Purdue University
Purdue e-Pubs

Open Access Dissertations | Theses and Dissertations

Spring 2015

Growth of low disorder GaAs/AlGaAs heterostructures by molecular beam epitaxy for the study of correlated electron phases in two dimensions

John D. Watson
Purdue University

Follow this and additional works at: https://docs.lib.purdue.edu/open_access_dissertations

Part of the Condensed Matter Physics Commons, Materials Science and Engineering Commons, and the Nanoscience and Nanotechnology Commons

Recommended Citation
Watson, John D., "Growth of low disorder GaAs/AlGaAs heterostructures by molecular beam epitaxy for the study of correlated electron phases in two dimensions" (2015). *Open Access Dissertations*. 585.
https://docs.lib.purdue.edu/open_access_dissertations/585

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

Graduate School Form 30
Updated 1/15/2015

PURDUE UNIVERSITY
GRADUATE SCHOOL
Thesis/Dissertation Acceptance

This is to certify that the thesis/dissertation prepared

By John Watson

Entitled

Growth of Low Disorder GaAs/AlGaAs Heterostructures by Molecular Beam Epitaxy for the Study of Correlated Electron Phases in Two Dimensions

For the degree of Doctor of Philosophy

Is approved by the final examining committee:

Michael Manfra (Chair)	John Reno
Chris H. Greene	
Gabor Csathy	
Ken Ritchie	

To the best of my knowledge and as understood by the student in the Thesis/Dissertation Agreement, Publication Delay, and Certification Disclaimer (Graduate School Form 32), this thesis/dissertation adheres to the provisions of Purdue University's "Policy of Integrity in Research" and the use of copyright material.

Approved by Major Professor(s): Michael Manfra

Approved by: Mark Haugan 3/10/2015

Head of the Departmental Graduate Program Date

GROWTH OF LOW DISORDER GAAS/ALGAAS HETEROSTRUCTURES BY MOLECULAR BEAM EPITAXY FOR THE STUDY OF CORRELATED ELECTRON PHASES IN TWO DIMENSIONS

A Dissertation

Submitted to the Faculty

of

Purdue University

by

John D. Watson

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

May 2015

Purdue University

West Lafayette, Indiana

To my parents, John and Irene.

ACKNOWLEDGMENTS

As I sit here working on the final revisions of my thesis, it's hard to believe that my time at Purdue is rapidly coming to a close. When I started, I don't think I really had a full understanding of what earning a Ph.D. entailed, but I'm glad I took the leap to begin pursuing a career in physics. Over the last seven years I have been fortunate to be exposed to a lot of different research as the lab has gotten up and running, and I have benefited from help from a lot of people at Purdue and elsewhere.

First and foremost, my adviser, Mike Manfra, deserves a huge thank you for hiring me in the first place and for entrusting a young, inexperienced graduate student with so much sensitive (and extremely expensive) lab equipment. His willingness to give me independence in the lab to pursue my own experiments was extremely beneficial to my development as a researcher. I also feel very fortunate for the many opportunities he gave me to interact with all our collaborators working on such a broad range of research. In addition, despite the often painfully slow pace of getting (and keeping) everything in the lab working, he was always up for keeping things light by debating the relative hipness of my favorite music compared to the glory days of '90s grunge rock.

I am also grateful to my Ph.D. committee for taking time out of their very busy schedules to advise me during the course of my work and read and give feedback on what has turned out to be a rather lengthy dissertation. In addition, Gabor Csáthy and John Reno have both gone above and beyond what is typical of committee members. Gabor has been an invaluable resource for cryogenic and low-level measurement techniques and has in many ways been a second adviser to me. John has been very generous with his time in agreeing to be my Sandia fellowship adviser and has always been more than willing to share his extensive MBE experience when we have run into new problems.

It is no understatement to say that my experiments would not have been possible without all the help from my labmates, particularly Geoff Gardner and Sumit Mondal. Geoff was always an excellent role model of patience and professionalism, and he was a great source of knowledge for everything from installing and maintaining equipment to navigating the maze of Purdue purchasing. His optimism was always encouraging, and his patience and attention to detail in training new students has been a huge benefit to the entire lab. Sumit and I faced the challenge of learning low temperature measurement and device fabrication together, and his perseverence in characterizing new wafers day after day was crucial to our progress as a group. The newer students in the lab, including Saeed Fallahi, Qi Qian, Naeem Islam, Mike Yannell, and Shih-Chieh Liu, all deserve thanks for being patient with me while I learned how to train people in the lab. They all brought fresh energy to the group, and it has been exciting seeing the group's impact in the field grow.

There were also several people outside the Manfra group that helped me along the way and made life in grad school more enjoyable. Nodar Samkharadze could be counted on to find humor in any situation, and he was always up for a trip to the gym to burn off some steam. I don't think I would have made it through Gabriele Giuliani's solid state courses without the many late-night study sessions with Abe Olson and Chandan Setty. Dan Whitenack was a driving force for getting people out of the physics building, and I have him to thank for getting me hooked on the music of Punch Brothers. Kurt Aikens could always be counted on to add enthusiasm to any event, and cheering on the Colts at his apartment was often the highlight of my week during football season. Emily and Dave Umulis have very kind and encouraging to me and everyone else in grad IV in opening up their home every week.

I also learned a lot about the fine points of experimental work from several of our collaborators. Bob Willett and Lisa Tracy were both very generous with their time in answering simple-minded questions about device fabrication and e-beam patterning, and Mikey Shulman and Shannon Harvey from the Yacoby group were also helpful in sharing processing recipes and providing feedback we used while optimizing our

heterostructure designs for spin qubits. While not technically a collaborator, Mark Jackson at Oxford Instruments was very gracious in providing detailed instructions for tuning up our dilution fridge.

A number of the staff at Purdue also helped me a lot along the way. The Birck engineering staff were very helpful in training me on all the equipment, and Kenny Schwartz and Dan Hosler were especially flexible in allowing me to change standard operating procedures to make my work with the back-gated devices possible. Keith Schmitter in the physics building could always be counted on to keep the helium liquefier producing liquid, even when it meant coming in early or staying late. Jim Corwin in the machine shop was very gracious in letting me borrow tools, and he was always a valuable source of advice for designing and machining parts. Linda Pacquay, Sandy Formica, Emjai Gregory, Carla Redding, and Nancy Black all deserve special thanks for taking care of paperwork, placing orders, and dealing with bureaucracy so I could focus on research.

My parents have been a constant source of love and support since my early days tinkering with erector sets and disolving egg shells with drano in the kitchen sink. I certainly would not have made it this far without their encouragement to keep pursuing my education. Last, but certainly not least, Enyi has been my number one fan and has been exceptionally patient of all the nights and weekends I have spent in the lab trying to coax some good results out of the experiments. Her love and encouragement have made West Lafayette a much brighter place.

TABLE OF CONTENTS

Page

LIST OF TABLES x

LIST OF FIGURES xi

ABSTRACT xxiv

1 Introduction 1
- 1.1 Dimensionality and Interactions 1
- 1.2 The Integer Quantum Hall Effect 2
- 1.3 The Fractional Quantum Hall Effect 6
- 1.4 The Incompressible State at $\nu = 5/2$ 10

2 Molecular Beam Epitaxy 17
- 2.1 Principles of Molecular Beam Epitaxy 17
  - 2.1.1 MBE Chamber Layout 17
  - 2.1.2 Pumping Methods 19
  - 2.1.3 Vacuum Analysis 21
  - 2.1.4 The RHEED Technique 24
- 2.2 Design Considerations for High Mobility MBE 30
- 2.3 Machine Setup 36
- 2.4 Lessons Learned During the First Growth Campaign 41
- 2.5 Heterostructure Design 47
  - 2.5.1 Smoothing and Gettering Layers 48
  - 2.5.2 Charge Transfer Fundamentals 49
  - 2.5.3 Silicon Dopant Incorporation 51
  - 2.5.4 Impact of Heterostructure Design on Scattering Mechanisms 53
  - 2.5.5 Doping Considerations for Amplifying 2^{nd} LL Physics 55
  - 2.5.6 Doping Considerations for Minimizing Charge Noise 56
  - 2.5.7 2D Holes in GaAs 59
- 2.6 Summary 60

3 Cryogenics for Electrical Transport Measurements 61
- 3.1 Janis Pumped ^{3}He Cryostat for Cooling to T = 300 mK 61
  - 3.1.1 Wiring of the ^{3}He System 64
- 3.2 Kelvinox 100 ^{3}He/^{4}He Dilution Refrigerator 64
  - 3.2.1 Basic Operating Principle 64
  - 3.2.2 Initial Construction 67
  - 3.2.3 Gas Handling System 68

Page

3.2.4 Wiring and Sample Mount 72
3.2.5 Thermometry 78

4 Scattering Mechanisms in a High-Mobility Low-Density Carbon-Doped (001) GaAs Two-Dimensional Hole System 83

5 Exploration of the Limits to Mobility in Two-Dimensional Hole Systems in GaAs/AlGaAs Quantum Wells 93

6 Impact of Heterostructure Design on Transport Properties in the 2nd Landau Level in In-Situ Back-Gated Two-Dimensional Electron Gases 111
6.1 Introduction 112
6.2 Device Growth and Fabrication 114
6.3 Low Temperature Transport 119
6.4 Conclusion 125

7 Summary and Future Work 126
7.1 High Mobility MBE 126
7.2 2D Hole Systems 126
7.3 The $\nu = 5/2$ FQHE State 127

A Computer Codes 128
A.1 Standard Structure Nextnano Input 128
A.2 In-Situ Back-Gated 2DEG Nextnano Input 148
A.3 Matlab Code for Calculating Transport Lifetimes 165

B MBE Standard Operating Procedure 178
B.1 Machine Check-Out 182
B.2 Beam Fluxes 184
B.3 RHEED Warm-up 187
B.4 RHEED Measurements 190
B.5 Loading New RHEED Wafers 196
B.6 Pre-Growth Transfer 198
B.7 Wafer Warm-Up 201
B.8 Growth Startup 203
B.9 Wrap-Up of Growth Setup and Checks During Growth 205
B.10 End of Growth 206
B.11 Unloading Wafers 208
B.12 Wafer Inspection 216
B.13 Nomarski Exam 221

C Wafer Characterization Standard Procedure 227
C.1 Sample Preparation 227
C.1.1 Initial Bookkeeping 227
C.1.2 Applying Contacts 228
C.1.3 Mixing Up Contact Alloy 228

Page

C.1.4 Annealing Contacts . . . 229
C.1.5 Mounting Samples on Header . . . 230
C.1.6 Room Temperature Checks . . . 230
C.2 4K Characterization . . . 230
C.2.1 Cooldown and Measurement . . . 230
C.2.2 Warm-Up . . . 231
C.2.3 Wrap-Up . . . 232

D Room Temperature Hall Effect System Standard Operating Procedure . . . 233
D.1 Sample Preparation . . . 233
D.2 Sample Measurement . . . 234

E Helium Transfer from Liquefier Standard Operating Procedure . . . 240
E.1 Pre-Transfer Bookkeeping . . . 240
E.2 Starting the Transfer . . . 241
E.3 Ending the Transfer . . . 247
E.4 Final Checks . . . 248

F ^{3}He Fridge Standard Operating Procedure . . . 250
F.1 Sample Loading . . . 250
F.2 Sample Cool Down . . . 253
F.3 Sample Warm Up . . . 255
F.4 Changing the Sample Mount . . . 256

G Device Fabrication Standard Operating Procedure . . . 260
G.1 Initial Preparations . . . 260
G.2 Tools and Tool Preparation . . . 261
G.3 Ga Removal . . . 263
G.4 Cleave and Initial Clean . . . 267
G.5 Etching . . . 272
G.6 Ohmic Contacts . . . 280
G.7 Dielectrics . . . 288
G.7.1 Deposition . . . 288
G.7.2 Dielectric Patterning . . . 291
G.8 Electron Beam Lithography . . . 292
G.8.1 Sample Preparation . . . 292
G.8.2 Raith Setup . . . 293
G.8.3 Sample Alignment . . . 295
G.8.4 Exposing and Developing . . . 298
G.8.5 Design Tips . . . 300
G.9 Optical Gates and Bond Pads . . . 303
G.10 Sample Mounting and Wire-Bonding . . . 304
G.11 Condensed Checklist for In-situ Back-gated Devices . . . 306
G.11.1 Ga Removal . . . 306

Page

G.11.2 Via Etch 307
G.11.3 Mesa Etch 308
G.11.4 Ohmics 309
G.11.5 E-beam Gates 310
G.11.6 Optical Gates and Bond Pads 311
G.11.7 Mounting and Wire Bonding 312

H Kelvinox Dilution Fridge Standard Operating Procedure 313
H.1 Introduction to the System 313
H.2 Preparing Samples 315
H.3 Preparing the Fridge 322
H.4 Cooling the Fridge to 4K with a Cold Magnet 330
H.5 Helium Transfer 333
H.6 Preparing to Condense 336
H.7 Sample Checks and Illumination 342
H.8 Condensing and Circulating 343
H.9 Cooling Down the Magnet 346
H.10 Warming Up the Insert 356
H.11 Condensed Checklist 359
H.11.1 Cooling Down Fridge (Magnet Cold) 359
H.11.2 Condensing and Cooling to Base 360
H.11.3 Cooling Down the Magnet 360
H.11.4 Warming Up 361

REFERENCES 363

VITA 371

LIST OF TABLES

Table | Page

5.1 Summary of structural parameters including δ-doping setback distance d, Al mole fraction around the dopants x_d, Al mole fraction surrounding the quantum well x_w, 2DHS density p, and $T = 300$ mK mobility after illumination μ. 96

LIST OF FIGURES

Figure Page

1.1 Change in 2D density of states under the application of an external, perpendicular magnetic field. The degeneracy of the Landau levels is set by the number of states in a range $\hbar\omega_c$ that is "compacted" into a single Landau level by the magnetic field. 2

1.2 Schematic of magnetotransport measurement setup. Shaded regions represent contacts to the Hall bar. 3

1.3 Magnetotransport in a 2DEG showing the Shubnikov - de Haas oscillations, the onset of spin splitting, and the development of the integer quantum Hall effect. 4

1.4 Landau level energy across a 2D Hall bar of width w. 5

1.5 2D DOS in the presence of a perpendicular magnetic field, finite temperature, and disorder. Cross-hatched regions indicate regions of charge localization in the 2D bulk. 7

1.6 Magnetotransport from a high quality 2DEG at low temperature showing a number of FQHE states. The filling fraction is indicated for a number of fractional quantum Hal states in the lowest LL as well as for the first few integer quantum Hall states. 8

1.7 Sketch of the geometry of a Fabry-Perot interferometer in the quantum Hall regime. The 2DEG is bound by the blue rectangle, the yellow regions represent Ohmic contacts, the black shapes represent surface gates used to electrostatically deplete the 2DEG underneath them, red lines represent the flow of edge states, and blue x's represent localized quasiparticles. As the side gate voltage is varied the number of quasiparticles encircled by the edge states changes, giving rise to oscillations in the conductance of the interferometer due to the change in phase of the edge states. . . . 15

2.1 Cross sectional sketch showing the functional form of a typical MBE chamber. 18

2.2 Schematic of the operating principle of an ion gauge. Electrons are produced by the filament and accelerated to the grid. Ions are accelerated to the collector where they produce a measurable current. 22

Figure Page

2.3 Sketch of an RGA. (a) Side view of a typical RGA showing the three main sections and direction of ion flow. (b) Illustration of the electrodes forming the quadrupole mass filter. Figure reproduced from Operating Manual and Programming Reference - Models RGA100, RGA200, and RGA300 Residual Gas Analyzer, copyright (1996), with permission from Stanford Research Systems. 23

2.4 RGA spectra. (a) RGA spectrum from a well-baked and leak-free chamber. (b) RGA spectrum from a leaking, unbaked chamber. The leak is evinced by the large peak at 32 and the fact that the 14 peak is larger than the 15 peak which together indicate the presence of O_2 and N_2 leaking into the chamber. 24

2.5 Sketch of the effect of increasing the lattice constant along the z direction. As the real space lattice constant increases, the corresponding reciprocal lattice constant decreases. In the limit of a 2D crystal in the x-y plane, the reciprocal lattice would consist of a series of thin rods parallel to the k_z axis. 25

2.6 Cross sectional views of the Ewald sphere construction used to determine momentum- and energy-conserving scattering processes from a 2D crystal. 26

2.7 Typical RHEED diffraction pattern in GaAs at growth temperature. (a) RHEED pattern showing the 2× reconstruction. (b) RHEED pattern showing the 4× reconstruction. 27

2.8 (a) RHEED pattern from a rough wafer immediately after desorbing the oxide at growth temperature. (b) RHEED pattern from a smooth wafer. 28

2.9 Sketch of crystal surface in cross section during growth. As the surface roughens due to the nucleation of a new monolayer, the specular intensity drops. Once the new atomic layer nears completion, the intensity recovers. 29

2.10 Typical oscillations in the intensity of the RHEED specular spot during growth of GaAs. The oscillations are damped due to roughening of the surface as islands nucleate on top of other islands. The oscillation period corresponds to the crystal growth rate. 30

2.11 The Manfra high mobility MBE machine as installed. The design and materials of construction play an important part in determining the ultimate crystal purity (see text). 32

Figure Page

2.12 RGA spectra showing the impact of replacing the rubber-sealed poppet valve on the chamber's cryopump with an all-metal burst disk. (a) RGA spectrum from chamber with poppet valve. The spectrum shows the tell-tale sign of a leak: the mass 14 peak is larger than the mass 15 peak and the 32 peak is large relative to the 28 peak. (b) RGA spectrum after replacing the poppet valve with an all-metal burst disk. The leak signature is gone, though the the chamber still required additional baking to improve the overall vacuum purity. 33

2.13 Plot of maximum 2DEG mobility as a function of temperature showing the result of improvements in MBE and heterostructure design and materials. Reprinted with permission from L. Pfeiffer, K. W. West, H. L. Stormer, and K. W. Baldwin. Appl. Phys. Lett. **55**, 1888 (1989). Electron mobilities exceeding 10^7 cm^2/Vs in modulation doped GaAs. Copyright (1989), AIP Publishing LLC. 34

2.14 MBE system with the baking system in place. 38

2.15 Homemade molybdenum tips for spot-welding MBE components. . . . 39

2.16 Summary of MBE clean-up early in first growth campaign. The first $\sim$ 12 heterostructures grown were electrically insulating at low temperature. Most of the growths between 12-17 were bulk, undoped GaAs used to measure the background impurity concentration. After extensive outgassing the mobility rapidly increased as we first worked on the design of a single heterojunction (SHJ) heterostructure and then switched to optimize the growth parameters before finally moving to a more complicated doping well type heterostructure. 42

2.17 Dedicated gallium outgassing chamber. The chamber featured a near-vertical port to enable the the crucible to be filled almost completely with gallium. The frame supporting the chamber could be rolled into position next to the MBE to allow sources to be passed into the MBE without exposure to air. 43

2.18 Wafer with so-called "island defects" after cleaving. Two defects (circled) away from the cleave lines as wells as a few defects along cleave lines are called out in the figure. 45

2.19 Nomarksi contrast optical micrograph of a wafer showing so-called cross hatch morphology. The field of view is $\sim$ 2.5 mm wide. 46

2.20 Sketch of the Γ band edge as a function of depth for a simple single interface heterojunction. The dashed line shows the position of the Fermi energy E_F. 49

Figure Page

2.21 Modern quantum well structure for studying 2^{nd} LL physics. (a) Sketch of the Γ band edge as a function of depth for a quantum well with a doping well doping scheme. The horizontal dashed line shows the position of the Fermi energy E_F while the vertical dashed lines show the position of the Si delta doping layers. An enlarged view of the dashed blue box is shown in the adjacent panel. (b) Close-up view of the doping well layer sequence. 55

2.22 Sketch of the Γ band edge as a function of depth for a heterostructure (a) cooled without a cooling bias, (b) during cooling with a forward bias, and (c) at low temperature after bias cooling. Reprinted figure with permission from M. Pioro-Ladrière, J. H. Davies, A. R. Long, A. S. Sachrajda, L. Gaudreau, P. Zawadzki, J. Lapointe, J. Gupta, Z. Wasilewski, and S. Studenikin, Phys. Rev. B, 72, 115331 2005. Copyright (2005) by the American Physical Society. 57

3.1 Sketch of the cross section of the Manfra group top-loading ^{3}He fridge. 62

3.2 Schematic of the main components of a dilution refrigerator. Black dashed lines represent cold plates in the fridge to which heat sinks, thermometers, or heaters can be attached. Red lines represent incoming mixture while blue lines represent outgoing mixture. Line thickness indicates the relative size of tubing in the fridge. 65

3.3 Lab space during construction of the pit. 68

3.4 Dewar, magnet, and stand in place at end of the intial construction phase. 69

3.5 Schematic of the homemade gas handling system. Line thickness denotes size of vacuum hoses. 70

3.6 Threaded rods (circled in red) that were adjusted to center the radiation shield inside the IVC. 71

3.7 Threaded insert and plug used to block thermal radiation from reaching the sample. 72

3.8 Homemade filter box that couples a D-SUB 25 cable to the Fischer connector on top of the fridge. 73

3.9 Home made heat sinks mounted on the 50 mK plate. 74

3.10 Tail after etching in 1:1 water:nitric acid. 75

3.11 Tail after cleaning with Brasso metal polish, toluene, acetone, and methanol. 76

3.12 Slurry of GE varnish and silver powder used to heat sink copper wires to tail. 76

3.13 Finished tail mounted on the mixing chamber. 77

Figure Page

3.14 Wiring diagram of the thermometers and heaters in the fridge. 79

3.15 Schematic of electronics setup and grounding. 80

4.1 Schematic of the device structure used in our experiments. 85

4.2 Effect of gate voltage on carrier density and leakage current at T = 300 mK. 86

4.3 Low field magnetoresistance of the backgated sample at T = 50 mK. . 87

4.4 Mobility as a function of the density at T = 300 mK (squares) and T = 50 mK (open circles). Straight lines are guides to the eye to the 300 mK data to illustrate 0.7 and 1.7 power laws. 88

4.5 (a) Theoretical density dependence of the exponent α in $\mu \propto p^{\alpha}$. (b) Comparison of experimental mobility data and theoretical results. Solid line represents RPA-Boltzmann calculation and dashed line represents RPA-Boltzmann calculation with the Hubbard approximation in both plots. 89

5.1 Layer structure of devices in this experiment. Note the use of two different Al mole fractions x_w and x_d in some of the devices as indicated in Table 5.1. 95

5.2 (Color online) Magnetotransport at $T = 300$ mK after illumination with a red LED of (a) peak mobility sample and (b) low density sample that exhibits many nascent fractional QHE features. 98

5.3 (Color online) (a) $T = 300$ mK mobility after illumination with a red LED as a function of density for various dopant setback distances d. Solid lines are guides to the eye. For fixed d the density was controlled by varying the Al mole fraction x. Samples were grown in random order to avoid continued machine clean-up from skewing the observed trend in mobility. Samples A-D were grown with varying x at fixed p to test the effect of alloy and interface roughness scattering on the mobility (see text). (b) Transport lifetime estimated as a function of density. Inset: Effective mass for our structures as a function of density extrapolated from refs. [108, 107]. 99

5.4 (Color online) Dashed lines show a comparison of the self-consistently calculated valence band edges (dashed lines) for the high density sample # 3 and and a single heterojunction sample with $x = 0.45$. Solid lines show a comparison of the self-consistently calculated wavefunction for sample # 3 and the Fang-Howard variational wavefunction. 102

Figure Page

5.5 (Color online) Comparison of $d = 80$ nm experimental data with mobility calculations. N_{BI} and Λ are used as free parameters to obtain a good fit to data. Pink star represents SHJ 2DES grown during this experiment. 105

5.6 (Color online) Comparison of $d = 80$ nm experimental data with mobility calculations. Δ and Λ are varied in order for mobility calculations to obtain agreement with an electron structure grown during this experiment (see text). 106

5.7 (Color online) Shubnikov-de Haas oscillations of high density $d = 80$ nm, $x = 0.45$ device. Inset (a): Sketch of the spin-split heavy hole and light hole ground states in a quantum well. Inset (b): Index of extrema in R_{xx} vs. B^{-1}. The high field slope gives the total density of 1.8×10^{11} cm^{-2}, and the low field slope gives the lighter sub-band density of 7×10^{10} cm^{-2} while the difference in the two gives the second sub-band density of 1.1×10^{11} cm^{-2}. 107

5.8 (Color online) (a) Estimate of the Hall density expected from Eq. 5.7 using the measured subband densities $p_1 = 1.1 \times 10^{11}$ cm^{-2} and $p_2 = 7 \times 10^{10}$ cm^{-2}. (b) Estimate of the expected measured mobility μ_{Hall} if the high density subband $p_1 = 1.1 \times 10^{11}$ cm^{-2} has a high mobility $\mu_1 = 2 \times 10^6$ cm^2/Vs. For the second subband density (dashed red line) $p_2 = 7 \times 10^{10}$ cm^{-2} measured in Fig. 5.7 the expected measured mobility $\mu_{Hall} \geq 1.75 \times 10^6$ cm^2/Vs. 108

6.1 Effect of Ohmic annealing temperature on device performance from wafer A. (a) Median 2-terminal resistance to ground of individual contacts measured in the dark at T = 4 K and $V_g = 0$ as a function of annealing temperature. (b) V_{leak}, defined as the voltage at which the gate leakage current reached 1 nA, as a function of annealing temperature. 113

6.2 Histogram of leakage turn-on voltage V_{leak} for devices fabricated with the optimized processing recipe and mask set. All the devices were annealed at 375 °C. The dashed line represents the voltage required to reach a 2DEG density of $\sim 3.2 \times 10^{11}$ cm^{-2}. 116

Figure Page

6.3 Magnetotransport in the lower spin branch of the 2^{nd} LL in device A. During the first cool down of the sample (black curves) the sample was mounted on a commercial ceramic chip carrier. At a gate leakage current (power) of $\sim$ 63 pA (76 pW) the electrons appear very warm as seen by the lack of RIQHE features, despite a low mixing chamber temperature T_{MC}. During the second cool down of the sample (blue curves), the device was mounted on a homemade header with a copper strip screwed onto the end of the cold finger on the mixing chamber. The electrons were obviously much colder even for a slightly higher T_{MC}. The green curve shows the transport around $\nu = 5/2$ during the second cooldown for T_{MC} = 51 mK. Comparing the green and black data, we conclude that the electron temperature was $\sim$ 50 mK for T_{MC} = 11 mK during the first cooldown. 118

6.4 Magnetotransport in device C after illumination with a red LED. The re-entrant states are labelled following the convention in reference [161]. Red data show the transport for the maximum strength in RIQHE states 2a and 2b while the blue data show transport at the highest density before the second subband became occupied. 120

6.5 Gap at $\nu = 5/2$ as a function of density for devices B and C. Inset shows the Arrhenius plot for device C at a density of 3.35×10^{11} cm^{-2} where the gap was measured to be 625 mK. 121

6.6 Cyclotron energy $\hbar\omega_c$ (red triangles) and spacing between E_F and the second sub-band (blue circles) overlaid with $\Delta_{5/2}$ for Wafer C. $\Delta_{5/2}$ drops suddenly at high density when the ground state is pushed into the lowest LL of the anti-symmetric sub-band. 122

6.7 Strength (as defined in the text) of the RIQHE in device C during the second cool-down as a function of density; the power dissipation from the gate leakage current is shown in the top panel. States 2a and 2b weaken over the measured density range while states 2c and 2d strengthen over the same range. 123

6.8 Comparison of the strength of the 2a RIQHE state in device C from two different cooldowns. The strength of the state is comparable between the two cooldowns despite the large change in gate power dissipation. . . . 124

A.1 Output of standard structure simulation. 130

A.2 Output of in-situ back-gated simulation. 150

B.1 The most important switch. 184

Figure Page

B.2 Typical beam equivalent pressures of each source during first growth campaign. Note that the first set of flux measurements has a larger percentage drop while the shutter is open, likely due to As build-up around the source. 186

B.3 Check the power output % as soon as you start ramping the CAR heater. 188

B.4 View of CAR control screen and As valve driver controller. 189

B.5 Main shutter with LED position indicator lights. 189

B.6 View of the scribe marks on the CAR index position. 190

B.7 RHEED controller and power supply. 191

B.8 Comparison of the 2× reconstruction (left) and 4× reconstruction (right) as seen on the RHEED screen. 191

B.9 Shutter control screen on the e-rack touch-screen computer. 193

B.10 Examples of good and bad GaAs and AlAs RHEED oscillations. 195

B.11 Slide the nose piece towards the growth chamber 9-10 inches before opening the gate valve. 198

B.12 The buffer chamber pressure must be sufficiently low, and the LL gate valve must be shut before the gate valve to the growth chamber can be opened. 199

B.13 The main shutter should also be down before the gate valve is opened (left). The gate valve is fully open when the second line on the arm is visible (right). 200

B.14 Double check that the transfer rod is all the way back before closing the gate valve. The gate valve is closed when the plastic sleeve reaches the top mark and the indicator is fully in the "closed" position. 202

B.15 Examples of "faint and hazy" morphology immediately after the wafer is loaded into the growth chamber. 203

B.16 Don't forget to double check the recipe in Molly line-by-line to make sure it matches what you intend to grow in the Excel sheet. 204

B.17 As valve setting during part of the first growth campaign. 207

B.18 Valves to open when warming up the gettering furnace. 208

B.19 LL vent line mechanical gauge and "N2 Vent" valve. 208

B.20 Prepare the LL ROMO and Ga beaker for re-loading the LL. 209

Figure Page

B.21 Ramp down the LL heater, turn off the LL ion gauge, and check that the sorption pumps are plugged before cooling down the sorption pumps. . 210

B.22 Watch for evidence of leaking gate valves when the LL is first vented. When the LL reaches atmospheric pressure, open the thumb screw to relieve the pressure. 211

B.23 Procedure for mounting a new wafer on a block. 213

B.24 Do not over-pressurize the LL. 215

B.25 Ga waste container. 216

B.26 Nomarski phase-contrast optical micrograph of slip lines at the edge of the wafer. 218

B.27 Nomarski phase contrast micrograph of a hazy wafer. The field of view is $\sim$ 2.5 mm wide. 221

B.28 Microscope setup. 222

B.29 Nomarski phase contrast micrograph showing a wafer with good morphology. 224

B.30 Nomarski phase contrast micrograph of a wafer with “cross hatch” morphology. The field of view is $\sim$ 2.5 mm wide. 225

D.1 A bulk doped sample after being measured. The scale shown is inches. 234

D.2 View of the inside of the MMR Hall effect system sample chamber. . . 235

D.3 Graph of data from the linearity check measurement. 237

E.1 View of the liquefier laptop when the liquefier is running smoothly. . . 241

E.2 Pressure gauge behind the computer showing the “P impure” reading. . 242

E.3 Recovery valve on the liquefier dewar. 243

E.4 Liquefier dewar pressure. Do not let this pressure exceed 6 psi if the liquefier is running. 244

E.5 Pressure building port. 245

E.6 Heat gun used to keep the gas meter from freezing. 246

F.1 (a) Gate valve. (b) Probe pump-out line. (c) Speedi-valve. (d) Lower Nupro valve (circled in red). 251

F.2 (a) Probe suspended above the gate valve by the winch. (b) Probe clamp. (c) Sample mount. (d) Tail. Make sure all set screws, stand-offs, and nuts are secure. 252

Figure Page

F.3 (a) Check that the heat-shrink tubing around the thermometer leads is not sticking out past the copper or brass. (b) Wait for the probe pressure to fall to $\sim$ 20 mTorr before proceeding. 252

F.4 Close the valve between the dump and the fridge. 253

F.5 (a) Start by prying the connectors apart with the tweezers sitting across multiple rows of connectors. (b) Move to the other end of the sample mount and pry the connectors apart a little more. (c) Do not try to pry the connectors apart with the tweezers parallel to rows of connectors. (d) Check that none of the pins are bent before trying to install the sample mount on the probe. 257

F.6 Support the sample mount and probe by holding either the circuit boards or metal disks. Do not touch the wires. 258

G.1 Gallium-covered backside of a wafer. 266

G.2 (a) Backside of wafer after gallium has been wiped off. (b) Approximate number of q-tips necessary to wipe all the gallium off a full wafer. . . . 266

G.3 Sketch of wire-bondable chip carrier and chip. The orange lines represent the chip carrier with the cross-hatched region representing the gold bond pad. The green rectangles represent the chip. The larger rectangle is the border of the chip and the smaller rectangle represents the area on the chip that will be free of significant photoresist edge-bead. All dimensions are in mm. 268

G.4 Sonicator setup used to clean the samples. 270

G.5 Homemade teflon inserts used to sonicate multiple samples in parallel. . 271

G.6 Surface of wafer after gallium removal and sonicated cleaning. Field of view is $\sim$ 2.5 mm wide. 272

G.7 Shallow-angle SEM view of mesa edge after etching using (a) AZ1518 resist and (b) S1805 resist. The thinner S1805 allows features to be defined more sharply and results in less edge roughness after etching. Plan-view SEM images (not shown) show that typical roughness in the plane of the original wafer surface is $\sim$ 150 nm for features defined with AZ1518 resist while the roughness drops to $\sim$ 50 nm for features defined with S1805 resist. 273

G.8 Mask design showing via and mesa layers with large features for aligning the sample to the mask. 275

Figure Page

G.9 Cross-sectional SEM views of different mesa etches. (a) Mesa etched with ammonium hydroxide piranha from [177]. (b) Mesa etched with phosphoric piranha from [176]. (c) Mesa etched with our standard phosphic piranha. 279

G.10 Shallow angle SEM images of resist profile for lift-off. (a) AZ1518 resist after patterning and hardening with chlorobenzene. (b) AZ1518 resist after metalization but before lift-off. Light regions are metal and dark regions are resist. 281

G.11 Drawings of different mesa and Ohmic designs. The black lines represent the mesa while the red lines represent the Ohmics. The Ohmics on the Hall bar on the left have well defined directions in which the electrons in the 2DEG would have to tunnel into each contact. The Hall bar on the right has Ohmics with “scallops” which allow the electrons to tunnel into the contacts along both crystallographic direction for all the contacts. The contact region of the voltage probe is shown enlarged for clarity. The bottom contact on the right has no overlap with the edge of the mesa. Such “interior” contacts were consistently insulating at low temperature. 285

G.12 Overlay of mesa (black) and Ohmic (red) designs. The large red rectangles at the top and bottom of the chip serve as initial orientation markers during the initial e-beam alignment. The Hall bars shown in the design are $\sim$ 1 mm long. The smaller L-shaped arrays of alignment marks are used for the actual 3-point alignment procedure to create the coordinate transformation between the mask and stage coordinates. In addition, the fact that these features are large enough to be seen with the naked eye enables one to mount the chip on the sample mount in a pre-determined orientation. The inset shows a magnified view of the alignent markers used for the 3-point alignment. These markers are arranged so that the array lacks rotational symmetry; this lack of symmetry makes errors during the alignment less likely. Dimensions shown are in microns. 296

G.13 Example of the influence of e-beam pattern design on exposure uniformity. (a) Pattern in which each arm of the multi-QPC device is “OR-ed” into a single large polygon. The black dashed box denotes the 100 μm write field used for the exposure. (b) The same device broken up into multiple polygons. (c) Optical micrograph of resist after exposing and developing using the design in (a). The device shows brighter lines at the boundaries of the write fields (where the design was exposed twice) as well as through the central arms of the device. (d) Optical micrograph of resist after exposing and developing using the design in (b). No excess exposure is seen in the central arms of the QPC structure. 300

Figure Page

G.14 Example of a successful dose test. Main panels show optical micrographs after metal lift-off while the insets show SEM micrographs of the center of the device. As the dose is increased from (a) to (d), smaller and smaller features should survive the lift-off procedure. 302

H.1 Top of the IVC with jack screw called out. The jack screws are used to break the indium seal. 315

H.2 Procedure for removing the radiation shield. (a) Align the cut-out in the radiation shield flange with the heat sink to avoid scratching the wires on the heat sink. (b) Once the flange is safely past the heat sink, rotate it slightly to align the cut-out with the corner of the connector. 316

H.3 Copper headers mounted on the tail. Care should be taken when removing the radiation shield to not bump the LEDs with the shield. 317

H.4 Homemade sandpaper tools for cleaning copper oxide off of small pieces. 2000 and 3000 grit sandpaper is glued to wooden sticks with quick set epoxy to aid in cleaning small copper pieces such as the headers. 319

H.5 Illustration of how to wire up processed samples on the dilution fridge headers. Start the bond on the chip carrier (left) and drag the wire over the solder fork before bonding to the sample (right). Note that a standard plastic header without the copper insert is shown here because the copper headers were both in the fridge during the preparation of this manuscript. 321

H.6 Manifold used for adding exchange gas to the IVC. Fill the 4-way cross to $\sim$ - 24 inHg with helium and then add this to the IVC to act as exchange gas. 328

H.7 Dewar exhaust port to helium recovery system. 331

H.8 Sliding seal exhaust valve #2. 332

H.9 Fridge ground lines. Both these ground lines must be securely connected to the fridge to keep the resistance bridge from showing overload at base temperature. 333

H.10 Sand-bucket manifold. 337

H.11 Fridge heater connections. 338

H.12 Primary gas handling manifold. 340

H.13 Roots blower manifold. 341

H.14 Dewar purge manifold. 349

H.15 1K pot manifold. 350

Figure Page

H.16 Lambda plate port. 350

H.17 High-tech LN_2-fridge coupling unit. 351

H.18 Primary lab helium recovery valve. 352

H.19 Magnet thermometers and persistent switch. 353

ABSTRACT

Watson, John D. Ph.D., Purdue University, May 2015. Growth of Low Disorder GaAs/AlGaAs Heterostructures by Molecular Beam Epitaxy for the Study of Correlated Electron Phases in Two Dimensions. Major Professor: Michael J. Manfra.

The unparalleled quality of GaAs/AlGaAs heterostructures grown by molecular beam epitaxy has enabled a wide range of experiments probing interaction effects in two-dimensional electron and hole gases. This dissertation presents work aimed at further understanding the key material-related issues currently limiting the quality of these 2D systems, particularly in relation to the fractional quantum Hall effect in the 2^{nd} Landau level and spin-based implementations of quantum computation.

The manuscript begins with a theoretical introduction to the quantum Hall effect which outlines the experimental conditions necessary to study the physics of interest and motivates the use of the semiconductor growth and cryogenic measurement techniques outlined in chapters 2 and 3, respectively. In addition to a generic introduction to the molecular beam epitaxy growth technique, chapter 2 summarizes some of what was learned about the material purity issues currently limiting the low temperature electron mobility. Finally, a series of appendices are included which detail the experimental methods used over the course of the research.

Chapter 4 presents an experiment examining transport in a low density two-dimensional hole system in which the hole density could be varied by means of an evaporated back gate. At low temperature, the mobility reached a maximum of 2.6×10^6 cm^2/Vs at a density of 6.2×10^{10} cm^{-2} which is the highest reported mobility in a two-dimensional hole system to date. In addition, it was found that the mobility as a function of density did not follow a power law with a single exponent. Instead, it was found that the power law varied with density, indicating a cross-over between dominant scattering mechanisms at low density and high density. At low

density the mobility was found to be limited by remote ionized impurity scattering, while at high density the dominant scattering mechanism was found to be background impurity scattering.

Chapter 5 details an experiment examining transport in a series of two-dimensional hole gases in which the dopant setback distance and the Al mole fraction in the barriers of the quantum well were varied. The hole density was tuned in this way from $0.18 - 1.9 \times 10^{11}$ cm^{-2}. Surprisingly, the mobility at T = 0.3 K was found to peak at 2.3×10^{6} cm^{-2} at an intermediate density of 6.5×10^{10} cm^{-2}. Self-consistent Schrödinger/Poisson calculations were performed for each wafer to examine the scattering rates due to a variety of potentials at low temperature. The drop in mobility at high density could be accounted for with the inclusion of interface roughness scattering, but using the same interface roughness scattering parameters for similar two-dimensional electron gases produced inconsistent results. This leaves open the possibility of contributions from other scattering mechanisms in the hole samples at high density.

Chapter 6 presents an in-depth study of in-situ backgated two-dimensional gases used for studying the fragile fractional quantum Hall states in the 2^{nd} Landau level. It was found that leakage currents as small as 4 pA could cause noticeable heating of the electron gas if the lattice was not properly thermally anchored to the cryostat. However, it was also found that when the heterostructure design and device fabrication recipe were properly optimized, gate voltages as large as 4 V could be applied before the leakage turned on, allowing the density to be tuned from full depletion to over 4×10^{11} cm^{-2}. As a result, heating effects at dilution refrigerator temperatures were negligible and the gap at $\nu = 5/2$ could be tuned continuously with density to a maximum value of 625 mK, the largest reported to date. An unusual evolution of the reentrant integer quantum Hall states as a function of density is also reported. Such devices should prove useful for the study of electron correlations in nanostructures in the 2^{nd} Landau level.

1. Introduction

1.1 Dimensionality and Interactions

The study of strongly interacting electrons is one of the central themes in modern research in condensed matter physics because, as Anderson succinctly put it, “More is different” [1]. A great deal of the effort expended in experimental work in condensed matter physics is therefore directed at finding or engineering systems in which interactions dominate the observed physics. One of the most fruitful methods of amplifying the effect of interactions is, as we shall see, to move to systems with reduced dimensionality.

Figure 1.1a shows the simple yet somewhat pedestrian case of the density of states of a two-dimensional (2D) system in the absence of any external potentials. If the system is instead subjected to an external magnetic field B perpendicular to the plane, the Hamiltonian describing the system can be written in the form of the familiar harmonic oscillator potential with the well-known energy eigenvalues $E_N = (N + 1/2)\hbar\omega_c$ [2]. As a result, the density of states is changed to a series of degenerate Dirac delta functions spaced apart in energy by the cyclotron energy (in the case of a negligible Zeeman spin splitting) as shown in figure 1.1b. The index N is referred to as the Landau level (LL) index and $\omega_c = eB/m^*$ where e is the electron charge and m^* is the electron effective mass. As figure 1.1 illustrates, the degeneracy of the LLs can be understood with a simple geometric argument. Under the application of an external magnetic field, the states become “compacted” into the individual LLs. The degeneracy d of a single spin-resolved LL is then simply given by

$$d = \frac{1}{2}\frac{m^*}{\pi\hbar^2}\hbar\omega_c = \frac{eB}{h} \tag{1.1}$$

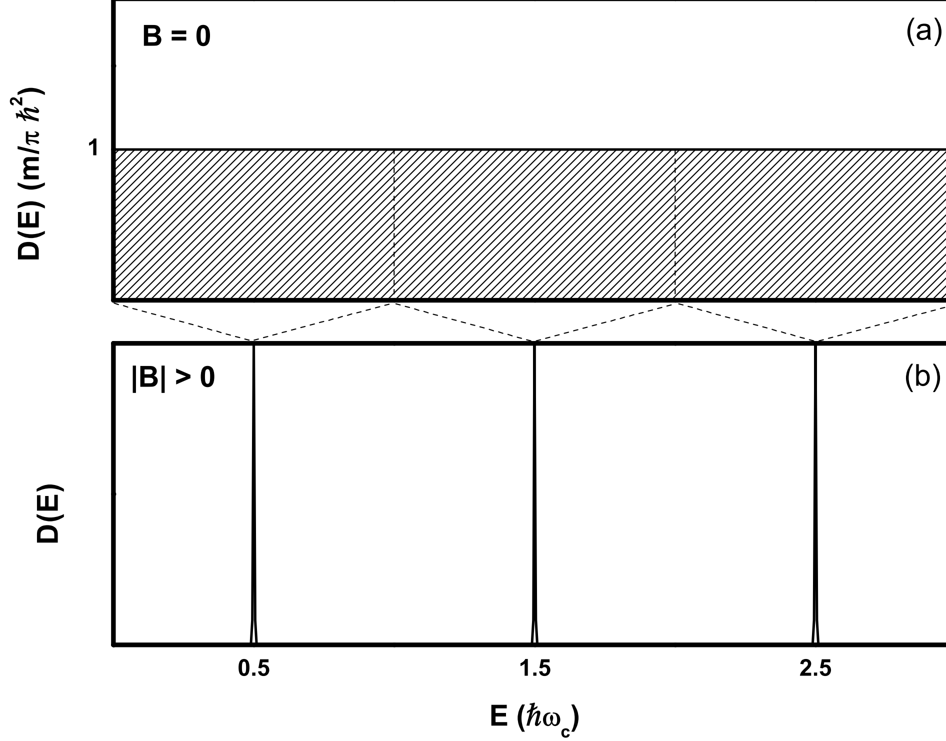

Figure 1.1. Change in 2D density of states under the application of an external, perpendicular magnetic field. The degeneracy of the Landau levels is set by the number of states in a range $\hbar\omega_c$ that is "compacted" into a single Landau level by the magnetic field.

The number of filled LLs (often referred to as the filling fraction ν) is thus

$$\nu = n/d = \frac{hn}{eB} \tag{1.2}$$

where n is the electron number density.

1.2 The Integer Quantum Hall Effect

Now suppose we pass a current I through a rectangular 2D system and measure the voltage drop V_{xx} parallel to the current flow as we increase B as shown in figure 1.2. Figure 1.3 shows the results of a representative measurement at low temperature. As the magnetic field is increased, the LLs will move up in energy and pass through

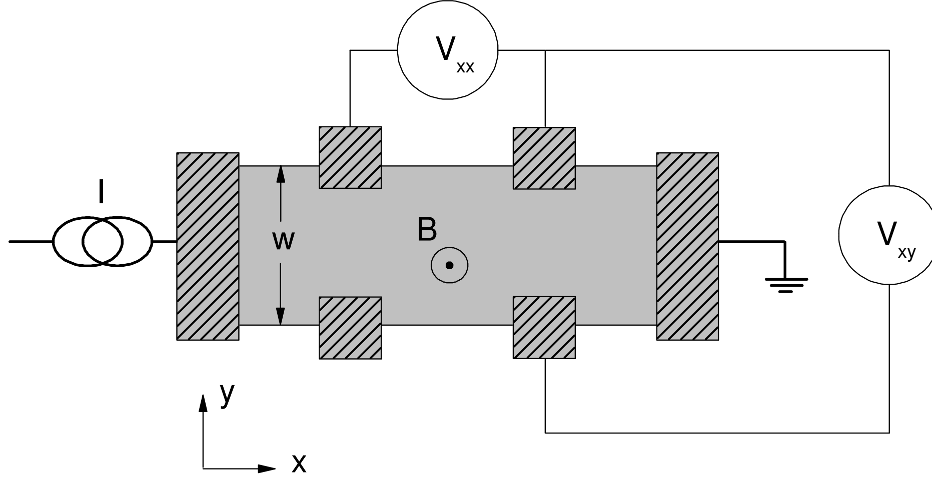

Figure 1.2. Schematic of magnetotransport measurement setup. Shaded regions represent contacts to the Hall bar.

the Fermi energy when the degeneracy increases sufficiently to accomodate all of the electrons in the next highest LL. The conductivity will fluctuate with the density of states at the Fermi surface resulting in an oscillating voltage V_{xx}. These fluctuations (usually plotted as a resistance R_{xx} vs. B) are known as Shubnikov de-Haas (SdH) oscillations and are periodic in inverse magnetic field. As the magnetic field is increased further, the LLs will become more separated and the density of states will become zero for certain values of the magnetic field. In this case since the density of states is zero, the conductivity σ_{xx} will vanish. In addition, the resistivity ρ_{xx} will also vanish since $\rho_{xx} \propto \sigma_{xx}$ in two dimensions [3]. This vanishing of the longitudinal resistance, accompanied by a plateau in the transverse resistance, is known as the integer quantum Hall effect (IQHE) and was first observed by von Klitzing et al. [4] in the inversion layer of a Si MOSFET at low temperature.

