs-GaAs-73.5% AlGaAs would take advantage of the extreme selectivity of HF.

C.3 Indium-tin oxide deposition

Indium-tin oxide (ITO) is a terrific material for terahertz-optical experiments. It is highly conducting below about 10 THz, but is practically lossless all the way up to 450 nm. It can be grown in the electron-beam vacuum deposition chamber for dielectrics in the cleanroom, named E-Beam 2. The ITO grown here should have a DC resistivity of about $400 \mu\Omega\text{cm}$ and a $\text{Im}(n) = k \sim 10^{-5}$ at 800 nm. Significant optical absorption should not exist until about 400 nm. Thanks to John Leonard from the GaN lasers group for teaching us this recipe.

Before signing up for time, prepare the sample and a test piece of undoped silicon wafer, about 1 cm square. Silicon has very high index contrast with ITO, $n_{\text{Si}} \sim 4$ versus $n_{\text{ITO}} \sim 1.7$, so it can be used to test the optical quality of the layer with the ellipsometer in the cleanroom, but be sure to ask for help with using that machine, analyzing the results can be complicated.

Depositing a high-quality film of ITO requires careful control of the wafer temperature, the O_2 pressure, and the deposition rate. All of these details are, again, courtesy of John Leonard, and meeting these requirements results in films that are extremely smooth (supposedly better than nanometer roughness, although that has not been measured by us), highly transparent, and very conducting. The best films grow when the sample is held at 290°C . Because of the complexities of ITO conduction, which are somehow related to the oxygen content of the material, the deposition chamber must be pumped down to about $6 \cdot 10^{-6}$ torr, then 30 sccm of O_2 is bled into the chamber, resulting in a

pressure of about $3 \cdot 10^{-4}$ torr. The ITO then needs to be deposited very slowly, 0.01 nm/s for the first 10 nm, then 0.05 nm/s until 50 total nm, then 0.1 nm/s until the film is complete. Once the film is deposited, the temperature must be allowed to drop slowly in the oxygenated atmosphere until it drops below the ITO glass transition temperature of about 140° C.

C.4 Dicing samples

The dicing saw run by the cleanroom staff in Engineering II is an extremely useful tool for making samples the exact size desired. Sapphire substrates, which have proven to be a very effective support material for GaAs epitaxial films, cannot be cleaved in rectangles. To cut 500 μm c-plane sapphire, we use 2.187-8C-54RU-3 diamond dicing blades. The width of the cuts are about 300 μm thick. We have also cut 1 mm sapphire wafers and 3 mm thick glass with this blade, but expect very colorful sparks and the blades to wear very quickly. The cleanroom staff can give specific advice for any given cut or material, including best practices to protect fragile processing. It is not recommended to dice samples post-epitaxial-transfer. The UV-release tape bonds stronger to the GaAs than the sapphire.

Bibliography

- [1] S. Arlt, U. Siegner, J. Kunde, F. Morier-Genoud, and U. Keller. Ultrafast dephasing of continuum transitions in GaAs and AlGaAs. *Technical Digest. Summaries of Papers Presented at the Quantum Electronics and Laser Science Conference*, 59(23):860–863, 1999.
- [2] Neil W. Ashcroft and N. David Mermin. *Solid State Physics*. Brooks/Cole, Belmont, CA, first edition, 1976.
- [3] Mary Lou P Bailey, Andrew T Pierce, Aaron J Simon, Devin T Edwards, Gerald J Ramian, Nikolay I Agladze, and Mark S Sherwin. Narrow-Band Water-Based Absorber With High Return Loss for Terahertz Spectroscopy. *IEEE Transactions on Terahertz Science and Technology*, 5(6):961–966, 2015.
- [4] Hunter B. Banks, Benjamin Zaks, Fan Yang, Shawn Mack, Arthur C. Gossard, Renbao Liu, and Mark S. Sherwin. Terahertz Electron-Hole Recollisions in GaAs Quantum Wells: Robustness to Scattering by Optical Phonons and Thermal Fluctuations. *Physical Review Letters*, 111(26):267402, dec 2013.
- [5] Robert W. Boyd. *Nonlinear Optics*. Academic Press, San Diego, CA, 1992.
- [6] S. G. Carter, V. Ciulin, Mark S. Sherwin, M. Hanson, A. Huntington, L. a. Coldren, and Arthur C. Gossard. Terahertz electro-optic wavelength conversion in GaAs quantum wells: Improved efficiency and room-temperature operation. *Applied Physics Letters*, 84(6):840, 2004.
- [7] P. Cavalie, J. Freeman, K. Maussang, E. Strupiechonski, G. Xu, R. Colombelli, L. Li, a. G. Davies, E. H. Linfield, J. Tignon, and S. S. Dhillon. High order sideband generation in terahertz quantum cascade lasers. *Applied Physics Letters*, 102(22):221101, 2013.
- [8] J. Cerne, J. Kono, Mark S. Sherwin, M. Sundaram, Arthur C. Gossard, and Ge Bauer. Terahertz Dynamics of Excitons in GaAs/AlGaAs Quantum Wells. *Physical review letters*, 77(6):1131–1134, aug 1996.