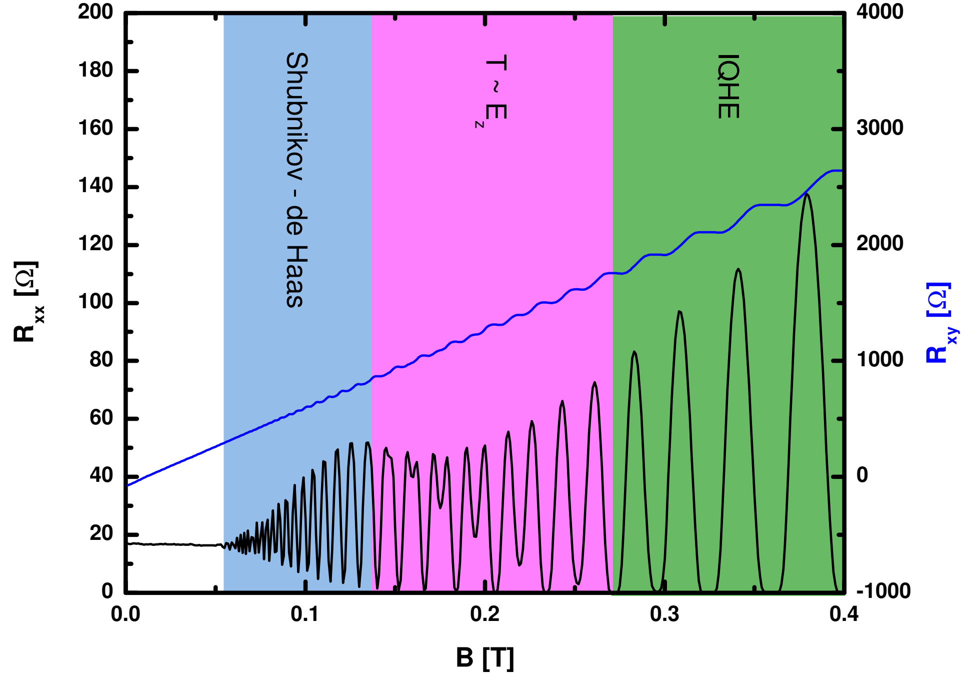

Figure 1.3. Magnetotransport in a 2DEG showing the Shubnikov - de Haas oscillations, the onset of spin splitting, and the development of the integer quantum Hall effect.

To better understand this peculiar situation of zero resitivity in a disordered system at finite temperature, it helps to consider the distribution of energy levels across the width of the Hall bar as shown in figure 1.4. When the Fermi energy E_F lies in between the Landau levels it is apparent that the only states at the Fermi surface (and hence the only states that can carry current at low temperature) are at the edge of the sample where the Landau levels cross the Fermi energy. In addition, these edge states are chiral (i.e. the edge states flow in opposite directions on opposite sides of the sample). The chiral nature of the edge states results from the presence of the edge confining potential $V(y)$ in concert with the magnetic field (see chapter 4 from reference [5] for a formal treatment). From a hand-waving, semi-classical view the edge states can be thought of as skipping orbits as the cyclotron orbits of the electrons collide with the edge of the sample. The chiral nature of the edge states results

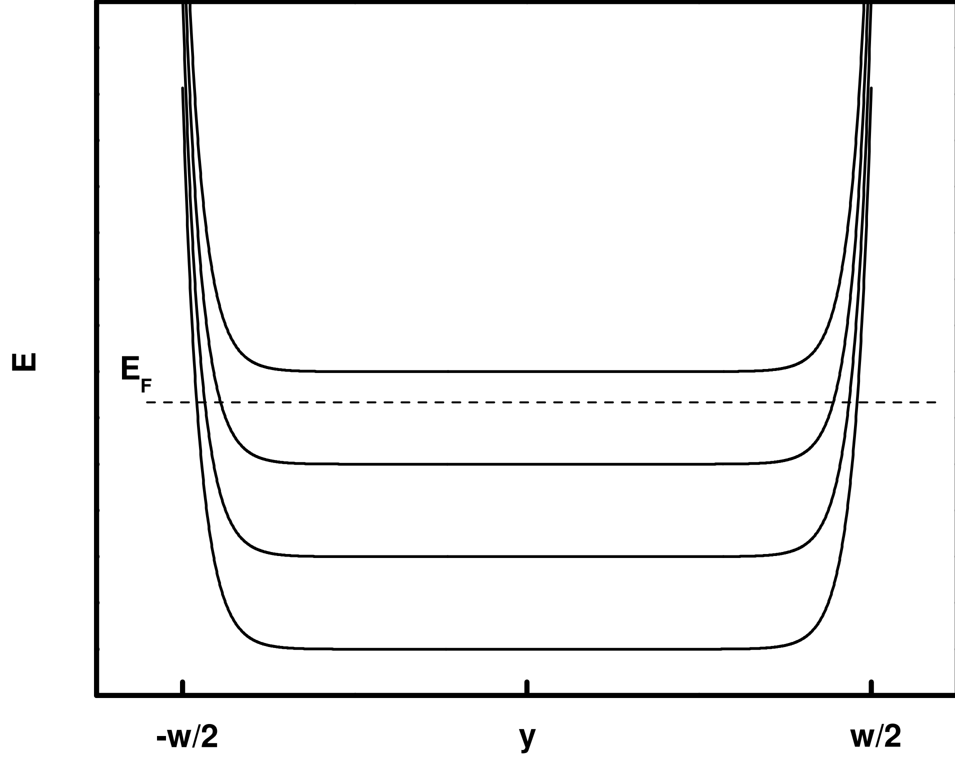

Figure 1.4. Landau level energy across a 2D Hall bar of width w.

in the suppression of backscattering (and hence resistance) so long as the edges of the sample are well-separated. For macroscopic samples the result is that R_{xx} is zero when the bulk is gapped.

Given the fact that the 4-terminal resistance R_{xx} is zero, one might expect that the 2-terminal source-drain resistance would also be zero apart, perhaps, from a small contact resistance associated with the semiconductor-metal interface in a real device. This, however, is not the case. As originally worked out by Landauer and Büttiker [6, 7, 8] and explained in a pedagogical manner in references [9] and [3], there is a finite conductance e^2/h associated with a ballistic, one-dimensional channel. This finite conductance arises due to the spatial quantization of modes in the channel. Thus, in the case shown in figure 1.4, the measured 2-terminal conductance would be $3e^2/h$ since there are 3 edge channels which cross the Fermi surface. Now due

to the fact that there is no voltage drop along the length of the sample ($R_{xx} = 0$), this means that each side of the sample is at the same chemical potential as the "upstream" source or drain contact. With one edge of the sample fixed at the source potential and the other fixed at the drain potential, the 2-terminal conductance can thus be measured in a 4-terminal arrangment by measuring the voltage drop across the sample (i.e. the Hall voltage V_{xy}). Taken together, this means that the Hall resistance R_{xy} will therefore be quantized at an integer multiple of h/e^2 whenever an integer number of LLs are completely filled and the Fermi energy lies in the gap in the bulk of the sample. An experimentally measured Hall resistance vs. B would still not look very exciting, though, since one would not be able to distinguish the quantized value of R_{xy} from the adjacent non-quantized, classical values of R_{xy}.

The somewhat surprising missing ingredient necessary to observe a universal, quantized conductance *plateau* in experiment is, in fact, disorder. In a real sample the disorder potential in the plane of the 2DEG will look like a landscape of hills and valleys. These hills and valleys serve as sources and sinks of electrons to keep an integer number of edge states filled as the LL degeneracy varies with magnetic field, thus keeping the conductance quantized over a finite range of B [10]. Said another way, the disorder gives rise to localized states in the tails of the LLs (broadened by temperature and disorder) as shown in figure 1.5. With many states localized, the measured conductance is set by the extended states. Therefore, as long as the Fermi energy lies in the region of localized states between adjacent regions of extended states, the conductance will remain quantized. As the sample disorder is increased or the temperature is decreased, the number of localized states will increase and R_{xy} will remain quantized over a larger range.

1.3 The Fractional Quantum Hall Effect

All of the features of the data in figure 1.3 can be explained using the relatively simple single-particle physics already discussed. Using this theoretical picture, one

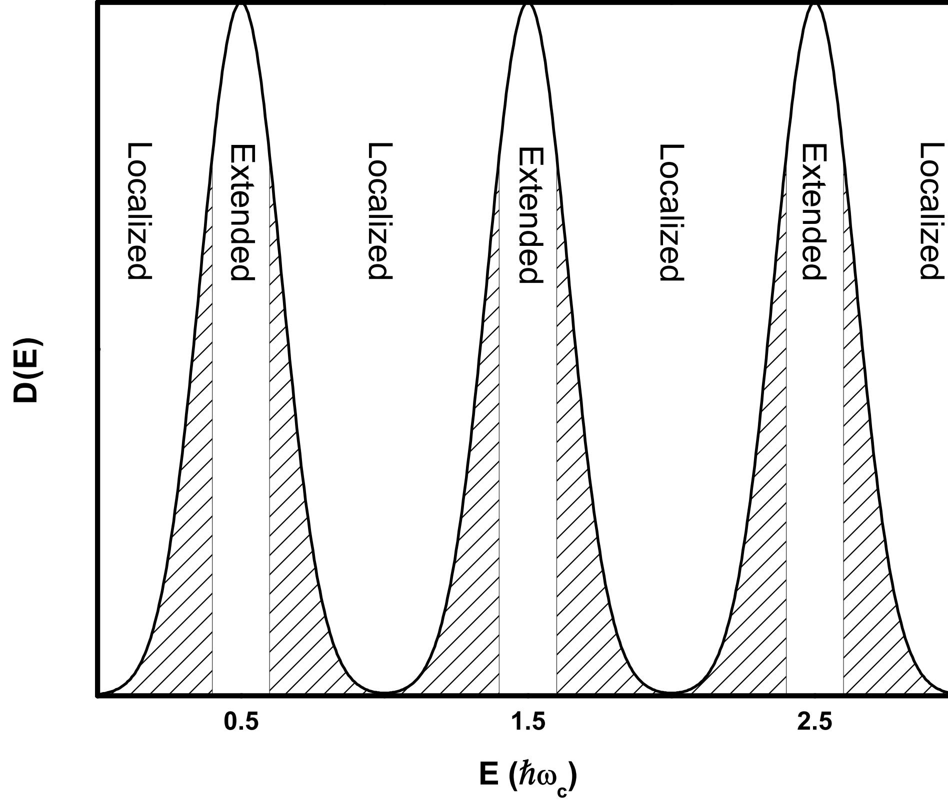

Figure 1.5. 2D DOS in the presence of a perpendicular magnetic field, finite temperature, and disorder. Cross-hatched regions indicate regions of charge localization in the 2D bulk.

would (incorrectly) expect not much interesting to happen at high magnetic field where all of the electrons occupy the lowest LL (LLL). In the LLL the kinetic energy of all the electrons is constant and equal to $\hbar\omega_c/2$ and as such is said to be "frozen". As a result, the only terms left in the electrons' Hamiltonian are interactions between the electrons and their environment (i.e. temperature and disorder) and each other. Thus, in a sufficiently clean 2D system at sufficiently low temperature one can expect interaction effects to dominate. It was this expectation that led Störmer and Tsui in 1982 to search for a so-called Wigner solid of electrons at high magnetic field. Instead, they discovered a plateau in R_{xy} and strong dip in R_{xx} at a filling fraction of 1/3 [11]. This effect, known as the fractional quantum Hall effect (FQHE), is due to

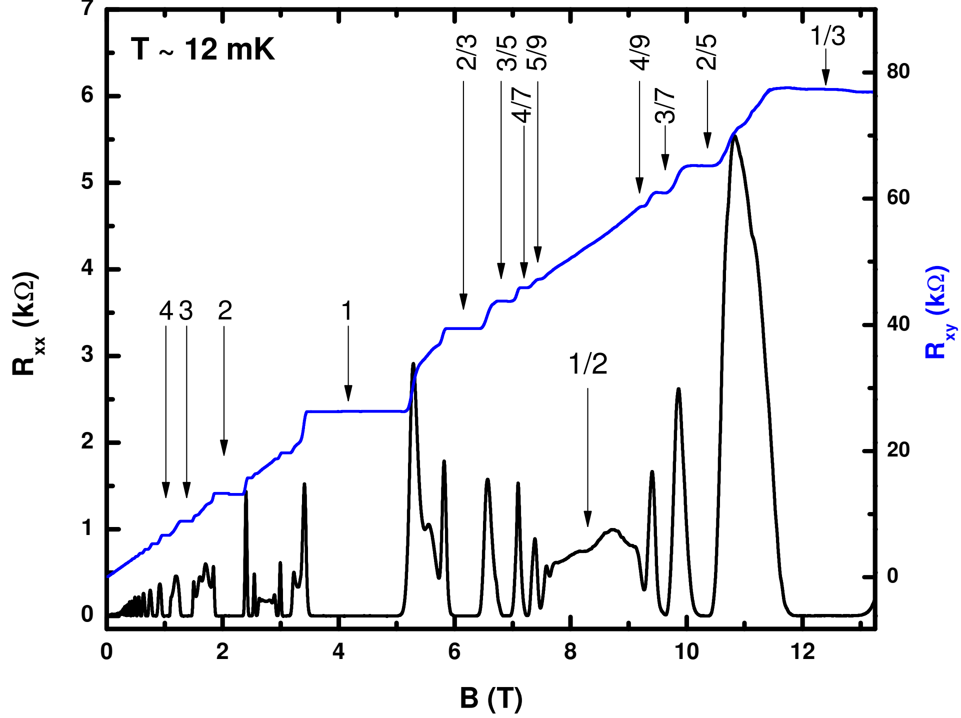

Figure 1.6. Magnetotransport from a high quality 2DEG at low temperature showing a number of FQHE states. The filling fraction is indicated for a number of fractional quantum Hal states in the lowest LL as well as for the first few integer quantum Hall states.

the opening of a new energy gap in the bulk brought about by strong electron-electron interactions as explained by Laughlin [12].

Laughlin's famous wave function describing the FQHE at primary fractions such as $\nu = 1/3,\ 1/5$, etc. is given by

$$\psi_\nu = \prod_{j<k} (z_j - z_k)^{1/\nu} \exp\left[-\frac{1}{4} \sum_l |z_l|^2 \right] \tag{1.3}$$

where $1/\nu$ is an odd integer and $z_j = x_j - iy_j$ is the position of the j^{th} electron. The requirement that $1/\nu$ be an odd integer is necessary for the wave function to be anti-symmetric under particle exchange. While this explanation of Störmer and Tsui's data was a great success (which earned Laughlin a share of the 1998 Nobel prize), it cannot by itself account for the menagerie of FQHE states observed in high

quality samples as shown in figure 1.6. In order to understand the multitude of other FQHE states in the LLL, it is necessary to invoke the composite-Fermion arguments proposed by Jain [13].

In Jain's theory, the energy of the many-body system is minimized when electrons pair with magnetic flux quanta to form new particles known as composite Fermions (CFs). These new composite particles then move through the residual, reduced magnetic field and form a new set of LLs often referred to as lambda levels (ΛLs). As their name implies, these CFs obey Fermionic statistics which requires that an electron bind with an even number of flux quanta. The simplest case is therefore an electron binding with two magnetic flux quanta. At $\nu = 1/2$ this means that all of the flux quanta will be attached to electrons and there will be no residual field left. At exactly half filling of the LLL one would therefore expect these CFs to move in straight lines analogous to electrons at zero field. This theoretical picture has in fact been confirmed experimentally through experiments (among others) involving resonant transport through anti-dot arrays [14] and extraction of the CF wavevector through surface acoustic wave transport measurements [15]. Away from half-filling, the ΛLs move through the Fermi level and R_{xx} and R_{xy} form zeros and plateaus, respectively. Indeed, if one examines R_{xx} and R_{xy} in figure 1.6 around $\nu = 1/2$, a remarkable similarity with the low field SdH oscillations and IQHE of electrons can be seen. In contrast, the extension of the edge state picture to the FQHE is not *a priori* obvious but was nonetheless explained by Wen [16, 17] who showed that the edge states in the FQHE regime are described by a one-dimensional, non-Fermi liquid known as a chiral Luttinger liquid (CLL). In this way fractions in the sequence $\nu = p/(2np + 1)$ with n and p integers can be accounted for in the CF model. This accounts for nearly all FQHE states that have been observed to date with the notable exception, among others, of the state at $\nu = 5/2$ in the 2^{nd} LL.

1.4 The Incompressible State at $\nu = 5/2$

From the first observation of an incompressible state at $\nu = 5/2$ by Willett et al. [18], it was clear that it was an "exotic" state because its even-denominator filling immediately precluded its explanation by the standard CF theory. The appearance of an incompressible state at half filling in a single quantum well is unique to the 2^{nd} LL. In the LLL at half-filling, a compressible composite-Fermi sea is formed while the higher LLs are dominated at half-filling by the formation of charge density wave states[19, 20]. Evidently in the 2^{nd} LL there is a strong competition between different families of many-body ground states with the incompressible, isotropic FQHE state favored at half filling. Indeed, in very close proximity to $\nu = 5/2$ there are re-entrant integer quantum Hall (RIQH) states [21, 22, 23] due to a bubble or liquid crystal phase in addition to more conventional odd-denominator incompressible states such as $\nu = 7/3$ and $\nu = 8/3$ and other more exotic states such as $\nu = 12/5$ [22, 24, 25] and $\nu = 2 + 6/13$ [25]. In addition to its exotic nature, the prediction that a non-Abelian state at $\nu = 5/2$ could be used as a topologically-protected qubit [26, 27] motivated intense research by a large number of groups.

Given the rather unique nature of the $\nu = 5/2$ state, it is useful to outline theoretical predictions about its nature and compare these predictions with experimental results. The most exciting candidate wave-function for $\nu = 5/2$ is the so-called Moore-Read state, named after its creators [28]. This wavefunction, often referred to as the Pfaffian, and its particle-hole conjugate the anti-Pfaffian became leading contenders for the state at $\nu = 5/2$ after being shown to have strong overlaps with exact diagonalization calculations for 2D systems with a few electrons [29, 30, 31, 32]. In the absence of LL mixing, these states are degenerate [33], but it is not clear which state is favored in the experimentally relevant case of finite LL mixing [34]. Both the Pfaffian and Anti-Pfaffian are expected to carry an effective charge $e^* = e/4$ and result from a fully spin-polarized ensemble of electrons [28]. The salient difference between the two wavefunctions is their edge structure; the Pfaffian has only down-stream propa-

gating charge modes while the Anti-Pfaffian has upstream neutral modes in addition to the downstream charge modes [34, 33]. One quantity that could be used to distinguish among different proposed wavefunctions is the Luttinger interaction parameter g. Theoretical models for the 2-terminal conductance as a function of temperature of a quantum point contact in the weak back-scattering regime in principle allow g to be extracted from experiment [35, 33, 36]. The Pfaffian state is predicted to have $g = 1/4$ while the anti-Pfaffian should have $g = 1/2$ (see reference [37] for a summary of g values for various wavefunctions). In addition, one of the easiest parameters to measure and use for comparison with theoretical models is the state's energy gap. In the absence of disorder, LL mixing, and finite width effects the gap at $\nu = 5/2$ has been calculated using exact diagonalization and density matrix renormalization group methods for small systems [29, 38]. The results of these calculations predict that the gap at $\nu = 5/2$ should be 0.03 - 0.05 in units of the Coulomb energy $e^2/4\pi\epsilon\ell_0$. For a 2DEG density of 3×10^{11} cm^{-2} this corresponds to an energy of $\sim$ 3-6 K. Finally, perhaps the most exciting experiment related to $\nu = 5/2$ would be a measurement of its quantum mechanical statistics using a two point-contact Fabry-Perot interferometer as this would definitively reveal whether $\nu = 5/2$ is a non-Abelian state of matter [39, 40].

Measurement of the effective charge e^* of the quasiparticles at $\nu = 5/2$ has been explored in measurements of the shot noise of tunneling through a quantum point contact (QPC) in the 2DEG and been found consistent with $e^* = e/4$ [41] where e is the electron charge in vacuum. Later measurements by the same group [42], though, found that the effective charge could be larger than $e/4$ for certain QPC transmission probabilities. Examination of the tunneling conductance of a QPC with a local filling fraction of $\nu = 5/2$ as a function of temperature has also been used to infer e^*. By fitting theoretical models for weak tunneling to the experimentally measured data, the best fit e^* was found to be $e^* = 0.17$ [37] in a narrow QPC, while a later measurement on the same device with a different set of depletion gates and different gate annealing parameters resulted in best fit values of $e^* = 0.25e$ and $e^* =$

0.22e for a modified QPC geometry and a quantum long contact (QLC), respectively [43]. A more recent experiment that examined tunneling at $\nu = 5/2$ found e^* to vary from 0.18e to 0.25e depending on the measurement setup [44]. Thus, while there seems to be general agreement with $e^*/e \sim 1/4$ there is still considerable variation from experiment to experiment which may suggest that the theoretical models [45, 36, 33] neglect important sample parameters that affect the value of the extracted charge.

A number of different experiments have been conducted to probe the degree of spin polarization at $\nu = 5/2$. The earliest was a tilted-field study conducted Eisenstein et al. [46]. As the sample is tilted at fixed filling fraction (i.e. fixed field perpendicular to the 2DEG), the total field and thus Zeeman energy increase. Thus, if the state at $\nu = 5/2$ required the co-existence of opposite spins, increasing the Zeeman energy would weaken and eventually destroy the state. In this early experiment, Eisenstein et al. were unable to directly measure the energy gap at 5/2 due to the relatively low-quality of their sample. Instead, they used the "deepness" of the minima in the longitudinal resistance R_{xx} as a proxy for the gap and found that the state appeared to weaken with increasing tilt angle, therby concluding that the state at 5/2 was unpolarized. However, it was later shown that the in-plane field caused a phase transition to an anisotropic stripe phase [47, 20], and thus tilted-field measurements could not be used to reliably determine the spin polarization at $\nu = 5/2$. Optical measurements utilizing resonant inelastic light scattering at $\nu = 5/2$ suggested a lack of spin-polarization, though accompanying resonant Rayleigh scattering measurements gave evidence of the coexistence of sub-micron domains of spin-polarized and spin-unpolarized fluids [48]. By contrast, measurements of the Knight shift via resistively-detected nuclear magnetic resonance [49, 50] suggested fully spin-polarized states throughout the lower spin branch of the 2$^{\text{nd}}$ LL. In addition, the measurements performed by Tiemann et al. [50] showed evidence for full spin-polarization at $\nu = 5/2$ across the entire measured density range (2.4-4.2 $\times 10^{11}$ cm^{-2}). Gap measurements at $\nu = 5/2$ as a function of density performed by Nübler et al. [51] and the Manfra group (see chapter 6) showed a monotonic dependence of activation energy on

density over the density range 1.3-3.35 $\times 10^{11}$ cm^{-2} in 30 nm quanum wells. This monotonic dependence appears to rule out a spin transition in this density range. While there is obviously still some controversy regarding the polarization of $\nu = 5/2$, it appears likely that at least in the density range producing the largest energy gaps ($\sim 2.7 - 3.3 \times 10^{11}$ cm^{-2}) the state is fully polarized.

Another theoretical prediction for some proposed wave functions at $\nu = 5/2$ is the existence of upstream neutral modes (see, for instance, reference [52] for a summary of predictions for various wavefunctions). Evidence for upstream neutral modes at $\nu = 5/2$ using shot noise measurements was first reported in reference [53], but later work by the same group showed a more nuanced picture as they found similar signatures for upstream modes at filling factors not initially anticipated to have upstream modes such as $\nu = 1/3$, $2/5$, and $4/3$ [54]. Experiments to measure the edge state temperature from the width of Coulomb blockade peaks in quantum dots located in close proximity to the edge states [55] were evidently unable to couple both the heater and thermometer to edges other than the outermost one. As a result, they were unable to make any claims about heat carried by upstream neutral modes at $\nu = 5/2$. However, Venkatachalam et al. [55] did find evidence for upstream neutral modes on the high-field side of $\nu = 1$ for gate-defined edges which points to edge state reconstruction for the shallow confining potential from the gate. Such edge reconstructions could explain the presence of edges modes seen in reference [54] in states not expected to have upstream modes for sharp confining potentials.

Returning to the tunneling experiments discussed earlier, yet another measurable quantity that differs among the various proposed wavefunction is the Coulomb interaction paramter g, also known as the Luttinger interaction parameter. The first edge state tunneling experiment at $\nu = 5/2$ by Radu et al. [37] found the best agreement with the predictions of the non-Abelian anti-Pfaffian and $U(1) \times SU_2(2)$ states. Later tunneling experiments with different gate configurations and preparations [43] on the same Hall bar and experiments by a different group [44] found results consistent with the Abelian (331) and (113) states. However, the fact that Radu et al. [37] and Lin

et al. [43] found different answers when using different gate preparations and configurations on the same Hall bar seems to suggest that such details may need to be included in the tunneling theory to extract meaningful results from the measurements. It could potentially be interesting to compare such tunneling experiments with similar experiments with the QPC gates deposited in trenches to create a sharper confining potential in an attempt to minimize edge state reconstruction.

One of the most frequently measured quantities in quantum Hall experiments is the energy gap of the incompressible states. As mentioned previously, theoretical predictions for the gap at $\nu = 5/2$ based on numerical calculations result in relatively large gaps of a few Kelvin. However, the largest gaps measured to date at $\nu = 5/2$ are < 600 mK [56]. The discrepancy between the predicted and measured gap is ascribed to a phenomenological disorder broadening Γ such that $\Delta_{int} = \Delta_{meas} + \Gamma$ where Δ_{int} is the intrinsic gap in the zero-disorder limit (i.e. the theoretically calculated gap) and Δ_{meas} is the experimentally measured gap. Due to the large discrepancy between the experimental and theoretical gaps, Γ is evidently quite large ($\mathcal{O}(1K)$ [57]). However, by examining the gap at filling fractions 5/2, 7/2, 7/3, and 8/3, Samkharadze et al. [57] were able to estimate the disorder broadening in a series of samples. Further, by extrapolating the intrinsic gap to the limit of zero LL mixing, they found an intrinsic gap of ~ 0.032 in units of the Coulomb energy which compares very well with the numerical results discussed previously [29, 38]. Nübler et al. [51] also found that inclusion of LL mixing in theoretical calculations of the gap at $\nu = 5/2$ had a large impact on the intrinsic gap. Therfore, LL mixing is evidently an important parameter to include in trying to reconcile theoretical predictions and experimental observations at $\nu = 5/2$.

Far and away the most convincing demonstration of the non-Abelian nature of the $\nu = 5/2$ state would be a measurement of braiding statistics in a quasiparticle interferometer. Figure 1.7 shows a sketch of the layout of a so-called Fabry-Perot interferometer. In such a setup the 2DEG in a Hall bar is depleted under electrostatic gates to define two constrictions. Tunneling of quasiparticles between the edges at

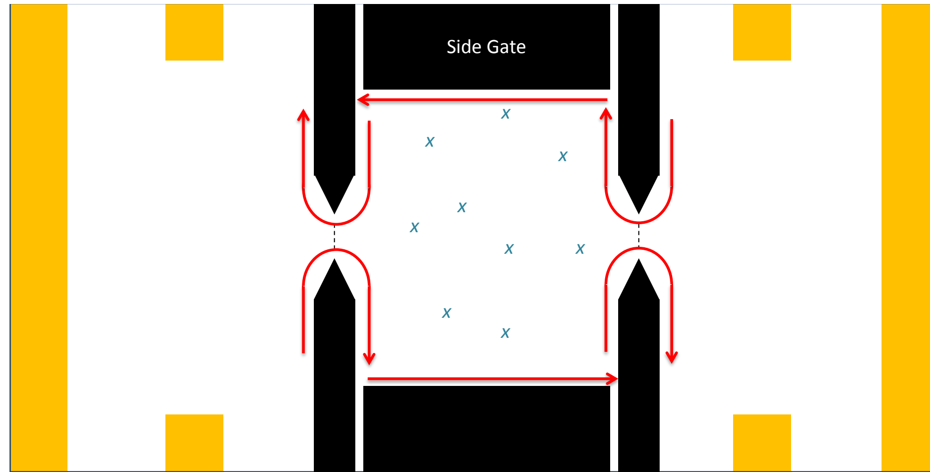

Figure 1.7. Sketch of the geometry of a Fabry-Perot interferometer in the quantum Hall regime. The 2DEG is bound by the blue rectangle, the yellow regions represent Ohmic contacts, the black shapes represent surface gates used to electrostatically deplete the 2DEG underneath them, red lines represent the flow of edge states, and blue x's represent localized quasiparticles. As the side gate voltage is varied the number of quasiparticles encircled by the edge states changes, giving rise to oscillations in the conductance of the interferometer due to the change in phase of the edge states.

these constrictions allows the edge quasiparticles to interfere with themselves. This then modifies the 2-terminal conductance of the device from its quantized value. The interference is determined both by the Aharanov-Bohm phase resulting from the encircling of magnetic flux quanta in the interior of the device as well as a phase resulting from the encircling (i.e. exchange) of localized quasiparticles in the bulk of the interferometer. If the area of the interferometer is changed (e.g. by tuning the side gate voltage), both the number of encircled quasiparticles and the number of encircled flux quanta change. It has been theoretically predicted [33] that if the state at $\nu = 5/2$ is either of the non-Abelian Pfaffian or Anti-Pfaffian wavefunctions, the 2-terminal conductance of the device should oscillate if an even number of quasiparticles are encircled while the interference effect should be lost if an odd number of quasiparticles is encircled. Thus, for sufficiently large side gate voltage sweeps,

one would expect to find an alternation between oscillations and no oscillations in the 2-terminal conductance as a function of the side gate voltage. Experiments with Fabry-Perot interferometers aimed at examining $\nu = 5/2$ have been conducted by a number of groups [58, 59, 60, 61] but have had mixed results. The experiments by Willett et al. [58, 59] showed an alternation of the frequency of oscillation as a function of side gate voltage with the two frequencies suggesting effective charges of $e/4$ and $e/2$. The results from the Chicago group [60] showed "phase slips" (possibly due to the entrance/exit of quasiparticles from the interferometer) at $\nu = 5/2$ and $\nu = 7/3$ consistent with a non-Abelian state at 5/2 and an abelian state at 7/3. However, the lack of quantization in the diagonal resistance through the device and the very poor quality of the bulk transport away from the interferometer cast some doubt on whether the observed phase slips were indeed due to quasiparticle interference. Finally, the Harvard group [61] was not able to observe oscillations at $\nu = 5/2$, though they did find oscillations consistent with a an effective charge of e at integer filling and $e/3$ at fractions $\nu = 1/3$, $2/3$, $4/3$, and $5/3$. At this point more work, both theoretical and experimental, is likely needed to understand the impact of factors such as device design and confining potential strength on the observed interference patterns.

2. Molecular Beam Epitaxy

2.1 Principles of Molecular Beam Epitaxy

Molecular beam epitaxy (MBE) is a highly controllable physical vapor deposition technique performed in an ultra-high vacuum (UHV) environment. The MBE growth technique is relatively simple conceptually. Atomic and/or molecular fluxes are generated thermally in furnaces known as effusion cells or Knudsen cells and are modulated by shutters. Due to the UHV enivronment and relatively low beam equivalent pressures (BEPs) of the sources, there is no gas-phase reaction as there would be in a chemical vapor deposition (CVD) process. The crystal growth rate is set by the BEP of each source, and the BEP is controlled by the effusion cell temperature. Growth rates can be set low enough to allow superlattice structures to be grown with thicknesses as small as a few monolayers (ML). Studying the physics of low dimensional systems requires spatial confinement on the scale of a few to tens of nanometers in addition to minimal disorder from crystal defects and impurities. The tight control of layer thickness coupled with the high crystal purity made possible by the UHV environment thus make MBE an ideal growth method for examining mesoscopic physics.

2.1.1 MBE Chamber Layout

Figure 2.1 illustrates a cross section of the basic design of an MBE growth chamber. The vacuum vessel is typically constructed from 316L or 304 stainless steel and electropolished to minimize the surface area and potential adsorption sites for gases. In a UHV environment the chamber surface can be thought of as a sponge that adsorbs gases when the chamber is vented and subsequently slowly releases these gases

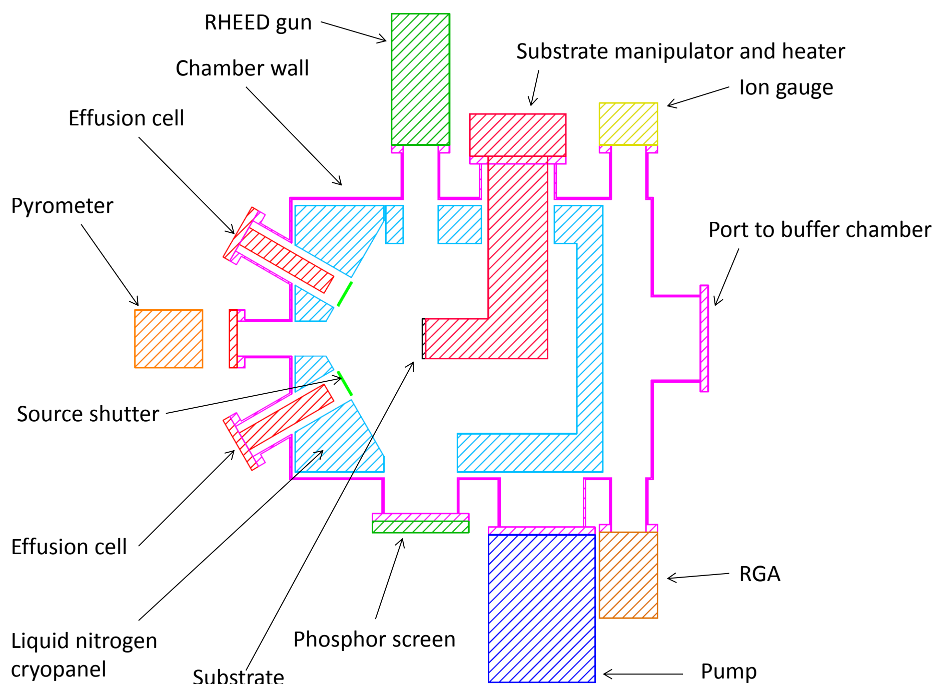

Figure 2.1. Cross sectional sketch showing the functional form of a typical MBE chamber.

when the chamber is under vacuum. Thus, minimizing the chamber surface area through electropolishing and minimizing the total size of the chamber is important for achieving a low base pressure. The desorption of gases from the surface can be greatly accelerated by elevating the temperature, and thus baking the chamber is a commonly used technique to quickly drive gases into the pumps to more quickly reach the UHV regime upon cooling the chamber.

Due to the difficulty in achieving UHV, it is necessary to include a load lock (LL) chamber and sometimes an additional buffer chamber so that the growth chamber is not exposed to air each time wafers are loaded and unloaded. New wafers are typically loaded into a small LL chamber which can be pumped down quickly, and the wafers are subsequently transferred into a buffer chamber where they can be heated to desorb

resdiual adsorbed gases. This approach minimizes the number of gaseous impurities introduced into the growth chamber by daily operation.

2.1.2 Pumping Methods

In addition, achieving UHV conditions requires specialized vacuum pumps. The most commonly types of pumps used in MBE systems are turbo pumps, ion pumps, titanium sublimation pumps (TSPs), and cryo pumps. Turbo pumps compress and exhaust gases with a set of fan blades spinning at high speed ($\sim$ 20,000 RPM). The fan blades and housing unit are machined to very tight tolerances to enable the pump to function without an oil seal between the blades and housing. The lack of oil is critical in UHV applications since any oil backstreaming from the pump into the chamber would compromise the vacuum quality. Turbo pumps have very high pumping speeds and exhaust the pumped gases to the room. This makes them especially useful for pumping LL chambers since the large gas load arising from frequent venting is continuosly removed from the pump. However, the use of turbo pumps in deposition chamber requires oil-free, magnetic-levitation designs to reduce the chance of oil back-streaming. Even with such oil-free designs, however, having such a large number of high-precision, moving parts internal to the vacuum raises reliability concerns. To get around this problem, deposition chambers are often pumped with ion pumps, TSPs, and cryo pumps which have no moving parts internal to the vacuum.

Ion pumps function by ionizing gas particles with magnetically confined, energetic electrons. The ions then either react with a titanium getter or are buried on the wall of the pump by sputtered titanium. Ion pumps are thus a type of entrapment pump since the gases are not exhausted into the room. However, if the pump loses power the pumped gases are not released back into the chamber. This means that there is no requirement for an expensive, high conductance valve between the pump and the chamber. Ion pumps, however, typically have a fairly low pumping speed ($<$ 300-500

L/s N_2 at 10^{-6} mbar [62]) and are thus not suitable as the only pump on a chamber with a large gas load.

Titanium sublimation pumps are another simple type of entrapment pump. TSPs operate by heating a filament or charge of titanium to sublimate pure titanium onto a surface typically cooled by water or liquid nitrogen. The fresh titanium is extremely reactive and thus getters reactive gases effectively; hydrogen, for instance, can be pumped at a rate > 1000 L/s throughout the UHV regime for a TSP with a 1-inch^2 water-cooled pumping surface; cooling the pumping surface with liquid nitrogen more than triples this pumping speed [62]. Like ion pumps, the pumped gases cannot escape back into the vacuum chamber, so valves are not necessary. TSPs do not pump all gases well, though. Noble gases in particular are not pumped effectively by a TSP, so a TSP is not an effective stand-alone pump.

Finally, cryo pumps act by cooling a surface with a large surface area (typically activated charcoal) to ~ 10 K. If a gas particle lands on the surface of the cold head, it will not have sufficient thermal energy to desorb and will thus stay adsorbed on the surface. Cryo pumps can have very high pumping speeds; for instance a cryo-pump mounted on a 14-inch flange can have have water-vapor pumping speeds as high as 9000 L/s [63]. In addition, cryo pumps have no moving parts internal to the vacuum which make them ideal for deposition chambers. However, if the pump loses power, the adsorbed gases will desorb and contaminate the chamber. A high conductance valve is thus needed in between the pump and the chamber so that the pump can be valved off for maintenance and/or in the case of an emergency. Cryo pumps also tend to be expensive, the pump and requisite compressor have moving parts, and the compressor typically requires a high voltage power source and water cooling. These factors make cryo pumps somewhat prone to failure and make fail-safe mechanisms such as uninterruptible power supply (UPS) systems harder and more costly to implement. Due to the various advantages and disadvantages of each of these types of pumps, many MBE chambers utilize more than one type of vacuum pump.

Returning to figure 2.1, an MBE chamber also usually has cooling shrouds around the sources and the main body of the growth chamber. The shrouding around the sources is necessary to eliminate thermal cross-talk between the cells and also to prevent damage to the stainless steel wall that is in close proximity to the hot effusion cells. This source shroud can be cooled with liquid nitrogen or in some cases coolant circulated through a closed-circuit chiller. Some sources with particularly high power requirements may also have built-in water cooling jackets to prevent fragile parts of the source from overheating. The shroud around the body of the growth chamber, called the growth shroud, is cooled with liquid nitrogen and functions as a large cryo pump.

The mean free path λ of gases in vacuum is given by [64]

$$\lambda = \frac{1}{\sqrt{2}\pi d_0^2 n} \tag{2.1}$$

Where d_0 is the molecular diameter and n is is the gas density. Assuming a diameter ~ 5 Å and the upper pressure bound of the UHV regime $\sim 10^{-9}$ Torr [65] this correponds to a mean free path $\sim 3 \times 10^4$ m, somewhat larger than typical MBE chambers. As a result, gas particles follow straight trajectories when they desorb from the chamber wall until they collide with another surface in the chamber. Thus, in order for impurities to reach the substrate they must desorb from a surface with a direct line of sight to the substrate. Impurity incorporation into the growing film can, therefore, be greatly reduced by keeping surfaces close to the substrate cooled to liquid nitrogen temperatures since the desorption rate from a 77 K surface will be reduced exponentially compared to that of a room temperature surface.

2.1.3 Vacuum Analysis

Quantitative analysis of the vacuum quality is accomplished most frequenly with two types of instruments: ion gauges and residual gas analyzers (RGAs). Ion gauges are used to measure the total pressure in a chamber. As shown in figure 2.2, electrons are produced by thermionic emission from the hot filament and accelerated towards

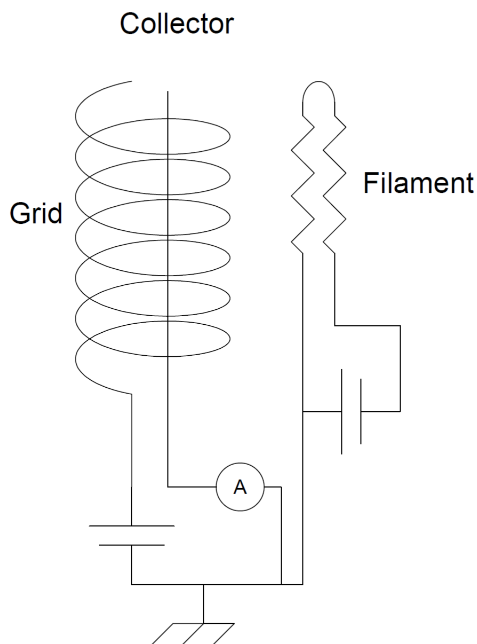

Figure 2.2. Schematic of the operating principle of an ion gauge. Electrons are produced by the filament and accelerated to the grid. Ions are accelerated to the collector where they produce a measurable current.

the grid by a large electric field. The energetic electrons ionize gas particles which are subsequently drawn to the collector electrode. The ion current is then easily measured and is proportional to the pressure of ionized gases. Since different gases have different ionization rates, the gauge must be calibrated for a specific gas (usually N_2); the ion gauge reading is, therefore, not an exact measure of the total pressure in the chamber. However, the pressure measured by an ion gauge is still a useful point of reference for qualitatively understanding the state of the vacuum. The minimum detectable pressure is set by the so-called x-ray limit which depends on the specific design of the gauge. The energetic electrons produce x-rays when they impact the grid, and these x-rays in turn produce a collector current even in the absence of any ionized gases. Standard ion gauges are usually x-ray limited in the low 10^{-11} Torr range, but specialized gauges[1] can measure the pressure as low as 10^{-13} Torr.

[1]For instance, the Ionivac Extraktor IE514 gauge used on the Manfra group GaAs MBE

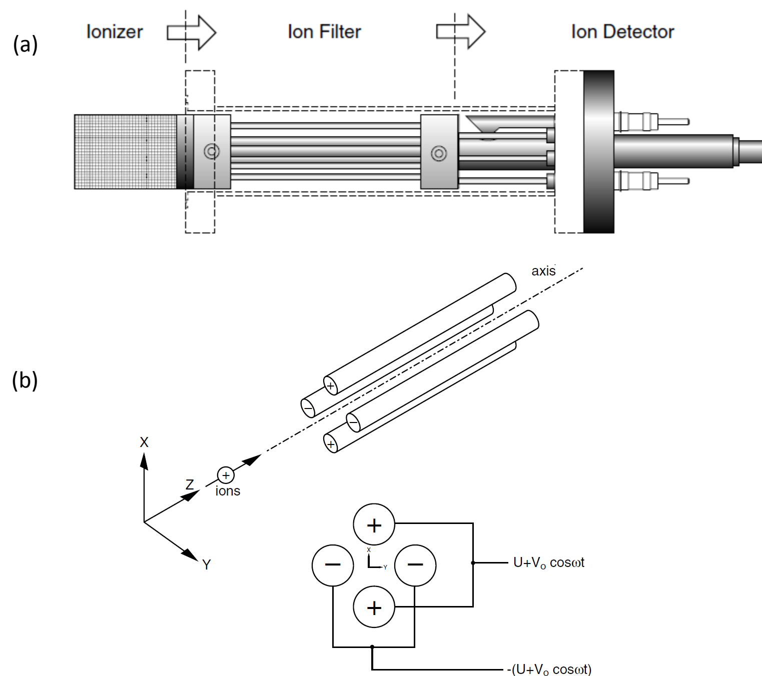

Figure 2.3. Sketch of an RGA. (a) Side view of a typical RGA showing the three main sections and direction of ion flow. (b) Illustration of the electrodes forming the quadrupole mass filter. Figure reproduced from Operating Manual and Programming Reference - Models RGA100, RGA200, and RGA300 Residual Gas Analyzer, copyright (1996), with permission from Stanford Research Systems.

By contrast, RGAs are used to map out a broad partial pressure spectrum of the gases in the vacuum; a schematic of an RGA is shown in figure 2.3. The gas particles are first ionized with energetic electrons, and the resulting ions are next accelerated down the center of four electrodes as shown in figure 2.3b. The electrodes are used to generate a quadrupole electric field which can be tuned in-situ to allow stable trajectories for a specific ion mass-to-charge ratio. The ion current is measured by a Faraday cup detector or amplified by an electron multiplier and converted to a partial pressure reading. The stable trajectory is subsequently swept, allowing the instrument to map out a large spectrum of partial pressures. Figure 2.4 shows RGA spectra in

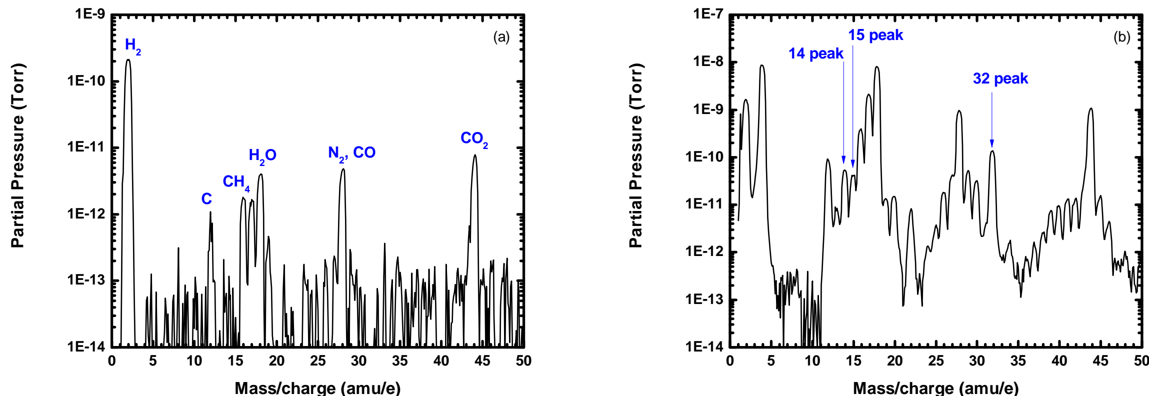

Figure 2.4. RGA spectra. (a) RGA spectrum from a well-baked and leak-free chamber. (b) RGA spectrum from a leaking, unbaked chamber. The leak is evinced by the large peak at 32 and the fact that the 14 peak is larger than the 15 peak which together indicate the presence of O_2 and N_2 leaking into the chamber.

two limiting cases. Figure 2.4a is taken from a well-baked, leak-tight chamber. The only residual gases in the chamber are called out, with hydrogen being the dominant partial pressure. Figure 2.4b, on the other hand, shows an RGA spectrum taken from a chamber shortly after assembly. There are many more peaks visible in figure 2.4b, and the tell-tale signature of a leak is present: the mass 14 peak is larger than the mass 15 peak, and the mass 32 peak is large. The large 14 peak relative to the 15 peak indicates the presence of nitrogen, and the 32 peak is due to atmospheric oxygen. Beyond revealing the presence of a leak, the RGA is vital to finding the leaking vacuum joint. To find a leak, the RGA is set to detect a mass:charge ratio of 4 and helium is sprayed around each seal. If the seal is leaking, the RGA will show an increased signal as the helium leaks into the chamber.

2.1.4 The RHEED Technique

In order to grow heterostructures for studying mesoscopic physics, it is important to have tight control over the crystal growth rate. The MBE growth method is at an

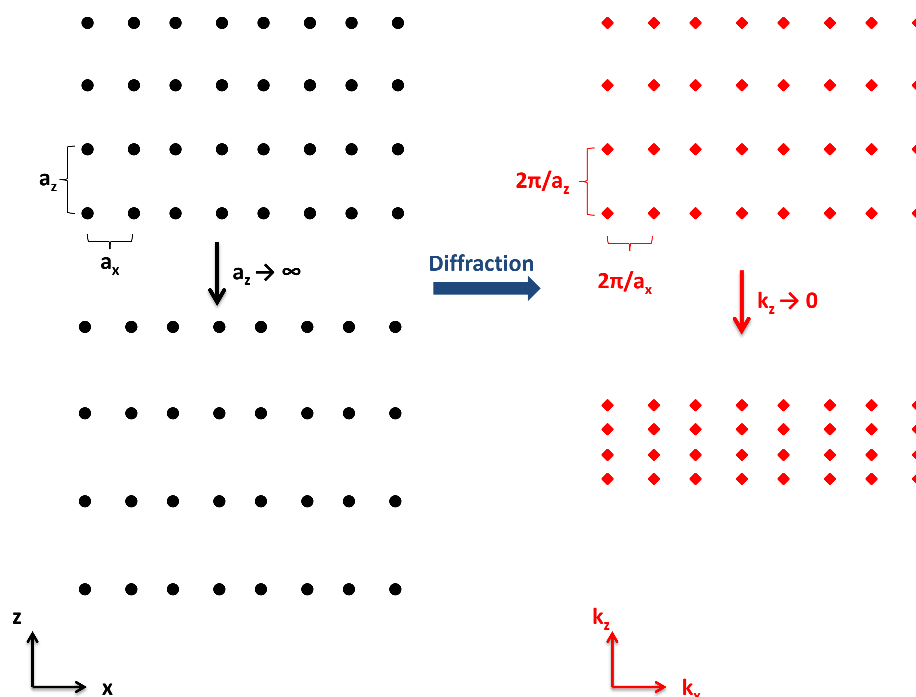

Figure 2.5. Sketch of the effect of increasing the lattice constant along the z direction. As the real space lattice constant increases, the corresponding reciprocal lattice constant decreases. In the limit of a 2D crystal in the x-y plane, the reciprocal lattice would consist of a series of thin rods parallel to the k_z axis.

advantage in this regard over growth processes occuring at higher pressures. Due to the UHV environment and the associated long mean free path, reflection high energy electron diffraction (RHEED) can be used to measure the crystal growth rate in-situ. In the RHEED technique a high energy beam of electrons ($\sim$ 10 keV) impinges on the substrate at a glancing angle. Due to this shallow angle, the electrons scatter off only the first few atomic layers, thereby generating a quasi-2D diffraction pattern which can be viewed on a phosphor screen. In trying to visualize the diffraction pattern from a 2D surface, it is useful to think of the 2D crystal as the limit of a 3D crystal with the spacing between atomic planes along the z direction tending to infinity. Since the diffraction pattern is essentially a Fourier transform of the real

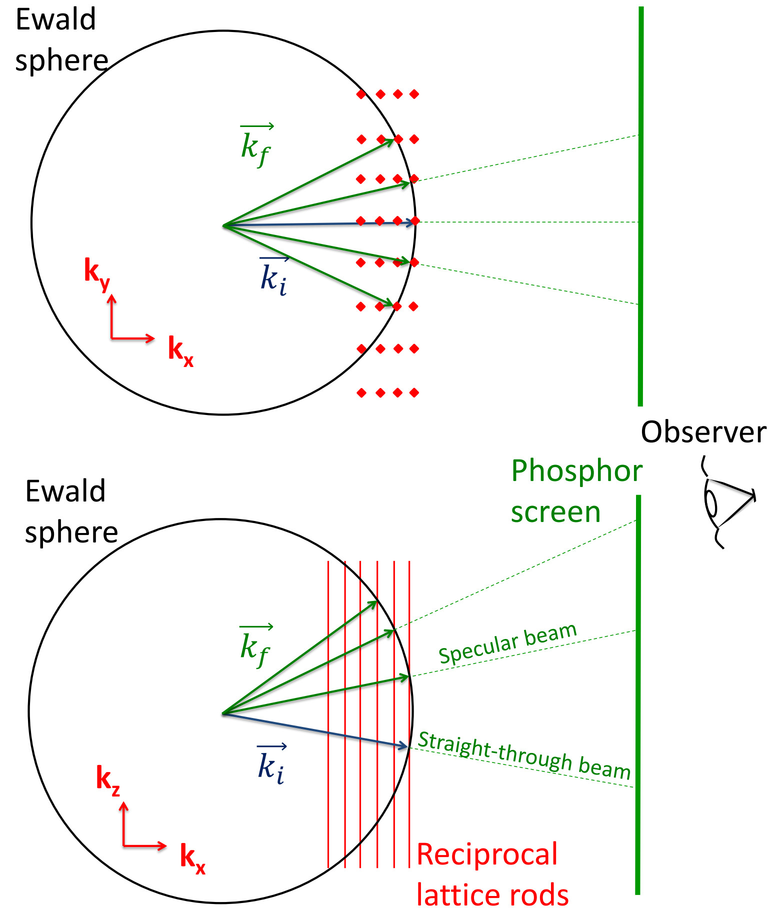

Figure 2.6. Cross sectional views of the Ewald sphere construction used to determine momentum- and energy-conserving scattering processes from a 2D crystal.

space crystal, the reciprocal lattice spacing along z will collapse to zero as the real space lattice spacing is taken to infinity as shown in figure 2.5. This means that the reciprocal lattice of a 2D crystal in the x-y plane would consist of a series of rods, broadened by disorder, parallel to the k_z axis. Since the electrons will scatter off the atoms elastically, their initial and final wave vectors $\vec{k}_i$ and $\vec{k}_f$ must be of equal magnitude. In addition, in order to conserve momentum, the initial and final wave vectors must differ by a reciprocal lattice vector $\vec{G}$. These conditions can be visualized using the Ewald sphere construction. Figure 2.6 shows 2D cross-sectional

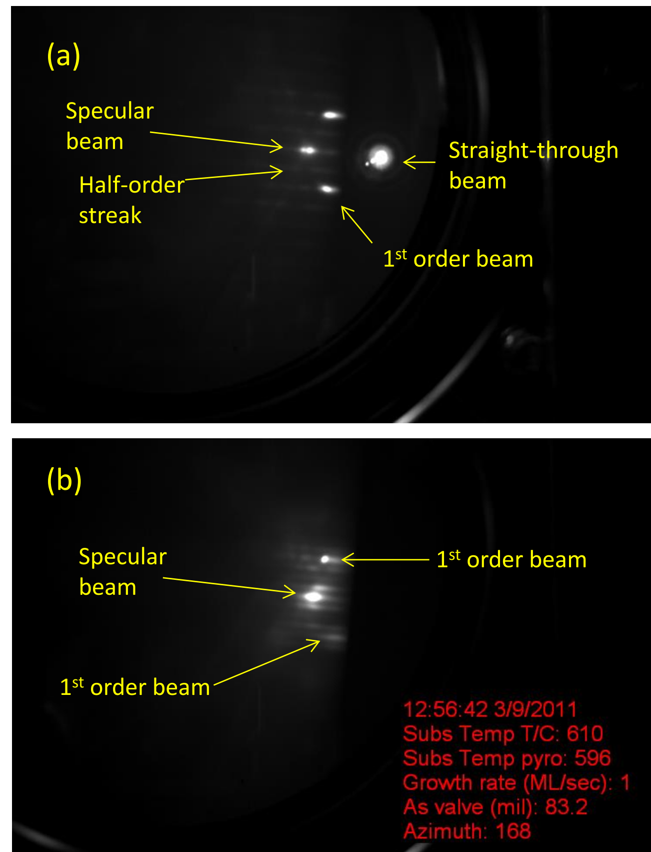

Figure 2.7. Typical RHEED diffraction pattern in GaAs at growth temperature. (a) RHEED pattern showing the 2× reconstruction. (b) RHEED pattern showing the 4× reconstruction.

views of the Ewald construction with a 2D crystal. For an incident momentum $\vec{k}_i$ the electron's kinetic energy and momentum will be conserved if its outgoing momentum vector $\vec{k}_f$ simultaneously lies on the surface of the Ewald sphere and intersects a reciprocal lattice rod. The trivial case of no scattering is also a possibility, and this results in the so-called straight-through beam. The brightest beam apart from the straight through beam is typically the specular beam which is simply the reflection of the incident beam from the surface. Figure 2.7a shows a typical diffraction pattern from GaAs showing the specular and straight-through beams along with the first- and half-order streaks in the first Laue zone. The half order streaks are due to the

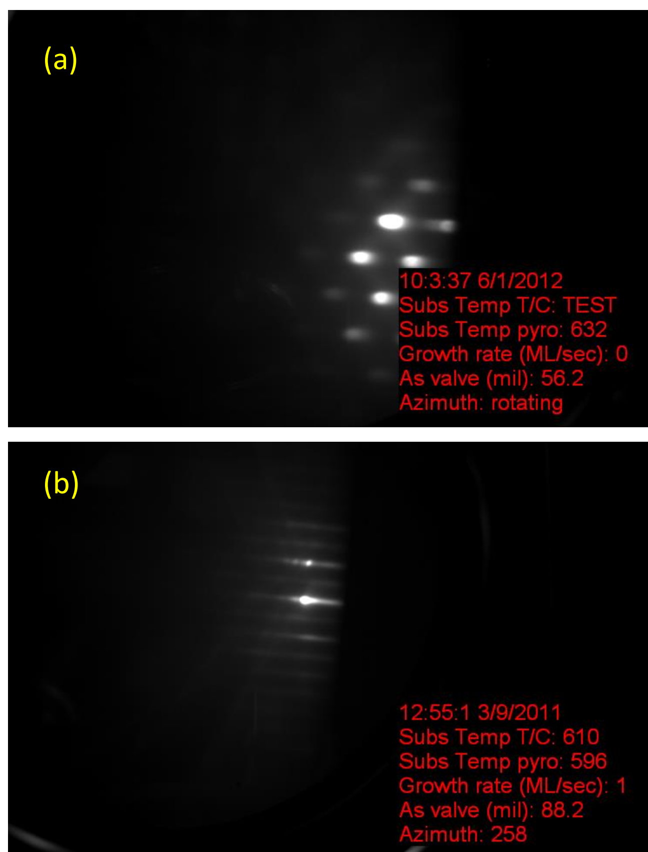

Figure 2.8. (a) RHEED pattern from a rough wafer immediately after desorbing the oxide at growth temperature. (b) RHEED pattern from a smooth wafer.

reconstruction of the surface which in this case doubles the periodicity along the RHEED beam. The orthogonal direction shown in figure 2.7b has a periodicity four times that of the bulk and thus has quarter-order streaks in addition to the primary streaks.

In addition, an important feature used for qualitative analysis of the wafer surface is the degree of “streakiness” of the diffraction spots. The diffraction pattern from a very smooth surface will exhibit long streaks perpendicular to the wafer surface on the phosphor screen. This is because for typical RHEED energies and scattering geometries, the finite width of the reciprocal lattice rods will intersect the Ewald

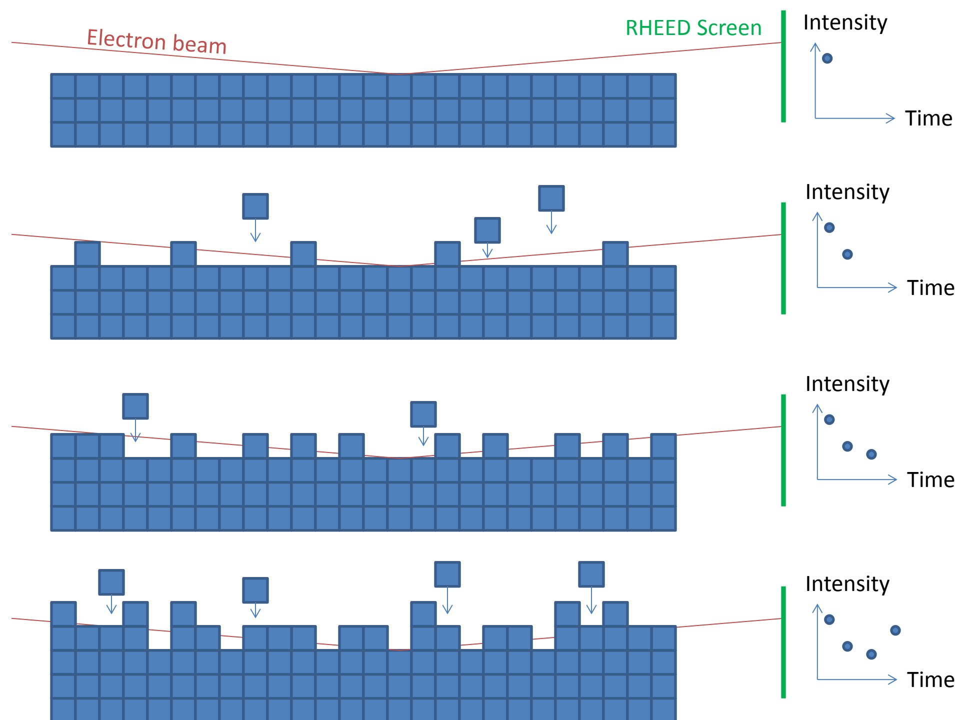

Figure 2.9. Sketch of crystal surface in cross section during growth. As the surface roughens due to the nucleation of a new monolayer, the specular intensity drops. Once the new atomic layer nears completion, the intensity recovers.

sphere over a large range along their length. A rough surface, which can be thought of as more “3D”, will instead show a diffraction pattern with well-defined spots as the spacing along z-direction of the reciprocal lattice planes becomes resolvable. Figure 2.8 contrasts the diffraction pattern from a rough surface in panel (a) to the diffraction pattern from a smooth surface in panel (b).

Finally, RHEED patterns can also be used to measure the growth rate by monitoring the intensity of the specular beam as a function of time. As figure 2.9 illustrates, the surface roughness of the growing crystal will oscillate in time as atomic monolayers nucleate and spread. The intensity of the specular beam is proportional to the surface roughness and, therefore, also oscillates, with each peak in reflected intensity corresponding to the completion of a single monolayer. Figure 2.10 shows intensity

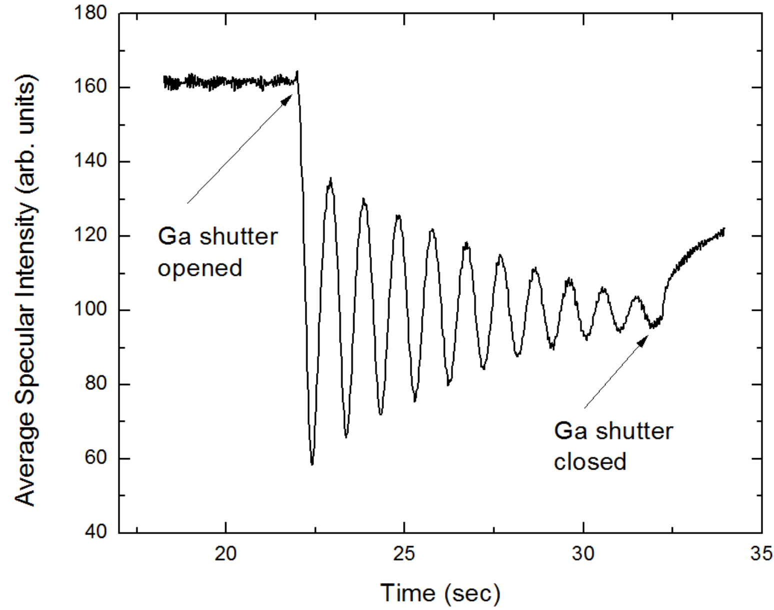

Figure 2.10. Typical oscillations in the intensity of the RHEED specular spot during growth of GaAs. The oscillations are damped due to roughening of the surface as islands nucleate on top of other islands. The oscillation period corresponds to the crystal growth rate.

oscillations as a function of time taken from a growing GaAs surface. By measuring the oscillation frequency one is able to accurately measure and adjust the growth rate as necessary prior to growing a heterostructure.

2.2 Design Considerations for High Mobility MBE

The preceding section described the basic features of a generic MBE system. In order to grow the low-disorder structures necessary for modern research in correlated electrons in GaAs, however, additional design constraints must be taken into account. Achieving and maintaining high electron mobilities requires that two guiding principles be considered during the machine design. First, every component in the growth chamber must be designed to maximize vacuum purity and minimize power input to the system. Second, the machine should be designed with as much redundancy and as

many back-up mechanisms as possible to extend the length of the growth campaign and minimize the impact of equipment failures on chamber purity.

The MBE design was based on the Varian High Mobility GenII system which was chosen as it is one of the few remaining 2 inch MBE systems commercially available. By contrast, most current research MBE machines are designed to grow on 3 or 4 inch wafers. However, it has been shown [66] that a larger substrate heater with its correspondingly larger power dissipation results in significantly lower electron mobilities. Furthermore, while reference [66] showed that a larger substrate heater in a given growth chamber decreases the maximum achievable mobility, this study did not take into account an additional source of mobility degradation from using a larger system. In order to achieve acceptable epilayer uniformity on a larger wafer, the substrate must be situated further from the effusion cells. This, in turn, requires larger effusion cells which dissipate more power to maintain the same growth rate. Therefore, from the standpoint of minimizing power input to the system, a smaller machine can be expected to perform better than a large machine.

Besides the 2 inch substrate, the High Mobility GenII design was chosen because it utilized all-metal components on the growth chamber to enable the entire growth chamber to be baked at 200 °C for extended periods. The primary difference between the standard and "High Mobility" designs is that the gate valves on the "High Mobility" growth chamber utilize metal sealing surfaces rather than commonly used rubber seals such as Viton. The rubber seal cannot be baked as hot as a metal seal can, particularly if the valve is closed, and this limits the clean-up of the machine. In addition to using all-metal gate valves, the gate valves on the Manfra high-mobility MBE were custom-designed to allow the full 10 inch opening of the CT-10 cryopump to see the growth chamber and take full advantage of the large pumping speed of the cryopump. Figure 2.11 shows the Manfra high-mobility MBE with its three cryopumps and custom gate valves. In addition to the growth chamber itself being bake-able, the cryopumps were also modified from the standard design. First, the rubber pressure relief valve on the back of each pump was replaced with an all-metal

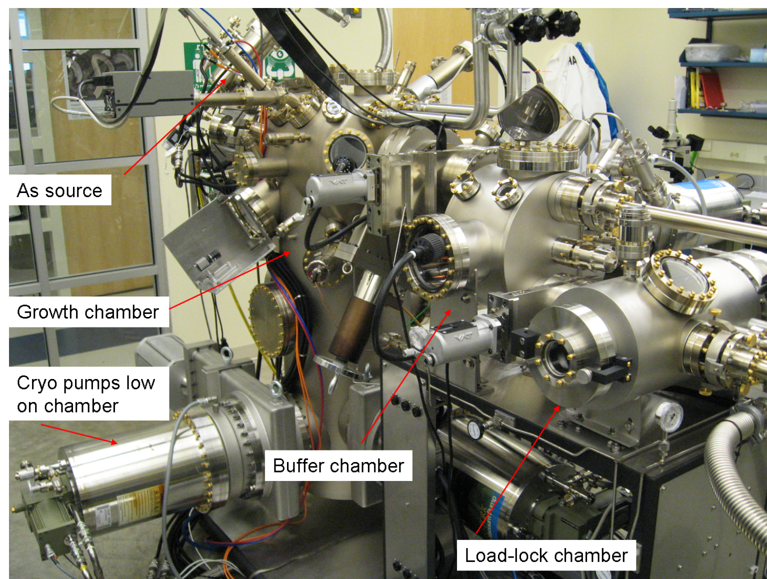

Figure 2.11. The Manfra high mobility MBE machine as installed. The design and materials of construction play an important part in determining the ultimate crystal purity (see text).

burst disk to improve the vacuum-tightness of the pump. Naively one might expect that a small leak on the backside of a fast pump such as a cryopump would not impact the vacuum quality or even be noticeable. However, figure 2.12 shows the RGA spectra from a different chamber in the lab before and after replacing the rubber poppet valve with an all-metal burst disk. This leak was very difficult to find since it required that the leak-checking first saturate the pump with helium before the helium signal was visible on the RGA, but removing the poppet valve nonetheless had a clear, positive impact on the vacuum quality. Second, the pumps themselves were modified to allow the outer vacuum can of the pump to be baked. This involved removing plastic components inside the pump and adding water cooling rings to the mouth of each pump to help manage the heat load on the pump during bake-outs. Baking the vacuum can of the pumps was shown to be important for achieving high mobilities by Pfeiffer and West [67]. As they pointed out, if as little as 1% of the vacuum surface is left unbaked, that surface will be the dominant source of outgassing in the chamber.

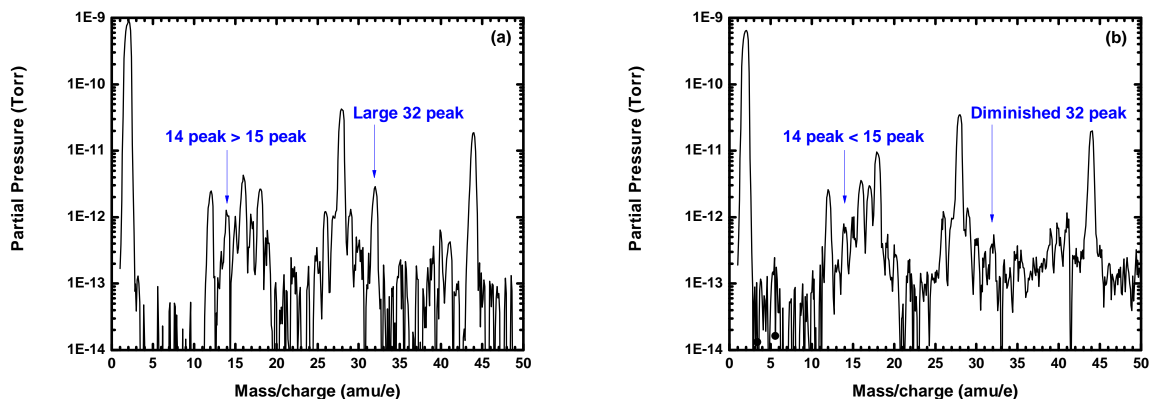

Figure 2.12. RGA spectra showing the impact of replacing the rubber-sealed poppet valve on the chamber's cryopump with an all-metal burst disk. (a) RGA spectrum from chamber with poppet valve. The spectrum shows the tell-tale sign of a leak: the mass 14 peak is larger than the mass 15 peak and the 32 peak is large relative to the 28 peak. (b) RGA spectrum after replacing the poppet valve with an all-metal burst disk. The leak signature is gone, though the the chamber still required additional baking to improve the overall vacuum purity.

Figure 2.13 [68] shows the result of various improvements to MBE technology, source material purity, and heterostructure design. The jump in maximum mobility between 1986 and 1988 ("English et al." data and "Our data" in plot) was a result of the installation of all-metal gate valves and bake-able cryopumps described in reference [67].

In order to achieve high mobilities, the choice of pumps was limited from what is often used in more standard MBE systems. First, no turbo pumps or other mechanical roughing pumps were installed on any of the MBE chambers due to the potential risk of particulate generation and/or oil back-streaming from nominally "dry" pumps. This left liquid nitrogen cooled sorption pumps as the only option for roughing pumps. In order to maintain high throughput, it was necessary to utilize cryo pumps (albeit non-bake-able designs) on the LL and buffer chambers because ion pumps would not have sufficient pumping speeds for chambers with such large gas loads. Finally, although the standard GenII design uses a combination of cryopumps, ion pumps, and

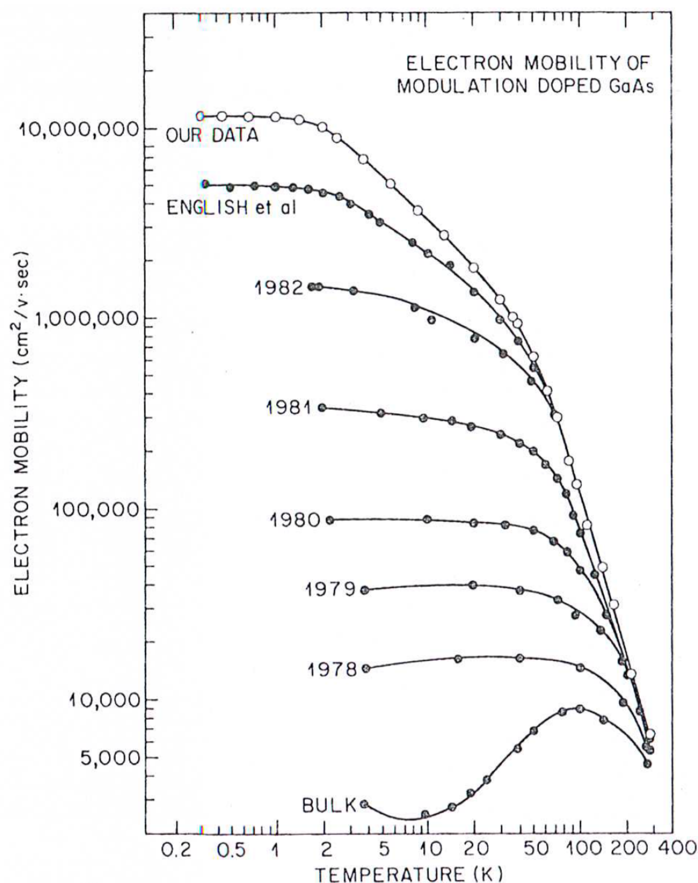

Figure 2.13. Plot of maximum 2DEG mobility as a function of temperature showing the result of improvements in MBE and heterostructure design and materials. Reprinted with permission from L. Pfeiffer, K. W. West, H. L. Stormer, and K. W. Baldwin. Appl. Phys. Lett. **55**, 1888 (1989). Electron mobilities exceeding 10^7 cm^2/Vs in modulation doped GaAs. Copyright (1989), AIP Publishing LLC.

TSPs to pump the growth chamber, there is anecdotal evidence that at the extremely low pressures needed for high mobility GaAs growth ion pumps may actually make the vacuum worse [67]. As a result, our design did not incorporate any ion pumps.

Next, the effusion cells themselves were optimized since, as the hottest part of the chamber, they represent the most likely source of impurity outgassing during growth. The cells were modified from the standard "high mobility" design to include additional heat shielding to reduce their power consumption. Moreover, wherever possible the

hot parts of the cells were specified to be made from high purity tantalum or from high purity pyrolitic boron nitride (pBN) where electrical insulation was required.

In order to further reduce the power input into the system, the dopant sources were home-made, resistively-heated filaments of silicon (for n-type doping) and carbon (for p-type doping). The filament design allowed the dopant sources to be ramped from zero current to their full operating temperature in as little as 30 seconds, thereby minimizing outgassing from the dopant sources and surrounding chamber walls during the rest of the growth. In addition, silicon and carbon were chosen as the dopant sources due to their low vapor pressure at room temperature. Carbon, in particular, is a significant improvement over other commonly used p-type dopants like berrylium and zinc which have large vapor pressures even at room temperature. Carbon, by contrast, must be heated to $\sim$ 2000 °C to develop an appreciable vapor pressure. As a result, whatever carbon leaves the filament will permanently stick to the first surface it impacts and thus not degrade the quality of subsequent growths.

The second consideration that was taken into accout during the design of the MBE was the incorporation of redundancy and fail-safe mechanisms. UHV equipment is usually quite fragile, and stopping a growth campaign to fix broken components can easily set a research team back 6 months or more. There are also a number of failure modes in a high mobility system that while not requiring any repairs can nevertheless significantly degrade the mobility for extended periods.

In terms of redundancy, the system was designed to use two sources for both gallium and aluminum, and each dopant source was designed with a back-up filament. The second gallium and aluminum sources act as back-ups in case the first source fails and also enable the growth of more complicated heterostructures with multiple alloy concentrations. Next, the system utilizes three cryopumps. This first creates a very large pumping capacity and second acts as a fail-safe in case one pump malfunctions.

Another failure mechanism that was designed around was a failure of the liquid nitrogen supply. With the lab being located in the Midwest, it is not difficult to imagine that a winter storm or other severe weather could prevent the building's

primary liquid nitrogen system from being re-supplied for a few days. If the liquid nitrogen supply were to run out, the cryopanels would warm up and release a large number of impurities into the chamber. What is more, if the cryo panel were to warm up, the stainless steel walls could potentially be heated to a dangerous level by the nearby effusion cells. As a safeguard against this failure mechanism, the MBE liquid nitrogen system was designed with a dedicated, continuously-filled 1000L back-up tank which would automatically start feeding the MBE if the house liquid nitrogen supply were to run out.

Power outages are another likely failure mode, and as such the liquid metal sources, ion gauges, control computer, cryo pump temperature monitors, cryo pump compressors, and liquid nitrogen control system were all powered by the building UPS power supply to ensure their continued operation. Next, while not originally designed into the system, our group has subsequently purchased a closed-cycle water chiller for cooling the cryopump compressors in the event of a failure in the building's process cooling water supply. Finally, if any of the likely failure mechanisms do occur, an auto-dialer system was setup to automatically alert all group members to the emergency situation.

2.3 Machine Setup

The MBE system was shipped to Purdue in pieces, so the initial installation was quite involved. One of the first steps was the installation of the high-purity argon lines used for venting the different vacuum chambers. The house argon was supplied from the boil-off of liquid argon and was routed to the lab via stainless steel lines. Inside the lab, the argon first passed through a gettering furnace[2]. The furnace operates by passing the gas to be purified over a hot charge of high purity titanium which reacts with impurities such as oxygen, carbon dioxide, hydrogen, and water. Our gettering furnace was specified by the manufacturer as being able to reduce an

[2]Centorr Vacuum Industries model 2G-100-SS inert gas gettering furnace.

initial oxygen concentration of 10 ppm to as low as 1×10^{-6} ppm. After leaving the gettering furnace, the argon was routed to the various chambers on the MBE by electropolished stainless steel lines. All connections were made with either butt-welds or metal-gasket-sealed VCR type connections. During the first growth campaign, the lines were simply flushed with argon; but prior to re-loading the source material at the start of the second growth campaign, all the lines were additionally baked under vacuum to remove residual water from the stainless steel.

One issue that arose during the initial leak-checking was that the helium background in the chamber was intially quite high. This was because the cryo pumps had been saturated with helium when the MBE was leak-checked at the factory. Despite the fact that the pumps were removed and sat in air for several months before the system was installed at Purdue, the helium level ($\sim 1.5 \times 10^{-8}$ Torr) was still large enough to make it difficult to find small leaks. Evidently, it is very difficult for the helium to find its way out of the "maze" of the charcoal on the cryo pump cold head, even at room temperature. In addition, this helium could not be simply flushed out of the charcoal with repeated pump/purge cycles. The only method we found that successfully removed the helium was to pump on the cryopumps with a turbo pump for $\sim$ 24 hours while heating the cryo pump to $\sim$ 50 C. During the repeated pump-downs of the MBE, the liquid nitrogen-cooled sorption pumps also became satured with helium. Once the sorption pumps were filled with helium they acted as a source of helium when used to rough down any helium-free chamber. As a result, we found it necessary to also regenerate the zeolite molecular sieve material as well. This was accomplished by baking the sorption pumps to $\sim$ 200C for a few days while the vacuum vessel was evacuated with a turbo pump. Once the helium was successfully removed from all the cryo pumps and sorption pumps, the helium background in the MBE was reduced to the mid-10^{-13} Torr range (i.e. the noise floor of the RGA).

Once the venting system was in the place and the initial leak checking was completed, the system was baked (without effusion cells) for 6 days with the growth chamber at 200 C for four of those days. This initial, short bake was meant to check

Figure 2.14. MBE system with the baking system in place.

that the system met the manufacturer's vacuum specification. Uniform baking of the chamber was facilitated by the convection bake-house (shown in figure 2.14) supplied by the manufacturer.

After the vacuum specification was met, the long work of preparing the sources began. Machined parts were de-greased in a multi-step process. First, the parts were soaked and wiped in trichloroethylene (TCE) to remove gross contamination. Second, the parts were sonicated in clean TCE in a new, clean beaker. Once the parts could be wiped on a clean, white wipe without leaving any dark residue, they were then sonicated in acetone followed by methanol followed by DI water; the cleaning process was finished with an extended rinse in running DI water. In addition to this de-greasing, de-greased tantalum parts were briefly etched in 1:1 $HF:HNO_3$ and then baked in an argon furnace at 850 °C for 4 hours to drive hydrogen out of the tantalum[3]. In order to spot-weld tantalum pieces without embedding any copper residue, a set of weld tips were made as shown in figure 2.15. The molybdenum wire was jammed

[3]It should be noted that the etchant used to etch tantalum is very aggressive and care should be taken to not over-etch the tantalum, particularly when dealing with small, threaded parts.

Figure 2.15. Homemade molybdenum tips for spot-welding MBE components.

into the copper housings while the copper was hot to create an interference fit with minimal electrical resistance. The copper housings then threaded into the existing electrodes on the spot welder.

The PBN crucibles were etched in aqua regia to remove residual metallic contaminants and thoroughly rinsed in DI water. Following the etch, the crucibles were baked in oxygen at 650 C for 4 hours to volatilize organic residues. Before the crucibles were outgassed in vacuum, the effusion cells were first heated to $\sim$ 1350 °C for 1 hour in an auxiliary UHV chamber. Care was taken to increase the power to the source slowly until the source was above $\sim$ 100 °C to allow heavy hydrocarbon molecules to desorb without cracking. The risk in cracking the hydrocarbons would be that atomic carbon could be left behind on the source and result in a p-type background that would be extremely difficult to remove due to the fact that carbon's vapor pressure does not rise rapidly until $\sim$ 2000 °C. After the source was outgassed and leak-checked, the crucible was installed in the source and outgassed to $\sim$ 1500 °C for $\sim$ 3 hours in the auxiliary chamber. In the case of sources and/or heater zones that could not reach

1500 °C, the source was outgassed to its maximum temperature as specified by the manufacturer.

Once each source was outgassed in the auxiliary chamber and deemed to be leak-tight, it was transferred to the MBE. The aluminum sources were filled at this point, but all other sources were left empty. The entire system was then baked at 200 °C for $\sim$ 5.5 days in an attempt to remove the small 32 amu signal (i.e. oxygen) seen in the RGA. As the 32 peak was removed by the baking, the chamber was (after extensive leak checking) deemed to be leak-tight and ready to receive the gallium and arsenic.

Great care was taken at this point to vent the system as cleanly as possible. The chamber was vented with the clean argon supplied by the gettering furnace as previously described. In addition, the sources were contained with glove bags sealed to the MBE so that the source and MBE would not be exposed to air. Furthermore, great care was taken to not handle the source material with anything other than its original packing to avoid introducing additional impurities. Once all the source material was loaded, the machine was baked at 200 °C for 2 weeks.

Following the extended bake, the cryo panels were flooded, and the TSP was run with a short duty cycle of 7 minutes on, 30 minutes off, for 3 days. Initially, however, it was necessary to use longer duty cycles to drive the hydrogen out of the titanium in the TSP. Once the TSP itself had cleaned up, the total pressure steadily dropped as the TSP pumped the hydrogen in the system (the dominant partial pressure). The total pressure in the MBE ultimately saturated $\sim 2.1 \times 10^{-12}$ Torr at which point the TSP was permanently powered down and initial wafers were grown to calibrate the dopant sources.

2.4 Lessons Learned During the First Growth Campaign

Over the course of the first growth campaign we learned a great deal about the finer points of growing high mobility GaAs. The first lesson we learned was about the great importance of source material purity. The first two months of attempting

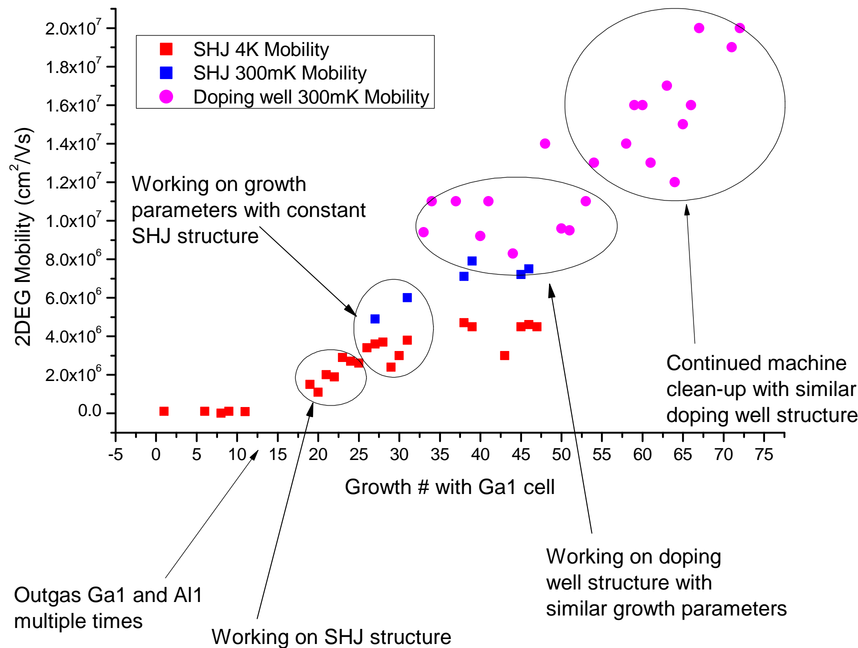

Figure 2.16. Summary of MBE clean-up early in first growth campaign. The first ∼ 12 heterostructures grown were electrically insulating at low temperature. Most of the growths between 12-17 were bulk, undoped GaAs used to measure the background impurity concentration. After extensive outgassing the mobility rapidly increased as we first worked on the design of a single heterojunction (SHJ) heterostructure and then switched to optimize the growth parameters before finally moving to a more complicated doping well type heterostructure.

to grow heterostructures resulted in failure after failure as every heterostructure we grew was electrically insulating at low temperature. After growing a nominally undoped epilayer of GaAs and measuring the Hall density at room temperature, it was determined that the GaAs had a p-type background $\sim 2.7 \times 10^{15}$ cm^{-3} even after the equivalent of ~ 150 μm of GaAs had been evaporated during previous growths and outgassings. After several additional outgassings and bulk GaAs growths which evaporated the equivalent of an additional ~ 190 μm of GaAs, we were able to grow our first working 2DEG with a mobility $\sim 1 \times 10^6$ cm^2/Vs. Following this extensive

outgassing of the gallium source, the mobility rose very rapidly as shown in figure 2.16 which charts the mobility as a function of growth number. This illustrates that the base pressure of the chamber (which was extremely low in our case) is, surprisingly, not a good predictor of resulting electron mobility. In addition, the RGA spectra were not particularly helpful in determining the source of contamination. The RGA spectra never showed a significant increase in impurities when the gallium cells were heated to growth temperature. In contrast, the RGA did show a noticeable rise in some impurities when the valve to the arsenic source was opened. Evidently, the impurities that readily incorporate into the growing film are not the same as the ones that are able to bounce around in the growth chamber and reach the RGA. Ultimately, then, the only feedback that can be reliably used to optimize the growth system is the electron mobility (or background impurity concentration in the case of very high impurity levels).

After all the initial outgassings and growths necessary to optimize our standard operating procedures and heterostructure designs, we were concerned that the length of the growth campaign would be limited by the amount of gallium we had loaded. As a result, we did not do any extended high temperature outgassings of the sources later in the campaign to try to push the mobility higher. However, it turned out that the length of the first growth campaign was limited not by the gallium but by the arsenic. Based on the amount of gallium remaining in each source at the end of the first growth campaign, we estimated that 1 μm of GaAs corresponded to $\sim$ 17 mg of gallium. Based on the total time each source was hot for each growth (including setup time) this would mean two gallium sources filled with a total of 200 g of gallium could grow $\sim$ 2000 wafers (roughly 10 years of work).

However, since we did not know that so little gallium would in fact be used until we removed the gallium cells at the end of the first campaign, we built a dedicated UHV chamber, shown in figure 2.17, for outgassing the gallium prior to loading the effusion cell into the MBE. The chamber was designed with a near-vertical port for the effusion cell so that the crucible could be filled with more gallium than the crucible could

Figure 2.17. Dedicated gallium outgassing chamber. The chamber featured a near-vertical port to enable the the crucible to be filled almost completely with gallium. The frame supporting the chamber could be rolled into position next to the MBE to allow sources to be passed into the MBE without exposure to air.

hold in the MBE. The intent was to then outgas the excess gallium in this auxiliary chamber and finally to transfer the source in an argon environment to the MBE. In addition to cleaning the gallium, outgassing the gallium in a separate chamber would minimize the risk of creating electrical shorts inside the growth chamber from all the evaporated gallium. The transfer into the MBE would be facilitated by the ability to roll the support stand to within a few feet of the MBE source flange where the outgassing chamber would mate with a custom made acrylic glove box [69] attached to the MBE. However, once we realized how much gallium was left after the first campaign, it was decided that the risks involved in moving a hot source through a

plastic glove bag and glove box outweighed the benefits of loading the MBE with pre-outgassed gallium.

Based on the lessons from the first growth campaign, we made some small changes to our loading procedure. The biggest change was to replace the glove bags used to load the sources with a homemade acrylic glovebox. This greatly increased visibility and dexterity when working on the vacuum components. In addition, this box combined with the outgassing chamber shown in figure 2.17 allowed the sources to be removed from the MBE, loaded with a new crucible, outgassed, loaded with material, and re-installed in the MBE without ever being exposed to air. The quality of the purge gas in the glovebox was monitored with a sensitive oxygen monitor in order to determine when the box was sufficiently purged to open the MBE. With appropriate purge techniques and flow rates it was possible to achieve oxygen concentrations as low as 50 ppm in the box. Oxygen contamination was a larger concern during the re-loading than it was during the initial machine setup due to the large amount of arsenic in the chamber. Large amounts of AsO would easily be formed by a bad vent and would likely result in a large n-type background in subsequent growths.

As a result of the improvements to our venting techniques, it was not necessary to bake the chamber after re-loading the source material. While the partial pressures of some gases were slightly higher after reloading than they were prior to removing the sources, we believed that while baking would improve the total vacuum quality, it would primarily move impurities from the walls of the chamber to the source material. Instead of baking the entire system, we simply heated the sources in the absence of any liquid nitrogen in the cryopanel to outgas the sources, shutters, and walls immediately surrounding the sources.

After re-assembling everything on the MBE, we immediately outgassed all the source material extensively prior to growing any heterostructures. After outgassing the equivalent of ~ 250 μm worth of GaAs from the first source, it began producing heterostructures with mobilities $> 20 \times 10^6$ cm^2/Vs. While this is not substantially less material than was evaporated at the beginning of the first growth campaign, it

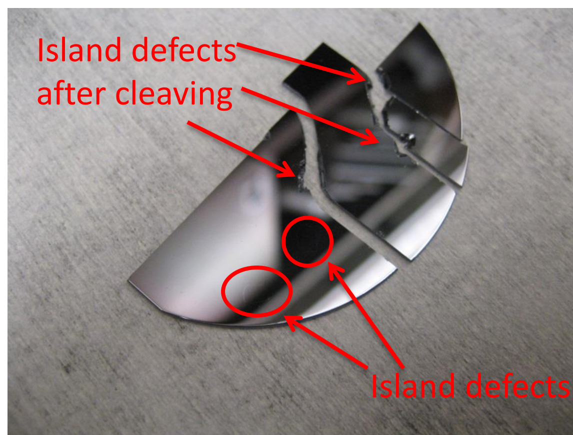

Figure 2.18. Wafer with so-called "island defects" after cleaving. Two defects (circled) away from the cleave lines as wells as a few defects along cleave lines are called out in the figure.

is worth noting that it took $\sim$ 6 months and 70 growths during the first campaign to reach this mobility. At the start of the second campaign, by contrast, it took 10 growths and less than a month to reach a mobility of 20×10^6 cm^2/Vs.

While the importance of outgassing the gallium was the most valuable lesson we learned during the first campaign, there were a number of other important details that we also worked out along the way. The most important of these were related to the formation of various kinds of defects in the wafers. The first prominent defect we encountered was what we termed "island defects". Figure 2.18 shows an image of a half-wafer with island defects after cleaving. It turned out that these defects formed when the gallium holding the wafer to the tantalum block froze in the LL chamber. The LL was designed in such a way that the wafers were situated in close proximity to the 80 K array of the cryo pump on the chamber. As a result, the inside of the chamber would radiatively cool to $\sim$ 13 °C which was cold enough to harden the gallium sufficiently to punch out chunks of the wafer. After we determined the cause of these defects, we eliminated them by simply keeping the LL heated to 50 °C with its built-in bake-out lamp.

Figure 2.19. Nomarksi contrast optical micrograph of a wafer showing so-called cross hatch morphology. The field of view is $\sim$ 2.5 mm wide.

In addition to these macroscopic defects, the first $\sim$ 12 months of our growths were plagued by a rough morphology we referred to as "cross hatching", shown in figure 2.19. This rough morphology was gradually improved by implementing several changes to our standard procedure. First, the initial smoothing sequence of each heterostructure was changed from 5 repeats of 100 nm of GaAs separated by 100 second smoothing pauses to 50 repeats of 10 nm of GaAs separated by 20 second pauses. Even though the total thickness of the smoothing sequence was not changed, the higher frequency of smoothing pauses evidently increased the rate at which the wafer surface smoothed out. In addition, the outgassing time at growth temperature in the growth chamber was reduced from 30 minutes to 10 minutes to reduce the period during which the substrate could roughen due to desorption. Finally, the tantalum blocks used to hold the wafers were gradually replaced with new, re-designed blocks. These blocks were modified in two ways. First, the block featured a short pedestal with a diameter $\sim$ 4 mm smaller than that of the wafer. This left the edge of the wafer free to move as the wafer and block expanded and contracted, thereby decreasing stress on the wafer during growth. In the absence of this pedestal, the gallium used to mount the wafer reacted with arsenic during the growth, and

this material would quickly build up on the edge of the block with the result that after several growths the edge of the wafer would typically be pushed up slightly by this build-up. This resulted in distortion lines appearing on the center of the wafer following a growth, and the blocks required regular etching to remove this build-up. In addition to eliminating this build-up, the pedestal on the block added to the total thickness of the block, thereby decreasing the radial thermal gradient in the wafer. This decreased thermal gradient helped reduce the prevalence of cross hatching in addition to reducing the occurence of slip line defects at the edge of the wafer.

2.5 Heterostructure Design

While the previous sections emphasized the work necessary to produce a high quality vacuum environment, the design of a heterostructure also has a large impact on the quality of low temperature electron transport. As a starting point, it is useful to ask why one would choose to study low-dimensional electrons in the GaAs/AlAs material system and not some other semiconductor system. The GaAs/AlAs system is unique in that the two materials are almost perfectly lattice-matched [56]. This means that heterolayers of GaAs, AlAs, and their alloy AlGaAs can be grown with arbitrary thickness without the introduction of significant strain or crystal defects. All other III-V and group IV semiconductor systems such as InGaAs/GaP, Si/Ge, GaN/AlN/InN, etc. require complicated strain-relieving buffer layers to achieve hiqh-quality heteroepitaxy. In addition, unlike the ubiquitous Si/SiO_2 system, the barrier material (AlAs) is crystalline and as such does not introduce the level of interface states or interface roughness found at a crystalline/amorphous interface. Moreover, the conduction band minimum in GaAs is at the Γ point which eliminates the complication of valley degeneracy that occurs in indirect gap materials such as Si. Finally, from a practical standpoint, the availability of low-cost, high-purity, semi-insulating GaAs substrates is an additional advantage over some other material systems such as GaN [56].

2.5.1 Smoothing and Gettering Layers

Having selected GaAs/AlAs as the semiconducting material of choice for studying electron correlations, there are a number of design considerations that must be taken into account when growing a heterostrcuture. First, it is necessary to smooth out the substrate prior to growing the active region of the device. This is necessary because when the native oxide is desorbed from the GaAs surface, it leaves the surface relatively rough (see figure 2.8a for an example of a RHEED pattern from a surface immediately after oxide desorption). The oxide layer consists of several different As- and Ga-oxides [70]. The As-oxides desorb below ~ 400 °C and Ga_2O desorbs between 400 - 500 °C. The surface roughens primarily when the final oxide is desorbed from the substrate (~ 600 °C) which proceeds according to the reaction [71] $Ga_2O_3 + 4Ga \rightarrow 3Ga_2O_3$. The oxide is believed to have defects which weaken it or expose GaAs, and these weak spots then act as sources of mobile Ga atoms which act to desorb the oxide via the aforementioned reaction. This then leads to pits on the surface, the lateral size of which grow with oxide thickness [71]. Thus the buffer layer, along with a series of growth interruptions, smooth out the substrate so that subsequent epitaxial layers can have the requisite smooth interfaces.