- [9] J. Cerne, A. G. Markelz, M. S. Sherwin, S. J. Allen, M. Sundaram, A. C. Gossard, P. C. Van Son, and D. Bimberg. Quenching of excitonic quantum-well photoluminescence by intense far-infrared radiation: Free-carrier heating. *Physical Review B*, 51(8):5253–5262, 1995.
- [10] E. Y. Chang, Yeong-Lin Lai, Y. S. Lee, and S. H. Chen. A GaAs/AlAs Wet Selective Etch Process for the Gate Recess of GaAs Power Metal-Semiconductor Field-Effect Transistors. *Journal of The Electrochemical Society*, 148(1):G4, 2001.
- [11] Ming-Che Chang and Qian Niu. Berry curvature, orbital moment, and effective quantum theory of electrons in electromagnetic fields. *Journal of Physics: Condensed Matter*, 20(19):193202, 2008.
- [12] Ming-Chang Chen, Christopher Mancuso, Carlos Hernández-García, Franklin Dolar, Ben Galloway, Dimitar Popmintchev, Pei-Chi Huang, Barry Walker, Luis Plaja, Agnieszka A Jaroń-Becker, Andreas Becker, Margaret M Murnane, Henry C Kapteyn, and Tenio Popmintchev. Generation of bright isolated attosecond soft X-ray pulses driven by multicycle midinfrared lasers. *Proceedings of the National Academy of Sciences of the United States of America*, 111(23):E2361–7, 2014.
- [13] NC Constantinou and BK Ridley. Electron energy relaxation via LO-phonon emission in free-standing GaAs wafers. *Journal of Physics: Condensed Matter*, 2(7465), 1990.
- [14] Paul B. Corkum. Plasma Perspective on Strong-Field Multiphoton Ionization. *Physical Review Letters*, 71(13):1994–1997, 1993.
- [15] Paul B. Corkum. Recollision physics. *Physics Today*, (March), 2011.
- [16] J. A. Crosse and Ren Bao Liu. Quantum-coherence-induced second plateau in high-sideband generation. *Physical Review B - Condensed Matter and Materials Physics*, 89(12):1–5, 2014.
- [17] John H. Davies. *The Physics of Low-dimensional Semiconductors: An Introduction*. Cambridge University Press, 1997.
- [18] J Feldmann, G Peter, EO Göbel, P. Dawson, K. Moore, C. T. Foxon, and R. J. Elliott. Dependence of Radiative Exciton Lifetimes in Quantum Wells. *Physical Review Letters*, 59(20):2337–2340, 1987.
- [19] M Ferray, Anne L’Hullier, XF Li, G. Mainfray, and C. Manus. Multiple-harmonic conversion of 1064 nm radiation in rare gases. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 21, 1988.
- [20] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich. Generation of optical harmonics. *Physical Review Letters*, 7(4):118–119, 1961.