Next, a GaAs/AlGaAs superlattice is grown for several reasons. First, it "exercises" the sources by repeatedly actuating the shutters; this is believed to shake loose contaminants that may have adsorbed on the moving parts after the previous growth. Second, the superlattice acts as a trapping barrier to impurities from the substrate. The exact method by which the superlattice traps impurities is not well agreed upon in the literature, but possible reasons are different solubilities of GaAs and AlGaAs to impurities [72] or strain-induced gettering [72, 73]. Third, the superlattice is believed to further smooth the growth front, particularly when compared to a thick AlGaAs cladding layer [72, 74].

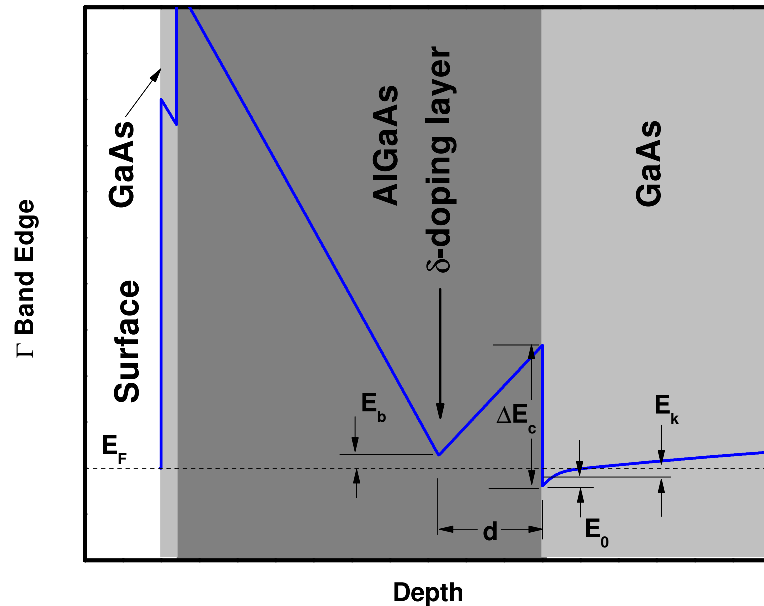

Figure 2.20. Sketch of the Γ band edge as a function of depth for a simple single interface heterojunction. The dashed line shows the position of the Fermi energy E_F.

2.5.2 Charge Transfer Fundamentals

Following the intial smoothing sequence and superlattice layers, the active region of the device is grown. Figure 2.20 shows the band edge as a function of depth for a simple single heterojunction with a single delta doping layer. This structure consists of (in the order of growth) a GaAs channel, an $Al_xGa_{1-x}As$ barrier with a delta doping layer, and a GaAs capping layer which forms a self-limiting oxide layer. Modern heterostrcutures, particularly those used in the study of the FQHE, are more sophisticated, but the single heterojunction is still useful for understanding the basics of band structure engineering. In order to create the electrostatic confinement of the 2DEG, barrier layers of $Al_xGa_{1-x}As$ or AlAs are needed. The conduction band offset ΔE_C between $Al_xGa_{1-x}As$ and GaAs is roughly linear in the aluminum mole fraction x and is $\sim 65\%$ of the band gap difference between $Al_xGa_{1-x}As$ and GaAs [75]. Larger values of x increase the conduction band offset until the conduction band minimum

crosses over from the Γ point to the X point of the Brillouin zone for x $\sim$ 0.45 [76] at which point the conduction band offset is $\sim$ 360 meV [3]. Most heterostructures used for electrical transport measurements, however, employ barriers with $x < 0.45$. A larger conduction band offset at the primary heterojunction will result in stronger spatial confinement and thereby increase the electric subband spacing.

In addition to spatial confinement, the conduction band offset affects the density of charge transferred from modulation doping layers to the 2DEG. Considering the energies sketched in figure 2.20, we can write down the following equation:

$$E_b + \frac{e^2}{\epsilon} nd = \Delta E_C - E_0 - E_k \tag{2.2}$$

where E_b is the donor binding energy, the second term on the left-hand side is the rise in energy due to the field between the doping layer and the 2DEG, e is the electron charge, ϵ is the permittivity, n is the 2DEG density, d is the dopant setback, ΔE_C is the conduction band offset, E_0 is the ground state energy of the roughly triangular confining potential, and E_k is the kinetic energy of the 2DEG in the plane. Note that it is assumed here that the temperature is essentially zero and the 2DEG density is sufficiently small so that only the ground electric subband is occupied. The kinetic energy in 2D can be written in terms of the density as

$$E_k = \frac{\hbar^2 k_F^2}{2m^*} = \frac{\pi \hbar^2}{m^*} n \tag{2.3}$$

where k_F is the Fermi wavevector and m^* is the electron effective mass. Next, the ground state energy of the confining potential can be approximated in closed form with the variational Fang-Howard solution [77].

$$E_0 = \frac{3}{2} (\frac{3}{2} \frac{e^2}{\epsilon} n \frac{\hbar}{\sqrt{m^*}})^{3/2} \tag{2.4}$$

Thus equation 2.2 can be re-written as

$$n(ad + c) + bn^{2/3} = \Delta E_C - E_b \tag{2.5}$$

where $a = e^2/\epsilon$, $b = \frac{3}{2}(\frac{3}{2}\frac{e^2}{\epsilon}\frac{\hbar}{\sqrt{m^*}})^{2/3}$, and $c = \pi\hbar^2/m^*$. For a dopant setback of the order of $\sim$ 50 nm and a 2DEG density $\sim 10^{11}$ cm^{-2}, $ad >> c$ and adn is $\sim$ 1 order of magnitude larger than $bn^{2/3}$ so to a rough approximation we can write 2.5 as

$$n \approx \frac{\Delta E_C - E_b}{ad} \tag{2.6}$$

The electron density is thus roughly linear in dopant setback as one would expect from treating the dopant layer and 2DEG as a parallel plate capacitor. The conduction band offset will be linear in x as mentioned previously, but for a dopant such as Si the binding energy will also change with x. This means that a more detailed microscopic understanding of the dopant binding is necessary to predict the charge transfer to the 2DEG.

2.5.3 Silicon Dopant Incorporation

Silicon in $Al_xGa_{1-x}As$ can incorporate as a shallow or deep donor depending on the Al concentration. For Al mole fractions $x < 0.22$ the Si exclusively incorporates as a shallow, hydrogenic donor with a binding energy $\sim$ 6 meV [77]. However, for $x > 0.22$, the majority of the Si atoms occupy interstitial sites which results in a deep donor level known as the DX center [78, 79, 80, 81]. The DX center is unusual in that it has a large barrier both to thermal emission and capture. Thus, if DX centers are ionized by illumination at low temperature the electrons cannot be re-captured by their parent ions and the 2DEG density will be increased essentially indefinitely after the illumination is removed. This so-called persistent photoconductivity can in some instances be used as a controllable way to change the 2DEG density in-situ [67]. The leading microscopic theory of DX centers, put forth by Chadi and Chang [80, 81], uses a negative U Hubbard model to describe the donors. This means that the ground state of the DX center is actually the negatively ionized DX^- state. The binding energy of the DX center is predicted by this model to be $E_b[\text{eV}] = 1.18x - 0.26$. It can be seen that the DX level lies above the conduction band edge for $x < 0.22$, though the DX

center can be induced in GaAs and low-x $Al_xGa_{1-x}As$ by the application of a large hydrostatic pressure [78].

Returning to equation 2.6, we can now predict the impact of Al concentration on the 2DEG density, at least in simple cases. For $x < 0.22$, the donor binding energy will be quite small compared to the conduction band offset, and thus

$$n \propto \frac{x}{d}, \quad x < 0.22 \tag{2.7}$$

Once x is large enough for DX centers to form, the electron density without illumination becomes more difficult to predict in closed form due to the presence of two donor energy levels. After illumination, however, the Fermi level should equilibrate with the shallow donor level for sufficiently high doping concentrations and the dependence of density on Al concentration will again be described by equation 2.7. In this case the electrons excited out of the DX centers will either populate the 2DEG or return to shallow donors. If the doping level is sufficiently high that the effective Bohr radii of the shallow, hydrogenic donors overlap, the electrons ejected from the DX centers that return to shallow donors will form a low-mobility, parallel-conducting channel that will interfere with transport measurements. Higher values of x will increase the maximum doping before parallel conduction sets in due to the larger conduction band offset.

Examining figure 2.20 again, it is clear that there is also a large electric field between the surface and the doping layer. This is due to the surface states arising from dangling bonds on the surface which steal charge from the dopant layer. These surface states pin the Fermi energy $\sim$ 0.7-0.8 eV below the GaAs conduction band edge at the surface. This means that for the 2DEG to be fully populated, the dopant concentration must be considerably larger than that necessary to populate the 2DEG. This effect becomes more pronounced as the depth from the surface is reduced and the field between the dopants and surface is increased. In practice, this makes it very difficult to grow doped heterostructures with a 2DEG depth less than $\sim$ 40-50 nm.

2.5.4 Impact of Heterostructure Design on Scattering Mechanisms

Armed with this basic understanding of heterostructure design and charge transfer, we can now turn to a discussion of the design considerations involved in growing a high mobility heterostructure. First, the four dominant electron scattering mechanisms at low temperature are background impurity (BI) scattering from the unintentional impurities in the crystal, remote impurity (RI) scattering from the ionized intential dopants, alloy disorder scattering where the electron wavefunction penetrates the Al_xGa_{1-x} barrier, and interface roughness (IR) scattering from variations in the width of the confining potential along the growth direction. Quantitative analysis of these scattering mechanisms is presented in chapter 5 in the context of 2D hole gases; for the purposes of this section I will focus on the qualitative description of these scattering mechanisms.

BI scattering is influenced by a number of factors in the heterostructure design. First, the growth rate influences impurity incorporation. At a slow growth rate vacuum impurities have more time to incorporate into the growing film thus resulting in higher BI scattering rates. Second, hot, reactive Al is known to be an effective getter of vacuum impurities [68, 56]. Therefore, reducing the Al concentration will reduce the incorporation of background impurities. However, for a fixed dopant setback this will also typically reduce the 2DEG density which results in reduced screening of the disorder potential by the 2DEG. To avoid this problem, one solution is to use a low Al concentration in the vicinity of the 2DEG and a higher Al concentration around the dopant layer. This results in efficient charge transfer while keeping the layers immediately surrounding the 2DEG as clean as possible.

RI scattering is, of course, due to the ionized donors, so reducing the concentration of ions and increasing their separation from the 2DEG is necessary to minimize RI scattering. Increasing the dopant setback, however, comes at the cost of reduced 2DEG density and reduced screening. In the highest mobility samples, though, it is possible to keep the electron density sufficiently high that RI scattering does not

limit the mobility [82]. In order to keep the RI concentration low, it is necessary to make the 2DEG sufficiently deep. As already discussed, decreasing the distance from the doping layer to the surface increases the dopant ionization due to surface compensation. Thus, the highest mobility samples typically have a 2DEG depth from the surface ~ 200 nm. However, this depth is rather large for experiments with nanostructures since the minimum feature that can be electrostatically defined in the 2DEG by the gates is of the order of the 2DEG depth. As a result, trade-offs between mobility and minimum feature size are often made.

Alloy scattering arises due to the random distribution of the Al and Ga atoms on the group-III lattice sites. Because both atoms have the same valence, however, the scattering is short ranged and therefore only occurs in regions where the wavefunction has an appreciable probability of being found. This is a major scattering mechanism in materials such as InGaAs where the channel is an alloy, but it is not typically very significant in GaAs/AlGaAs structures. The alloy scattering that does occur in the barrier, though, depends on both the degree of randomness of the alloy and the amplitude of the wavefunction. Maximal alloy disorder is achieved for $x = 0.5$, but the large conduction band offset results in very minimal penetration of the wavefunction into the barrier. In practice, alloy scattering increases as x is decreased; in other words as x is decreased, the increased penetration of the wavefunction into the barrier wins out over the decreased randomness of the alloy.

Finally, IR scattering is closely related to alloy scattering. It arises from variations in the width of the confining potential in the plane of the 2DEG which results in variation of the electric subband energy as the electron moves in the x-y plane. It can be reduced by adding short pauses at the quantum well interfaces to allow the surface to smooth out. However, the pauses can also increase background impurity incorporation, so it is necessary to empirically find the appropriate pause duration. The growth temperature and Al concentration also influence interface roughness. Higher Al concentrations and lower growth temperatures typically cause rougher growth because the optimal growth temperature for AlAs is significantly higher than that typically

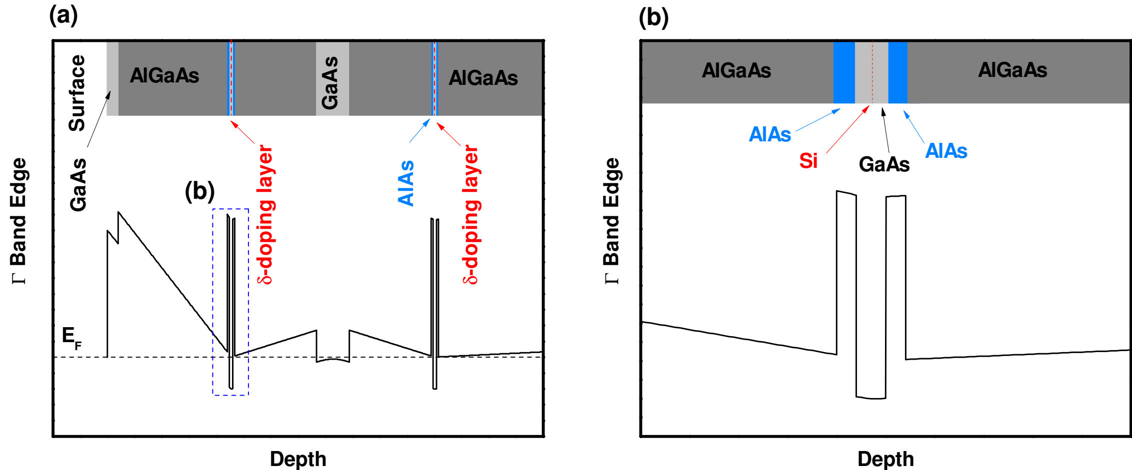

Figure 2.21. Modern quantum well structure for studying 2$^{\text{nd}}$ LL physics. (a) Sketch of the Γ band edge as a function of depth for a quantum well with a doping well doping scheme. The horizontal dashed line shows the position of the Fermi energy E_F while the vertical dashed lines show the position of the Si delta doping layers. An enlarged view of the dashed blue box is shown in the adjacent panel. (b) Close-up view of the doping well layer sequence.

used for GaAs/Al$_x$Ga$_{1-x}$As heterostructures. Finally, the growth rate can also impact the roughness with higher growth rates resulting in rougher interfaces. However, IR scattering is not usually the dominant scattering mechanism in GaAs except for very narrow quantum wells. By contrast, IR scattering is known to be an important scattering mechanism in Si MOSFETs where the 2DEG is pulled up against the amorphous oxide interface [83].

2.5.5 Doping Considerations for Amplifying 2$^{\text{nd}}$ LL Physics

The method of incorporating dopants also has a significant impact on the behavior of the 2DEG. In order to maximize the FQHE energy gaps in the 2$^{\text{nd}}$ LL, it is necessary to incorporate the donors in a so-called “doping well” scheme, also referred to as a short-period superlattice (SPSL) doping scheme[84, 85]. In this design, shown

in figure 2.21, the Si donors are deposited in a thin GaAs layer sandwiched between thin layers of AlAs. The strong spatial confinement created by the AlAs barriers ionizes the Si atoms and results in efficient charge transfer to the 2DEG. In addition, the layer thicknesses can be set such that the ground state for X-band electrons in the AlAs layer is below that of Γ-band electrons in the GaAs layer [84]. In this case, charge from excess Si atoms will reside as heavy electrons in the AlAs layer. Due to their large effective mass and close proximity to their parent ions, these electrons will have a very low mobility and not be visible as parallel conduction in magneto-transport measurements. However, they will still be sufficiently mobile to be very effective at screening their parent Si atoms. This strong screening effect evidently has a strong influence on transport in the 2^{nd} LL[85]; using this doping scheme we have achieved record energy gaps at $\nu = 5/2$ as large as ~ 600 mK [56, 86]. Creating this strong screening of the Si ions, however, comes at the expense of charge stability in nanostructures. Since there is so much loosely-bound charge in the doping layers, it is often necessary to wait several hours for the charge to stabiliize after changing the gate voltage [87]. This difficulty in controlling the gating of nanostructures on heterostructures with optimal bulk transport is one of the outstanding challenges in current research on the 2^{nd} LL.

2.5.6 Doping Considerations for Minimizing Charge Noise

In situations that require more stable gating, such as spin qubits, it is necessary to remove the doping well and instead deposit the Si directly in $\text{Al}_x\text{Ga}_{1-x}\text{As}$ with a large Al concentration such that the electrons will be tightly bound to DX centers at low temperature. In addition to the persistent photoconductivity effect already discussed, the freeze-out of the DX centers at low temperature can also be used to manipulate the low-temperature threshold voltage of a gated device. Figure 2.22 shows the effect of the so-called bias cooling technique in which the device is cooled from room temperature with a forward bias applied to the gate. Figure 2.22a shows

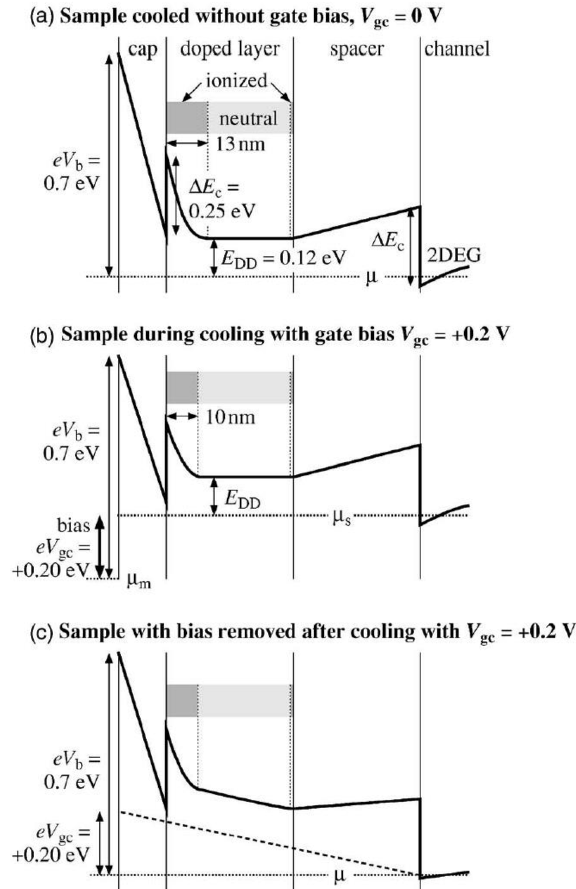

Figure 2.22. Sketch of the Γ band edge as a function of depth for a heterostructure (a) cooled without a cooling bias, (b) during cooling with a forward bias, and (c) at low temperature after bias cooling. Reprinted figure with permission from M. Pioro-Ladrière, J. H. Davies, A. R. Long, A. S. Sachrajda, L. Gaudreau, P. Zawadzki, J. Lapointe, J. Gupta, Z. Wasilewski, and S. Studenikin, Phys. Rev. B, 72, 115331 2005. Copyright (2005) by the American Physical Society.

the band structure after cooling without a gate bias. The bands bend and charge transfers to the interface and surface as previously shown in figure 2.20. The only significant difference from the heterostructure shown in figure 2.20 and the one shown in figure 2.22 is that the latter employs a uniform doping scheme in which the doping is spread out over a finite stretch of $Al_xGa_{1-x}As$ rather than being placed in a single

delta layer. As figure 2.22a shows, in the absence of a cooling bias the surface ionizes 13 nm of the doped region in order to satisfy the dangling bonds at the surface. Figure 2.22b shows the heterostructure during cooldown with a positive bias of 0.2V applied to a top gate. In this case the field from the gate reduces the total field between the surface and doping layer which results in a thinner depletion layer (10nm) in the doping layer. While the bias is maintained, the 2DEG density is the same as in figure 2.22a. Once the sample is below the freeze-out temperature of the DX centers ($\sim$ 100 K [88]) no charge can enter or leave the DX centers. Thus if the cooling bias is removed, the only free charge that can be stolen by the surface resides in the 2DEG itself, and the system behaves as if there is a built-in reverse bias of 0.2V as shown in figure 2.22c.

This so-called "bias-cooling" technique has been studied both as a way to control charge correlation in the donor layer [89] and also as a way to control charge noise in nanostructures [88, 90]. One dominant mechanism for generating charge noise in nanostructures is evidently charge leakage from the reverse biased gate to the 2DEG via an intervening trap [91, 88, 90]. When an electron tunnels from the gate to a trap in the vicinity of the nanostructure this alters the local electrostatic environment. Such local changes in the electrostatics results in variation in the conductance of QPC constrictions and is also believed to result in dephasing of spin qubits [92]. By using the bias cooling method the required gate voltage necessary to reach a given 2DEG depletion is lessened. This results in dramatically suppressed tunneling rates from the gate and thus more temporally stable electrostatic environments in nanostructures.

Lateral charge hopping in the doping layer, presumably between hydrogenic Si donors, may also contribute to charge noise. To combat this effect, spacing the dopants out by changing from a delta doping distribution to a uniform doping distribution would be expected to reduce this charge noise mechanism. There is some evidence supporting this ideal [90], so our group typically avoids the use of delta doping layers in wafers intended for spin qubit experiments.

2.5.7 2D Holes in GaAs

Thus far the discussion has been centered on 2D *electron* gases, but 2D hole gases (2DHGs) are also of interest for various experiments. Their large effective mass enhances interaction effects and makes them useful for studying phenomena such as a 2D metal-insulator transition [93]. In addition, their p-wave symmetry results in a reduced contact hyperfine interaction with the nuclear spin bath. As such, 2DHG-based qubits are expected to have longer coherence times than their electron counterparts [94, 95, 96]. Early work on high mobility 2DHGs in GaAs took advantage of the amphoteric nature of Si and used (311)A-oriented substrates to cause the Si to incorporate as an acceptor, though this came at the expense of a large in-plane mobility anisotropy due to interface corrugations [97]. With the advent of filament-type dopant sources[98, 99, 100, 101] it became possible to use carbon as an acceptor on the same high-symmetry (100) substrates use for high-mobility Si-doped 2DEGs and thereby eliminate the mobility anisotropy of (311)A 2DHGs. Design of carbon-based 2DHGs follows similar principles as previously discussed. One important difference, though, is that carbon is not known to incorporate as a DX center as evinced by the lack of a strong persistent photoconductivity effect [101, 102]. Instead of forming a DX center, carbon incorporates as a shallow acceptor with a binidng energy $\sim$ 26 meV [77]. The lack of dopant freeze-out is a likely cause of the difficulty encountered in fabricating stable nanostructure involving C-doped heterostructures [103]. By contrast, undoped p-type heterostructures in GaAs have shown promise as platforms for nanostructures [104, 105]. A second important difference between n- and p-type heterostructures in GaAs is that the mixing of the light- and heavy-hole bands result in an effective mass that is dependent on both the hole density and the shape of the confining potential [106, 107, 108]. This variable effective mass means that the hole mobility cannot be used directly to quantify the relative quality of two wafers unless the density and confining potential are fixed.

2.6 Summary

In summary, MBE is a highly controllable crystal growth technique ideally suited for the study of low-dimensional electron correlations in a variety of material systems. The GaAs/Al_xGa_{1-x} material system is particularly-well suited for studying electron interactions due to its low effective electron mass, lattice-matched barrier and channel, direct bandgap, and mature growth and processing technology. In order to reduce disorder to the levels necessary for state-of-the-art research, it is necessary to go to great lengths to achieve the necessary growth chamber vacuum quality. In addition to meticulous vacuum hygiene, a thorough understanding of growth mechanisms, material properties, and band structure engineering is necessary to produce high quality 2D electron and hole gases.

3. Cryogenics for Electrical Transport Measurements

In the previous chapters I have laid out the motivation for studying electrons in low dimensions and discussed the method the Manfra group uses to grow the GaAs/AlGaAs heterostructures used to cofine the electrons to two dimensions. The remaining requirement to study electron interactions is to cool the electrons sufficiently that the Coulomb energy dominates over the thermal energy. There are several methods for cooling below 4.2 K, the boiling point of ^{4}He, but I will focus on the two types of cryostats used in the Manfra lab, evaporatively cooled ^{3}He fridges and ^{3}He/^{4}He dilution refrigerators.

3.1 Janis Pumped ^{3}He Cryostat for Cooling to T = 300 mK

One commonly used method for achieving temperatures below $\sim$ 1.2K, the limit of evaporatively-cooled ^{4}He, is to evaporatively cool ^{3}He. Figure 3.1 shows a sketch of the cross section of the fridge used in the Manfra lab. Cooling is accomplished through a multi-step process. Starting from the top of the cryostat at room temperature and working down, heat is first removed by the cold ^{4}He gas and finally the liquid ^{4}He contained in the dewar. To further cool the system, the cold parts of the fridge must be insulated from the "warm" 4.2 K bath. This is accomplished by evacuating the inner vacuum can (IVC). The outer wall of the IVC can is plated with copper. The copper's high thermal conductivity serves to keep the temperature at the top of the IVC fixed near 4.2K even when the liquid level falls significantly below the top of the IVC. Next, a so-called 1K pot is used to evaporatively cool liquid ^{4}He to $\sim$ 1.5K. In our fridge's top loading design, the 1K pot has a square toroidal geometry to allow the sample probe to pass through its center. ^{4}He is drawn in through a thin "sipper" that extends into the main ^{4}He bath. Flow into the 1K pot is regulated by a needle valve;

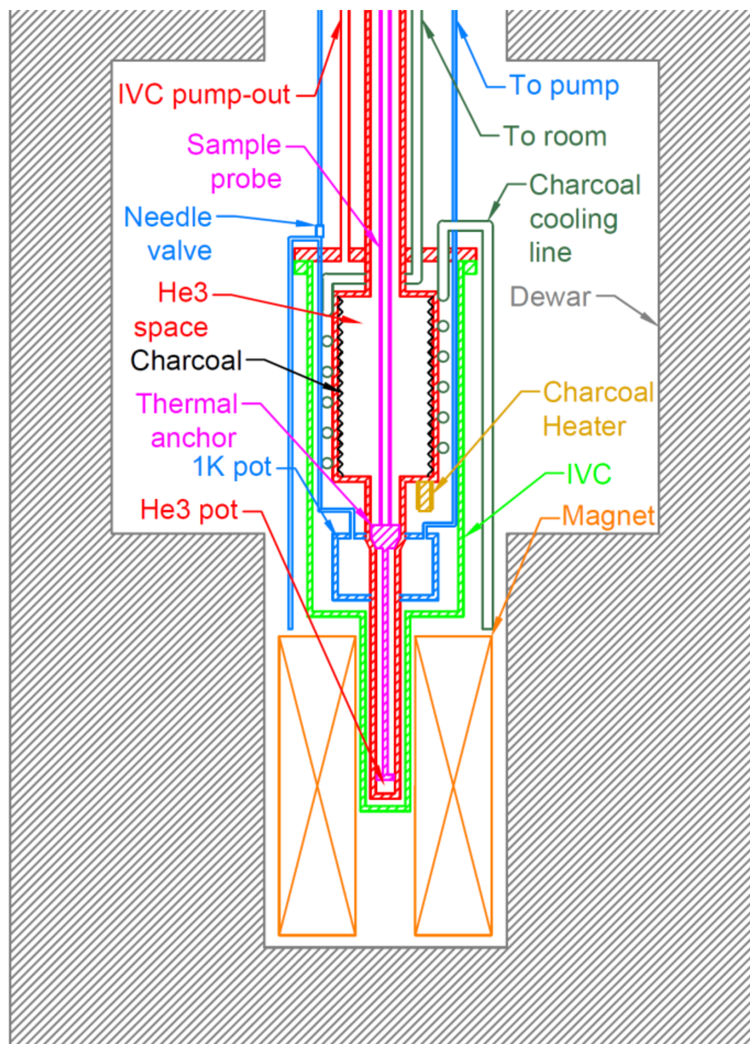

Figure 3.1. Sketch of the cross section of the Manfra group top-loading ^{3}He fridge.

since there is no regulation on the 1K pot pump, the conductance of the needle valve acts to control the pressure in the 1K pot. As the pressure in the 1K pot is lowered, the temperature of the remaining liquid will in turn drop as the hot He atoms boil off. The minimum temperature that can be achieved in a 1K pot is $\sim$ 1.3K [109], but in practice the temperature is often higher due to the extenal heat load.

To reach temperatures below 1.3 K, it is necessary to utilize ^{3}He. Like its heavier cousin ^{4}He, ^{3}He can also be cooled by evaporation. However, as the world's ^{3}He

supply comes primarily from the decay of the US nuclear weapons stockpile [110] it is rather difficult to obtain compared to ^{4}He. Thus, simply pumping liters of ^{3}He out into the lab on a daily basis is not a cost-effective cooling method. Instead, the ^{3}He is cooled by closed-cycle cryo-pumping. As shown in figure 3.1, a large portion of the ^{3}He space inside the IVC region is coated with activated charcoal. The charcoal is cooled by a continuous flow of cold ^{4}He gas from the main dewar space that flows through the charcoal cooling line after it evaporates. The desorption rate of ^{3}He from a surface at 4.2K is extremely slow, so the cold charcoal with its large surface area acts as a very effective entrapment pump of the ^{3}He. The liquid ^{3}He bath contained in the ^{3}He pot can thereby be cooled to 300 mK. The boil-off of ^{3}He is the final heat sink that absorbs heat from the sample and heat radiatively coupled from the walls of the IVC. This method of evaporatively cooling ^{3}He is, of course, a "one-shot" cooling method. Once all the ^{3}He boils off, the sample will be warmed above the temperature of the 1K pot by the thermal radiation from the IVC. In order to re-condense the ^{3}He, the charcoal is heated to $\sim$ 30 K by a non-inductively wound, resistive heater. The ^{3}He driven out of the charcoal will then condense on the inner wall of the 1K pot and drip down to the ^{3}He pot. Our system is designed to use 20 L of ^{3}He and can give a hold time $>$ 24 hours if a sufficient portion of the ^{3}He is condensed and if the overall heat load on the ^{3}He pot is minimized.

The hold time, however, is not the most important parameter for our group. As a group which is focused primarily on materials growth, the most important factor in the fridge design is the turn-around time necessary to change and cool new samples. Given that we typically grow $\sim$ 200 wafers a year and that most of these wafers require characterization at sub-K temperatures, a top-loading design is essential as it allows samples to be cooled down, measured, and warmed back up in one day. Returning again to figure 3.1, the sample is lowered into the fridge on the end of the long sample probe. This probe can be withdrawn into a load-lock on the top of the fridge so that the fridge itself can stay cold while samples are warmed up and replaced.

3.1.1 Wiring of the ^{3}He System

The wiring of the ^{3}He probe is relatively simple. The BNC break-out box on the electronics rack is connected by shielded, twisted pairs to Pi filters, with a 4 nF capacitance, on top of the probe. At the top of the sample probe the wires are fed into the probe through hermetically sealed Fischer connectors. Inside the probe, the wires to the sample, the sample thermometer (a calibrated RuO_2 resistor), and a red LED positioned over the sample are constantan twisted pairs which are threaded down the inside of the probe to the thermal anchor. At the thermal anchor they are varnished to a copper bobbin which acts as a 1.5 K heat sink when the heat sink is pressed tightly against the inner wall of the 1K pot. The wires are connectorized below the thermal anchor and then continue to the sample mount. The socket on the end of the probe has a 16-pin dual inline package (DIP) layout which is convenient for measuring cleaved Van der Pauw squares. The DIP socket can also be removed and replaced with an LCC socket which is convenient for characterizing processed samples requiring more connections.

3.2 Kelvinox 100 ^{3}He/^{4}He Dilution Refrigerator

3.2.1 Basic Operating Principle

The most commonly used cryostat for achieving temperatures below $\sim$ 300 mK is the ^{3}He/^{4}He dilution refrigerator. The basic operating principle of a dilution fridge is sketched in figure 3.2. The mixture first enters the fridge through the condenser line at a relatively high pressure ($\sim$ -28 inHg). It is first cooled to $\sim$ 4 K by the main ^{4}He bath in the dewar. It then enters the interior of the 1K pot (discussed in the previous section) where it is cooled to $\sim$ 1.7K and condenses. After leaving the 1K pot, the liquid flows through the primary impedance which is necessary to keep the pressure of the mixture in the 1K pot high enough to cause condensation. The liquid leaves the 1K pot and is cooled to $\sim$ 700-800 mK by the still (the operation of

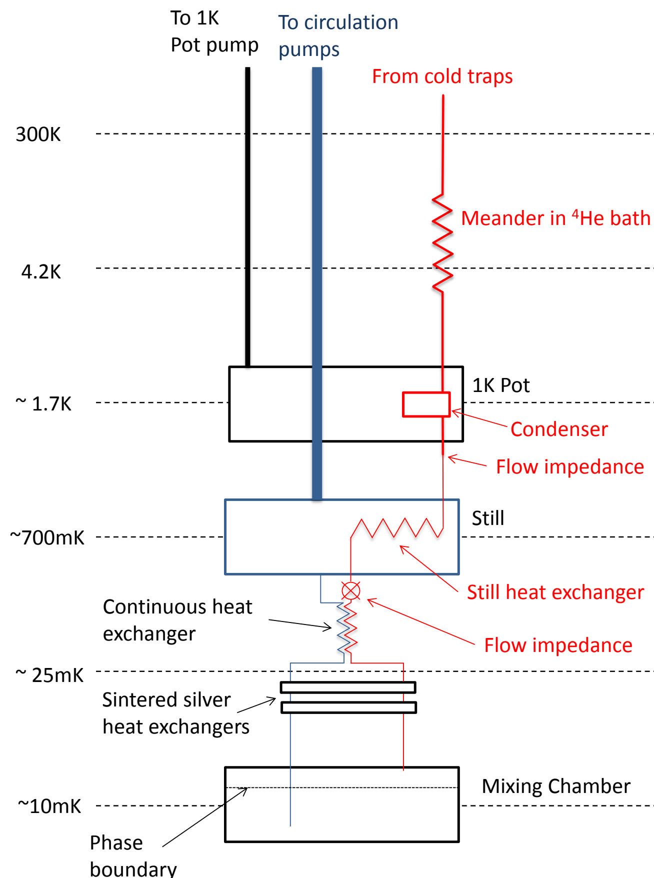

Figure 3.2. Schematic of the main components of a dilution refrigerator. Black dashed lines represent cold plates in the fridge to which heat sinks, thermometers, or heaters can be attached. Red lines represent incoming mixture while blue lines represent outgoing mixture. Line thickness indicates the relative size of tubing in the fridge.

the still will be discussed later). After being cooled in the still it encounters another impedance which is necessary to keep the mixture from evaporating at the relatively high temperature in the still [109]. Next, the mixture flows through a number of heat exchangers where it is cooled by the mixture leaving the mixing chamber. The first heat exchanger is the so-called counter-flow heat exchanger which is consists of a pair of coaxial tubes carrying the counter-flowing mixture. The final stage of heat exchangers are sintered silver exchangers. The sintered silver has an extremely large

surface area which, in addition to its high thermal conductivity, make it an ideal heat exchanger. The Manfra group fridge has only two sintered exchangers, but it is not uncommon for fridges to have several more. Finally, the mixture enters the coldest part of the fridge, the mixing chamber.

The key phenomena that makes the whole dilution fridge work is the finite solubility of ^{3}He in ^{4}He at low temperature. Below 0.87 K [109] a mixture of ^{3}He and ^{4}He will separate into two phases, one rich in ^{3}He and the other poor in ^{3}He. Due to its lower density, the ^{3}He-rich phase will float on top of the ^{3}He-poor phase. As the temperature is lowered, the molar fraction of ^{3}He in the ^{3}He-poor phase will decrease but saturate at 6.6% at zero temperature. At the lowest temperatures this means that if one ^{3}He atom were removed from the dilute phase another one would have to cross the phase boundary to take its place in order to maintain the minimum concentration of 6.6%. It turns out that the heat capacity of the dilute phase is larger than that of the concentrated phase [109], and as a result when the ^{3}He crosses the phase boundary the system is able to cool (or absorb heat if its temperature is fixed). A rough way of thinking of this process is to imagine the ^{3}He expanding into the "vacuum" of the inert, Bose-condensed ^{4}He background.

The key experimental challenge to take advantage of this cooling method is to design a way to extract ^{3}He from the dilute phase. This is accomplished by linking the still to the mixing chamber with a tube that extends below the phase boundary in the mixing chamber. At the high temperature in the still ($\sim$ 700 mK) the vapor pressure of ^{3}He is almost 3 orders of magnitude larger than that of ^{4}He [109]. Therefore, if the still is pumped, the gas removed from the still will be almost entirely ^{3}He. This will cause an osmotic pressure gradient to develop between the mixing chamber and still which will drive ^{3}He from the mixing chamber up to the still, thereby accomplishing the goal of removing ^{3}He from the mixing chamber and absorbing heat. As the circulation rate is increased, either by increasing the pumping speed of the pumps or by increasing the temperature of the still, the cooling power of the mixing chamber will increase. However, this will only reduce the temperature of the fridge to a point;

eventually the heat exchangers will not be able to sufficiently cool the incoming ^{3}He and further increases in circulation rate would cause the mixing chamber to warm. It is worth noting that the cooling power of the fridge, typically quoted at a temperature of 100 mK, is not a direct indication of the base temperature. In other words, a 400 μW fridge will not necessarily get four times colder than a 100 μW fridge.

Once the mixture leaves the fridge, it is compressed by the pumps and returned to the condenser. Unfortunately, the pumps and other room temperature vacuum connections can add impurities such as oil and air to the mixture. These impurities will quickly plug up the primary impedance in the condenser, so it is necessary to utilize cold traps to prevent the impurities from reaching the fridge. The first trap, cooled by liquid nitrogen, will catch the majority of the air and oil. However, the hot, oil-sealed pumps will also generate hydrogen over time, and this hydrogen will not be caught by the nitrogen trap. For this reason, it is also necessary to have a liquid helium cooled trap immediately upstream of the condenser. If the system is leak-tight it is possible in this manner to keep the fridge continuously cold (i.e. below 4 K) for many months.

3.2.2 Initial Construction

The Manfra group fridge came to Purdue after having been moth-balled for several years following its use in optical experiments. As a result, it required a bit of work to get it (and the lab) setup for making electrical transport measurements. Due to the system's large magnet (17T at 2.2K), Oxford recommended removing as much structural steel as possible from close proximity to the fridge. In addition, the low ceiling in the lab required that the dewar[1] be setup in a pit in order to extract the insert from the dewar. Thus setup of the system started with digging up an 8' × 8' section of floor around the fridge to remove the rebar and install a plastic insert to define the walls of the pit as shown in figure 3.3. The lab preparation phase also

[1]Custom made vapor-shielded dewar with 60 L belly from Precision Cryo, Indianapolis, IN.

Figure 3.3. Lab space during construction of the pit.

involved mounting a chain hoist[2] on the ceiling to simplify loading and unloading of the insert and the construction of an aluminum stand[3] for the dewar. The system at the end of the initial construction is shown in figure 3.4

3.2.3 Gas Handling System

The original gas handling system designed by Oxford did not make the trip to Purdue with the rest of the fridge, so we made our own manifolds to control the flow of the mixture. Figure 3.5 shows a schematic of the gas handling system (GHS). The GHS was designed with several features in mind. First, the system was designed such that the mixture could return to the storage dump via a 10 psi check valve in the event of a plug in the condenser or either of the cold traps. Second, the GHS can be setup to “wash” the mixture by bypassing the fridge while still circulating the

[2]McMaster-Carr PN 3287T61 trolley-mount chain hoist with 1100 lbs capacity. The large capacity of the hoist enabled us to use it to center the stand and dewar under the hoist during initial setup.
[3]Designed and built by Tom Halsmer in the physics machine shop. The assembly was simplified by the use of pre-formed 80/20® extruded aluminum beams.

Figure 3.4. Dewar, magnet, and stand in place at end of the intial construction phase.

mixture through the cold traps. This is useful as a way to check for leaks after the system has been sitting unused for an extended period. In addition, the pumping system was setup in such a way that the system could be pumped with the rotary pump alone or with the rotary pump in series with the high speed roots blower. This was necessary as we did not know initially what circulation rate would give the lowest base temperature; it has since been determined that running the two pumps in series without putting any heat into the still minimizes the base temperature. Moreover, in order to maximize the circulation rate, the vacuum lines from the still to the pumps were all large, NW-50 flanged hoses, and the valves between the still and pumps were all high-conductance gate valves. In order to minimize the long-term potential for leaks, all the connections used in the dashed box in figure 3.5 utilized metal compression fittings instead of rubber-gasket sealed NW connections which could dry out and crack over time. The final vacuum connections to the fridge were all made with plastic clamps and plastic centering rings to electrically isolate the fridge from the pumps. In addition, the vacuum lines from the still and 1K pot were clamped

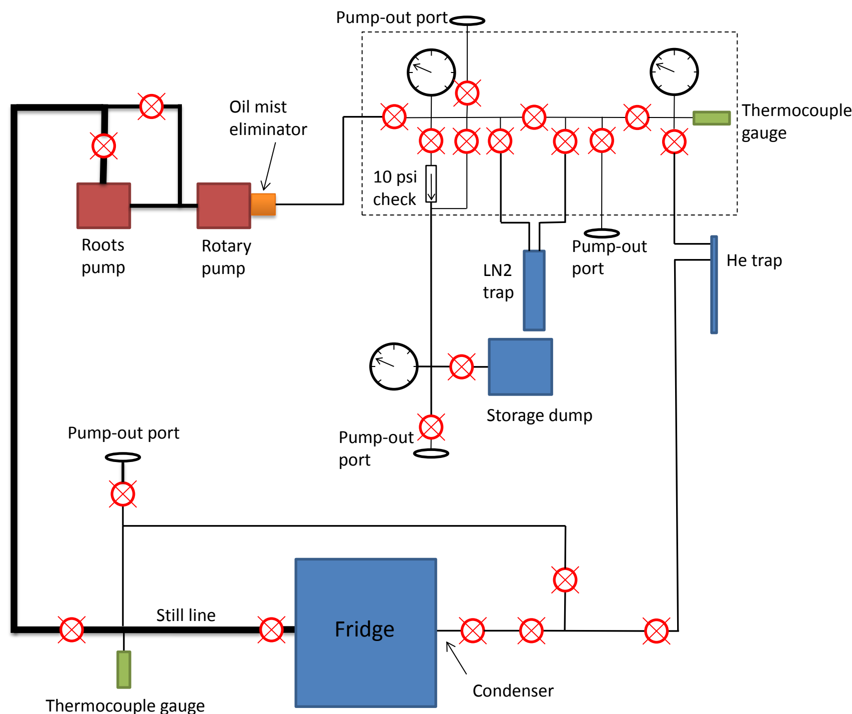

Figure 3.5. Schematic of the homemade gas handling system. Line thickness denotes size of vacuum hoses.

to a bucket of sand to minimize the propagation of vibrations from the pumps. A better design, however, would have been to bury the pumping lines in the sand since the original design did not significantly attenuate the vibrations in the pumping lines. That being said, the vibration level thus far has been sufficiently low that it has not hampered our measurements.

After assembling and leak checking the GHS, it was necessary to adjust the amount of ^{3}He and ^{4}He in the mixture as some mixture had evidently leaked out of the dump during the fridge's several-year-long hibernation. To optimize the circulation rate, mixture volume, and still power we first cooled the fridge without a tail to eliminate the chance of any thermal shorts between the tail and the radiation shield. In addition, this allowed us to test the effectiveness of the magnet's cancellation coil. The ratio of ^{3}He to ^{4}He was determined by one-shotting the fridge with the still temperature

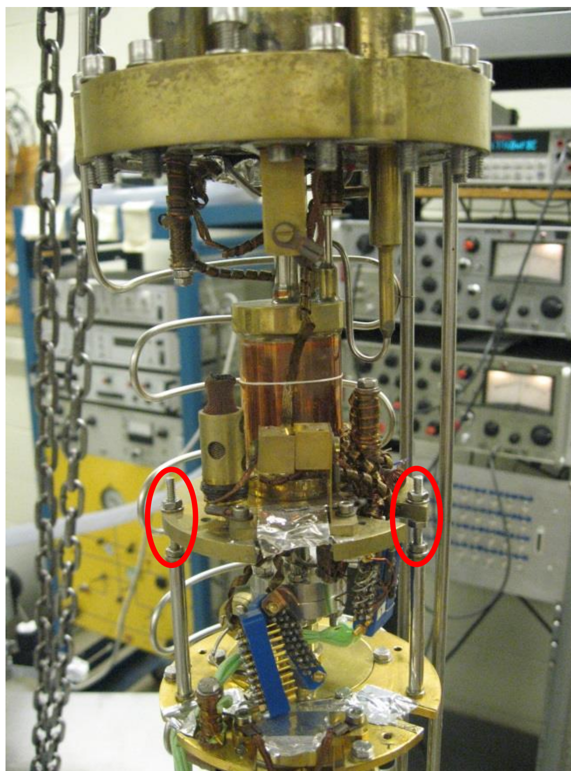

Figure 3.6. Threaded rods (circled in red) that were adjusted to center the radiation shield inside the IVC.

< 1 K. At this temperature the ^{3}He vapor pressure in the still is, of course, much higher than that of the ^{4}He, so once the still pressure dropped to zero it was safe to assume that the majority of the mixture returned to the dump was ^{3}He. The fridge plus GHS as originally designed by Oxford was specified to use 38 L of ^{4}He and 7 L of ^{3}He; however, we found the cooling power to be maximized using ~ 33 L ^{4}He and ~ 10 L ^{3}He. During the initial cool-down it was also determined that the radiation shield was touching the IVC as the base temperature would not drop below ~ 30 mK and the cold plate thermometer was stuck ≥ 160 mK. This was likely due to the fridge being stored horizontally for several years. To fix this, the two threaded rods shown in figure 3.6 were adjusted to center the radiation shield in the IVC. This was made much easier by the fact that the IVC had a window in the bottom of it

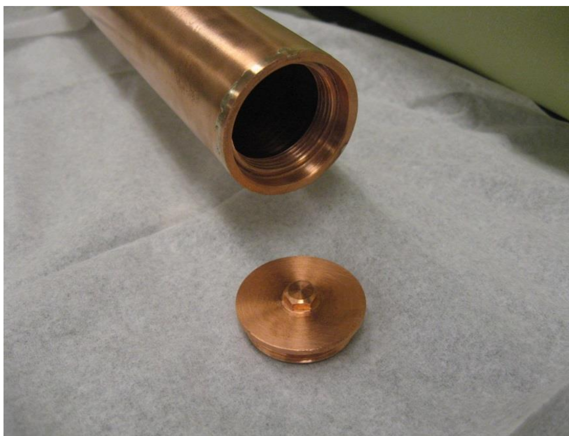

Figure 3.7. Threaded insert and plug used to block thermal radiation from reaching the sample.

(originally to allow laser illumination of the sample). In addition to re-centering the radiation shield, we machined a threaded insert and matching plug to seal the end of the radiation shield as shown in figure 3.7. The insert was hard-soldered in place[4], and the removable plug allowed the tail to be inspected for thermal shorts to the radiation shield. After removing the thermal short and adding helium to the mixture, the base temperature of the fridge was minimized ($\sim$ 10 mK) for zero power input into the still. Finally, during the initial cool-down, the mixing chamber temperature did not change for fields up to 10 T, indicating that the cancellation coil was quite effective at eliminating eddy current heating in the mixing chamber.

3.2.4 Wiring and Sample Mount

Starting from the BNC breakout box, the sample leads consisted of copper twisted pairs in two separate shielded cables with D-subminiature (D-SUB) style 25-pin connectors. On top of the fridge the cables were connected to so-called data transfer switches[5]. The data transfer switches allowed all of the leads going to the sample to

[4]Harris "Safety-Silv 56" 56% silver brazing alloy with Harris "Stay-Silv" flux.

[5]Manhattan DB25 data switch, www.qualitycables.com part number 150460

Figure 3.8. Homemade filter box that couples a D-SUB 25 cable to the Fischer connector on top of the fridge.

be grounded simultaneously to the fridge during cool downs when the cables to the breakout box would be too short to reach the top of the fridge. Following the data transfer switches, a short, shielded cable with twisted pairs connected the transfer switches to home made filter boxes (see figure 3.8). These filter boxes then plugged directly into the Fischer connectors on top of the fridge[6]. The filters were Pi-filters built into a D-sub connector for ease of fabrication[7].