- [21] D. Gammon, B. V. Shanabrook, and D. S. Katzer. Excitons, phonons, and interfaces in GaAs/AlAs quantum-well structures. *Physical Review Letters*, 67(12):1547–1550, 1991.
- [22] Shambhu Ghimire, Anthony D. DiChiara, Emily Sistrunk, Pierre Agostini, Louis F. DiMauro, and David A. Reis. Observation of high-order harmonic generation in a bulk crystal. *Nature Physics*, 7(2):138–141, 2010.
- [23] Dennis H. Goldstein. *Polarized Light*. CRC Press, Boca Raton, FL, third edition, 2011.
- [24] Ronald L Greene, Krishan K. Bajaj, and Dwight E. Phelps. Energy levels of Wannier excitons in GaAs/GaAlAs quantum-well structures. *Physical Review B*, 29(4):1807, 1984.
- [25] M. A. Herman, D. Bimberg, and J. Christen. Heterointerfaces in quantum wells and epitaxial growth processes: Evaluation by luminescence techniques. *Journal of Applied Physics*, 70(2), 1991.
- [26] FA Ilkov, JE Decker, and SL Chin. Ionization of atoms in the tunnelling regime with experimental evidence using Hg atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 25:4005–4020, 1992.
- [27] P. U. Jepsen, D. G. Cooke, and M. Koch. Terahertz spectroscopy and imaging - Modern techniques and applications. *Laser and Photonics Reviews*, 5(1):124–166, 2011.
- [28] C. J. Joachain, N. J. Kylstra, and R. M. Potvliege. *Atoms in Intense Laser Fields*. Cambridge University Press, Cambridge, UK, 2012.
- [29] R A Kaindl, M A Carnahan, D. Hägele, R. Lövenich, and Daniel S. Chemla. Ultrafast terahertz probes of transient conducting and insulating phases in an electron-hole gas. *Nature*, 423:734–738, 2003.
- [30] J. P. Kaminski, J. S. Spector, C. L. Felix, D. P. Enyeart, D. T. White, and G. Ramian. Far-infrared cavity dump coupling of the UC Santa Barbara free-electron laser. *Applied Physics Letters*, 57(26):2770–2772, 1990.
- [31] D. S. Katzer, D. Gammon, and B. V. Shanabrook. Modification of the microroughness of molecular-beam epitaxially grown GaAs/AlAs interfaces through changes in the growth temperature. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 10(2):800, mar 1991.
- [32] LV Keldysh. Ionization in the field of a strong electromagnetic wave. *Sov. Phys. JETP*, 20(5):1307, 1965.

- [33] Dai Sik Kim, Jagdeep Shah, J. E. Cunningham, T. C. Damen, Wilfried Schäfer, Michael Hartmann, and Stefan Schmitt-Rink. Giant excitonic resonance in time-resolved four-wave mixing in quantum wells. *Physical Review Letters*, 68(7):1006–1009, 1992.
- [34] J. Kono, M.Y Su, J. Cerne, Mark S. Sherwin, S. J. Allen, Jr., T. Inoshita, T Noda, and H. Sakaki. Terahertz dynamics in confined magnetoexcitons. *Physica B: Condensed Matter*, 249-251:527–533, jun 1998.
- [35] R. F. Kopf, E. F. Schubert, T. D. Harris, and R. S. Becker. Photoluminescence of GaAs quantum wells grown by molecular beam epitaxy with growth interruptions. *Applied Physics Letters*, 58(6):631, 1991.
- [36] JL Krause, Kenneth Schafer, and KC Kulander. High-order harmonic generation from atoms and ions in the high intensity regime. *Physical Review Letters*, 68(24):3535–3538, 1992.
- [37] G. Lambert, T. Hara, D. Garzella, T. Tanikawa, M. Labat, B. Carre, H. Kitamura, T. Shintake, M. Bougeard, S. Inoue, Y. Tanaka, P. Salieres, H. Merdji, O. Chubar, O. Gobert, K. Tahara, and M.-E. Couprie. Injection of harmonics generated in gas in a free-electron laser providing intense and coherent extreme-ultraviolet light. *Nature Physics*, 4(4):296–300, 2008.
- [38] F. Langer, M. Hohenleutner, C. P. Schmid, C. Poellmann, P. Nagler, T. Korn, C. Schüller, M. S. Sherwin, U. Huttner, J. T. Steiner, S. W. Koch, M. Kira, and R. Huber. Lightwave-driven quasiparticle collisions on a subcycle timescale. *Nature*, 533(7602):225–229, 2016.
- [39] M. Lewenstein, Ph. Balcou, Misha Yu Ivanov, Anne L’Huillier, and Paul B. Corkum. Theory of high-harmonic generation by low-frequency laser fields. *Physical Review A*, 49(3):2117, 1994.
- [40] Anne L’Huillier and Ph Balcou. High-order harmonic generation in rare gases with a 1-ps 1053-nm laser. *Physical Review Letters*, 70(6):774–777, 1993.
- [41] Renbao Liu and Bang-Fen Zhu. High-order THz-sideband generation in semiconductors. *AIP Conference Proceedings*, 893:1455–1456, 2007.
- [42] JM Luttinger and W. Kohn. Motion of Electron and Holes in Perturbed Periodic Fields. *Physical Review*, 173(4):869, 1955.
- [43] JJ Macklin, JD Kmetec, and CL Gordon III. High-order harmonic generation using intense femtosecond pulses. *Physical Review Letters*, 70(6):766–769, mar 1993.
- [44] S.W. McKnight, K.P. Stewart, H. D. Drew, and K. Moorjani. Wavelength-independent anti-interference coating for the far-infrared. *Infrared Physics*, 27(5):327–333, sep 1987.