Inside the fridge, twisted pairs of constantan wire connected the hermetically sealed Fischer connectors to connectors on the 1 K plate. The heat sinking at this stage and at the 4 K plate was part of the wiring originally supplied by Oxford and appeared to consist of wires wrapped around copper posts and varnished in place. The heat sinking below the 1 K plate was not deemed to be satisfactory, so we made our own heat sinks. The wires consisted of constantan twisted pairs in a cotton

[6]24 pin Fischer series 105 connector part number S 105 A093-80+ with clamp set part number E31 105.2/10.7+B, both purchased from Kensington Electronics, Austin, TX, www.keiconn.com

[7]Spectrum-Control part number 56-721-013 1 nF Pi filter with 3 dB point of 3.2 MHz and $>$ 70 dB attentuation above 1 GHz, purchased from Newark/Element i4 www.newark.com

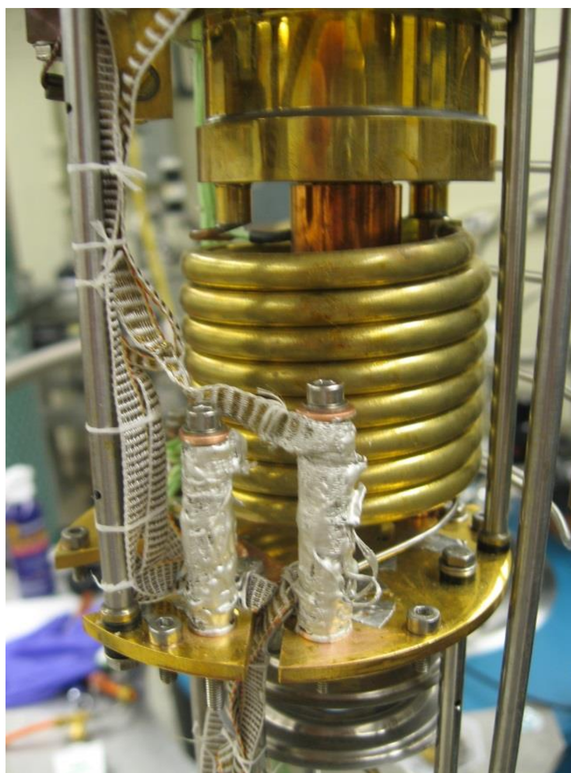

Figure 3.9. Home made heat sinks mounted on the 50 mK plate.

loom[8]. The heat sinks at the still plate, 50 mK plate, and mixing chamber plate all consisted of loom wrapped tightly around copper posts after being coated with silver epoxy[9]. The copper posts were polished with fine grit sandpaper and cleaned with isopropanol immediately prior to applying the epoxy to minimize thermal resistance from oxidation and debris. Figure 3.9 shows the heat sinks on the 50 mK plate. At the mixing chamber the constan wire was connectorized[10] and transistioned to copper twisted pairs in a cotton loom[11]. Care was taken during construction to ensure that the same sets of twisted pairs were maintained from the BNC breakout box all the way to the sample mount on the end of the tail.

[8]Part number A8-312 from Oxford Instruments www.cryospares.com

[9]Part number EJ2189-LV from Epoxy Technology www.epotek.com

[10]Cinch Connector Inc. Lombard, IL. Mating part numbers DCDM25PSB and DCDM25SSB.

[11]Oxford Instruments part number A8-311 www.cryospares.com

Figure 3.10. Tail after etching in 1:1 water:nitric acid.

The first tail that we fabricated to hold the samples was a simple design consisting of an OFHC copper rod that threaded into the bottom of the mixing chamber with a copper stud. The wires were epoxied to the side of the tail with the same silver epoxy used for the heat sinks at higher stages. However, this design had two flaws. First, due to the soft copper threads and the small diameter of the rod, the rod could not be fastened to the mixing chamber well enough to completely thermalize the bottom of the tail with the mixing chamber. Second, upon cooling, the epoxy contracted significantly more than the copper and, as a result, pulled away from the copper rod.

In order to solve these problems, several changes were made. Construction of a new tail began by designing a large copper plate that could be tightened against the mixing chamber with stainless steel screws passing through clearance holes rather than tapped copper holes. This allowed the stainless steel screws to be tightened with stainless steel nuts and eliminated the risk of stripping soft copper threads. This plate was then welded to the rest of the tail to ensure a strong thermal link.

In order the clean the copper after welding, the copper was etched in 1:1 water:nitric acid for 3 minutes. This removed the heavy, black oxide left from the welding process, but a different oxide started to form during rinsing as shown in figure 3.10. This oxide was then cleaneded off with Brasso metal polish followed by toluene,

Figure 3.11. Tail after cleaning with Brasso metal polish, toluene, acetone, and methanol.

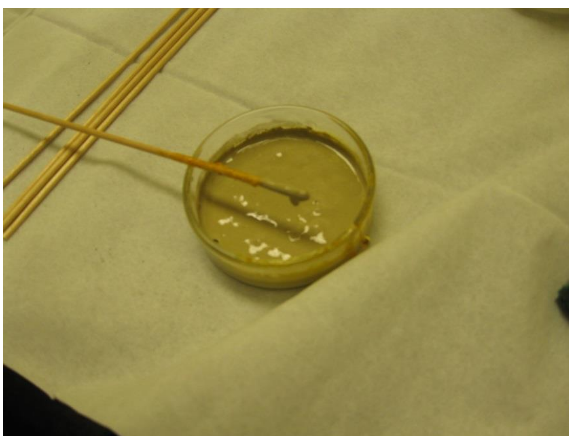

Figure 3.12. Slurry of GE varnish and silver powder used to heat sink copper wires to tail.

acetone, and methanol. Immediately following this cleaning, the copper wires were pasted in place with GE Varnish[12] infused with silver powder[13]. The varnish was diluted with 1:1 toluene:methanol in order to extend its working life. The slurry of varnish and silver powder, shown in figure 3.12 had the consistency of molasses when

[12]VGE-7031 varnish from Lakeshore Cryotronics www.lakeshore.com.

[13]Part number 61-310 from Ted Pella, Inc. www.tedpella.com.

Figure 3.13. Finished tail mounted on the mixing chamber.

first mixed and an electrical resistance of $\sim$ 1 MΩ aross $\sim$ 1 inch of the petri dish shown in figure 3.12. The resistance dropped to $\sim 0.5\ \Omega$ across 1 inch after the slurry dried. As soon as the wires were pasted in place and covered with more varnish, the top plate (shown in the background of figure 3.11) was placed on top of the wires and screwed in place with stainless steel screws. Figure 3.13 shows the completed tail mounted on the fridge. Prior to mounting the tail on the fridge, both the mixing chamber plate and the mounting plate on the tail were polished with 3000 grit sandpaper and cleaned with acetone and methanol. The sandpaper was mounted on a smooth wooden block to ensure that the polishing did not de-planarize either surface.

This new and improved tail design significantly improved the thermal contact of the samples to the mixing chamber. However, the large amount of copper in the magnetic field resulted in a large nuclear demagnetization effect. While this was in

some sense useful for reaching lower temperatures, it made it difficult to assign a single temperature to data from large field sweeps. What is more, we have found that the coldest looking data typically come from up-field sweeps that closely follow down-field sweeps. We speculate that this is due to the long thermalization time constants present at base temperature. In other words, during the down-sweep the tail demagnetizes, but the electron temperature does not change significantly until near the end of the sweep. If the field is then swept back up, the electrons stay at this colder temperature for much of the up-sweep. Future tail designs would probably benefit in this regard by using silver as in reference [111] due to its low nuclear magnetic moment.

3.2.5 Thermometry

The temperatures at the various cold stages on the fridge were measured with resistance thermometers. Uncalibrated RuO_2 resistors with typical resistance vs. temperature profiles were used to estimate the temperature at the 1K plate, still plate, and 50 mK plate. The temperature at the mixing chamber and on the end of the tail were measured with calibrated RuO_2 thermometers.[14] The mixing chamber thermometer was mounted using the original copper package supplied by Oxford. The tail thermometer was glued to a small copper tab with cigarette paper soaked in GE varnish. The copper tab was then screwed into the end of the tail close to the sample holder. The resistor was thermally cycled between 77 K and room temperature four times and was additionally cycled between 4 K and room temperature four times prior to calibration to minimize the potential changes in calibration during subsequent cool-downs.

The wiring of the thermometers in the fridge was not modified from what was originally supplied by Oxford. Heat sinking at each cold plate was accomplished by varnishing the wires (presumably constantan twisted pairs) to copper posts. In order

[14] A special thank you to Gabor Csáthy and Ethan Kleinbaum for calibrating our tail thermometer in their fridge.

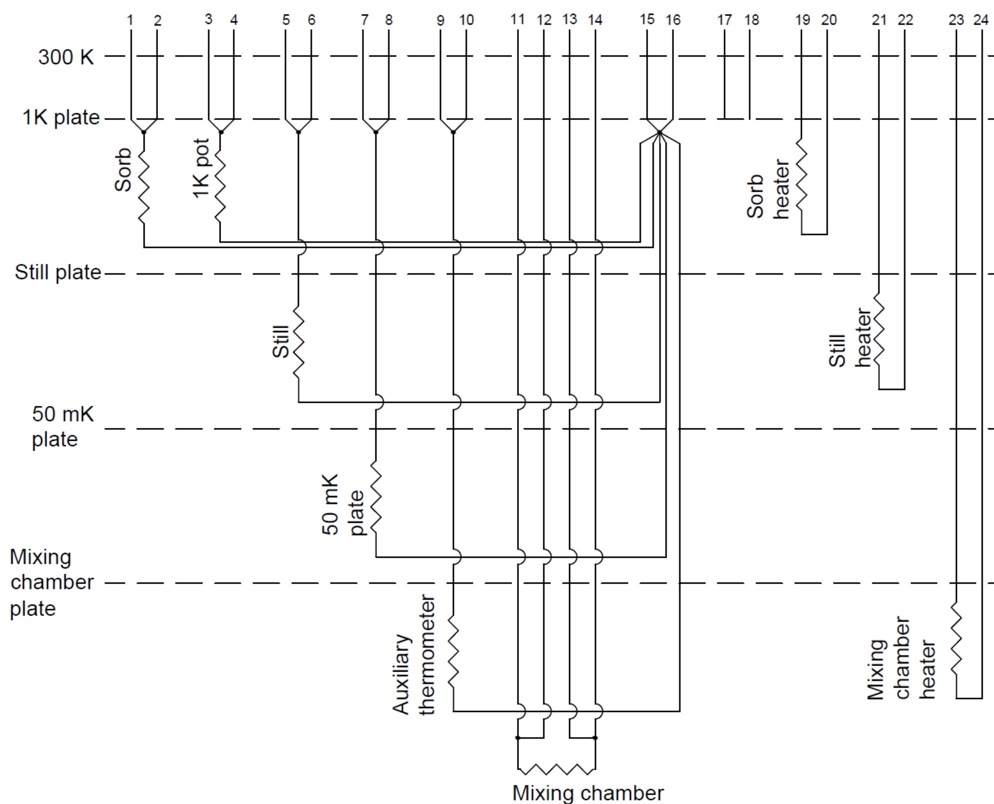

Figure 3.14. Wiring diagram of the thermometers and heaters in the fridge.

to fit all the wires in a 24 pin Fischer connector, the uncalibrated thermometers were wired with a common ground as shown in figure 3.14. The tail thermometer is not shown in figure 3.14; it was wired using extra wires on the tail which were heat sunk and filtered as described earlier.

The thermometers shown in figure 3.14 were measured using an AVS-47 resistance bridge with its associated pre-amplifier connected to the fridge with a shielded cable while the tail thermometer was measured with a PAR 124A lock-in amplifier. To avoid self-heating at base temperature, the mixing chamber thermometer was measured with a 3 μV excitation from the resistance bridge and the tail thermometer was

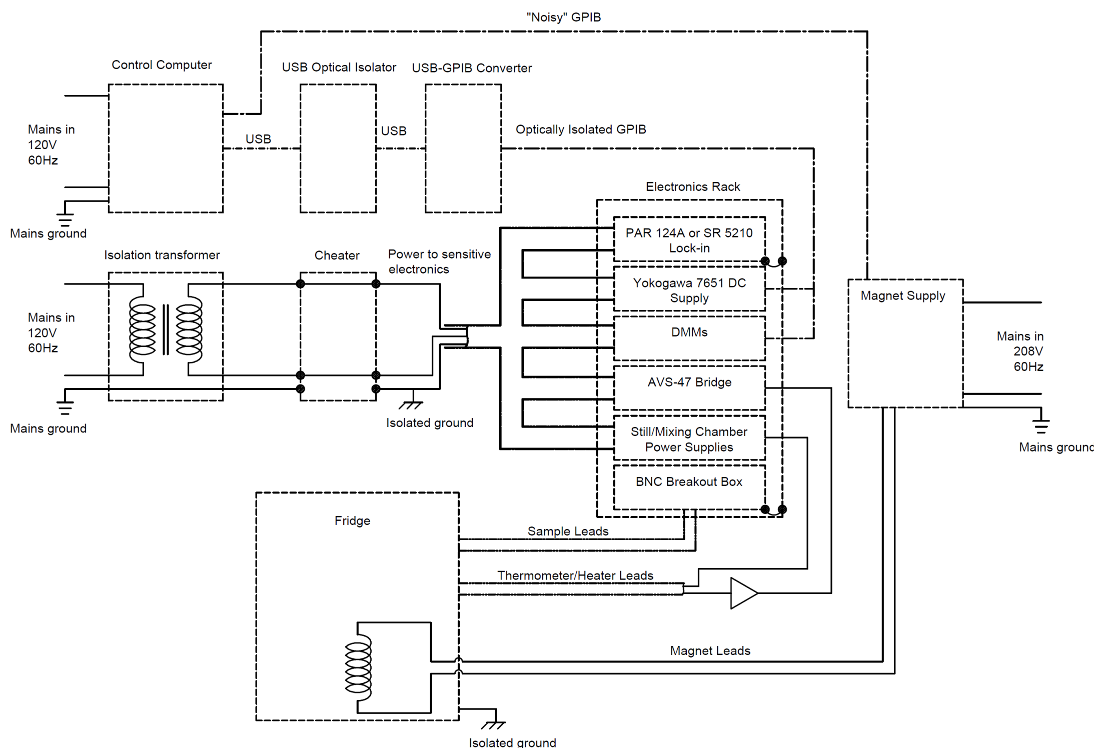

Figure 3.15. Schematic of electronics setup and grounding.

measured with a 100 pA excitation from the lock-in amplifier. This corresponded to power dissipations of 100 aW and 1 fW, respectively.

It turned out that the resistance bridge was a very sensitive detector of ground loops and interference. If the mixing chamber thermometer was measured with a 3 μV excitation, the bridge would easily overload in the presence of grounding issues or noise from other electronics. In addition, the monitor output channel on the bridge could be viewed with an oscilloscope to check for the presence of noise signals that were present but too small to overload the bridge. This served as a convenient diagnostic for setting up the electronics. The first step in achieving a low-noise setup was to electrically isolate the fridge from the pumps by using plastic centering rings and plastic clamps on all the pumping lines. The second adjustment that improved the

noise level in the resistance bridge was to isolate the bridge, lock-in amplifiers, fridge power supplies, and multimeters (used to measure the lock-ins and bridge) from the mains ground. This was accomplished by using an isolation transformer and breaking the ground on the secondary as shown in figure 3.15. We also optically isolated the GPIB communications cables to the sensitive electronics using a combination of a USB optical isolator[15] and a USB-GPIB converter[16]. Finally, since the magnet leads were not grounded to the fridge, we did not install an isolation transformer and cheater combination to isolate the magnet power supply[17]. Consequently, it was necessary to run the magnet supply off a different, non-isolated GPIB bus to keep the grounds separated.

While this setup eliminated problems due to ground loops, it was also necessary to eliminate noisy electrical instruments from the system. The first two methods we used in trying to isolate the GPIB actually made the noise level in the bridge worse. The first GPIB isolator we tried[18] caused the bridge to overload as soon as the isolator was turned on, even if it was not connected to any electronics in the rack. The first USB optical isolator we tried[19] in conjunction with the GPIB-USB adapter also behaved as a source of high frequency noise. The bridge was also quite sensitive to noise from the original fridge power supply[20]; as soon as the still power was turned on the bridge would overload. Despite many attempts at changing the grounding and location of the power supply, the problem persisted, likely due to capacitive coupling between the heater leads and the thermometer leads in the cabling leading to the shared Fischer connection on top of the fridge. As a result, we were forced to switch to powering the still and mixing chambers with small, analog power supplies[21]. Finally, the first DC power supply[22] we tried to use with gated samples also resulted in a large high

[15]Sealevel Hub7i optically isolated USB hub from www.sealevel.com.
[16]National Instruments GPIB-USB-HS adapter.
[17]Oxford IPS 120-10 power supply.
[18]National Instruments GPIB-120B bus isolator/expander.
[19]Keterex 751-KXUSB-150
[20]Oxford PS2603 fridge power supply.
[21]HP 6218A DC power supply.
[22]Keithley 2612B source-meter.

frequency noise signal on the bridge and significant heating of the thermometers. Once again, changes to the location and grounding of the power supply did not improve the interference, and we were forced to switch to a different DC supply[23].

[23]Yokogawa 7651 programmable DC supply.

4. Scattering Mechanisms in a High-Mobility Low-Density Carbon-Doped (001) GaAs Two-Dimensional Hole System

J. D. Watson[1], S. Mondal[1], G. A. Csáthy[1], M. J. Manfra[1,2,3,], E. H. Hwang[4], S. Das Sarma[4], L. N. Pfeiffer[5], K. W. West[5]

[1]*Department of Physics and Birck Nanotechnology Center, Purdue University West Lafayette, IN 47907, USA*

[2]*School of Electrical and Computer Engineering, Purdue University West Lafayette, IN 47907, USA*

[3]*School of Materials Engineering, Purdue University West Lafayette, IN 47907, USA*

[4]*Condensed Matter Theory Center and Department of Physics University of Maryland, College Park, MD, 20742, USA*

[5]*Department of Electrical Engineering Princeton University, Princeton, NJ 08544, USA*

[1] *Abstract:* We report on a systematic study of the density dependence of mobility in a low-density Carbon-doped (100) GaAs two-dimensional hole system (2DHS). At T=50 mK, a mobility of 2.6 × 10^6 cm^2/Vs at a density p=6.2 × 10^{10} cm^{-2} was measured. This is the highest mobility reported for a 2DHS to date. Using a back-gated sample geometry, the density dependence of mobility was studied from 2.8 × 10^{10} cm^{-2} to 1 × 10^{11} cm^{-2}. The mobility vs. density cannot be fit to a power law dependence of the form $\mu \sim p^{\alpha}$ using a single exponent α. Our data indicate a continuous evolution of the power law with α ranging from $\sim$ 0.7 at high density and increasing to $\sim$ 1.7 at the lowest densities measured. Calculations specific to our structure indicate a crossover of the dominant scattering mechanism from uniform background impurity scattering at high density to remote ionized impurity scattering at low densities. This is the first observation of a carrier density-induced transition from background impurity dominated to remote dopant dominated transport in a single sample.

[1]This chapter is adapted with permission from Phys. Rev. B **83**, 241305(R) (2011). Copyright (2011) American Physical Society.

The two-dimensional hole system (2DHS) offers an attractive platform for the study of strong carrier interactions parameterized by r_s : the ratio of the Coulomb energy to the Fermi energy. $r_s = E_c/E_f \propto m^*/\sqrt{p}$, where p is the hole density and m^* is the effective mass. Recent developments in the growth of Carbon-doped (100) GaAs heterostructures by molecular beam epitaxy (MBE) have resulted in 2DHSs of unprecedented quality [112, 101]. Such structures have been utilized in the study the metal-to-insulator transition (MIT) [93], fractional quantum Hall physics in the 2nd Landau level (LL)[113], spin-orbit coupling in Aharonov Bohm rings [114] and charge density wave formation in partially filled LL's[115],[116]. These initial experiments and the prospect of studying strong correlations in the presence of tunable spin-orbit coupling provide strong motivation to understand the scattering processes presently limiting mobility in the highest quality samples. Here we present mobility vs. density data on an unprecedently high mobilty 2DHS. One of the most exciting avenues for future research is the investigation of ultra-low density 2DHSs at very large r_s. Thus, our data and calculations will inform the design of new hole heterostructures of ever increasing quality.

Carbon doping[101] of 2DHSs offers advantages over the more commonly used acceptor dopants Beryllium and Silicon. Carbon diffuses and surface segregates much less at typical MBE growth temperatures (T∼630 °C) than Beryllium[77]. Additionally Carbon can be incorporated as an acceptor on multiple crystallographic orientations, including on the high-symmetry (100) face of GaAs. Silicon can also act as an acceptor to produce high quality 2DHSs but so far high mobility ($\mu \sim 10^6$ cm^2/Vs) Silicon-doped 2DHSs have only been realized on (311)A face[97]. The (311)A face has a well known mobility anisotropy between the [$\bar{2}$33] and [01$\bar{1}$] directions [97] whereas Carbon-doped structures on the (100) face have a significantly lower anisotropy between the [011] and [0$\bar{1}$1] directions [101]. Furthermore, the high symmetry of the (100) orientation dramatically alters the nature of spin-orbit interactions in 2DHSs as compared to quantization along the (311)A direction. Indeed, further experimental work is needed to fully exploit the potential benefits of Carbon-doped (100) 2DHSs.

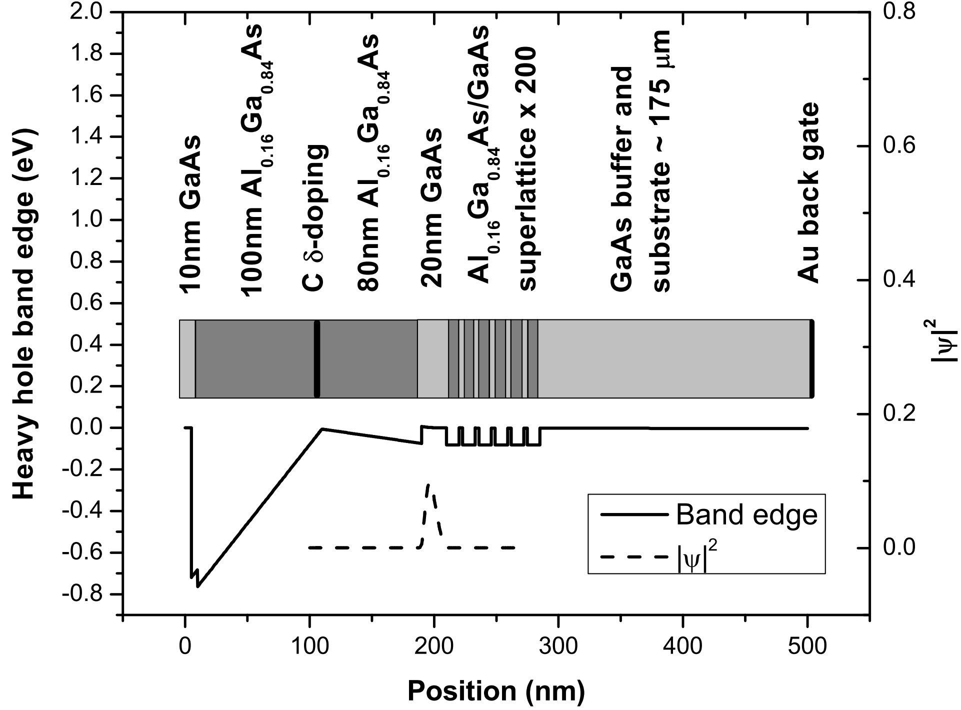

Figure 4.1. Schematic of the device structure used in our experiments.

Our sample consists of a 20 nm $Al_{0.16}Ga_{0.84}As/GaAs/Al_{0.16}Ga_{0.84}As$ quantum well asymmetrically modulation doped with Carbon at a density of 1×10^{12} cm^{-2} above the quantum well at a setback of 80 nm. FIG. 4.1 shows a sketch of the device along with the numerically calculated [117] band structure and heavy hole ground state wavefunction (normalized to unity). For simplicity in simulation the superlattice and buffer regions are truncated. In order to modulate the density in the quantum well, we utilized a back-gate geometry. The sample was first thinned to approximately 150 μm and then cleaved into a Hall bar approximately 2 mm $\times$ 9 mm. Ohmic contacts consist of In/Zn dots positioned approximately 1 mm apart along the length of the Hall bar and annealed at T = 430 °C. The hall bar was subsequently fixed to a gold backgate evaporated on an undoped GaAs substrate. The carrier density was measured from minima in the longitudinal magnetoresistance, and the conductivity was obtained from four-terminal zero field measurements using standard lock-in techniques. As

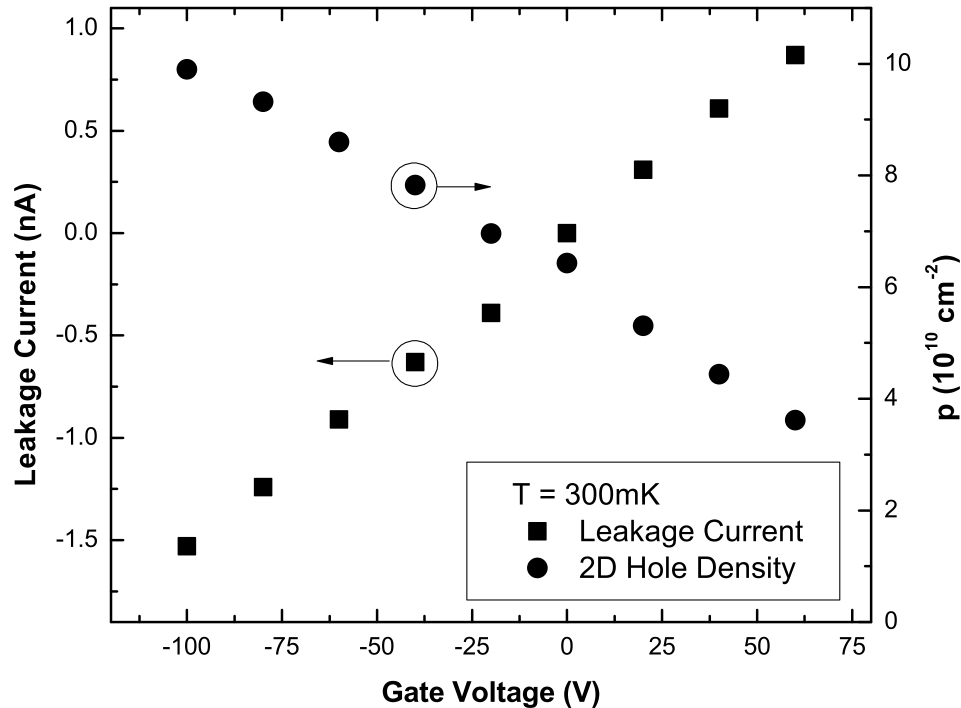

Figure 4.2. Effect of gate voltage on carrier density and leakage current at T = 300 mK.

shown in FIG. 4.2, the 2DHS density depended linearly on voltage over the range measured. Modeling the structure as a parallel plate capacitor with one plate being the 2DHS and the other being the backgate we estimate the gate to be situated 175 μm from the well. The peak mobility μ at low temperature (T = 50 mK) was measured to be 2.6×10^6 cm^2/Vs at a density of 6.2×10^{10} cm^{-2} in an as-grown sample.

FIG. 4.2 also shows the leakage current as a function of the gate voltage. The linear dependence of the leakage current on the gate voltage and its small magnitude ($<$ 1.5 nA as compared to an excitation current of 50 nA) suggest that the observed leakage represents parasitic current passing through the measurement circuit and not hard breakdown in the GaAs. In addition, the linear dependence of the density on the gate voltage also suggests that sharp breakdown did not occur. FIG. 4.3 shows a

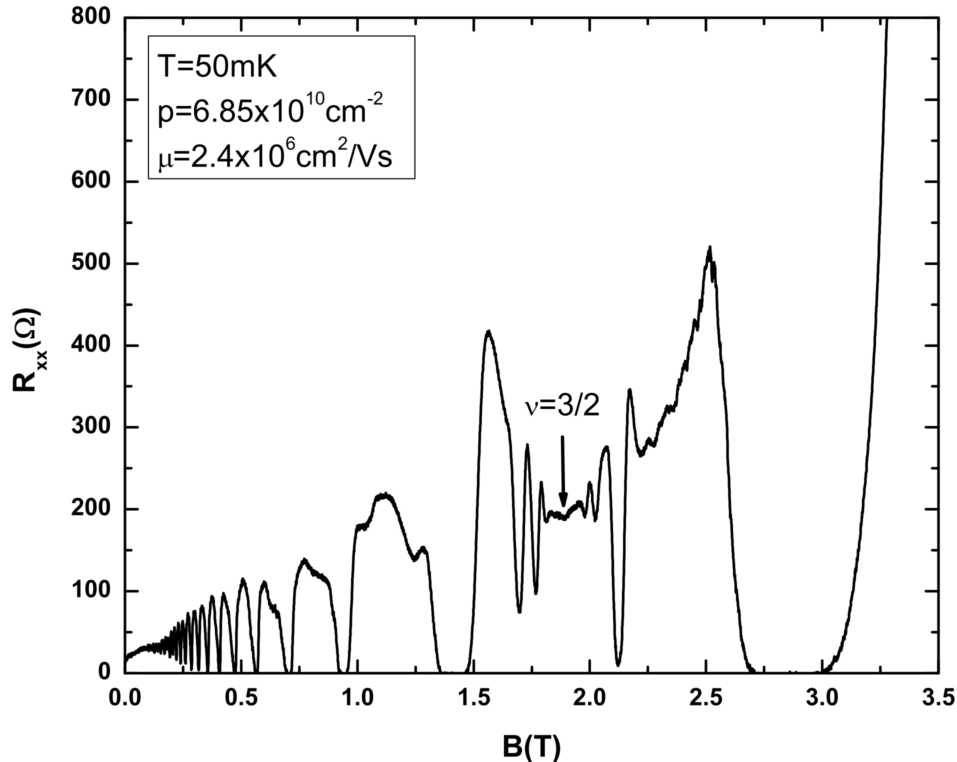

Figure 4.3. Low field magnetoresistance of the backgated sample at T = 50 mK.

representative trace of the magnetoresistance at T = 50 mK. The deep minima in the fractional quantum Hall states around filling factor $\nu = \frac{3}{2}$ illustrate the high quality of the processed sample. We note that this sample has also been studied at ultra-low temperatures (T $\leq$ 10 mK) in which the first evidence of a fully formed fractional quantum Hall state at $\nu = 8/3$ in the 2nd Landau level in a 2DHS was observed [113].

In order to examine the scattering mechanisms limiting mobility in our system, we measured the dependence of the mobility on the 2D hole density modulated by the backgate as shown in FIG. 4.4. As can be clearly seen on this log-log plot the data points do not fall on a straight line as would be expected for a single dominant scattering mechanism. The mobility vs. density cannot be fit to a power law depen-

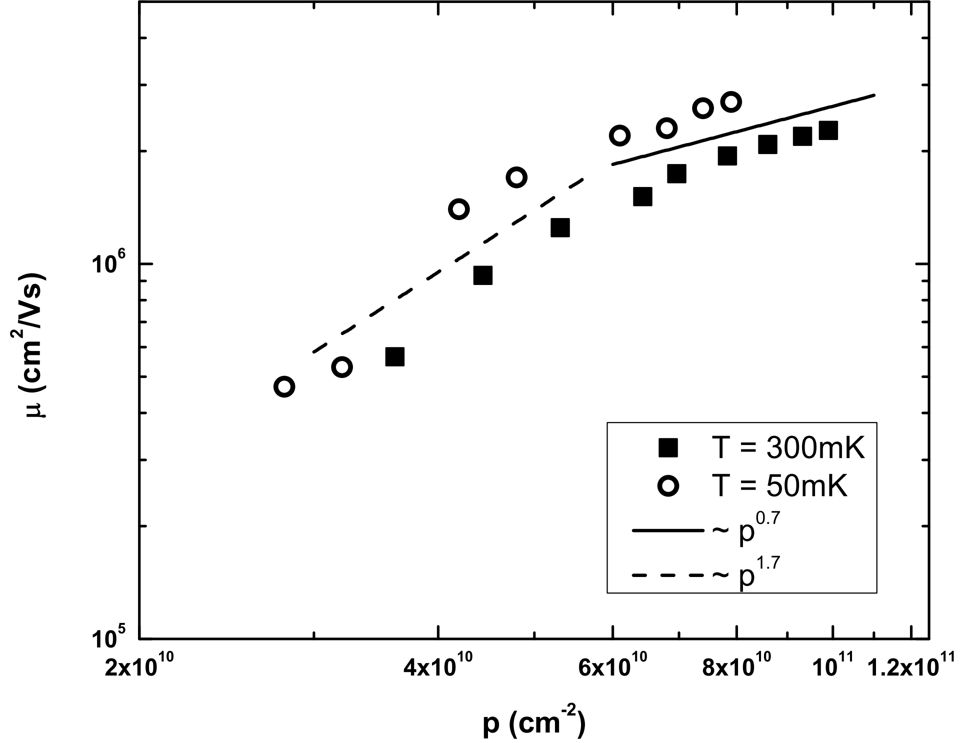

Figure 4.4. Mobility as a function of the density at T = 300 mK (squares) and T = 50 mK (open circles). Straight lines are guides to the eye to the 300 mK data to illustrate 0.7 and 1.7 power laws.

dence of the form $\mu \sim \mathrm{p}^{\alpha}$ using a single exponent α. Our data indicate a continuous evolution of the power law with α ranging from ~ 0.7 at high density and increasing to $\sim$1.7 at the lowest densities measured. Thus the data indicate the presence of multiple dominant scattering mechanisms over the range of density tested. Indeed at the lowest densities measured, the mobility decreases rapidly indicating that the system will eventually approach a finite density MIT[93]. We emphasize, however, that $k_F l$, the product of the Fermi wavevector and the mean free path, remains larger than 50 over the entire range of density tested. It can be seen at high density that the mobility follows a power law behavior $\mu \propto p^{0.7}$ which is indicative of uniformly distributed charged background impurity (BI) scattering [118], [66] in 2D carrier systems. However, the power law continuously shifts towards higher α of approximately

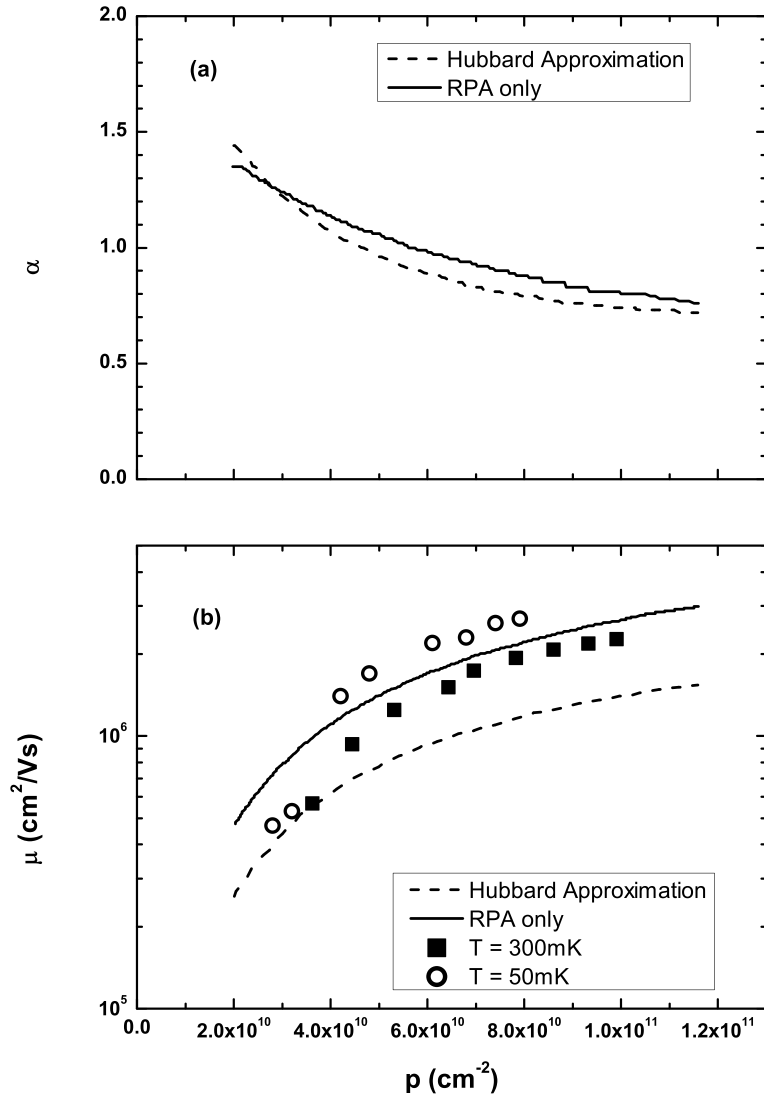

Figure 4.5. (a) Theoretical density dependence of the exponent α in $\mu \propto p^{\alpha}$. (b) Comparison of experimental mobility data and theoretical results. Solid line represents RPA-Boltzmann calculation and dashed line represents RPA-Boltzmann calculation with the Hubbard approximation in both plots.

1.7 at low density. Typically, an exponent $\sim 1.5 - 1.7$ is taken as an indication of the dominance of remote ionized (RI) impurity scattering[118],[119] originating in the remote doping layer in 2D systems. We emphasize that it is unusual that such a crossover behavior can be seen in a single sample while remaining in the high mobility

(or equivalently large $k_F l$) regime. It is interesting to note that our setback distance is 80 nm, a distance at which remote ionized impurity scattering usually makes a relatively minor contribution to the total scattering in samples with density > 10^{11} cm^{-2}. Nevertheless, it is clear that for the material parameters of our structure it dominates scattering at lower density. A similar transition to remote ionized impurity scattering in low density 2D electron samples was report by Jiang et al. [120], but in these samples $k_F l$ was substantially lower and the samples were approaching a conduction threshold. Unusual transport behavior was also recently reported in the 2DHS in an undoped electron-hole bilayer [121] and subsequently explained by carrier inhomogeneities resulting from strong carrier-carrier interactions and non-linear screening [122]. It could be argued that T = 300 mK data is only marginally outside the range where phonons could be playing some role, but the fact that our T = 50 mK data displays the same functional dependence seems to rule out any critical scattering contributions from phonons. Interface roughness (IR) scattering should also be considered, but significant IR scattering should also manifest itself in a significant mobility anisotropy [66] which was not observed in our sample. Nevertheless we cannot rule out possible contributions from IR and alloy scattering as the sample is gated to higher densities. Finally, scattering between the light and heavy hole bands[112] as well as changes in the effective mass [108] have also been reported. However, both of these effects should decrease as the density is lowered and the Fermi level moves towards the top of the valence band and away from light hole band.

To explain our experimental observation we calculated the Coulomb scattering rate due to background charged impurity scattering and remote charged impurity scattering using a Boltzmann transport method[123], [124]. Screening was taken into account using the random phase approximation (RPA). The calculation was performed both with and without the inclusion of correlation effects via the Hubbard approximation. The analysis assumed a 3D background impurity concentration of $n_{i3D} = 3 \times 10^{13}$ cm^{-3}, dopant setback from the center of the quantum well $d = 90$ nm, hole effective mass $m^* = 0.4 m_e$ where m_e is the free electron mass, quantum well

width $a = 20$ nm, and a remote ionized impurity concentration $n_i = 2 \times 10^{11}$ cm^{-2}. We note that m^* has been measured for our sample by cyclotron resonance to be $m^* = 0.4m_e$[107]. FIG. 4.5(a) shows a crossover of the mobility exponent changing from $\alpha \sim 0.7 - 0.8$ to $\alpha \sim 1.5 - 1.7$ which has the same qualitative behavior as seen in the data in FIG. 4.4. The results of the calculations are compared with the experimental data in FIG. 4.5(b). Qualitatively, as the hole density is lowered, screening of the remote dopants by the 2D hole gas becomes less effective, and the dominant scattering mechanism transitions from being dominated by uniform background impurities to being dominated by remote dopants. Such a transition point in the density is governed entirely by n_{i3D}, d, and n_i, and thus it is not surprising that this is the first time (to our knowledge) that such a transition has been observed in a single sample within the high mobility regime.

To understand the transition in the transport mechanism observed in our data, it is important to realize[122] that the impurity scattering strength in the transport theory carries the form-factor $\exp(-2k_F d)$ at low temperatures where $2k_F$ is the typical momentum transfer for resistive scattering by impurities, and d is the typical separation of the impurities from the 2D carrier layer. Since $k_F \sim p^{1/2}$, it is clear that lowering the carrier density would lead to stronger scattering by remote impurities which is exponentially suppressed at higher values of $k_F d$. For a screened Coulomb potential with two impurity contributions (remote and background charged impurity) we can derive the approximate qualitative formula for the mobility

$$\mu \propto \frac{(k_F d)^3 q_{TF}}{n_i + A n_{i3D} (k_F d)^3 q_{TF} / (2k_F + q_{TF})^2} \tag{4.1}$$

where A is a density independent constant and $q_{TF} = 2/a_B$ is the Thomas Fermi screening wave vector with the effective Bohr radius $a_B = \hbar^2/m^* e^2$. In the high-density limit, $k_F d \gg 1$ the mobility is proportional to the square of the sum of two wave vectors, i.e., $\mu \propto (2k_F + q_{TF})^2$. However, as $k_F d$ (or, density) decreases the mobility behaves approximately $\mu \propto (k_F d)^3$. Thus, as long as strong localization does not set in, which is the usual situation for lower mobility samples[120], lowering carrier density should always lead to a continuous increase of the exponent α as

scattering from the remote dopant impurities becomes important. This is exactly the theoretical behavior predicted in the theory (Fig. 4.4), and experimentally observed in our extremely high-mobility hole sample.

In conclusion, we measured the density dependence of mobility in a very high quality 2DHS. The 50 mK mobility was found to be 2.6×10^6 cm^2/Vs at a density of 6.2×10^{10} cm^{-2} in a pristine sample. The mobility appears to be limited by background charged impurity scattering at high density but in the low density regime is a stronger function of the density indicating an increasingly important scattering contribution from remote impurities. From this data, we can surmise that increased 2DHS mobility at low density can be realized by significantly increasing the spacer thickness. Our work also demonstrates that in samples of sufficiently high quality, where the 2D MIT transition is pushed down to very low carrier densities, the theoretically predicted continuous transition from background impurity scattering limited transport to remote dopant scattering limited transport can be quantitatively verified by decreasing the carrier density in a single sample. We are currently exploring the structural parameter space of Carbon-doped 2DHSs in order to optimize mobility.

Acknowledgment

JDW is supported by a Sandia Laboratories/Purdue University Excellence in Science and Engineering Fellowship. MJM acknowledges support from the Miller Family Foundation. The work by Maryland (EH and SDS) is supported by Microsoft Q. GAC acknowledges support from the NSF DMR-0907172.

5. Exploration of the Limits to Mobility in Two-Dimensional Hole Systems in GaAs/AlGaAs Quantum Wells

J. D. Watson[1,2,], S. Mondal[1,2], G. Gardner[2,3], G. A. Csáthy[1], and M. J. Manfra[1,2,3,4]

[1] *Department of Physics, Purdue University, West Lafayette, IN 47907, USA*

[2] *Birck Nanotechnology Center, Purdue University, West Lafayette, IN 47907, USA*

[3] *School of Materials Engineering, Purdue University, West Lafayette, IN 47907, USA*

[4] *School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA*

[1] *Abstract:* We report on the growth and electrical characterization of a series of two-dimensional hole systems (2DHSs) used to study the density dependence of low temperature mobility in 20 nm GaAs/AlGaAs quantum wells. The hole density was controlled by changing the Al mole fraction and the setback of the delta-doping layer. We varied the density over a range from 1.8×10^{10} cm^{-2} to 1.9×10^{11} cm^{-2} finding a nonmonotonic dependence of mobility on density at T = 0.3 K. Surprisingly, a peak mobility of 2.3×10^{6} cm^2/Vs was measured at a density of 6.5×10^{10} cm^{-2} with further increase in density resulting in reduced mobility. We discuss possible mechanisms leading to the observed non-monotonic density dependence of the mobility. Relying solely on interface roughness scattering to explain the observed drop in mobility at high density requires roughness parameters which are not consistent with measurements of similar electron structures. This leaves open the possibility of contributions from other scattering mechanisms at high density.

Two-dimensional hole systems (2DHSs) on (001) oriented GaAs offer an interesting alternative to the more widely studied two-dimensional electron systems (2DESs). 2DHSs on (001) GaAs have effective masses roughly 4.5 to 7.5 times larger[108, 107, 106] than that in corresponding 2DESs which increases the importance of Coulomb

[1]This chapter is adapted with permission from Phys. Rev. B **85**, 165301 (2012). Copyright (2012) American Physical Society.

interactions relative to the kinetic energy resulting in enhancement of importance of many-body effects. In addition, the p-wave symmetry of the valence band in GaAs leads to a much reduced hyperfine coupling of hole spins to the atomic nuclei which makes them an exciting alternative to electrons for quantum dot spin-based qubits[94, 95, 96]. The presence of spin-orbit coupling and light/heavy hole mixing in the valence band of GaAs also allows extensive band structure engineering[125, 126, 127]. This feature has been exploited to alter the nature of groundstates in the quantum Hall regime.[116, 115]

Here we describe our efforts to understand the limits to low temperature mobility for (001) 2DHSs. Continued improvement in 2DHS quality is motivated by the well-established paradigm for 2DESs that increased low-temperature mobility often leads to the observation of new correlated groundstates[68]. Historically, improvement to the low temperature mobility of 2DHSs has lagged behind that of 2DESs due to the lack of a p-type dopant in GaAs that does not diffuse or segregate significantly at typical molecular beam epitaxy (MBE) growth temperatures ~ 635 °C. Si can act as a low-diffusivity acceptor on the (311)A face of GaAs, but subsequent transport experiments are known to be complicated by a significant mobility anisotropy due to surface corrugation[97]. However, recent use of low diffusivity carbon doping (C-doping)[100, 101, 112] has rapidly led to low temperature mobilities $> 10^6$ cm^2/Vs without the accompanying transport anisotropy. Purely from a growth standpoint, then, there does not appear to be any reason why low temperature hole mobilities should not approach that of electrons once scaled by the appropriate effective mass. Presently it is widely believed that uniformly distributed ionized background impurities limit the mobility in the best 2DESs [118]. However, the highest hole mobility reported to date[128] of 2.6×10^6 cm^2/Vs is still about a factor of two lower than record mobility 2DESs grown in the same MBE chamber[68] once the heavy hole to electron effective mass ratio of $0.4\text{m}_e : 0.067\text{m}_e$ is taken into account. The question then remains, if sufficiently reducing background impurities[118] is the main obstacle

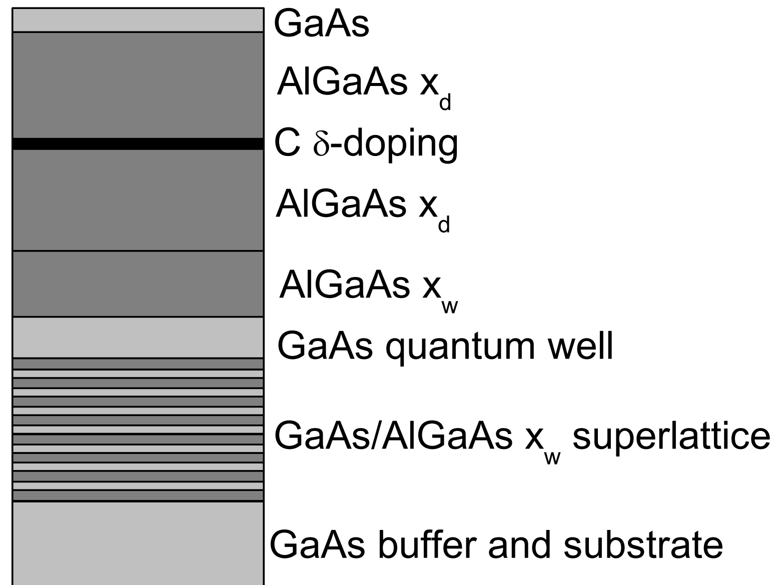

Figure 5.1. Layer structure of devices in this experiment. Note the use of two different Al mole fractions x_w and x_d in some of the devices as indicated in Table 5.1.

to reaching an electron mobility of 100 $\times$ 10^6 cm^2/Vs, what are the key ingredients to a hole mobility of 15 $\times$ 10^6 cm^2/Vs?

In order to answer this question, we have begun to explore the impact of varying structural parameters on the resulting mobility. Samples in this work were grown in a customized Veeco GenII MBE which has recently achieved *electron* mobilities > 20 $\times$ 10^6 cm^2/Vs and extremely large excitation gaps for the fragile ν=5/2 fractional quantum Hall state. C-doping was performed with a carbon filament capable of producing a doping rate of 2.8 $\times$ 10^{10} cm^{-2}/sec at a total power (including parasitic dissipation) of $\sim$ 150 W [129]. In this experiment, we utilized a 20 nm quantum well situated 190 nm below the surface and asymmetrically δ-doped from above at a setback d of 80, 110, or 140 nm. The Al mole fraction x was also varied between 0.07 and 0.45 to allow further tuning of the 2DHS density. Table 5.1 summarizes the structures grown in the experiment, and Fig. 5.1 shows the epilayer design. Square

Table 5.1.
Summary of structural parameters including δ-doping setback distance d, Al mole fraction around the dopants x_d, Al mole fraction surrounding the quantum well x_w, 2DHS density p, and $T = 300$ mK mobility after illumination μ.

Sample	d	x_d	x_w	p	μ
	nm			10^{11}cm^{-2}	10^6 cm^2/Vs
1	80	0.24	0.24	1.1	1.2
2	80	0.24	0.24	0.98	1.4
3	80	0.45	0.45	1.9	0.55
4	80	0.10	0.10	0.32	1.8
5	80	0.35	0.35	1.4	.80
6	80	0.20	0.20	0.80	1.6
7	80	0.07	0.07	0.18	1.3
8	110	0.10	0.10	0.29	1.5
9	140	0.10	0.10	0.23	1.4
10	110	0.24	0.24	0.70	1.6
11	80	0.16	0.16	0.65	2.3
12	110	0.13	0.13	0.36	1.8
A	80	0.45	0.16	1.7	0.73
B	80	0.45	0.24	1.5	0.78
C	80	0.35	0.16	1.34	1.3
D	80	0.35	0.24	1.30	1.1

samples were prepared using InZn contacts annealed at 430 °C for 15 minutes in H_2/N_2 forming gas. Characterization was performed in the dark and after illumination with a red LED at $T = 300$ mK using standard lock-in techniques, and the density was determined from quantum Hall effect (QHE) minima. Illumination typically resulted in $\sim$ 3-5% increase in density and as much as a 27% increase in mobility for low density samples. Transport data also showed a qualitative improvement after illumination, indicating that illumination increases the homogeneity of the 2DHS and has a favorable impact on the screened disorder potential. Figure 5.2 shows transport data of the highest mobility sample and a low density sample; the number of nascent fractional QHE features attest to the sample quality.

Fig. 5.3a shows the measured mobility as a function of density for various values of d. We note that remote ionized impurity (RI) scattering does not play a significant factor in limiting the mobility since within experimental uncertainty there is no meaningful difference between the mobility at different values of d for the same density. However, increased d should allow these samples to be gated to ultra-low densities before RI scattering begins to cause the mobility to rapidly drop off with further decreased density[128]. The most interesting feature of the data in Fig. 5.3a, however, is the strongly non-monotonic dependence of the mobility on density. For 2DESs in this density range with such a large value of d, one would expect the mobility to monotonically increase with density[68, 118, 66, 67] following a power law dependence $\mu \propto p^{\alpha}$ where $\alpha \sim 0.6$ - 0.8 with ionized background impurity (BI) scattering being the dominant scattering mechanism. In analyzing our results, we first note that the effective mass is known to vary throughout the density range of our samples due to the valence band non-parabolicity arising from light- and heavy-hole band mixing. By performing a linear fit to cyclotron resonance data on 2DHSs in (001) 20 nm quantum wells in refs. [108, 107] and assuming the cyclotron mass plateaus at 0.5m_e at high density we estimate the transport lifetime for our structures as shown in Fig. 5.3b. The transport lifetime, however, follows the same non-monotonic behavior as

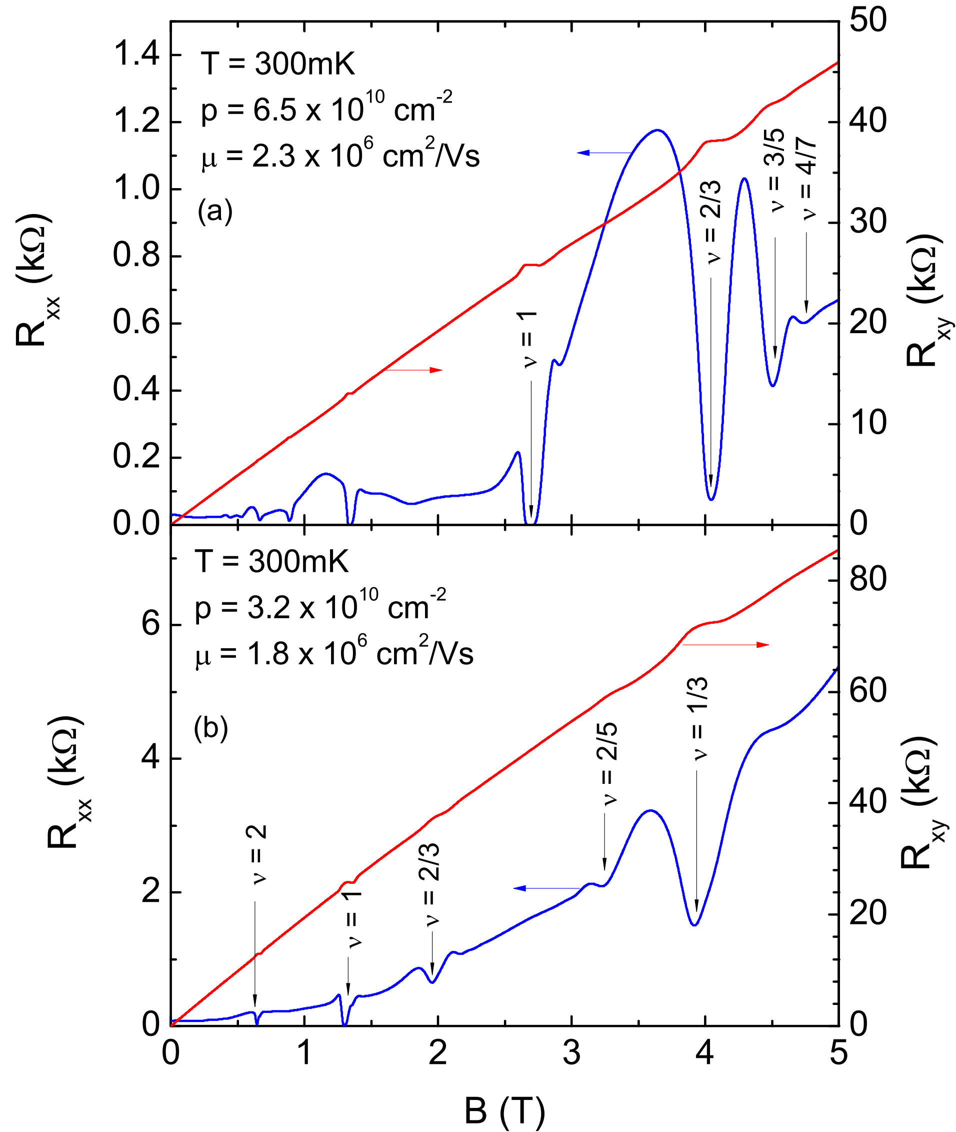

Figure 5.2. (Color online) Magnetotransport at T = 300 mK after illumination with a red LED of (a) peak mobility sample and (b) low density sample that exhibits many nascent fractional QHE features.

the mobility which indicates a competition between different scattering mechanisms throughout the density range of our experiment in addition to the changing mass.

To shed further light on possible scattering mechanisms, we have performed a series of scattering calculations including the effects of BI, RI, alloy, and interface

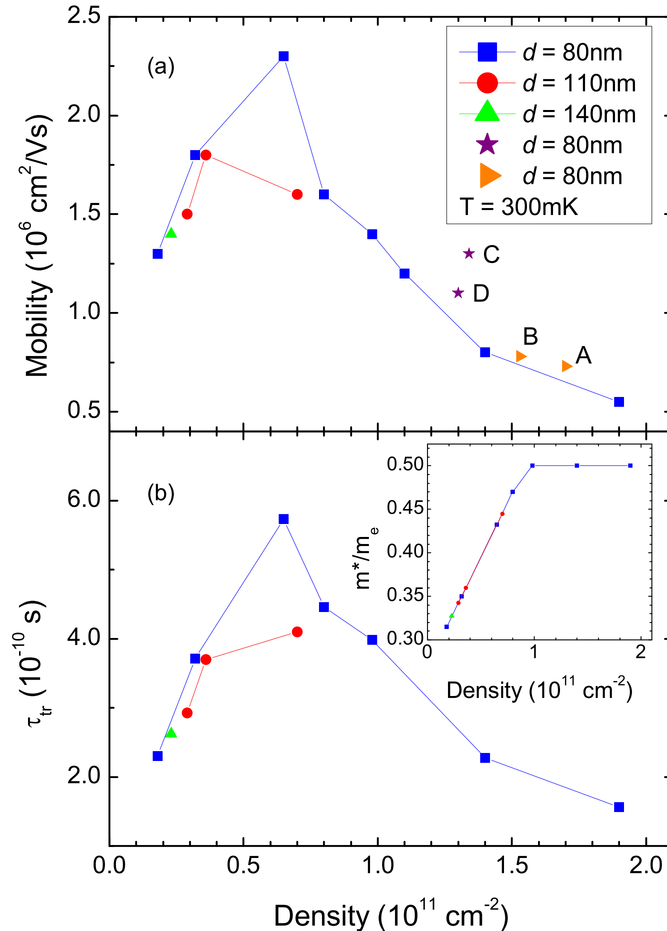

Figure 5.3. (Color online) (a) $T = 300$ mK mobility after illumination with a red LED as a function of density for various dopant setback distances d. Solid lines are guides to the eye. For fixed d the density was controlled by varying the Al mole fraction x. Samples were grown in random order to avoid continued machine clean-up from skewing the observed trend in mobility. Samples A-D were grown with varying x at fixed p to test the effect of alloy and interface roughness scattering on the mobility (see text). (b) Transport lifetime estimated as a function of density. Inset: Effective mass for our structures as a function of density extrapolated from refs. [108, 107].

roughness (IR) scattering. We follow the derivation of the transport relaxation time in [130] which assumes $T = 0$ and neglects intersubband scattering, multiple scattering events, and correlation between ionized impurities. This simple calculation is

intended to elucidate the expected trend of the mobility as the density is increased and determine if scattering mechanisms dominant in 2DESs can qualitatively explain our observations. More sophisticated calculations have been made by S. Das Sarma and coworkers.[123, 124, 118] Transport relaxation times are calculated individually and then the total mobility is calculated using Mathiessen's rule. For BI and RI scattering the transport lifetime is given by

$$\frac{1}{\tau_{tr}(\epsilon_F)} = \frac{m^*}{\pi\hbar^3}\sum_i \int_0^{\pi} \mathrm{d}\theta(1-\cos(\theta)) \times \left[\frac{2\pi e Z_i e_i}{4\pi\epsilon\left(q + q_{TF}g_s(\mathbf{q})\right)}\right]^2 \int_{-\infty}^{\infty} \mathrm{d}z N_i(z) g_{imp}^2(q,z) \quad (5.1)$$

where m^* is the hole effective mass (as estimated in Fig. 5.3b), $\hbar$ is the reduced Planck constant, θ is the scattering angle, $Z_i e_i$ is the impurity charge, ϵ is the dielectric constant of the semiconductor, $\mathbf{q}$ is the scattering vector, k_F is the Fermi wavevector, q_{TF} is the Thomas-Fermi screening wavevector, $N_i(z)$ is the i$^{\text{th}}$ impurity distribution, and the form factors are given by

$$g_s(\mathbf{q}) = \int \chi^2(z)\chi^2(z')\exp(-q|z-z'|)\mathrm{d}z\mathrm{d}z' \quad (5.2)$$

$$g_{imp}(q,z) = \int \chi^2(z')\exp(-q|z'-z|)\mathrm{d}z' \quad (5.3)$$

where $\chi(z)$ is the self-consistently calculated[117] envelope function in the effective mass approximation. For the BI calculation we use a three-dimensional impurity concentration N_{3D} as a fitting parameter and find the best agreement with the experimental data for $N_{3D} = 2 \times 10^{13}$ cm^{-3}. We use a remote impurity sheet concentration N_{RI} equal to the hole concentration p. A more realistic value of N_{RI} could also include some of the ionized impurities due to the surface compensation; however, we assume a simple parallel-plate capacitor model of the surface-δ-layer charge and thus neglect the surface compensation contribution to N_{RI}. This neglect of charge due to surface compensation is typical in these types of calculations.[118, 119, 131] For our purposes, though, the exact value of N_{RI} is not important since it will not change the qualitative dependence of the RI-limited mobility as p is varied.

To calculate alloy scattering we use the virtual crystal approximation with a square well potential limited over a spherical range[132] which is independent of temperature in 2D systems[133]. The alloy limited relaxation lifetime is unscreened due to its short range nature and given by[130, 134]

$$\frac{1}{\tau_{alloy}(\epsilon_F)} = \frac{4\Omega^2 m^* U^2 x(1-x)}{a^3\hbar^3} \int\limits_{barrier} \chi^4(z)\mathrm{d}z \tag{5.4}$$

where $a = 0.565$ nm is the lattice constant of the compound semiconductor, Ω is the volume of the scattering potential given by $\Omega = (4/3)\pi r^3$, and $r = (\sqrt{3}/4)a$ is the nearest-neighbor separation. There is a broad range of estimates of the magnitude of the scattering potential U in the literature[134], ranging from 0.12 to 1.56 eV. We take U= 1 eV (as suggested in ref. [130]) as a rough estimate.

To examine the possible effect of interface roughness scattering, we employ a simple model which makes use of the Fang-Howard variational wavefunction and associated potential[135] which takes the distortion of the wavefunction with increased density into account. In this model, the IR scattering rate is given by[83, 136]

$$\frac{1}{\tau_{IR}(\epsilon_F)} = \left(\frac{\Delta\Lambda e^2 p}{2\epsilon}\right)^2 \frac{m^*}{\hbar^3} \int_0^\pi \left(\frac{q}{q + g_s(\mathbf{q})q_{TF}}\right)^2 \times (1-\cos\theta)\exp(-\Lambda^2 q^2/4)\mathrm{d}\theta \tag{5.5}$$

where the wavefunction used to calculate g_s is

$$\chi(z) = \begin{cases} \frac{1}{\sqrt{2}} b^{3/2} z e^{-bz/2} & z > 0 \\ 0 & z \leq 0 \end{cases}$$

where the variational parameter is[3, 83]

$$b = \left(\frac{33 m^* e^2 p}{8\hbar^2 \epsilon}\right)^{1/3} \tag{5.6}$$

We take one monolayer roughness height to be a reasonable estimate and thus set Δ = 0.2825 nm and use Λ as a fitting parameter with the result that Λ = 6 nm.

As a justification for using the Fang-Howard wavefunction to model our asymmetric quantum well system, we show in Fig. 5.4 a comparison of the self-consistently

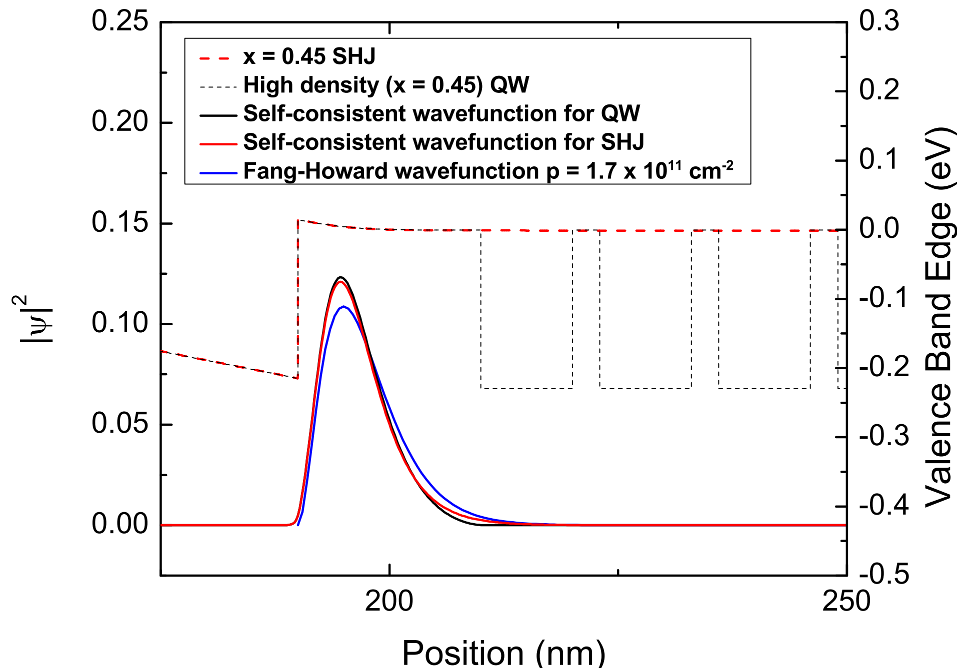

Figure 5.4. (Color online) Dashed lines show a comparison of the self-consistently calculated valence band edges (dashed lines) for the high density sample # 3 and and a single heterojunction sample with $x = 0.45$. Solid lines show a comparison of the self-consistently calculated wavefunction for sample # 3 and the Fang-Howard variational wavefunction.

calculated valence band edges for quantum well (QW) sample # 3 and a single heterojunction (SHJ) along with the self-consistently calculated wavefunction for the QW structure and the Fang-Howard wavefunction which is often taken as an approximation of the wavefunction in SHJ structures. The band edges show that the bottom barrier of the QW changes the confining potential very little, and the high density samples (where IR scattering could be important) can therefore be approximated by the Fang-Howard model.

The results of our calculations are compared with the d = 80 nm experimental results in Fig. 5.5. It is clear that even with a changing effective mass and wave function profile the BI and RI limited mobilities steadily increase with increasing

density and therefore cannot account for the drop in mobility at high density. The exact contributions of alloy and interface roughness scattering are initially less clear, however. We will address alloy disorder first.

If U is large enough, alloy scattering could conceivably contribute to the drop in mobility seen in the experimental data. Before continuing, it should be noted that the slight increase in the calculated alloy-limited mobility at high density is simply due to the saturation of the effective mass as shown in the inset of Fig. 5.3b. We have repeated the calculations (not shown) without forcing the mass to plateau at 0.5m$_e$, but even with a mass as high as 0.7m$_e$ at high density the alloy limited mobility does not appear to be limiting the total mobility. To test the contribution of alloy scattering we grew a series of four test structures (labeled A-D in Fig. 5.3a) in which x_d, the Al concentration starting 25 nm above the quantum well (e.g. around the δ-doping layer), was kept fixed to keep the density constant while x_w, the Al concentration around the quantum well, was varied between $x_w = 0.16$ and $x_w = 0.24$. Samples A, B, and 3 ($x_d = 0.45$) suggest that x_w has no impact on the mobility, though there is scatter in the resulting density which we attribute to wafer-to-wafer variation and possible variation in the illumination. Samples C, D, and 5 ($x_d = 0.35$), however, suggest that increased x_w does cause the mobility to decrease somewhat. Most importantly, this is the opposite trend that would be expected if alloy scattering per se were limiting the mobility. Our calculations for the test structures (not shown) and ref. [136] predict that the alloy-limited mobility would increase for increased x_w since as x_w is increased for fixed density the wavefunction is more confined. This in turn causes the integral of χ^4 to decrease faster than the $x(1-x)$ term increases in equation (5.4) resulting in a decrease in the alloy scattering rate for increased x_w. The results from this set of structures is consistent, however, with the theory that Al getters impurities[68], thus an increase in x_w would locally increase N_{BI} and the associated scattering. Regardless, samples C, D, and 5 show that the negative side effects of increasing x_w are not enough to explain the data of Fig. 5.3. If the increase in x_w was dominating the mobility we would expect test structures A and C to have

significantly higher mobilities than the peak mobility sample #11 due to the higher hole density of the test structures. This, however, is clearly not the case.

Finally, our fit seems initially to indicate that IR scattering is limiting the mobility at high density. However, whenever parameters can be freely adjusted caution must of course be exercised to obtain physically meaningful results. The dashed pink line in Fig. 5.5 shows the IR limited mobility for a SHJ 2DES in the Fang-Howard model while the pink star shows a 2DES SHJ structure with $x = 0.35$ grown during the course of this experiment. Evidently the Fang-Howard calculation overestimates the IR scattering by at least a factor of four. Repeating our self-consistent calculation for BI, RI, and alloy scattering in this 2DES SHJ using the impurity concentrations and alloy potential listed in the inset of Fig. 5.5 we find that the IR-limited mobility at a density of 2.4×10^{11} cm^{-2} would have to be 86×10^6 cm^2/Vs to fit the measured total mobility of 7.9×10^6 cm^2/Vs. To get such a high IR-limited mobility we are forced to set $\Delta = 0.1$ nm and $\Lambda = 2.2$ nm. Figure 5.6 shows the result of our calculation for the hole structures using these smaller roughness parameters. With these reduced roughness parameters there is no longer a good fit to the hole data at high density as the IR term makes almost no contribution to the total mobility, though we still obtain a good fit at low to medium density. We conclude that our crude model of interface roughness scattering cannot simultaneously account for our experimental data in both electrons and holes and are thus hesitant to conclude that interface roughness scattering is the dominant source of our drop in mobility at high density. Similar discrepancies between electron and hole data have been noted in ref. [137].

Another possible scattering mechanism that must be kept in mind at high density is scattering between the electric subbands of the quantum well which is known to degrade the mobility in high density 2DESs.[138] To estimate the possibility of such scattering, we use a finite square well with a barrier height of 230 meV and an effective mass of 0.5m$_e$, which corresponds to our highest density sample. This estimate results in an energy spacing of 5.0 meV between the heavy hole ground and first excited state. Assuming a light hole mass along the (001) direction[3, 126] of 0.090m$_e$ the spacing

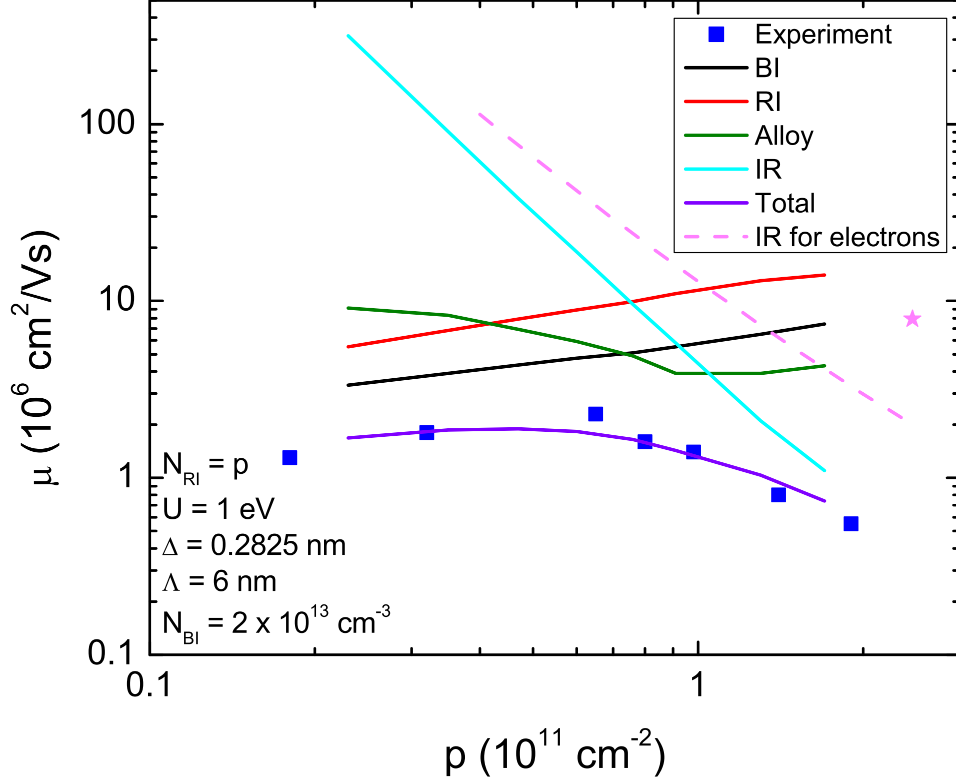

Figure 5.5. (Color online) Comparison of $d = 80$ nm experimental data with mobility calculations. N_{BI} and Λ are used as free parameters to obtain a good fit to data. Pink star represents SHJ 2DES grown during this experiment.

between the heavy hole and light hole ground states is 6.4 meV. In both cases this energy spacing is significantly larger than the Fermi energy $E_F = \frac{\pi\hbar^2 p}{m*} \sim 0.9$ meV which precludes a significant contribution from intersubband scattering between the electric subbands.

Next, we note the presence of beating in the Shubnikov-de Haas oscillations in Fig. 5.7 which is indicative of B = 0 spin-splitting. Such spin-splitting is known to occur in structurally-asymmetric devices[125, 127, 139] due to Rashba spin-orbit coupling[126]. We sketch the qualitative effect of this splitting in inset (a) of Fig. 5.7. As the 2DHS density is increased, the electric field (and hence spin splitting) in the well is also increased. Furthermore, it is know that the presence of a parallel channel can result

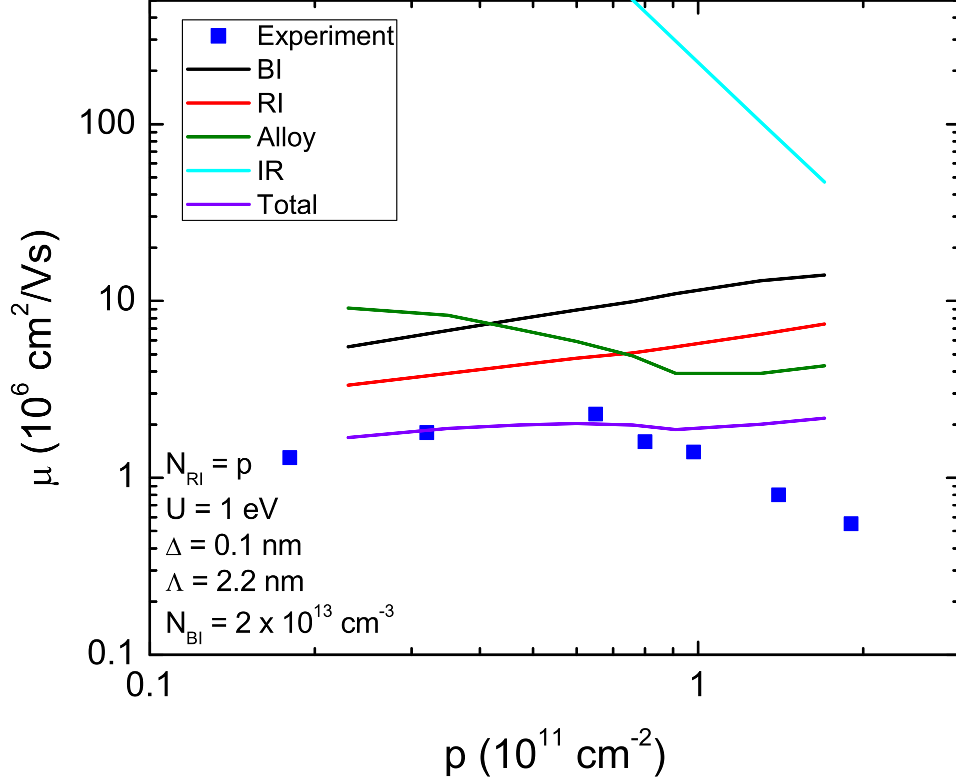

Figure 5.6. (Color online) Comparison of $d = 80$ nm experimental data with mobility calculations. Δ and Λ are varied in order for mobility calculations to obtain agreement with an electron structure grown during this experiment (see text).

in a Hall density different from the sum of the subband densities and a measured mobility different from that of either subband even in the absence of intersubband scattering. In our case, we assume that the two parallel channels are non-interacting B = 0 spin-split subbands of the heavy hole ground state. The measured Hall density p_{Hall} and mobility μ_{Hall} in the absence of intersubband scattering are given by[3]

$$p_{Hall} = \frac{(p_1\mu_1 + p_2\mu_2)^2}{p_1\mu_1^2 + p_2\mu_2^2} \tag{5.7}$$

$$\mu_{Hall} = \frac{p_1\mu_1^2 + p_2\mu_2^2}{p_1\mu_1 + p_2\mu_2} \tag{5.8}$$

where $\mu_{1(2)}$ and $p_{1(2)}$ are the mobility and density, respectively, of the first (second)

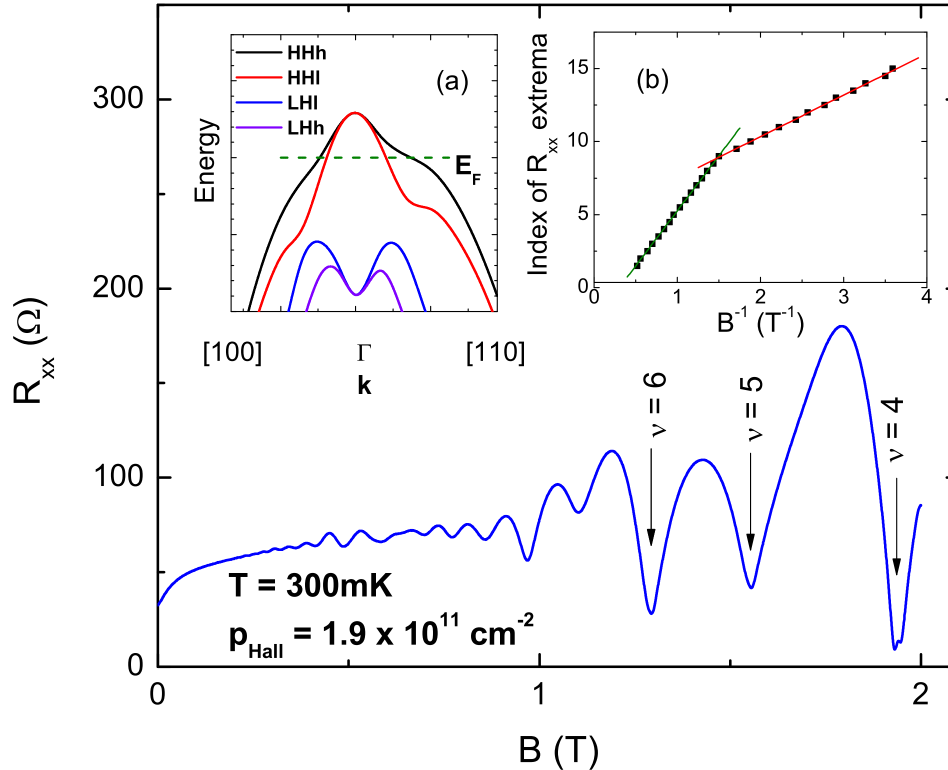

Figure 5.7. (Color online) Shubnikov-de Haas oscillations of high density d = 80 nm, x = 0.45 device. Inset (a): Sketch of the spin-split heavy hole and light hole ground states in a quantum well. Inset (b): Index of extrema in R_{xx} vs. B^{-1}. The high field slope gives the total density of 1.8 $\times$ 10^{11} cm^{-2}, and the low field slope gives the lighter sub-band density of 7 $\times$ 10^{10} cm^{-2} while the difference in the two gives the second sub-band density of 1.1 $\times$ 10^{11} cm^{-2}.

subband. Figure 5.8a illustrates the Hall density as a function of the subband mobilities in our peak density sample (sample # 3) predicted by Eq. 5.7 using the subband densities extracted in Fig. 5.7. It is clear from Fig. 5.8a that in order to measure a Hall density $\sim 1.8 \times 10^{11}$ cm^{-2} the subband mobilities should be comparable, though the high density subband should have a slightly higher mobility. In order to estimate the effect of the presence of two subbands on the measured mobility, we therefore assume that the high density subband is also the high mobility subband. In order to determine if the presence of the lower mobility subband could by itself account for

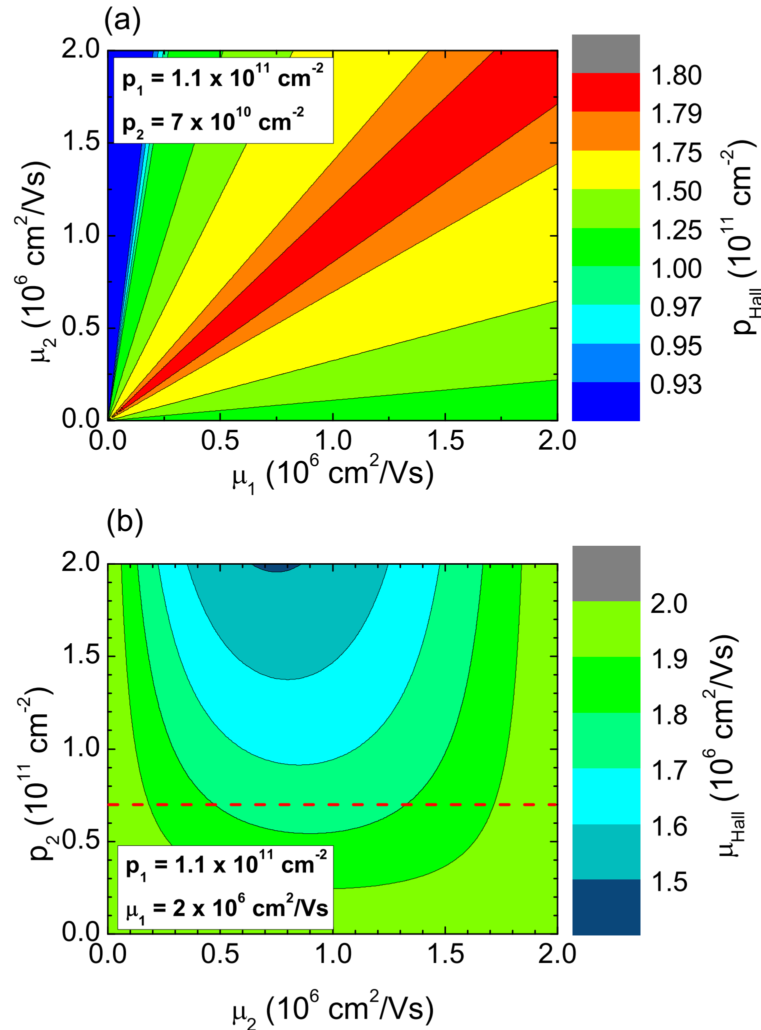

Figure 5.8. (Color online) (a) Estimate of the Hall density expected from Eq. 5.7 using the measured subband densities $p_1 = 1.1 \times 10^{11}$ cm^{-2} and $p_2 = 7 \times 10^{10}$ cm^{-2}. (b) Estimate of the expected measured mobility μ_{Hall} if the high density subband $p_1 = 1.1 \times 10^{11}$ cm^{-2} has a high mobility $\mu_1 = 2 \times 10^6$ cm^2/Vs. For the second subband density (dashed red line) $p_2 = 7 \times 10^{10}$ cm^{-2} measured in Fig. 5.7 the expected measured mobility $\mu_{Hall} \geq 1.75 \times 10^6$ cm^2/Vs.

the drop in mobility seen in Fig. 5.3, we assume that the high mobility subband is unchanged from the peak total mobility value ($\sim 2 \times 10^6$ cm^2/Vs) at low density. Figure 5.8b shows what we would thus expect to measure as a function of density

and mobility in the low mobility subband if the high mobility subband has a density $p_1 = 1.1 \times 10^{11}$ cm^{-2} as we measure in Fig. 5.7. For the measured second subband density of $p_2 = 7 \times 10^{10}$ cm^{-2} (dashed red line) we see that this parallel subband effect would not decrease the measured mobility below $\sim 1.75 \times 10^6$ cm^2/Vs. We therefore conclude that the presence of a second, possibly low mobility B = 0 spin-split subband cannot explain our observed drop in mobility at high density in the absence of intersubband scattering.

A final possible mechanism for the observed drop in mobility at high density is intersubband scattering between the spin-split subbands of the heavy hole ground state of the quantum well. The question remains, however, whether or not there exists a potential capable of coupling the spin-split sub-bands and causing back-scattering. Such scattering is typically neglected in theoretical calculations of the mobility due to the assumed lack of a significant spin-flip mechanism[140], though intersubband hole-hole scattering in inversion-asymmetric structures is not without precedent.[141] At this time more theoretical work is needed to resolve the relative contributions of the different scattering mechanisms.

In conclusion, we have performed an experimental study of the density dependence of mobility in C-doped (001) GaAs/AlGaAs quantum wells by varying the dopant setback d and Al mole fraction x. The mobility was seen to depend non-monotonically on the density. At low density the mobility increased with density. The $T = 300$ mK mobility was found to peak at a value of 2.3×10^6 cm^2/Vs at a density of 6.5×10^{10} cm^{-2}. This 2DHS mobility is among the highest ever reported. Increasing the density further, however, resulted in a sharp drop in mobility. Scattering calculations indicate that background ionized impurities and remote ionized impurities will not lead to a decrease in mobility at high density even with a changing effective mass, and alloy scattering cannot account for all of our experimental results from various test structures. Interface roughness scattering contributions remain unclear due to the difficulty in obtaining physically reasonable roughness parameters that predict both electron and hole mobilities. Beating in the Shubnikov-de Haas oscillations in

our high density samples is indicative of zero-field spin-splitting which leaves open the possibility of an intersubband scattering contribution to the mobility. Further theoretical work is needed to determine the mechanism and magnitude of such a contribution.

Acknowledgment

JDW would like to thank S. Birner, C. Rössler, and T. Feil for helpful discussions regarding calculations with Nextnano3. JDW is supported by a Sandia Laboratories/Purdue University Excellence in Science and Engineering Fellowship. MJM acknowledges support from the Miller Family Foundation. The MBE growth and transport measurements at Purdue are supported by the U.S. Department of Energy, Office of Basic Energy Sciences, Division of Materials Sciences and Engineering under Award DE-SC0006671.

6. Impact of Heterostructure Design on Transport Properties in the 2$^{\mathrm{nd}}$ Landau Level in In-Situ Back-Gated Two-Dimensional Electron Gases

J. D. Watson[1,2,], G. A. Csáthy[1], and M. J. Manfra[1,2,3,4]

[1] *Department of Physics, Purdue University, West Lafayette, IN 47907, USA*

[2] *Birck Nanotechnology Center, Purdue University, West Lafayette, IN 47907, USA*

[3] *School of Materials Engineering, Purdue University, West Lafayette, IN 47907, USA*

[4] *School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA*

[1] *Abstract:* We report on transport in the 2$^{\mathrm{nd}}$ Landau level in in-situ back-gated two-dimensional electron gases in GaAs/Al$_x$Ga$_{1-x}$As quantum wells. Minimization of gate leakage is the primary heterostructure design consideration. Leakage currents resulting in dissipation as small as a few pW can cause noticeable heating of the electrons at 10 mK, limiting the formation of novel correlated states. We show that when the heterostructure design is properly optimized, gate voltages as large as 4V can be applied with negligible gate leakage, allowing the density to be tuned over a large range from depletion to over 4×10^{11} cm^{-2}. As a result, the strength of the $\nu = 5/2$ state can be continuously tuned from onset at n $\sim 1.2 \times 10^{11}$ cm^{-2} to a maximum $\Delta_{5/2} = 625$ mK at n $= 3.35 \times 10^{11}$ cm^{-2}. An unusual evolution of the reentrant integer quantum Hall states as a function of density is also reported.

[1]J. D. Watson, G. A. Csáthy, and M. J. Manfra (2015). "Impact of heterostructure design on transport properties in the 2$^{\mathrm{nd}}$ Landau level in in-situ back-gated two-dimensional electron gases". Manuscript submitted for publication.

6.1 Introduction

Since the discovery of the fractional quantum Hall effect (FQHE) at $\nu = 5/2$ over 25 years ago[18], this state has drawn intense scrutiny. The well-known Laughlin wave-function and extensions enabled by composite Fermion theory [13] cannot explain the existence of incompressible states with even-denominator filling. Many potential wavefunctions have been proposed (see references [142, 52] for a summary of candidate states), but the exact nature of the ground state at 5/2 is still controversial. The exact ground state realized in experiment may depend on sample parameters such as electron density, well width, edge confining potential, etc. Numerical work[29, 30, 31, 32] has shown strong overlaps with the Pfaffian wavefunction[28] and its particle-hole conjugate state, the so-called anti-Pfaffian[34, 143]. This is a tantalizing prospect as both these states give rise to non-Abelian quasiparticle excitations which have been proposed as a route to fault-tolerant quantum computing[26, 27].

There have also, however, been theoretically proposed wavefunctions for the $\nu = 5/2$ state that exhibit Abelian statistics[142, 52]. To date, the experimental tests to determine the nature of the ground state at $\nu = 5/2$ have failed to agree on the identity of the wavefunction. Experiments probing the temperature dependence of tunneling between the edge states at $\nu = 5/2$ have been proposed[144] and conducted[37, 43] as a way to measure the quasiparticle effective charge e^* and Luttinger liquid interaction parameter g in order to discriminate between proposed wavefunctions. These experiments, however, were inconclusive as tunneling experiments performed on the same Hall bar mesa but with different electrostatic confinement potentials gave results consistent with the non-Abelian anti-Pfaffian and $U(1) \times SU_2(2)$ states[37] and the Abelian 331 state[43]. As it is unclear how possible edge reconstruction[145] due to shallow confining potentials might influence the interpretation of these experiments, it is possible that the different confinement parameters in the previously studied devices could be responsible for this apparent discrepancy. It would therefore be desirable to examine transport in nanostructures in the quantum Hall regime more systemati-

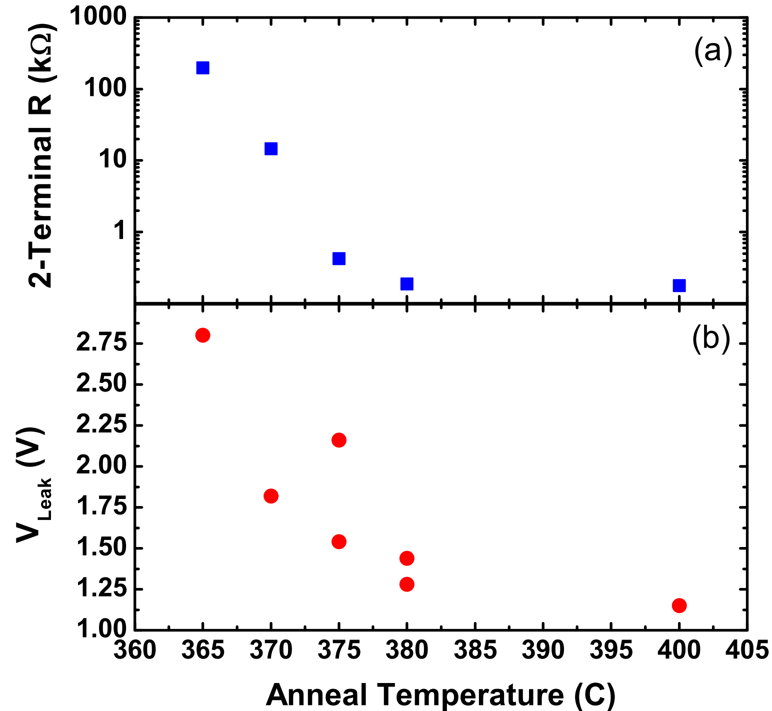

Figure 6.1. Effect of Ohmic annealing temperature on device performance from wafer A. (a) Median 2-terminal resistance to ground of individual contacts measured in the dark at T = 4 K and $V_g = 0$ as a function of annealing temperature. (b) V_{leak}, defined as the voltage at which the gate leakage current reached 1 nA, as a function of annealing temperature.

cally in samples in which the electron density and confining potential could be tuned simultaneously in a single structure. A variable density would also allow for direct comparisons between experimental results in the N = 1 Landau level (LL) and the more well-understood N = 0 LL in a single device.

In order to undertake such experiments, however, it is necessary to have a thorough understanding of how heterostructure design and device fabrication parameters affect device yields and the quality of transport in the 2^{nd} LL. Towards this aim, we have grown and processed a series of high quality, in-situ backgated two-dimensional electron gases (2DEGs). The processing of similar devices of lower mobility has been reported[146, 147, 148, 149, 150, 151, 152, 153] and a similar high mobility device has

been used to examine the energy gaps of FQHE states in the 2^{nd} LL[51]. However, to our knowledge there has not been a published systematic study of heterostructure design and processing conditions and their impact on the visibility of states in the 2^{nd} LL.

6.2 Device Growth and Fabrication

We studied three wafers utilizing two heterostructure designs to study the impact of the heterostructure on the gate leakage and the low temperature transport. Both designs feature a 2DEG located approximately 200 nm from the surface in a 30 nm GaAs quantum well flanked by $Al_{0.24}Ga_{0.76}As$ barriers modulation doped from the top side only at a setback of $\sim$ 70 nm. The dopants were incorporated in a so-called doping well scheme (also known as a short-period superlattice) [21, 68, 85, 56] which has been found empirically to maximize the FQHE energy gaps in the 2^{nd} LL. The in-situ gate consisted of an N+ GaAs layer situated 850 nm below the bottom interface of the quantum well. The key difference between the two designs was that design #1 used a 200 nm $Al_{0.24}Ga_{0.76}As$ barrier to separate the quantum well from a GaAs/AlAs (2/2 nm) superlattice while design #2 decreased this superlattice setback to 20 nm while keeping the gate setback fixed at 850 nm. Wafers A and B utilized design #1 while wafer C utilized design #2.

Device fabrication began by etching via holes to the gate layer using an etchant consisting of 50:5:1 water:phosphoric:peroxide followed by a second, $\sim$ 160 nm deep etch to define 1 mm Van der Pauw square mesas. Ohmic contacts consisted of a 8/80/160/36 nm stack of Ni/Ge/Au/Ni and were annealed for 1 min in forming gas at a variety of temperatures. Following the annealing, large TiAu pads off of the mesa were deposited in order to facilitate wirebonding.

Figure 6.1 shows the effect of annealing temperature on the quality of the contacts and the gate leakage measured in the dark at T = 4 K on devices fabricated from wafer B. The lead resistance of the measurement setup was $\sim 1\ \Omega$, so the 2-terminal

resistance values quoted here are reasonable proxies for the true contact resistance. At an annealing temperature of 360 °C, the contacts were electrically open at low temperature, and the contact morphology was extremely smooth, indicating that the metal did not melt or diffuse significantly during the anneal. Figure 6.1b displays the effect of the annealing on the gate leakage. To quantify the leakage from our devices, we defined V_{leak} as the gate voltage V_g at which the gate leakage current reached 1 nA; thus high values of V_{leak} are expected for a high quality gate insulating layer. Both the 2-terminal resistance and V_{leak} decrease monotonically as the annealing temperature is increased and the NiAuGe diffuses further into the semiconductor.

To further study the impact of mask design and processing parameters on the gate leakage and contact resistance, we fabricated a set of test structures (not shown) which gave evidence that the gate leakage was primarily through the annealed contacts and not through the bulk of the mesa. In addition, the test structures gave evidence that the leakage current density through annealed metal in etched regions was larger than that through annealed metal in un-etched regions. The increased electric field due to the decreased gate-contact separation in the etched regions was insufficient to account for this increase in leakage density. This observation appears to imply that the etching procedure enhances the subsequent diffusion of the contacts. With this in mind, we designed our lithographic mask sets to minimize the total Ohmic area, particularly in the region off of the etched mesa. In our final design the total Ohmic area was $< 1.5 \times 10^4$ μm^2 per device, and the total Ohmic overhang off each mesa was $\sim$ 6000 μm^2. By minimizing the total time the etched sidewall of the mesa was exposed to air between the etch step and the metallization (typically $\sim$ 3-4 hours) and optimizing the geometry of the Ohmic contacts to include 45° scallops, we were able to produce devices with acceptably low contact resistances in the range of a few hundred Ohms while minimizing the gate leakage.