- [45] Eun-A Moon, Jong-Lam Lee, and Hyung Mo Yoo. Selective wet etching of GaAs on $[\text{Al}_{1-x}\text{Ga}_x]_{1-x}\text{As}$ for AlGaAs/InGaAs/AlGaAs pseudomorphic high electron mobility transistor. *Journal of Applied Physics*, 84(7):3933, 1998.
- [46] KB Nordstrom, K Johnsen, S. J. Allen, Jr., AP Jauho, B. Birnir, J. Kono, T. Noda, Hidefumi Akiyama, and H. Sakaki. Excitonic dynamical Franz-Keldysh effect. *Physical Review Letters*, 81(2):457–460, 1998.
- [47] Andrew T Pierce. *Development of a Frequency Domain THz Vector Reflectometer*. Bachelor honor's thesis, University of California, Santa Barbara, 2015.
- [48] Gerald Ramian. The new UCSB free-electron lasers. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 318(1-3):225–229, 1992.
- [49] S. Rudin and T. Reinecke. Exciton-acoustic-phonon linewidths in GaAs bulk and quantum wells. *Physical Review B*, 65(12):9–12, 2002.
- [50] K Miyazaki Sakai and H. High-order harmonic generation in rare gases with intense subpicosecond dye laser pulses. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 25(3):L83–L83, 1992.
- [51] O. Schubert, M. Hohenleutner, F. Langer, B. Urbanek, C. Lange, U. Huttner, D. Golde, T. Meier, M. Kira, S. W. Koch, and R. Huber. Sub-cycle control of terahertz high-harmonic generation by dynamical Bloch oscillations. *Nature Photonics*, 8(2):119–123, 2014.
- [52] S Takahashi, L C Brunel, D T Edwards, J van Tol, G Ramian, S Han, and M S Sherwin. Pulsed electron paramagnetic resonance spectroscopy powered by a free-electron laser. *Nature*, 489(7416):409–413, 2012.
- [53] Susumu Takahashi, Gerald Ramian, and Mark S. Sherwin. Cavity dumping of an injection-locked free-electron laser. *Applied Physics Letters*, 95(23):2007–2010, 2009.
- [54] K L Teo, J S Colton, P Y Yu, E R Weber, and M F Li. An analysis of temperature dependent photoluminescence line shapes in InGaN An analysis of temperature dependent photoluminescence line shapes in InGaN. *Applied Physics Letters*, 1697(1998):10–13, 2012.
- [55] G Vampa, T J Hammond, N Thiré, B E Schmidt, F Légaré, C R McDonald, T Brabec, and P B Corkum. Linking high harmonics from gases and solids. *Nature*, 522(7557):462–464, 2015.
- [56] G. Vampa, T. J. Hammond, N. Thiré, B. E. Schmidt, F. Légaré, C. R. McDonald, T. Brabec, D. D. Klug, and P. B. Corkum. All-Optical Reconstruction of Crystal Band Structure. *Physical Review Letters*, 115(19):1–5, 2015.