Next, we examined the impact of heterostructure design on device performance. Using our optimized fabrication recipe and mask set, we fabricated devices on both wafers B and C, using an annealing temperature of 375 °C. Figure 6.2 is a histogram

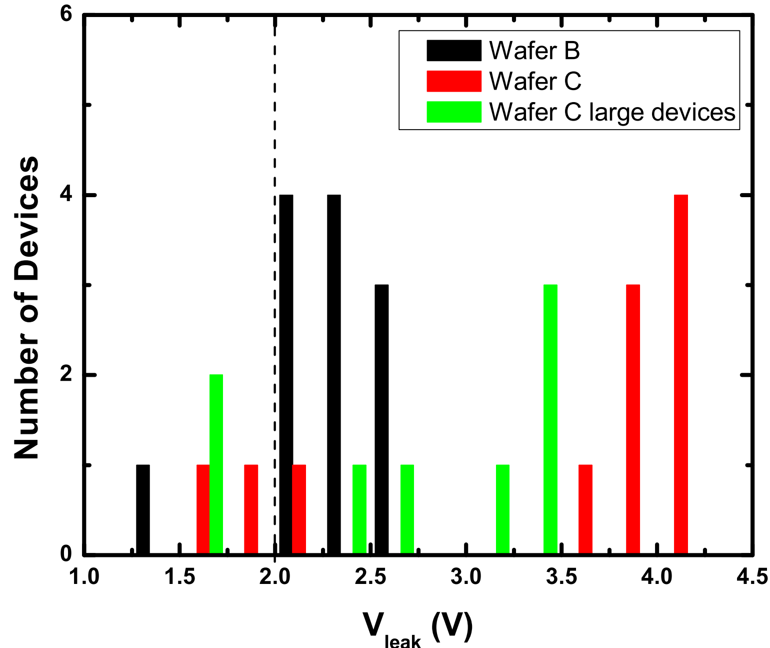

Figure 6.2. Histogram of leakage turn-on voltage V_{leak} for devices fabricated with the optimized processing recipe and mask set. All the devices were annealed at 375 °C. The dashed line represents the voltage required to reach a 2DEG density of $\sim 3.2 \times 10^{11}$ cm^{-2}.

of the leakage turn-on V_{leak} for devices from each wafer. The leakage in the majority of devices from wafer B (black bars) turned on around 2.2 V while the leakage in devices from wafer C (red bars) typically turned on around 3.8 V. Evidently, the proximity of the superlattice to the quantum well has a pronounced effect on the gate leakage.

The dashed line in Fig. 6.2 represents the gate voltage required in our geometry to reach a 2DEG density of $\sim 3.2 \times 10^{11}$ cm^{-2}, roughly twice the zero-bias density and the approximate electron density of 2DEGs exhibiting state-of-the-art energy gaps in the 2nd LL (see for instance references [154, 86, 56]). As the devices from wafer C clearly could be biased well beyond the point necessary to study the FQHE of the 2nd LL, we fabricated Hall bar devices with larger contacts on wafer C to check how much less stringent the device design and fabrication requirements were for this wafer

to exhibit acceptable gate leakage. These devices were based on a design[87] known to both exhibit high quality transport in the 2^{nd} LL and allow the incorporation of nanostructures. The total Ohmic area per device was 3.0×10^5 μm^2 with 4.6×10^4 μm^2 overhanging the edge of the mesa. Even though the Ohmic area in the etch field increased by a factor of ~ 8 and the total Ohmic area increased by a factor of ~ 20 from our optimized mask design, the leakage turn-on in most devices was still beyond 2.5 V, further highlighting the importance of proper heterostructure design.

We speculate that the large reduction in gate leakage in wafer C is due to two effects. First, the alternating layers of the superlattice act as a diffusion barrier[72, 73] to the metal from the Ohmic contacts; thus, by moving the superlattice closer to the quantum well, less metal is able to diffuse towards the gate, thereby reducing the shorting of the Ohmics to the gate. In addition, Fowler-Nordheim tunneling[155, 156, 157, 158, 159] from the bulk of the 2DEG to the gate can be expected to be reduced by moving the tall AlAs barriers of the superlattice closer to the 2DEG. The fact that there was such a minimal decrease in V_{leak} when the contact area was increased on wafer C (see Fig. 6.2) may suggest that with our optimized fabrication procedure Fowler-Nordheim tunneling in wafer C is comparable to leakage from the contacts.

While moving the superlattice closer to the quantum well has the benefit of dramatically increasing the maximum achievable density, it also has the undesirable consequence of placing a significant amount of AlAs close to the quantum well. It is known that Al is an effective getter of vacuum impurities during MBE growth[68], and thus moving the superlattice closer to the 2DEG may degrade the quality of the FQHE states. Indeed, the average maximum electron mobility in devices from wafer B was $\sim 15 \times 10^6$ cm^2/Vs while that from wafer C was $\sim 11 \times 10^6$ cm^2/Vs. Wafers B and C were grown on the same day, so it appears likely that the decrease in mobility can be attributed to the change in heterostructure design. That being said, it has become clear in recent years that the zero field mobility is not a good predictor of energy gaps in the 2^{nd} LL[51, 154, 160, 56]. Consequently, it was necessary to examine

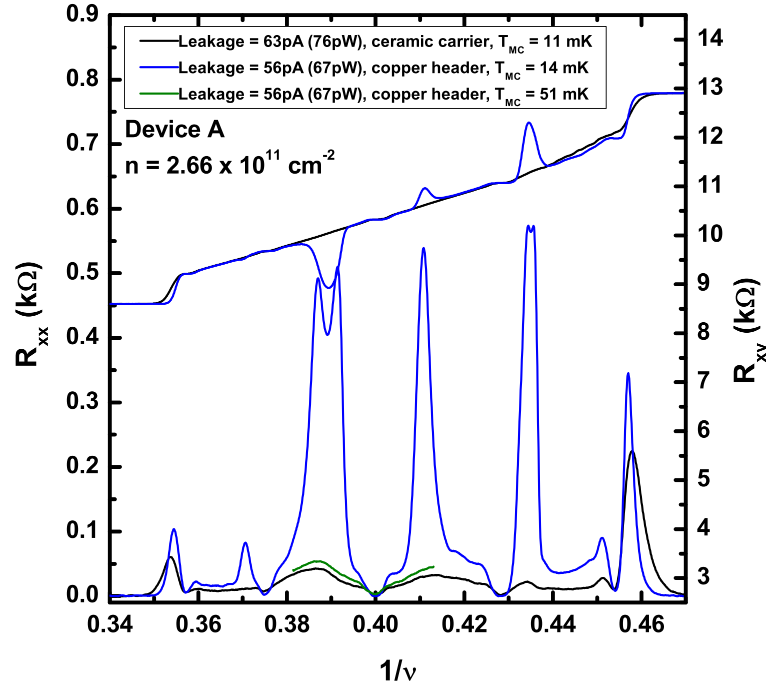

Figure 6.3. Magnetotransport in the lower spin branch of the 2nd LL in device A. During the first cool down of the sample (black curves) the sample was mounted on a commercial ceramic chip carrier. At a gate leakage current (power) of $\sim$ 63 pA (76 pW) the electrons appear very warm as seen by the lack of RIQHE features, despite a low mixing chamber temperature T_{MC}. During the second cool down of the sample (blue curves), the device was mounted on a homemade header with a copper strip screwed onto the end of the cold finger on the mixing chamber. The electrons were obviously much colder even for a slightly higher T_{MC}. The green curve shows the transport around $\nu = 5/2$ during the second cooldown for $T_{MC} = 51$ mK. Comparing the green and black data, we conclude that the electron temperature was $\sim$ 50 mK for $T_{MC} = 11$ mK during the first cooldown.

the magnetotransport at low temperature to make any definitive statement on the potential negative impact of moving the superlattice closer to the 2DEG.

6.3 Low Temperature Transport

Figure 6.3 illustrates the importance of minimizing the gate leakage and properly heat sinking the sample in order to study the 2nd LL at low temperatures (T < 25 mK). The data shown were taken from an early device from wafer A which was fabricated prior to the final optimization of our processing recipe. During the first cool down of the device, the Joule heating of the electrons due to the gate leakage current evidently caused the electron temperature to depart from the mixing chamber temperature T_{MC} for a gate leakage current (power) $\sim$ 4 pA ($\sim$ 3.5 pW) as evinced by the weakening of the reentrant integer quantum Hall effect (RIQHE) features (data not shown). By contrast, the excitation current of 2.1 nA contributed a neglible power dissipation of $\sim$ 45 fW at $\nu = 5/2$. In order to facilitate wire bonding, we mounted the device on a commercial bondable ceramic chip carrier during the first cool-down. This meant, however, that the sample was only cooled through the 18 μm thick Au bond wires. To improve the heat sinking, we re-wired the same device on a homemade header. In this design the sample was mounted to a strip of Cu with Ag paint, and the Cu strip was screwed directly onto the Cu cold finger of the mixing chamber resulting in a continuous metal connection between the mixing chamber and sample. With this improved heat sinking, heating of the electrons was not evident until a gate leakage current (power) $\sim$ 56 pA (67pW). Figure 6.3 illustrates the vast improvement in electron temperature achieved by improving the heat sinking of the sample. For a fixed density and approximately constant gate leakage current, the data taken with the Cu strip header show strong RIQHE features while the data taken with the ceramic chip carrier shows no RIQHE features. To quantify $T_{electron}$ during the first cool-down, we show data (green curve in Fig. 6.3) taken at T_{MC} = 51 mK during the second cooldown. The insulating peaks in R_{xx} in the vicinity of $\nu = 5/2$ at T_{MC} = 51 mK during the second cool-down are comparable to those seen at T_{MC} = 11 mK during the first cool-down. This allows us to estimate $T_{electron} \sim$ 50 mK for the black curves in Fig. 6.3.

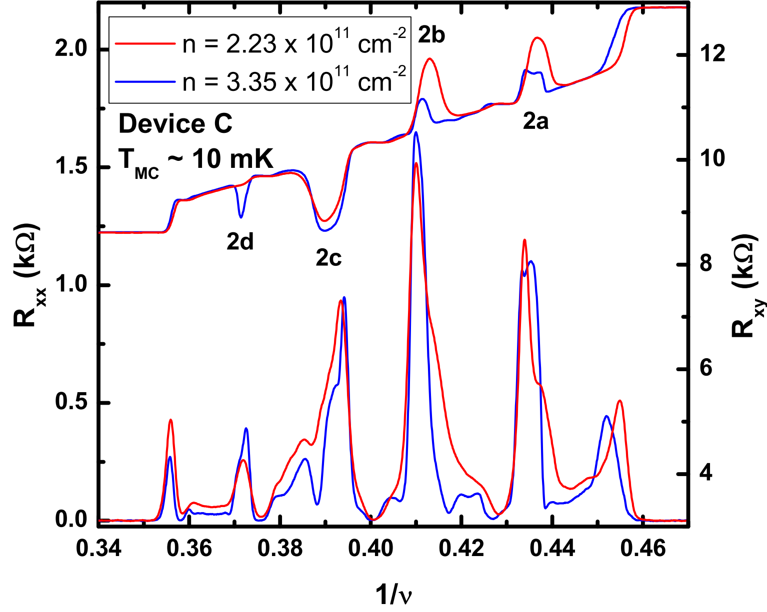

Figure 6.4. Magnetotransport in device C after illumination with a red LED. The re-entrant states are labelled following the convention in reference [161]. Red data show the transport for the maximum strength in RIQHE states 2a and 2b while the blue data show transport at the highest density before the second subband became occupied.

After screening devices from each wafer at T = 4 K and T = 300 mK, we cooled one exemplary device each from wafers B (device B) and C (device C) to low temperature (< 25 mK) to examine the transport as a function of density at low temperature. Figure 6.4 shows low temperature transport ($T_{MC} \sim 10$ mK) at two different densities for device C after illumination with a red LED. The device shows excellent transport with all four RIQHE states present and well developed FQHE states at $\nu = 14/5$, 8/3, 5/2, and 7/3. In addition, nascent states at $\nu = 12/5$ and $\nu = 2 + 6/13$ begin to develop at high density. This is, to our knowledge, the first time these states have been observed in a back-gated device, and their presence in spite of their extreme fragility[22, 25, 162] further points to the high quality of the 2DEG.

We examined the strength of the FQHE in each device quantitatively by measuring the gap at $\nu = 5/2$ ($\Delta_{5/2}$). Figure 6.5 displays the gap at $\nu = 5/2$ as a function of

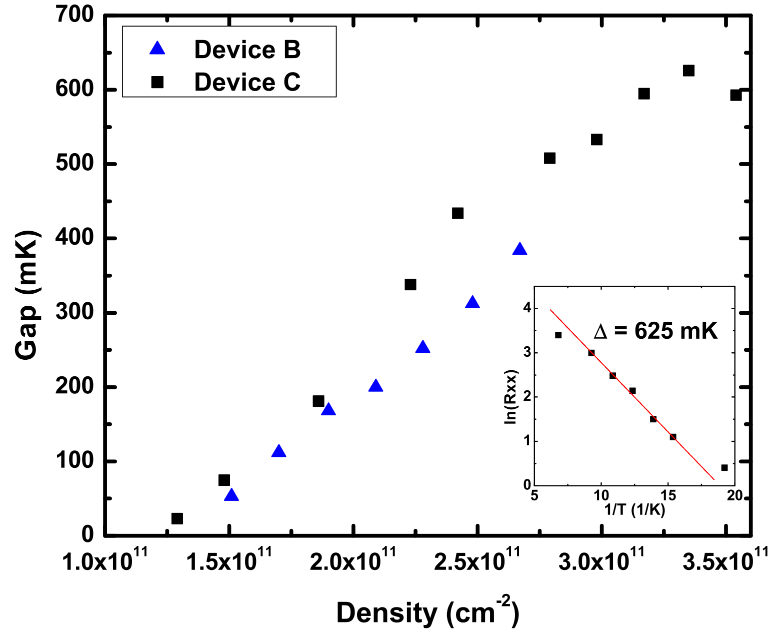

Figure 6.5. Gap at $\nu = 5/2$ as a function of density for devices B and C. Inset shows the Arrhenius plot for device C at a density of 3.35×10^{11} cm^{-2} where the gap was measured to be 625 mK.

density for devices B and C. It is clear that, within the experimental resolution, the gaps are nearly identical for both devices at low density (n $< 2.5 \times 10^{11}$ cm^{-2}). Evidently, neither the day-to-day variation in the MBE growth conditions nor the uncontrolled sample degradation from device fabrication nor the proximity of the superlattice to the 2DEG significantly affect the gap at $\nu = 5/2$. Device C, however, allows investigation of much higher 2DEG densities. Moreover, the magnitude of the gaps are very large with the gap in device C reaching a maximum value of 625 mK, the highest reported to date, at a density of 3.35×10^{11} cm^{-2}.

One noticeable feature of the data from device C in Fig. 6.5 is that at the highest density measured the gap shows a pronounced drop. It has been previously reported[163] that the gap at $\nu = 5/2$ drops suddenly when the energy difference between the Fermi energy E_F and the first excited electric sub-band in the quantum well equals the cyclotron energy. In this case, there is a level crossing and the ground

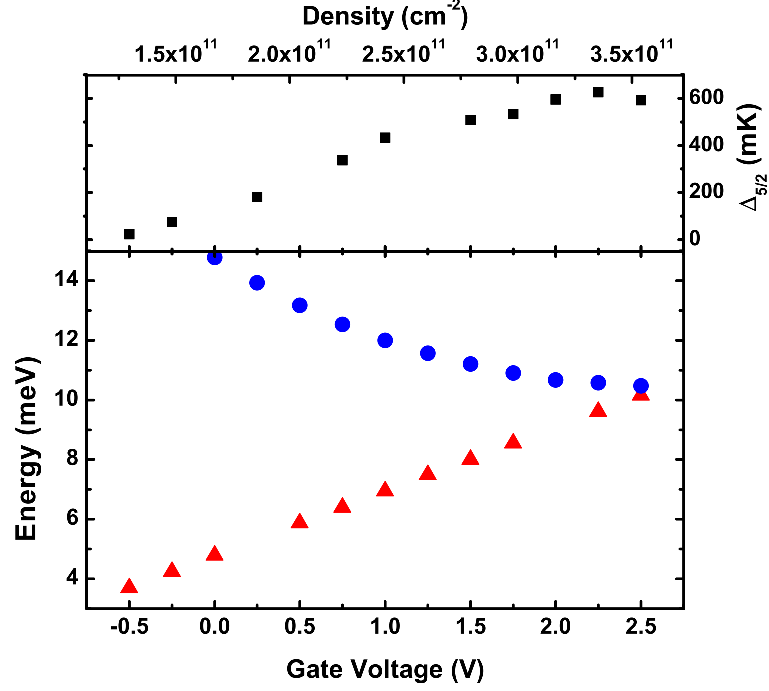

Figure 6.6. Cyclotron energy $\hbar\omega_c$ (red triangles) and spacing between $\mathrm{E_F}$ and the second sub-band (blue circles) overlaid with $\Delta_{5/2}$ for Wafer C. $\Delta_{5/2}$ drops suddenly at high density when the ground state is pushed into the lowest LL of the anti-symmetric sub-band.

state is pushed into the lowest LL of the anti-symmetric sub-band. Figure 6.6 shows the calculated[117] energy spacing along with the cyclotron energy as a function of density. As expected, the experimentally measured gap at $\nu = 5/2$ is seen to drop suddenly when the cyclotron energy becomes approximately equal to the gap between $\mathrm{E_F}$ and the second sub-band.

Finally, the RIQHE states in device C (from wafer C) showed an interesting evolution with density. In order to quantitatively compare the states, we defined the strength S of the RIQHE states as

$$S \equiv \frac{|R^c_{xy} - R_{xy}|}{|R^c_{xy} - R^i_{xy}|} \tag{6.1}$$

where R^c_{xy} is the the classical Hall resistance at the filling fraction of interest, R_{xy} is the actual Hall resistance at the peak position, and R^i_{xy} is the resistance of the

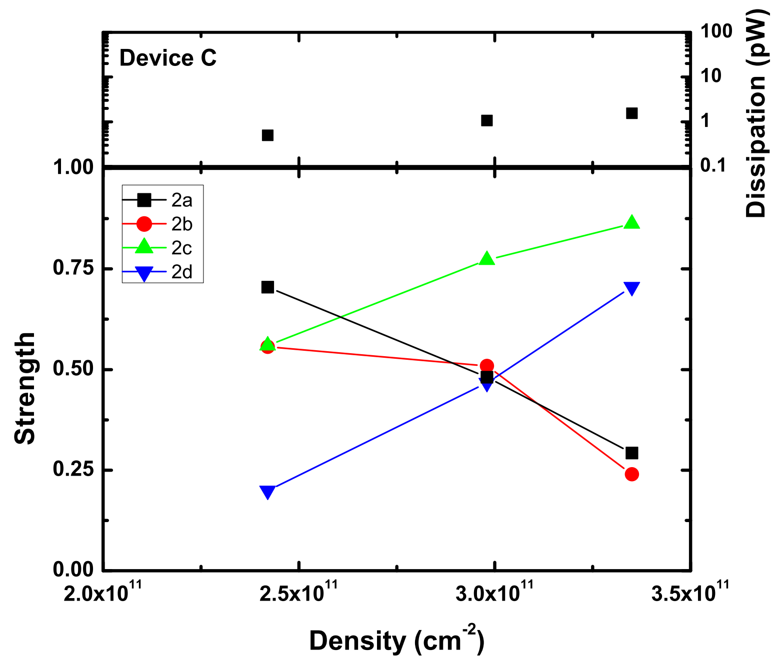

Figure 6.7. Strength (as defined in the text) of the RIQHE in device C during the second cool-down as a function of density; the power dissipation from the gate leakage current is shown in the top panel. States 2a and 2b weaken over the measured density range while states 2c and 2d strengthen over the same range.

nearest integer Hall plateau. Using this definition, a fully quantized RIQHE state has a strength of 1 while a completely absent state has a strength of 0. Figure 6.7 shows the evolution of the RIQHE states in device C during its second cool-down. The states on the high field side of $\nu = 5/2$, 2a and 2b, are seen to weaken over the measured density range while states 2c and 2d continue to strengthen. Figure 6.8 shows a comparison of the evolution of state 2a in device C as a function of density for two different cooldowns. Even though the power dissipation from the gate leakage varied by ~ 1 order of magnitude between the two cool-downs, the data show the same trend. Comparisons between the other three states for the two cool-downs show similar agreement. This appears to indicate that the observed evolution in strength is driven primarily by the 2DEG density and not by heating from the

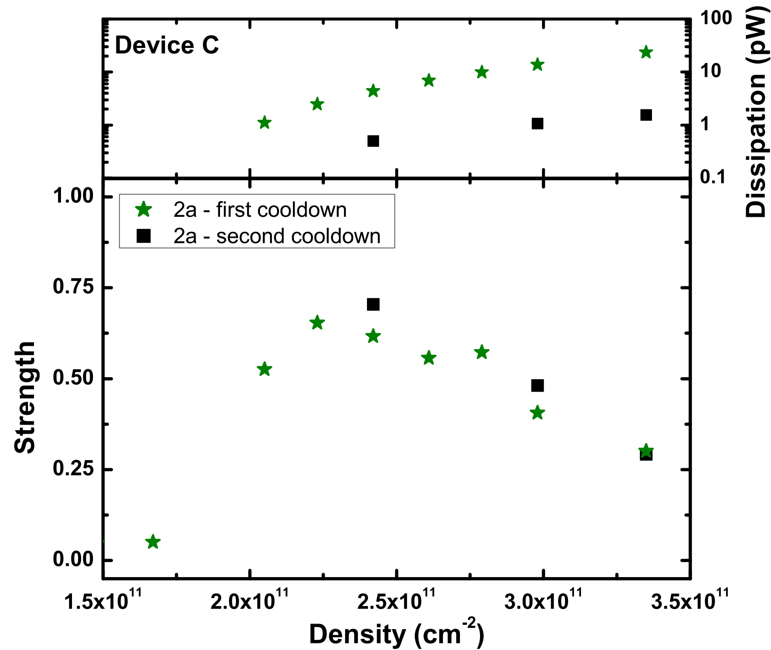

Figure 6.8. Comparison of the strength of the 2a RIQHE state in device C from two different cooldowns. The strength of the state is comparable between the two cooldowns despite the large change in gate power dissipation.

gate leakage. At present, the origin of this behavior is not understood. Regardless of the mechanism that causes states 2a and 2b to weaken with increasing density, this behavior is qualitatively different than that seen in states 2c and 2d and may point to a difference in the underlying localization mechanisms. In contrast, the strength of all the RIQHE states in device B (from wafer B) with the larger superlattice setback (data not shown) were seen to increase with density up to a density (power dissipation) of 2.67×10^{11} cm^{-2} (6.4 pW) after which all the RIQHE states weakened. While we cannot identify the mechanisms causing localization, it appears that the different proximity of the superlattice to the 2DEG in wafers B and C has a significant impact. Based on our measurements from all three devices, we estimate that ~ 10 pW is the maximum acceptable power dissipation from the gate leakage before the RIQHE states begin to weaken given the cooling power of our dilution refrigerator.

This extremely low power level serves to highlight the necessity of minimizing the gate leakage in order to study the 2^{nd} LL.

6.4 Conclusion

To summarize, we have examined the effect of heterostructure design and device processing on the performance of in-situ back-gated 2DEGs in the 2^{nd} LL. We found that the position of the GaAs/AlAs superlattice barrier relative to quantum well has a large impact on the leakage characteristics of the device due to its effectiveness in blocking the diffusion of the Ohmic contacts towards the gate and minimizing Fowler-Nordheim tunneling. Moving the superlattice closer to the 2DEG greatly increases the range of low-leakage gating without significantly degrading the strength of the gap at $\nu = 5/2$ or other correlated states in the 2^{nd} LL. In addition, we found that gate leakage dissipation powers as small as a few pW are sufficient to cause electronic heating that impacts transport in the 2^{nd} LL. By improving the heat sinking of the lattice, the acceptable power dissipation is increased to $\sim$ 10 pW. Moreover, it is likely that the FQHE gaps would continue to rise at higher density beyond what we report here if the electric sub-bands were spaced sufficiently far apart. Thus, examining gaps as a function of density in narrower quantum wells could potentially yield important results on the density dependence of the gap at more exotic fractions such as $\nu = 12/5$. As we have demonstrated a robust recipe for these structures and as the FQHE states in the 2^{nd} LL are very strong over a wide range of density, these devices should prove an interesting platform for studying transport in nanostructures as a function of density in the 2^{nd} LL.

Acknowledgment

This work was supported by the US DOE Office of Basic Energy Sciences, Division of Materials Sciences and Engineering Award DE-SC0006671. J.D.W. thanks L.A. Tracy and R.L. Willett for helpful discussions regarding device processing.

7. Summary and Future Work

7.1 High Mobility MBE

Our group's experience during the first two growth campaigns has consistently pointed to the purity of the Ga source material as the primary limit on the resulting 2DEG mobility. Further improvements to the gallium quality, however, will likely be a very involved endeavor as material suppliers are typically not interested in changing their procedures without a clear need in industrial settings. Perhaps at some point in the future the Ga outgassing chamber I built will in fact be necessary to make further improvements in Ga quality.

That being said, our group has also done extensive work to show that the zero field mobility is in fact a poor indicator of the quality of 2^{nd} LL physics. In addition, it seems unlikely that major progress in the field (e.g. an experiment which decisively reveals the identity of the ground state at $\nu = 5/2$) will result from bulk transport measurements alone. Future high-impact experiments will very likely require more complex nanostructures. Due to the poor gate-ability of the doping well heterostructures that typically give the best 2^{nd} LL transport, however, it appears that improvements in heterostructure design which could allow more stable gating while preserving the strength of the states in the 2^{nd} LL would have a large impact on the field.

7.2 2D Hole Systems

Due to the combination of large effective mass and high mobility, 2D hole systems in GaAs offer a unique platform for studying strongly interacting systems. In addition, the tuneability of spin-orbit coupling and effective mass allow for extra experimental

knobs not present in the conduction band of GaAs. Moreover, the reduced hyperfine coupling of holes to nuclear spins presents a possible avenue to more robust spin qubits in GaAs.

Given that our results in chapter 5 indicate that interface roughness scattering may limit the mobility at higher density, it would be interesting to more thoroughly explore symmetrically doped quantum wells as well as other heterostructure designs intended to minimize interface roughness scattering. By contrast, the mobility in low density samples appears to be limited by remote impurity scattering, and, as such, growing samples with even larger dopant setbacks could be important for work with metal-insulator transitions, for instance. The relative importance of remote impurity scattering, however, is still not firmly established experimentally since the zero field mobility had not saturated by 300 mK in the wafers measured in chapter 5. Characterizations at dilution refrigerator temperatures should, therefore, shed some light on routes to higher quality low density hole systems.

7.3 The $\nu = 5/2$ FQHE State

With the potential for discovering a non-Abelian state of matter at $\nu = 5/2$ there is significant motivation for continuing to study this system. In addition, there is still much that is unclear regarding the true nature of the ground state. Given the incomplete and sometimes even contradictory results of various experiments, it seems clear that systematic studies of the influence of heterostructure and device designs are worth pursuing. In particular, a clearer understanding of the tuneability of the edge state structure could provide motivation for more theoretical work including the impact of edge reconstructions on transport in nanostructures. The density-tuneable back-gated heterostructures presented in chapter 6 are now at a position to be useful in these experiments. Given the extremely high quality of the transport and the widely tuneable density, these structures will serve as an ideal platform for studying nanostructure transport with varying energy scales.

APPENDICES

A. Computer Codes

In this section I've included some input files for the Nextnano3 Schrodinger/Poisson solver as well as the Matlab code I used to calculate the scattering rates in [102]. For the Nextnano simulations of gated devices sometimes small, seemingly insignificant changes can cause the simulation to not converge. The scattering calculations were intended for estimating the transport lifetime resulting from various scattering mechanisms as a function of wafer design. The calculations primarily follow those in reference [130]; they can be easily adjusted to calculate the quantum lifetime as well.

A.1 Standard Structure Nextnano Input

The code shown below is what I have been using for our so-called "standard structure" which is a doping-well type heterostructure with a 30 nm quantum well. This structure has been found to maximize the mobility and $\nu = 5/2$ energy gap, and we commonly use it as a test of the cleanliness of the MBE. The main issue in simulating these structures is that Nextnano calculates the dopant ionization prior to doing any Schrodinger calculation. This means that if you naively input the binding energy of Si in GaAs ($\sim$ 6 meV), Nextnano will show flat bands. To get the code to show physically reasonable charge accumulation in the quantum well it is necessary to first calculate the binding energy with the spatial quantization of the doping well taken into account and then insert this by hand into the code. To first order the dopant level E_d which you should input into Nextnano is given by $E_d = E_0 - E_b$ where E_0 is the ground state energy of the GaAs/AlAs finite square well referenced to the bottom of the well (remember this is backwards from many elementary quantum mechanics textbooks which usually define the energy levels relative to the top of the well) and E_b is the binding energy of Si in GaAs ($\sim$ 6 meV). Assuming you insert enough Si into

the doping layer, E_0 will define the position of the Fermi level (Nextnano defines this as $E = 0$ in the band structure output) and thus cause charge to accumulate in the main quantum well. Since Nextnano defines a positive E_d as below the conduction band edge, you need to stick a minus sign in front of the energy you calculate for E_d to put the Fermi level above the conduction band edge in the GaAs layer in the doping wells. I've inserted comments where appropriate to document the flow of the code below; the Nextnano3 website (http://www.nextnano.com/nextnano3/) has a lot of documentation, though, and the "Keywords" section is particularly useful for understanding the various parts of the input files. This input file was run successfully on a 64-bit Windows 7 PC with the "nn3_Intel_32bit.exe" executable compiled on 6/9/2011. Note that newer (or older) versions of Nextnano may or may not run this file successfully. I would recommend saving all versions of Nextnano you use in their own folders in case a new release fixes a bug for one input file but creates a new bug for a different input file. Figure A.1 shows the result from the input file given below.

```
!****BEGIN SIMULATION********!
!Exclamation points denote comments
!--------------------------------------------------------------!
$simulation-dimension
dimension   = 1  !Can also do 2 & 3, but that is more complicated
orientation = 0 0 1
$end_simulation-dimension
!--------------------------------------------------------------!

%FunctionParser = yes !Allows the user to use variables - very
                      !useful for creating input files that update
                      !automatically when layer thicknesses are
                      !changed


!--------------------------------------------------------------!
```

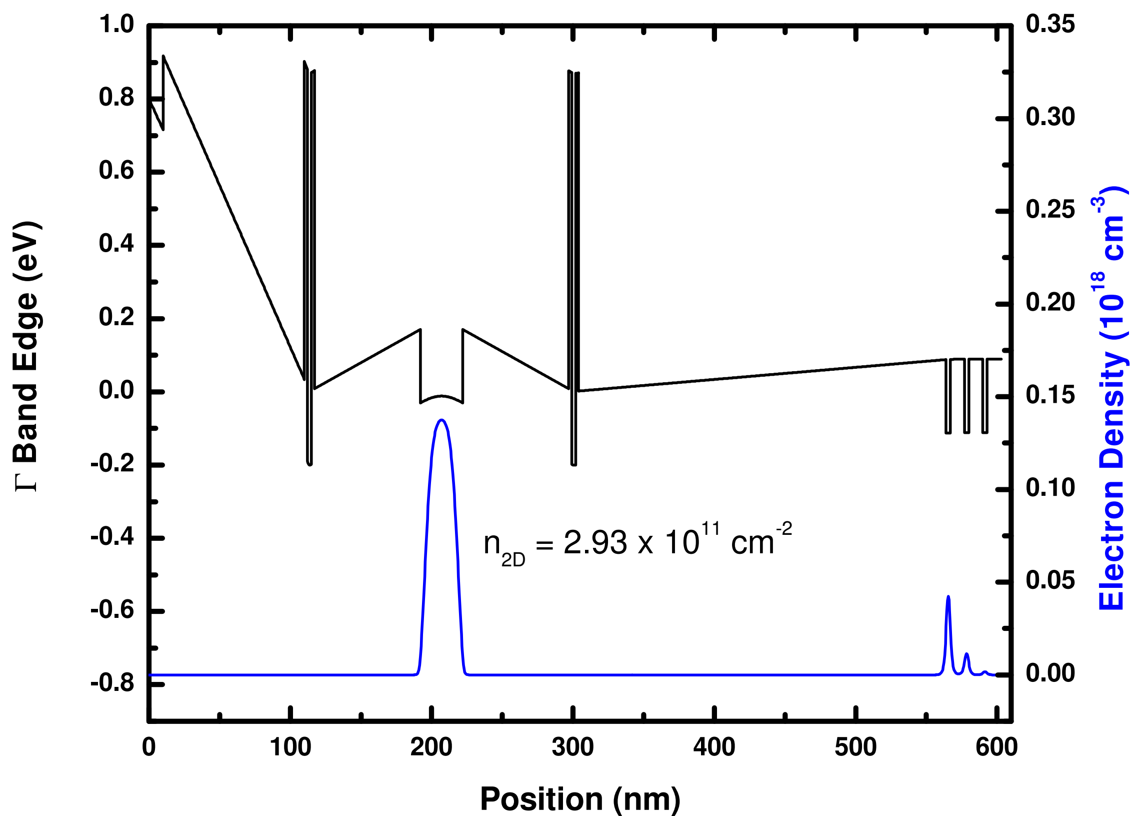

Figure A.1. Output of standard structure simulation.

```
$warnings

warnings = .FALSE. !This doesn't make much difference

$end_warnings

!---------------------------------------------------------------!


!---------------------------------------------------------------!

$global-parameters

lattice-temperature = 1d0 !Kelvin

$end_global-parameters

!---------------------------------------------------------------!
```

```
!-------------------------------------------------------------!
$simulation-flow-control
flow-scheme         = 2 !Calculates Schrodinger/Poisson
                        !self-consistently
raw-directory-in    = raw_data1\ !Not used here, but can be used
                                 !with other flow schemes
raw-potential-in    = no !Used with other flow schemes
strain-calculation  = zero-strain-amorphous !All you need for GaAs
$end_simulation-flow-control
!-------------------------------------------------------------!

!Variables to define simulation region
!Variables denoted by "%"
!-------------------------------------------------------------!
!Layer thicknesses I will use
!Because I used variables the code will update the coordinates
!of the boundaries automatically.  If I hadn't used variables,
!I would have to update all the coordinates following the
!changed layer by hand (very time consuming)

%Cap = 10d0 !Thickness in nm.  "d0" denotes floating point number
%AlGaAsTopSpacer = 100d0
%AlAsBarrier = 2d0
%GaAsDope = 3d0
%TopBarrier = 75d0
%QW = 30d0
%BotBarrier = 75d0
%AlGaAsBotSpacer = 250d0
%AlGaAsSL = 10d0
```

```
%GaAsSL = 3d0

!Define sequence of layers (starting from surface)
!Each variable represents  the coordinate of a layer boundary
%R2 = %Cap !Start with region 2 - region 1 is used for a Poisson
           !boundary condition defined later
%R3 = %R2 + %AlGaAsTopSpacer
%R4 = %R3 + %AlAsBarrier
%R5 = %R4 + %GaAsDope
%R6 = %R5 + %AlAsBarrier
%R7 = %R6 + %TopBarrier
%R8 = %R7 + %QW
%R9 = %R8 + %BotBarrier
%Rten = %R9 + %AlAsBarrier !NN doesn't like "R10"
%Releven = %Rten + %GaAsDope
%Rtwelve = %Releven + %AlAsBarrier
%Rthirteen = %Rtwelve + %AlGaAsBotSpacer
%Rfourteen = %Rthirteen + %AlGaAsSL
%Rfifteen  = %Rfourteen + %GaAsSL
%Rsixteen = %Rfifteen + %AlGaAsSL
%Rseventeen = %Rsixteen + %GaAsSL
%Reighteen = %Rseventeen + %AlGaAsSL
%Rnineteen = %Reighteen + %GaAsSL
%Rtwenty = %Rnineteen +%AlGaAsSL

%TopDeltaHigh = %R5 - 1.45d0 + .5d0 !Define a delta doping layer by
                                    !a 1 nm thick layer of doped
                                    !material (1 nm makes the math
                                    !easier). Center the doping in
```

```
                                    !the middle of the doping well
%TopDeltaLow  = %R5 - 1.45d0 - .5d0
%BotDeltaHigh = %Releven - 1.45d0 + .5d0
%BotDeltaLow  = %Releven - 1.45d0 - .5d0

!--------------------------------------------------------------!
$domain-coordinates
domain-type            = 0 0 1          !Do not change
z-coordinates          = -.5d0 %Rtwenty !Start at -.5d0 for the
                                        !Poisson boundary.  First
                                        !real layer will start at
                                        !0d0
growth-coordinate-axis = 0 0 1
pseudomorphic-on       = GaAs           !Do not change
$end_domain-coordinates
!--------------------------------------------------------------!

!--------------------------------------------------------------!
!Now assign the coordinates you created in variables to
!"region-numbers" (i.e. what Nextnano uses internally to define
!different materials, grid points, etc.)
$regions
 region-number  = 1    base-geometry = line   region-priority = 1
 z-coordinates = -.5d0 0d0

 region-number  = 2    base-geometry = line   region-priority = 1
 z-coordinates = 0d0 %R2

 region-number  = 3    base-geometry = line   region-priority = 1
```

```
z-coordinates = %R2 %R3

region-number  = 4    base-geometry = line   region-priority = 1
z-coordinates = %R3 %R4

region-number  = 5    base-geometry = line   region-priority = 1
z-coordinates = %R4 %R5

region-number  = 6    base-geometry = line   region-priority = 1
z-coordinates = %R5 %R6

region-number  = 7    base-geometry = line   region-priority = 1
z-coordinates = %R6 %R7

region-number  = 8    base-geometry = line   region-priority = 1
z-coordinates = %R7 %R8

region-number  = 9    base-geometry = line   region-priority = 1
z-coordinates = %R8 %R9

region-number  = 10    base-geometry = line   region-priority = 1
z-coordinates = %R9 %Rten

region-number  = 11    base-geometry = line   region-priority = 1
z-coordinates = %Rten %Releven

region-number  = 12    base-geometry = line   region-priority = 1
z-coordinates = %Releven %Rtwelve
```

```
 region-number  = 13    base-geometry = line   region-priority = 1
 z-coordinates = %Rtwelve %Rthirteen

 region-number  = 14    base-geometry = line   region-priority = 1
 z-coordinates = %Rthirteen %Rfourteen

 region-number  = 15    base-geometry = line   region-priority = 1
 z-coordinates = %Rfourteen %Rfifteen

 region-number  = 16    base-geometry = line   region-priority = 1
 z-coordinates = %Rfifteen %Rsixteen

 region-number  = 17    base-geometry = line   region-priority = 1
 z-coordinates = %Rsixteen %Rseventeen

 region-number  = 18    base-geometry = line   region-priority = 1
 z-coordinates = %Rseventeen %Reighteen

 region-number  = 19    base-geometry = line   region-priority = 1
 z-coordinates = %Reighteen %Rnineteen

 region-number  = 20    base-geometry = line   region-priority = 1
 z-coordinates = %Rnineteen %Rtwenty
$end_regions
!-------------------------------------------------------------!

%AlAsFactor = 1d0 !Other values can be used for non-uniform grid
%GaAsFactor = 1d0
%AlAsGrid = 7 !Will give 0.25nm spacing
```

```
%GaAsGrid = 5 !Will give 0.25nm spacing
%INT(AlAsGrid) = %AlAsGrid !Convert to integer
%INT(GaAsGrid) = %GaAsGrid

!-------------------------------------------------------------!
!Note z-grid-lines, z-nodes, z-grid-factors should each be on their
!own line.  These sections extend to multiple lines for formatting
!purposes in this dissertation
$grid-specification
 grid-type        = 0 0 1 !Don't change
!Grid line at boundary of each region
 z-grid-lines     = -.5d0 0d0 %R2 %R3 %R4 %R5 %R6 %R7 %R8 %R9 %Rten
                        %Releven %Rtwelve %Rthirteen %Rfourteen
                        %Rfifteen %Rsixteen %Rseventeen %Reighteen
                        %Rnineteen    %Rtwenty

!z-nodes = # of grid points.  Set this to a uniform value.  As a
!check with new input files, always set the donor binding energy to
!-100d0 (i.e. fully ionized) and make sure ionized impurity
!concentration matches what you told Nextnano (poor grid spacing can
!cause NN to get this wrong)

z-nodes          = 1 19 199 %INT(AlAsGrid) %INT(GaAsGrid)
                    %INT(AlAsGrid) 149 59 149 %INT(AlAsGrid)
                    %INT(GaAsGrid) %INT(AlAsGrid) 499 19 5 19 5 19
                    5 19
z-grid-factors   = 1d0 1d0 1d0 %AlAsFactor %GaAsFactor %AlAsFactor 1d0
                   1d0 1d0 %AlAsFactor %GaAsFactor %AlAsFactor 1d0 1d0
                   1d0 1d0 1d0 1d0 1d0 1d0
```

```
$end_grid-specification
!--------------------------------------------------------------!


!--------------------------------------------------------------!
!Assign regions to clusters.  You will assign each cluster to a
!material and/or Poisson boundary condition later
$region-cluster
cluster-number = 1   region-numbers = 1
cluster-number = 2   region-numbers = 2 5 8 11 15 17 19 21
cluster-number = 3   region-numbers = 3 7 9 13 14 16 18 20
cluster-number = 4   region-numbers = 4 6 10 12
$end_region-cluster
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$material !See Keywords on website for allowed materials
!Cluster 1 will be a Poisson boundary condition
material-number = 1
material-name   = GaAs
cluster-numbers = 1

material-number = 2
material-name   = GaAs
cluster-numbers = 2

material-number = 3
material-name   = Al(x)Ga(1-x)As
cluster-numbers = 3
alloy-function  = constant
```

```
material-number = 4
material-name   = AlAs
cluster-numbers = 4
$end_material
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$alloy-function
material-number      = 3
function-name         = constant
xalloy                       = 0.24d0
$end_alloy-function
!--------------------------------------------------------------!

!--------------------------------------------------------------!
!Define coordinates of doped regions.  "only-region" defines min/max
!z coordinates of that doping function.
$doping-function
doping-function-number     = 1
impurity-number            = 1
doping-concentration       = 10d0 !Units: 10^18 cm^-3 - i.e.
                                  !10*10^18 cm^-3 * 1 nm = 1E12 cm^-2
                                  !sheet density
only-region                = %TopDeltaLow %TopDeltaHigh

doping-function-number     = 2
impurity-number            = 1
doping-concentration       = 8d0
```

```
only-region                = %BotDeltaLow %BotDeltaHigh
$end_doping-function
!--------------------------------------------------------------!


!--------------------------------------------------------------!
!Specify properties of each dopant used
$impurity-parameters
impurity-number              = 1
impurity-type                = n-type ! n-type, p-type
number-of-energy-levels      = 1      !Do not change
energy-levels-relative       = -0.2d0 !Binding energy of dopant [eV].
                                      !Negative value means dopant
                                      !level above conduction band
                                      !edge (or below valence band
                                      !edge)
degeneracy-of-energy-levels  = 2  !2 for n-type, 4 for p-type
$end_impurity-parameters                                                 !
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$poisson-boundary-conditions
!Specify Poisson boundary conditions.  Here this is used to set Fermi
!level pinning due to surface states
poisson-cluster-number     = 1 !Can use this for defining current
                               !calculations or voltage sweeps applied
                               !at different points in the
                               !heterostructure
region-cluster-number      = 1
applied-voltage            = 0.0d0
```

```
boundary-condition-type    = schottky !See Keywords for options
contact-control            = voltage
schottky-barrier           = .8d0 !Height of barrier [eV]
$end_poisson-boundary-conditions
!--------------------------------------------------------------!


!--------------------------------------------------------------!
!Define variables for position of quantum regions
%Quant1low = %TopDeltaHigh - 20d0
%Quant1high = %TopDeltaHigh + 20d0
%Quant2low = %R7 - 30d0
%Quant2high = %R8 + 30d0
%Quant3low = %BotDeltaHigh - 20d0
%Quant3high = %BotDeltaHigh + 40d0
%Quant4low = %Rfourteen - 20d0
%Quant4high = %Rtwenty

!--------------------------------------------------------------!
$quantum-regions
!Define regions where Schrodinger will be solved
region-number          = 1 !Top doping region
base-geometry          = line
region-priority        = 1
z-coordinates          = %Quant1low %Quant1high

region-number          = 2 !2DEG
base-geometry          = line
region-priority        = 1
z-coordinates          = %Quant2low %Quant2high
```

```
region-number          = 3 !Bottom doping region
base-geometry          = line
region-priority        = 1
z-coordinates          = %Quant3low %Quant3high

region-number          = 4 !Superlattice
base-geometry          = line
region-priority        = 1
z-coordinates          = %Quant4low %Quant4high
$end_quantum-regions
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-cluster
!Group quantum regions into cluster. Will later specify unique quantum
!calculations for each cluster
 cluster-number         = 1 !Quantum cluster number
 region-numbers         = 1 !Quantum region number
 deactivate-cluster     = no

 cluster-number         = 2
 region-numbers         = 2
 deactivate-cluster     = no

 cluster-number         = 3
 region-numbers         = 3
 deactivate-cluster     = no
```

```
 cluster-number         = 4
 region-numbers         = 4
 deactivate-cluster     = no
$end_quantum-cluster
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$quantum-model-holes
model-number                            =  1
model-name                              =  effective-mass !8x8kp,
                                           !6x6kp or effective-mass
cluster-numbers                         =  1
valence-band-numbers                    =  1 2 3  !i.e. HH, LH,
                                                  !Split-off
number-of-eigenvalues-per-band          =  3  3  3  !Array size must
                                          !match array size of
                                          !valence-band-numbers
boundary-condition-001                  =  Dirichlet !See Keywords
boundary-condition-010                  =  Dirichlet
boundary-condition-100                  =  Dirichlet
$end_quantum-model-holes
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$quantum-model-electrons
model-number                            =  1
model-name                              =  effective-mass !8x8kp or
                                          !effective-mass
cluster-numbers                         =  1
```

```
conduction-band-numbers                 =  1 2 3 !i.e. Gamma, L, X
number-of-eigenvalues-per-band          =  3 3 3 !Array size must
                                          !match array size of
                                          !conduction-band-numbers
boundary-condition-001                  =  Dirichlet !See Keywords
boundary-condition-010                  =  Dirichlet
boundary-condition-100                  =  Dirichlet
$end_quantum-model-electrons
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-model-holes
model-number                            =  2
model-name                              =  effective-mass !8x8kp,
                                           !6x6kp or effective-mass
cluster-numbers                         =  2
valence-band-numbers                    =  1 2 3  !i.e. HH, LH,
                                                  !Split-off
number-of-eigenvalues-per-band          =  3  3  3  !Array size must
                                          !match array size of
                                          !valence-band-numbers
boundary-condition-001                  =  Dirichlet !See Keywords
boundary-condition-010                  =  Dirichlet
boundary-condition-100                  =  Dirichlet
$end_quantum-model-holes
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-model-electrons
```

```
model-number                          =  2
model-name                            =  effective-mass !8x8kp or
                                        !effective-mass
cluster-numbers                       =  2
conduction-band-numbers               =  1 2 3 !i.e. Gamma, L, X
number-of-eigenvalues-per-band        =  3 3 3 !Array size must
                                        !match array size of
                                        !conduction-band-numbers
boundary-condition-001                =  Dirichlet !See Keywords
boundary-condition-010                =  Dirichlet
boundary-condition-100                =  Dirichlet
$end_quantum-model-electrons
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-model-holes
model-number                          =  3
model-name                            =  effective-mass !8x8kp,
                                         !6x6kp or effective-mass
cluster-numbers                       =  3
valence-band-numbers                  =  1 2 3  !i.e. HH, LH,
                                                !Split-off
number-of-eigenvalues-per-band        =  3  3  3  !Array size must
                                        !match array size of
                                        !valence-band-numbers
boundary-condition-001                =  Dirichlet !See Keywords
boundary-condition-010                =  Dirichlet
boundary-condition-100                =  Dirichlet
$end_quantum-model-holes
```

```
!----------------------------------------------------------------!

!----------------------------------------------------------------!
$quantum-model-electrons
model-number                            =  3
model-name                              =  effective-mass !8x8kp or
                                          !effective-mass
cluster-numbers                         =  3
conduction-band-numbers                 =  1 2 3 !i.e. Gamma, L, X
number-of-eigenvalues-per-band          =  3 3 3 !Array size must
                                          !match array size of
                                          !conduction-band-numbers
boundary-condition-001                  =  Dirichlet !See Keywords
boundary-condition-010                  =  Dirichlet
boundary-condition-100                  =  Dirichlet
$end_quantum-model-electrons
!----------------------------------------------------------------!

!----------------------------------------------------------------!
$quantum-model-holes
model-number                            =  4
model-name                              =  effective-mass !8x8kp,
                                           !6x6kp or effective-mass
cluster-numbers                         =  4
valence-band-numbers                    =  1 2 3  !i.e. HH, LH,
                                                  !Split-off
number-of-eigenvalues-per-band          =  3  3  3  !Array size must
                                          !match array size of
                                          !valence-band-numbers
```

```
boundary-condition-001                  =  Dirichlet !See Keywords
boundary-condition-010                  =  Dirichlet
boundary-condition-100                  =  Dirichlet
$end_quantum-model-holes
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-model-electrons
model-number                            =  4
model-name                              =  effective-mass !8x8kp or
                                          !effective-mass
cluster-numbers                         =  4
conduction-band-numbers                 =  1 2 3 !i.e. Gamma, L, X
number-of-eigenvalues-per-band          =  3 3 3 !Array size must
                                          !match array size of
                                          !conduction-band-numbers
boundary-condition-001                  =  Dirichlet !See Keywords
boundary-condition-010                  =  Dirichlet
boundary-condition-100                  =  Dirichlet
$end_quantum-model-electrons
!--------------------------------------------------------------!

!--------------------------------------------------------------!
!Used for sequential calculations with different flow schemes (i.e.
!not important in this example)
$output-raw-data
 destination-directory                   = raw_data1\
 potential                               = yes
 fermi-levels                            = yes
```

```
 kp-eigenstates                           = no
$end_output-raw-data
!--------------------------------------------------------------!

!****IMPORTANT**************************!
!NOTE: CHANGING WHAT NEXTNANO OUTPUTS CAN AFFECT THE CALCULATIONS
!THEMSELVES. BE CAREFUL WHEN CHANGING PARAMETERS IN THIS SECTION.
!ON THE FLIP SIDE, IF A CALCULATION IS NOT CONVERGING, TRY CHANGING
!THE OUTPUT SECTIONS - SOMETIMES THIS HELPS
!****************************************!

!--------------------------------------------------------------!
$output-1-band-schroedinger
 destination-directory                    = sg_1band1\
 sg-structure                             = yes
 conduction-band-numbers                  = 1 2 3
 cb-min-ev                                = 1
 cb-max-ev                                = 10
 valence-band-numbers                     = 1 2 3
 vb-min-ev                                = 1
 vb-max-ev                                = 10
$end_output-1-band-schroedinger
!--------------------------------------------------------------!

!--------------------------------------------------------------!

!--------------------------------------------------------------!
$output-bandstructure
 destination-directory                    = band_struc1\
```

```
 conduction-band-numbers                = 1 2 3
 valence-band-numbers                   = 1 2 3
 potential                              = yes
$end_output-bandstructure
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$output-densities
 destination-directory                  = densities1\
 electrons                              = yes
 holes                                  = yes
 charge-density                         = yes
 intrinsic-density                      = yes
 ionized-dopant-density                 = yes
 piezo-electricity                      = yes
 pyro-electricity                       = yes
 interface-density                      = yes
 integrated-density                     = yes
 subband-density                        = yes
$end_output-densities
!--------------------------------------------------------------!


!*****END OF SIMULATION*****************!
```