- [57] G. Vampa, C. R. McDonald, G. Orlando, D. D. Klug, P. B. Corkum, and T. Brabec. Theoretical Analysis of High-Harmonic Generation in Solids. *Physical Review Letters*, 113(7):073901, 2014.
- [58] Carlo Vicario, Balazs Monoszlai, and Christoph P. Hauri. GV/ m single-cycle terahertz fields from a laser-driven large-size partitioned organic crystal. *Physical Review Letters*, 112(21):1–5, 2014.
- [59] N. Q. Vinh, S. James Allen, and Kevin W. Plaxco. Dielectric spectroscopy of proteins as a quantitative experimental test of computational models of their low-frequency harmonic motions. *Journal of the American Chemical Society*, 133(23):8942–8947, 2011.
- [60] Martin Wagner, Harald Schneider, Dominik Stehr, Stephan Winnerl, Aaron M. Andrews, Stephan Schartner, Gottfried Strasser, and Manfred Helm. Observation of the intraexciton Autler-Townes effect in GaAs/AlGaAs semiconductor quantum wells. *Physical Review Letters*, 105(October):1–4, 2010.
- [61] Martin Wagner, Harald Schneider, Stephan Winnerl, Manfred Helm, T. Roch, Aaron M. Andrews, S. Schartner, and Gottfried Strasser. Resonant enhancement of second order sideband generation for intraexcitonic transitions in GaAs/AlGaAs multiple quantum wells. *Applied Physics Letters*, 94(24):241105, 2009.
- [62] Di Xiao, Ming Che Chang, and Qian Niu. Berry phase effects on electronic properties. *Reviews of Modern Physics*, 82(3):1959–2007, 2010.
- [63] Xiao Tao Xie, Bang Fen Zhu, and Ren Bao Liu. Effects of excitation frequency on high-order terahertz sideband generation in semiconductors. *New Journal of Physics*, 15, 2013.
- [64] E. Yablonovitch, D. M. Hwang, T. J. Gmitter, L. T. Florez, and J. P. Harbison. Van der Waals bonding of GaAs epitaxial liftoff films onto arbitrary substrates. *Applied Physics Letters*, 56(1990):2419–2421, 1990.
- [65] Eli Yablonovitch, T. Gmitter, J. P. Harbison, and R. Bhat. Extreme selectivity in the lift-off of epitaxial GaAs films. *Applied Physics Letters*, 51(1987):2222–2224, 1987.
- [66] Jie-Yun Yan. Theory of excitonic high-order sideband generation in semiconductors under a strong terahertz field. *Physical Review B*, 78(7):075204, aug 2008.
- [67] Fan Yang and Rb Liu. Nonlinear optical response induced by non-Abelian Berry curvature in time-reversal-invariant insulators. *Physical Review B*, 90:245205, 2014.
- [68] Fan Yang and Renbao Liu. Berry phases of quantum trajectories of optically excited electronhole pairs in semiconductors under strong terahertz fields. *New Journal of Physics*, 15(11):115005, nov 2013.

- [69] Fan Yang, Xiaodong Xu, and Ren Bao Liu. Giant Faraday rotation induced by the Berry phase in bilayer graphene under strong terahertz fields. *New Journal of Physics*, 16, 2014.
- [70] K. L. Yeh, M. C. Hoffmann, J. Hebling, and Keith A. Nelson. Generation of 10 μJ ultrashort terahertz pulses by optical rectification. *Applied Physics Letters*, 90(17):1–3, 2007.
- [71] Benjamin Zaks, Hunter B. Banks, and Mark S. Sherwin. High-order sideband generation in bulk GaAs. *Applied Physics Letters*, 102(1):012104, 2013.
- [72] Benjamin Zaks, James Heyman, Dominik Stehr, Dan Allen, Nelson Coates, and Mark Sherwin. Single shot high resolution THz upconversion spectrometer. *33rd International Conference on Infrared and Millimeter Waves and the 16th International Conference on Terahertz Electronics, 2008, IRMMW-THz 2008*, pages 1–2, 2008.
- [73] Benjamin Zaks, Renbao Liu, and Mark S. Sherwin. Experimental observation of electron-hole recollisions. *Nature*, 483(7391):580–3, mar 2012.
- [74] Benjamin Zaks, Dominik Stehr, Tuan-Anh Truong, Pierre M Petroff, Stephen Hughes, and Mark S. Sherwin. THz-driven quantum wells: Coulomb interactions and Stark shifts in the ultrastrong coupling regime. *New Journal of Physics*, 13(8):083009, aug 2011.
- [75] Ph. Zeitoun, G. Faivre, S. Sebban, T. Mocek, A. Hallou, M. Fajardo, D. Aubert, Ph. Balcou, F. Burgy, D. Douillet, S. Kazamias, G. de Lache‘ze-Murel, T. Lefrou, S. le Pape, P. Merce‘re, H. Merdji, A. S. Morlens, J. P. Rousseau, and C. Valentin. A high-intensity highly coherent soft X-ray femtosecond laser seeded by a high harmonic beam. *Nature*, 431(1):426, 2004.