A.2 In-Situ Back-Gated 2DEG Nextnano Input

This input file was used to simulate the band structure and electron density of an in-situ back-gated 2DEG under voltage bias. It can be quite tricky to get these gated simulations to converge, so do not be surprised if even minor changes to the

following input file results in a simulation that doesn't converge. In particular, overlap between the Schottky regions and the quantum regions can cause problems. Changes to the output of the current calculation can also cause problems; as called out in the previous input file, changing whether or not Nextnano outputs data from the current calculation can affect the simulation (i.e. cause it to stop converging). Note that in order to simplify the input file I only included the first few superlattice layers closest to the quantum well. After that I simply specified a 50% AlGaAs layer to represent the superlattice. I wasn't too worried about inaccuracies this might cause since I was mostly interested in the electric field in the quantum well, the wavefunction symmetry, and when the second sub-band became occupied. As with the previous input file, this simulation was run with a 64-bit Windows 7 PC with the "nn3_Intel_32bit.exe" execeutable compiled on 6/9/2011. For the sake of time, I have not included as many comments in this input file as compared with the standard structure input file; I would recommend familiarizing yourself with that input file first before attempting the gated simulation. Figure A.2 shows the output generated from the input file below.

```
!****BEGIN SIMULATION********!
!Exclamation points denote comments
!--------------------------------------------------------------!
$simulation-dimension
 dimension   = 1  !Can also do 2 & 3, but that is more complicated
 orientation = 0 0 1
$end_simulation-dimension
!--------------------------------------------------------------!

%FunctionParser = yes !Allows the user to use variables - very
                      !useful for creating input files that update
                      !automatically when layer thicknesses are
                      !changed
```

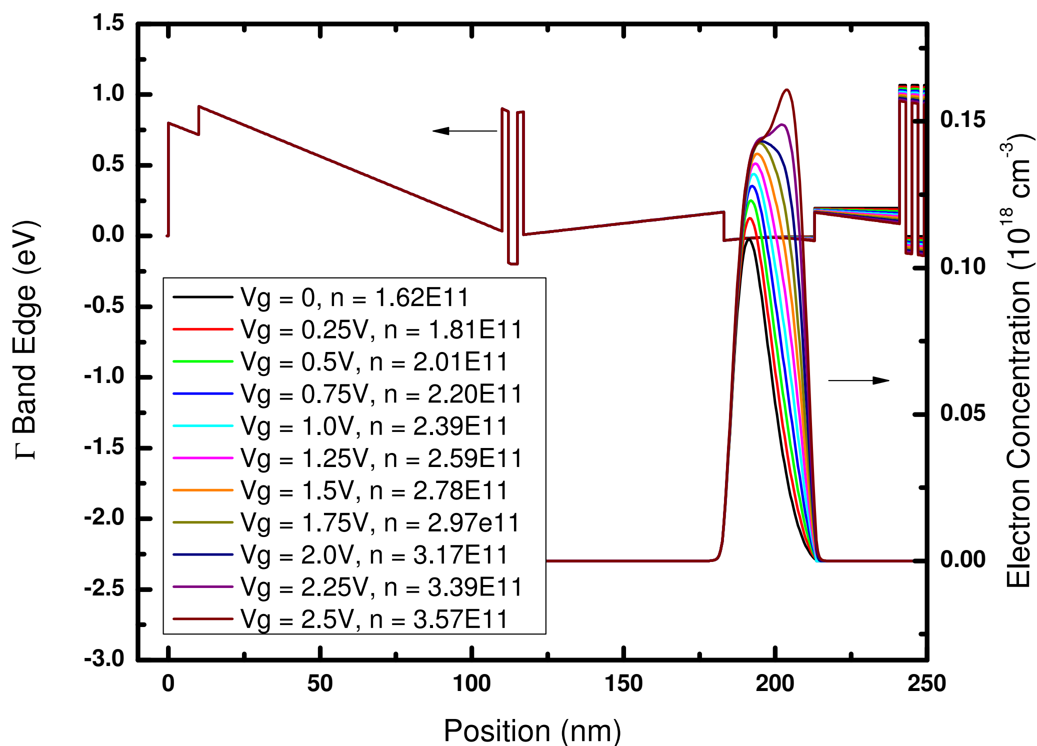

Figure A.2. Output of in-situ back-gated simulation.

```
!--------------------------------------------------------------!
$warnings
warnings = .FALSE. !This doesn't make much difference
$end_warnings
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$global-parameters
 lattice-temperature = 1d0 !Kelvin
$end_global-parameters
!--------------------------------------------------------------!
```

```
!--------------------------------------------------------------!
$numeric-control
 simulation-dimension            = 1
 newton-method                   = Newton-3
 current-poisson-method          = couple-all-false
 current-problem                 = solve-for-fermi-cg
 current-problem-iterations      = 100
 current-problem-residual        = 1d-11
 current-block-relaxation-Fermi  = 1d0
$end_numeric-control
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$simulation-flow-control
 flow-scheme         = 1
 raw-directory-in    = C:\Users\John\Documents\nextnano
                         \Output\ISBG103121\raw_data1\
 raw-potential-in    = no !Used with other flow schemes
 strain-calculation  = zero-strain-amorphous !All you need for GaAs
$end_simulation-flow-control
!--------------------------------------------------------------!


%BotBarrier = 20d0
%AlGaAsBotSpacer = 10d0
%AlAsSL = 2d0
%GaAsSL = 2d0
%AvgSL = 758d0
```

```
%Buffer = 50d0
%Gate = 500d0

%R2 = %Cap
%R3 = %R2 + %AlGaAsTopSpacer
%R4 = %R3 + %AlAsBarrier
%R5 = %R4 + %GaAsDope
%R6 = %R5 + %AlAsBarrier
%R7 = %R6 + %TopBarrier
%R8 = %R7 + %QW
%R9 = %R8 + %BotBarrier
%Rten = %R9 + %AlAsSL
%Releven = %Rten + %GaAsSL
%Rtwelve = %Releven + %AlAsSL
%Rthirteen = %Rtwelve + %GaAsSL
%Rfourteen = %Rthirteen + %AlAsSL
%Rfifteen  = %Rfourteen + %GaAsSL
%Rsixteen = %Rfifteen + %AlAsSL
%Rseventeen = %Rsixteen + %GaAsSL
%Reighteen = %Rseventeen + %AlAsSL
%Rnineteen = %Reighteen + %GaAsSL
%Rtwenty = %Rnineteen +%AlAsSL
%Rt1 = %Rtwenty + %AvgSL
%Rttwo = %Rt1 + %Buffer
%Rtthree = %Rttwo + %Gate

%TopDeltaHigh = %R5 - 1.45d0 + .5d0
%TopDeltaLow  = %R5 - 1.45d0 - .5d0
```

```
!--------------------------------------------------------------!
$domain-coordinates
 domain-type             = 0 0 1
 z-coordinates           = -.5d0 %Rtthree
 growth-coordinate-axis = 0 0 1
 pseudomorphic-on        = GaAs
$end_domain-coordinates
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$regions
 region-number  = 1    base-geometry = line   region-priority = 1
 z-coordinates = -.5d0 0d0
 region-number  = 2    base-geometry = line   region-priority = 1
 z-coordinates = 0d0 %R2
 region-number  = 3    base-geometry = line   region-priority = 1
 z-coordinates = %R2 %R3
 region-number  = 4    base-geometry = line   region-priority = 1
 z-coordinates = %R3 %R4
 region-number  = 5    base-geometry = line   region-priority = 1
 z-coordinates = %R4 %R5
 region-number  = 6    base-geometry = line   region-priority = 1
 z-coordinates = %R5 %R6
 region-number  = 7    base-geometry = line   region-priority = 1
 z-coordinates = %R6 %R7
 region-number  = 8    base-geometry = line   region-priority = 1
 z-coordinates = %R7 %R8
 region-number  = 9    base-geometry = line   region-priority = 1
 z-coordinates = %R8 %R9
```

```
 region-number  = 10    base-geometry = line   region-priority = 1
 z-coordinates = %R9 %Rten
 region-number  = 11    base-geometry = line   region-priority = 1
 z-coordinates = %Rten %Releven
 region-number  = 12    base-geometry = line   region-priority = 1
 z-coordinates = %Releven %Rtwelve
 region-number  = 13    base-geometry = line   region-priority = 1
 z-coordinates = %Rtwelve %Rthirteen
 region-number  = 14    base-geometry = line   region-priority = 1
 z-coordinates = %Rthirteen %Rfourteen
 region-number  = 15    base-geometry = line   region-priority = 1
 z-coordinates = %Rfourteen %Rfifteen
 region-number  = 16    base-geometry = line   region-priority = 1
 z-coordinates = %Rfifteen %Rsixteen
 region-number  = 17    base-geometry = line   region-priority = 1
 z-coordinates = %Rsixteen %Rseventeen
 region-number  = 18    base-geometry = line   region-priority = 1
 z-coordinates = %Rseventeen %Reighteen
 region-number  = 19    base-geometry = line   region-priority = 1
 z-coordinates = %Reighteen %Rnineteen
 region-number  = 20    base-geometry = line   region-priority = 1
 z-coordinates = %Rnineteen %Rtwenty
 region-number  = 21    base-geometry = line   region-priority = 1
 z-coordinates = %Rtwenty %Rt1
 region-number  = 22    base-geometry = line   region-priority = 1
 z-coordinates = %Rt1 %Rttwo
 region-number  = 23    base-geometry = line   region-priority = 1
 z-coordinates = %Rttwo %Rtthree
$end_regions
```

```
!------------------------------------------------------------!

%AlAsFactor = 1d0
%GaAsFactor = 1d0
%AlAsGrid = 7
%GaAsGrid = 5
%INT(AlAsGrid) = %AlAsGrid
%INT(GaAsGrid) = %GaAsGrid

!------------------------------------------------------------!
$grid-specification
 grid-type        = 0 0 1
 z-grid-lines     = -.5d0 0d0 %R2 %R3 %R4 %R5 %R6 %R7 %R8 %R9 %Rten
                    %Releven %Rtwelve %Rthirteen %Rfourteen %Rfifteen
                    %Rsixteen %Rseventeen %Reighteen %Rnineteen
                    %Rtwenty %Rt1 %Rttwo %Rtthree
 z-nodes          = 1 19 199 %INT(AlAsGrid) %INT(GaAsGrid)
                    %INT(AlAsGrid) 131 59 39 3 3 3 3 3 3 3 3 3 3 3
                    378 24 249
 z-grid-factors   = 1d0 1d0 1d0 %AlAsFactor %GaAsFactor %AlAsFactor
                    1d0 1d0 1d0 %AlAsFactor %GaAsFactor %AlAsFactor
                    1d0 1d0 1d0 1d0 1d0 1d0 1d0 1d0 1d0 1d0 1d0
$end_grid-specification
!------------------------------------------------------------!

!------------------------------------------------------------!
$region-cluster
 cluster-number = 1   region-numbers = 1 !Schottky barrier
 cluster-number = 2   region-numbers = 2 5 15 17 19 24 !GaAs
```

```
 cluster-number = 3   region-numbers = 3 7 9 10 11 12 13 !AlGaAs
 cluster-number = 4   region-numbers = 4 6 14 16 18 20 !AlAs
 cluster-number = 5   region-numbers = 21 !Avg Superlattice
 cluster-number = 6   region-numbers = 8 !Quantum well w/Poisson BC
 cluster-number = 7   region-numbers = 22 23 !Backgate w/Poisson BC
$end_region-cluster
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$material
 material-number = 1
 material-name   = GaAs
 cluster-numbers = 1

 material-number = 2
 material-name   = GaAs
 cluster-numbers = 2

 material-number = 3
 material-name   = Al(x)Ga(1-x)As
 cluster-numbers = 3
 alloy-function  = constant

 material-number = 4
 material-name   = AlAs
 cluster-numbers = 4

 material-number = 5
 material-name   = Al(x)Ga(1-x)As
```

```
 cluster-numbers = 5
 alloy-function  = constant

 material-number = 6    !Quantum well Fermi contact
 material-name   = GaAs
 cluster-numbers = 6

 material-number = 7    !Backgate
 material-name   = GaAs
 cluster-numbers = 7
$end_material
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$alloy-function
 material-number        = 3
 function-name          = constant
 xalloy                 = 0.24d0

 material-number        = 5
 function-name          = constant
 xalloy                 = 0.5d0
$end_alloy-function
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$doping-function
 doping-function-number     = 1
 impurity-number            = 1
```

```
 doping-concentration       = 10d0
 only-region                = %TopDeltaLow %TopDeltaHigh

 doping-function-number     = 2 !Backgate
 impurity-number            = 2
 doping-concentration       = 1d0
 only-region                = %Rttwo %Rtthree
$end_doping-function
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$impurity-parameters
 impurity-number              = 1
 impurity-type                = n-type
 number-of-energy-levels      = 1
 energy-levels-relative       = -0.2d0
 degeneracy-of-energy-levels  = 2

 impurity-number              = 2
 impurity-type                = n-type
 number-of-energy-levels      = 1
 energy-levels-relative       = .006d0 !Si in GaAs
 degeneracy-of-energy-levels  = 2
$end_impurity-parameters
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$poisson-boundary-conditions
 poisson-cluster-number     = 1
```

```
 region-cluster-number      = 1
 applied-voltage            = 0.0d0
 boundary-condition-type    = schottky
 contact-control            = voltage
 schottky-barrier           = .8d0

 poisson-cluster-number     = 2 !2DEG
 region-cluster-number      = 6
 applied-voltage            = 0.0d0
 boundary-condition-type    = Fermi !i.e. QW = ground
 contact-control            = voltage

 poisson-cluster-number     = 3 !backgate
 region-cluster-number      = 7
 applied-voltage            = 0.0d0
 boundary-condition-type    = schottky
 contact-control            = voltage
 schottky-barrier           = .364d0
$end_poisson-boundary-conditions
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$current-cluster
 cluster-number                          = 1
 region-numbers                          = 1
 deactivate-cluster                      = no
$end_current-cluster
!--------------------------------------------------------------!
```

```
!--------------------------------------------------------------!
$current-models
 model-number                              = 1
 transport-model-name                      = simple-drift-model
 cluster-numbers                           = 1
$end_current-models
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$simple-drift-models
 model-number                              = 1
 mobility-model                            = mobility-model-simba-1
 current-model-numbers                     = 1
! charge-carriers                           = holes-only
$end_simple-drift-models
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$voltage-sweep
 sweep-number              = 1
 sweep-active              = yes
 poisson-cluster-number    = 3
 step-size                 = .05d0
 number-of-steps           = 50
 data-out-every-nth-step   = 5
$end_voltage-sweep
!--------------------------------------------------------------!

%Quant1low = %TopDeltaHigh - 20d0
```

```
%Quant1high = %TopDeltaHigh + 20d0
%Quant2low = %R7 - 30d0
%Quant2high = %R8 + 10d0
%Quant3low = %Rfourteen - 8d0
%Quant3high = %Rtwenty + 20d0

!------------------------------------------------------------!
$quantum-regions
 region-number          = 1
 base-geometry          = line
 region-priority        = 1
 z-coordinates          = %Quant1low %Quant1high

 region-number          = 2
 base-geometry          = line
 region-priority        = 1
 z-coordinates          = %Quant2low %Quant2high

 region-number          = 3
 base-geometry          = line
 region-priority        = 1
 z-coordinates          = %Quant3low %Quant3high
$end_quantum-regions
!------------------------------------------------------------!

!------------------------------------------------------------!
$quantum-cluster
 cluster-number         = 1
 region-numbers         = 1
```

```
 deactivate-cluster     = no

 cluster-number         = 2
 region-numbers         = 2
 deactivate-cluster     = no

 cluster-number         = 3
 region-numbers         = 3
 deactivate-cluster     = no
$end_quantum-cluster
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-model-electrons
 model-number                                =  1
 model-name                                  =  effective-mass
 cluster-numbers                             =  1
 conduction-band-numbers                     =  1 3 !Gamma and X
 number-of-eigenvalues-per-band              =  3 3
 boundary-condition-001                      =  Neumann
 boundary-condition-010                      =  Neumann
 boundary-condition-100                      =  Neumann
$end_quantum-model-electrons
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-model-electrons
 model-number                                =  2
 model-name                                  =  effective-mass
```

```
 cluster-numbers                          =  2
 conduction-band-numbers                  =  1 3 !Gamma and X
 number-of-eigenvalues-per-band           =  3 3
 boundary-condition-001                   =  Neumann
 boundary-condition-010                   =  Neumann
 boundary-condition-100                   =  Neumann
$end_quantum-model-electrons
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$quantum-model-electrons
 model-number                             =  3
 model-name                               =  effective-mass
 cluster-numbers                          =  3
 conduction-band-numbers                  =  1 3 !Gamma and X
 number-of-eigenvalues-per-band           =  3 3
 boundary-condition-001                   =  Neumann
 boundary-condition-010                   =  Neumann
 boundary-condition-100                   =  Neumann
$end_quantum-model-electrons
!--------------------------------------------------------------!

!--------------------------------------------------------------!
$output-raw-data
 destination-directory                    = raw_data1\
 potential                                = yes
 fermi-levels                             = yes
 kp-eigenstates                           = no
$end_output-raw-data
```

```
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$output-1-band-schroedinger
 destination-directory                        = sg_1band1\
 sg-structure                                 = yes
 conduction-band-numbers                      = 1  3
 cb-min-ev                                    = 1
 cb-max-ev                                    = 10
$end_output-1-band-schroedinger
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$output-bandstructure
 destination-directory                        = band_struc1\
 conduction-band-numbers                      = 1 2 3
 potential                                    = yes
$end_output-bandstructure
!--------------------------------------------------------------!


!--------------------------------------------------------------!
$output-densities
 destination-directory                        = densities1\
 electrons                                    = yes
 charge-density                               = yes
 ionized-dopant-density                       = yes
 integrated-density                           = yes
$end_output-densities
!--------------------------------------------------------------!
```

```
!--------------------------------------------------------------!
$output-current-data
 destination-directory                    = current1\
 current                                  = no
 fermi-levels                             = yes
 mobility-out                             = no
$end_output-current-data
!--------------------------------------------------------------!

!*********END OF SIMULATION***************!
```

A.3 Matlab Code for Calculating Transport Lifetimes

This section has the code I used in reference [102] to calculate the transport lifetime from different scattering mechanism for each wafer I grew. The code opens the wavefunction (from the Schrodinger output folder) produced from the appropriate Nextnano simulation. In order to keep the coordinate systems consistent, the Nextnano input file must be setup to start the wavefunction at z = 0 (i.e. the surface). In other words, use a single quantum region that extends from the surface to below any scattering sites (i.e. the bottom doping layer). Once the Nextnano simulation is done, specify the file path of the wavefunction data in the Matlab code (i.e. "A=importdata();"). One useful check on your understanding of the code is to try to reproduce the results on ionized impurity scattering from reference [118], digitize their plots, and overlay your results; I was able to get essentially identical results using my code below.

```
% Denotes comment
clear;
%Summary of 2D hole density from simulations
```

```
%d80 = 80nm dopant setback
%x07 = 7% AlGaAs barrier
 %d80x07 p = 2.3e10
 %d80x10 p = 3.5E10
 %d80x13 p = 4.7E10
 %d80x16 p = 6.0E10
 %d80x20 p = 7.6E10
 %d80x24 p = 9.1E10
 %d80x35 p = 1.3E11
 %d80x45 p = 1.7E11

 Holes = 1; %1 for holes, 0 for electrons
 N_S = 1.7e11*1E4; %2DEG concentration in m^-2
 %V_surf = 0.6; %Surface potential in Volts
 A=importdata('H:\My Documents\MATLAB\d80x45alloy.dat');
 N_ac = 2e13*1e6; %5E13*1E6; %acceptor concentration [m^-3]
 x = .45; %Al mole fraction
 N_RI = N_S; %Assume RI conc = 2DEG/2DHG density
%This is necessary to get reasonable results since the scattering
%model does not take correlation in doping layer into account
 delta_V = 1.0*1.6e-19; %alloy scattering potential [J]
 Delta = .1e-9; %Interface roughness height [m]
 Lambda = 22e-10; %Interface roughness lateral size [m]
 L = 20e-9; %Finite square well width [m] - for IR calculation

 %****Calculation to Perform*****
 Remote = 0; %yes = 1, no = 0
 Background = 0;
 Alloy = 1;
```

```
IR_FNTSQWL = 0;
IR_FNGHWRD = 0;
%*****************************


hbar=1.05e-34; %in J.s
chg=1.6e-19; %electronic charge in C
epsilonnot=8.85e-12; %in C^2/N.m^2
kappa=12.9*4*pi; %note that Bastard calls
                 %kappa = 4*pi*12.9*epsilonnot
DopePos = 110e-9; %110nm for 80nm setback
kf = sqrt(2*pi*N_S);

a = .566e-9; %lattice constant for alloy scattering calculation
r = sqrt(3)*a/4; %alloy scattering potential range
Omega_not = (4/3)*pi*r^3; %alloy spherical potential well volume



if(Holes == 1)
 if(N_S < 9E14)
    m = (N_S*1e-4*2.5e-12+.27)*9.1e-31;
 else
    m = .5*9.1e-31;
 end
else
    m = .067*9.1e-31;
end
```

```
 %N_Surf = V_surf*kappa*epsilonnot/(DopePos*chg); %Surface charge
                                                       %stolen from
                                                       %delta layer
                                                       %[m^-2]
 %N_RI = N_S+N_Surf; %Net RI concentration in delta layer [m^-2]

 %Import mod square of wave function from nextnano
 RawData=A.data;
 z=RawData(:,1)*1e-9;
 Usqr=RawData(:,2);
 scatter(z,Usqr);

 %interpolate envelope function and establish real space grid
 zi=min(z):.5*1e-9:max(z);
 Usqri=1e9*interp1(z,Usqr,zi,'pchip');

 scatter(zi,Usqri);

 area=(zi(2)-zi(1))*trapz(Usqri);

%
% %Begin scattering calculation
% %
%
  q_not=2*m*chg^2/(epsilonnot*kappa*hbar^2);
%T.F. screening wavevector (=2/bohr radius)
%
%calculate form factor g(theta), theta array must have
%same length as zi array
```

```
 theta=0:pi/size(zi,2):(pi-pi/size(zi,2));

if(Remote == 1)
 G_imp = zeros(1,size(theta,2));

 for index = 1:size(theta,2)
   G_imp(index) = N_RI*((zi(2)-zi(1))*trapz(Usqri.*exp(-2*kf*
                   sin(theta(1,index)/2)*abs(DopePos-zi))))^2;
%Note previous equation is split to two lines to fit dissertation
%format
 end
end
% % %calculate g_s(q) screening form factor

if((Remote == 1) || (Background == 1))
  z_prime_integrand=zeros(1,size(zi,2));
  g_s = zeros(1,size(theta,2));

 for index = 1:size(theta,2)
   for index2 = 1:size(zi,2)
             z_prime_integrand(index2) = Usqri(index2)*
          (zi(2)-zi(1))*trapz(Usqri.*exp(-2*kf*
          sin(theta(1,index)/2)*abs(zi-zi(1,index2))));
%Note previous equation is split to two lines to fit dissertation
%format
   end
     g_s(index) = (zi(2)-zi(1))*trapz(z_prime_integrand);
 end
end
```

```
% %  %g_s = 1;

if(Remote == 1)
%  %scattering rate
 RI_rate = (m/(pi*hbar^3))*(theta(2)-theta(1))*trapz((1-cos(theta))
          .*G_imp.*(2*pi*chg^2./(epsilonnot*kappa*(2*kf*sin(theta/2)
         +q_not*g_s))).^2);
 %Note previous equation is split to two lines to fit dissertation
%format

 mu_RI = 1e4*chg/(RI_rate*m); %mu in cm^2/Vs
end
%  %BI impurities

if(Background == 1)
 clear c;
 clear G_imp;

 c = zeros(1,size(zi,2));
 for index = 1:size(zi,2)
     if((zi(1,index)>0)); % && (zi(1,index)>-10.4e-9-7.8e-9))
         c(index) = N_ac;
     else
         c(index) = 0;
     end
 end

 g_imp = zeros(size(theta,2),size(zi,2));
 for index = 1:size(theta,2)
```

```
   for index2 = 1:size(zi,2)
             g_imp(index,index2) = ((zi(2)-zi(1))*trapz(Usqri
                                   .*exp(-2*kf*sin(theta(1,index)/2)
                                   *abs(zi-zi(1,index2)))))^2;
%Note previous equation is split to two lines to fit dissertation
%format

   end
 end

 G_imp = zeros(1,size(theta,2));
 for index =1:size(theta,2)
     G_imp(index) = (zi(2)-zi(1))*trapz(c.*g_imp(index,:));
 end

 %scattering rate - BI in channel
 BI_rate = (m/(pi*hbar^3))*(theta(2)-theta(1))*trapz((1-cos(theta))
          .*G_imp.*(2*pi*chg^2./(epsilonnot*kappa*(2*kf*sin(theta/2)
         +q_not*g_s))).^2);
%Note previous equation is split to two lines to fit dissertation
%format

 mu_BI = 1e4*chg/(BI_rate*m); %mu in cm^2/Vs
end

% %Alloy scattering estimate
if(Alloy == 1)
  clear zi;
  clear Usqri;
```

```
zi=min(z):.01*1e-9:max(z);
 Usqri=1e9*interp1(z,Usqr,zi,'linear');

alloylower = zeros(1,size(zi,2));
for index=1:size(zi,2)
    if(zi(1,index) < 190e-9)
        alloylower(index) = 1;
    else
        alloylower(index) = 0;
    end
end

alloyupper = zeros(1,size(zi,2));
for index=1:size(zi,2)
    if(zi(1,index) > 210e-9)
        alloyupper(index) = 1;
    else
        alloyupper(index) = 0;
    end
end

alloy = zeros(1,size(zi,2));
for index=1:size(zi,2)
    if((zi(1,index) < 190e-9) || (zi(1,index) > 210e-9))
        alloy(index) = 1;
    else
        alloy(index) = 0;
    end
end
```

```
  Int_lower = (zi(2)-zi(1))*trapz(alloylower.*(Usqri.^2));
  Int_upper = (zi(2)-zi(1))*trapz(alloyupper.*(Usqri.^2));
  Int_both = (zi(2)-zi(1))*trapz((alloylower+alloyupper)
             .*(Usqri.^2));
  Int_sum = Int_lower + Int_upper;
  Alloy_Rate = (4*m*Omega_not^2/(a^3*hbar^3))*delta_V^2*x*(1-x)
               *(zi(2)-zi(1))*trapz(alloy.*(Usqri.^2));
  mu_alloy = 1e4*chg/(Alloy_Rate*m);
end

%IR Scattering calculation

if(IR_FNTSQWL == 1)
clear zi;
zi = -100e-9:.5e-9:100e-9; %new grid - QW from -L/2 to L/2
clear theta;
theta=0:pi/size(zi,2):(2*pi-pi/size(zi,2));
q = 2*kf*sin(theta/2);

if(Holes == 1)
    V = .35*(1.087*x+.438*x^2)*1.6e-19;
%Use band offset according to Winkler's book
else
    V = .65*(1.087*x+.438*x^2)*1.6e-19;
end

alpha = sqrt(m*L^2/(2*hbar^2));
```

```
f = @(T)sqrt((V-T)/T) - tan(alpha*sqrt(T));
Energy = fzero(f,2.5e-22);
%for holes in 20nm well start looking around .0017*1.6e-19
%for electrons look around 1.6e-21
figure(2);
T = 0:.0001*1.6e-19:5e-22;
scatter(T,tan(alpha*sqrt(T)));
hold all;
scatter(T,sqrt((V-T)./T));
hold off;

wv_vctr1 = sqrt(2*m*(V-Energy)/hbar^2);
wv_vctr2 = sqrt(2*m*Energy/hbar^2);
H = 1/sqrt(L/2+1/wv_vctr1);
G = exp(wv_vctr1*L/2)*cos(wv_vctr2*L/2)/sqrt(L/2 + 1/wv_vctr1);

psi = zeros(1,size(zi,2));
for index = 1:size(zi,2)
    if((zi(1,index) < -L/2) || (zi(1,index) > L/2))
        psi(index) = G*exp(-wv_vctr1*abs(zi(index)));
    else
        psi(index) = H*cos(wv_vctr2*zi(index));
    end
end

area_fntsqrwll = (zi(2)-zi(1))*trapz(psi.^2);

clear z_prime_integrand;
clear g_s;
```

```
 z_prime_integrand=zeros(1,size(zi,2));
 g_s = zeros(1,size(theta,2));

  for index = 1:size(theta,2)
    for index2 = 1:size(zi,2)
             z_prime_integrand(index2) = (psi(index2)^2)
                                        *(zi(2)-zi(1))
       *trapz((psi.^2).*exp(-2*kf*sin(theta(1,index)/2)
       *abs(zi-zi(1,index2))));
%Note previous equation is split to two lines to fit dissertation
%format
    end
      g_s(index) = (zi(2)-zi(1))*trapz(z_prime_integrand);
  end
% g_s = 1;
F = (1/(2*pi))*(theta(2)-theta(1))*trapz(((q./(q+g_s*q_not)).^2)
    .*exp(-Lambda^2*q.^2/4).*(1-cos(theta)));
%Note previous equation is split to two lines to fit dissertation
%format



IR_rate = 4*pi*m*Energy^2*Delta^2*Lambda^2*F/(hbar^3
          *(sqrt(2*hbar^2/(m*(V-Energy)))+L)^2);
%Note previous equation is split to two lines to fit dissertation
%format
mu_IR = 1e4*chg/(IR_rate*m);
end

if(IR_FNGHWRD == 1)
```

```
    clear zi;
    zi = 0:.25e-9:75e-9; %new grid - QW from -L/2 to L/2
    clear theta;
    theta=0:pi/size(zi,2):(pi-pi/size(zi,2));
    q = 2*kf*sin(theta/2);

    b = (33*m*chg^2*N_S/(8*hbar^2*12.9*epsilonnot))^(1/3);

    psi_FH = zeros(1,size(zi,2));
    for index = 1:size(zi,2)
        %if((zi(1,index) < -L/2) || (zi(1,index) > L/2))
        psi_FH(index) = (1/sqrt(2))*b^(1.5)*zi(index)
                        *exp(-b*zi(index)/2);
%Note previous equation is split to two lines to fit dissertation
%format
    %     else
    %         psi(index) = H*cos(wv_vctr2*zi(index));
    %     end
    end

    area_FH = (zi(2)-zi(1))*trapz(psi_FH.^2);

    clear z_prime_integrand;
    clear g_s;
    z_prime_integrand=zeros(1,size(zi,2));
    g_s = zeros(1,size(theta,2));

    for index = 1:size(theta,2)
```

```
        for index2 = 1:size(zi,2)
                  z_prime_integrand(index2) = (psi_FH(index2)^2)
                                         *(zi(2)-zi(1))
                  *trapz((psi_FH.^2).*exp(-2*kf*sin(theta(1,index)/2)
                  *abs(zi-zi(1,index2))));
%Note previous equation is split to two lines to fit dissertation
%format
        end
          g_s(index) = (zi(2)-zi(1))*trapz(z_prime_integrand);
    end
    % g_s = 1;
    J = (theta(2)-theta(1))*trapz(((q./(q+g_s*q_not)).^2)
       .*exp(-Lambda^2*q.^2/4).*(1-cos(theta)));
%Note previous equation is split to two lines to fit dissertation
%format

    IR_FH_rate = (Delta*Lambda*chg^2*N_S/(2*epsilonnot*12.9))^2
                  *m*J/hbar^3;
    mu_IR_FH = 1e4*chg/(IR_FH_rate*m);
end

%**********END OF CALCULATION***********%
```

B. MBE Standard Operating Procedure

An exhaustive (or maybe just exhausting) treatise

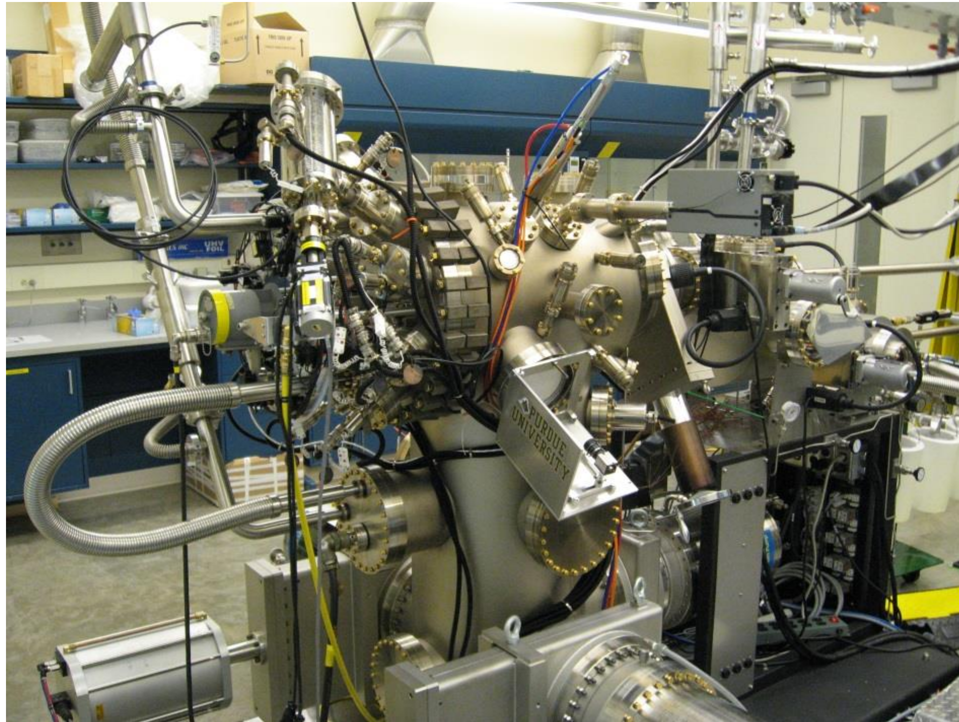

By John Watson

Last updated 2/18/2013

A brief comment

The goal of this document is to describe in detail how to run the GaAs MBE and to solidify the standard procedures so that everyone running the machine is running it consistently (i.e. to try to turn random errors in the growth into systematic errors). The standard procedures I outline here are just what I have been using for the past two years; they are by no means the only (or best) way to do things. It is my hope that this will be a working document that will be updated when improvements to the standard procedures are found. To this end, the editable version of this will be stored in the directory "Manfra MBE\GaAs growth data\Standard Operating Procedures" along with the other standard operating procedures for the growth and low temperature labs. This document will likely be quite long which will make it inconvenient for reference during the growth setup, so to any new users I would suggest copying off/making bullet points of sections that you have a hard time remembering. Ultimately, you need to have all these procedures memorized so that setting up a growth is controlled by muscle memory, so that you do the correct things automatically. As an example, think of how you don't have to consciously think about checking your blind spot before changing lanes in your car; checking that the transfer arm is all the way back before closing the gate valves should be a similar habit you always do without having to think about it. The pictures I've included are the things you should be seeing when you make these safety checks.

The Golden Rules (i.e. lessons we've already learned the hard way)

GOLDEN RULE #1: SEE SOMETHING, SAY SOMETHING. If something appears to be changing over time or is suddenly different, clearly communicate this to everyone else in the group so that we are all aware of how the machine is behaving/changing. For instance, if the liquid nitrogen panel is flashing "Emergency Fill Condition" at 5pm on Friday, don't wait until Monday afternoon to mention it to the boss.

GOLDEN RULE #2: IF SOMETHING IS OPERATING NORMALLY, WRITE IT DOWN. IF SOMETHING IS NOT OPERATING NORMALLY, WRITE IT DOWN AND TELL EVERYONE ELSE. As an example, this would include things like the wafer morphology. Currently the morphology is consistently good (a vast improvement from a year or two ago), so it may seem pointless to check all the wafers with the Nomarski microscope. However, it has been shown[1] that as soon as you stop recording "superfluous" information like this, something will change in the machine that will affect the morphology. Tracking down the source of this degradation in film quality will be much easier if you can reference the database and say, "The morphology started degrading as soon as we started using wafers from this new ingot." It will be much harder to fix problems if you look through the database and say, "Well, the morphology was good the last time we checked it 6 months ago, but since then we have used new wafers, reloaded source material, and changed our outgassing procedure."

GOLDEN RULE #3: FOLLOW THIS STANDARD OPERATING PROCEDURE, EXACTLY, EVERY TIME, EVEN IF "IT SHOULDN'T MAKE A DIFFERENCE". IF THINGS NEED TO BE CHANGED, DISCUSS THIS WITH THE GROUP AND UPDATE THIS S.O.P. IF A CHANGE IS AGREED UPON. This rule could also be titled "There is no such thing as an absolute measurement in MBE." Temperatures, pressures, even growth growth rates (I suspect to some degree) are relative measurements. Comparing these numbers between different MBEs is essentially meaningless

[1] Murphy, et al.

because of the importance of the exact placement of thermocouples, ion gauges, and beam flux profiles across the RHEED sample, respectively. The idea behind this rule is to keep systematic errors from becoming random errors. This is even more important if two different people are growing wafers for the same project.

GOLDEN RULE #4: DO NO HARM. The MBE is the most complicated and easy to break piece of equipment in the building. If you are ever unsure of what to do, ask someone before proceeding. All of the group member cell phone numbers and lab numbers are posted on the wall, so you should always be able to contact someone who will know what to do.

Growth Setup

B.1 Machine Check-Out

Let's assume you're getting ready to start a growth, coming in bright-eyed and bushy-tailed in the morning. The first thing to do is check that the machine is running normally. Assuming there was one growth the previous day which finished mid-afternoon, the reactor pressure should be in the low 10^{-10} Torr range, the buffer chamber should be in the low 10^{-11} Torr range, and the load-lock (LL) will vary depending on when it was last opened. If it was opened last night, and the LL outgassing recipe was run, it will probably be $\sim 1 \times 10^{-8}$ Torr. The cells should all be at their idle temperatures, the viewport shutters closed, the LN2 phase separator should be $\sim$ 75%, the gettering furnace should be off, the valves on the Ar lines should be closed, and all the pressure-relief holes in the sorption pumps should be plugged with their corks.

Next, check that the RHEED block is securely held in the CAR. This is not as big an issue since we cleaned the Ga build-up out of the transfer mechanisms, but it is still possible to drop the block if the pins in the block are not fully locked in the "V"s in the As shield. While you are running the machine, everything that happens is ultimately your responsibility, so don't just trust that that lazy John guy did a good job getting the block transferred last night before you start moving the CAR around.

Before you forget, start warming up the wafer you will be using. The wafer should have already been outgassed in the buffer chamber for 3.5 hours at 350C and should now be sitting at 100C. Send the heated station setpoint to 210C at 2-3 degrees/min. This will clean the wafer again a little bit (in case anything else condensed on it overnight) and get it warm for transferring. I assume by the time it gets taken off the heater and transferred to the CAR it is pretty close to the CAR temperature $\sim$ 150C.

At this point, write down the purpose of the growth in the notebook. This will later get copied into the other/comment section of the sample database. For this document, I will be explaining what I did during growth 2-14-13.1 (so you can reference the notebook if necessary). The notes I have for 2-14-13.1 are "Repeat uniformly doped SHJ 1-26-13.1 with two Ga cells to get faster growth rate up until top barrier to try to improve mobility." The important thing to note here is I state what I am growing (uniformly doped SHJ), why I am growing it (to try to improve the mobility of 1-26-13.1), and I list the growth name from which I am iterating (1-26-13.1). Including the name of the previous growth is important because it makes it much easier to go back through many months' worth of growths and summarize all the work done on a given project. Also include the time you do each step in the margin of the notebook so you can come back later and compare Epitrend data with what you did (this is useful in the case of failures/emergencies to understand what actions caused the problem).

Before you do anything else you should also get the recipe more or less written out in the Excel recipe time calculator so you at least know what growth rates you should be shooting for. Finally, make sure the most important switch is in the "ON" position.

Figure B.1. The most important switch.

B.2 Beam Fluxes

The next thing to do is to heat up the sources you will use. Look through the notebook and use the temperatures that were last used to give the growth rate you want. If you plan to take beam flux measurements, do not open the As valve yet. The purpose of taking the beam flux is two-fold. First, it gives us a record over time of how the flux is changing which will hopefully give us some warning of when the source is running out. Second, it allows us to correct for the shutter transients. Since our shutters are oriented perpendicular to the crucible opening, the shutter reflects a lot of heat back into the cell. The thermocouple, however, is in the back of the cell at the bottom of the crucible. When the shutter opens, the top part of the crucible cools off because the shutter is no longer reflecting heat back in, but the thermocouple's temperature hasn't really changed. This causes the flux (and thus the growth rate) to drop a few percent over the course of about 5 minutes. The problem is that when we measure RHEED oscillations, we are measuring the average growth rate over the first $\sim$ 30-45 seconds the shutter is open. When we are growing, however, we are doing most of the deposition when the shutter has been open for a long time, so there

is less flux than there was during the RHEED measurement. The crude work-around is to just set the RHEED-measured growth rate a few percent higher than what you want and then assume the growth rate during the bulk of the growth is what you used to calculate the time each shutter should be open. In other words, if you want 1 ML/s growth rate and the flux droped 4.5% over the course of 5 minutes, adjust the cell temperature so that the RHEED growth rate is 1.045 ML/s. Then the growth rate should be 1.0 ML/s 5 minutes after the shutter is opened.

If you are taking beam fluxes, the “Take Beam Fluxes” recipe (in the “Frequently Used Recipes” folder) is convenient to use. You will first have to “resume” the “Outgas Cracker Recipe” which has been paused since the cracker cooled down. Resuming the recipe will ramp the current through the doping filaments to zero. Be sure to wait until the Stepper shows the “end of run” message before loading the beam flux recipe. While the dopant sources are ramping down, edit the beam flux recipe. I typically keep the “layers” that I'm not using in the “unused recipe step” drop down menu so that I can just drag and drop the layers into the recipe sequence if I need to change which shutters are being opened. Set the initial pause so that the Al cell will be hot for 10 minutes before its shutter is opened (always take the Al flux first since it heats up/stabilizes before the Ga cells). The standard sequence for measuring a source's flux is to open the shutter for 5 minutes and then wait at least 5 minutes before taking its flux the second time (this allows the material to warm back up due to having the shutter closed). Save the recipe, load it in the stepper, and start the recipe. The beam fluxes need to be taken with the CAR index at 180, but the CAR is probably still at 179 from when the RHEED block got transferred onto it last night. First rotate the index to 150, then to 180. The CAR can't accurately do small changes in position, hence why you have to first send it to 150. Open the main shutter so the beam flux gauge is looking at the sources.

When the shutter opens, look at the source. You should do this every time you grow so that you know what the source (and source material) normally look like. Write down the state of the cell. If it is unchanged from the last growth, just write

"Ga1 looks normal". If it looks like a lot of Al has crept up in Al2, write some thing like "Large puddle of Al at opening of crucible in Al2, does not appear to be in danger of spilling out yet." This is a pain but worth doing (see golden rule #2).

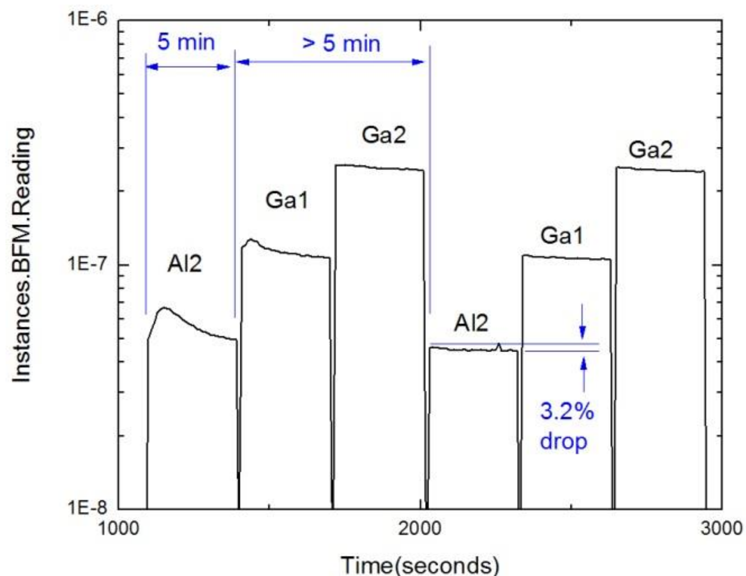

Figure B.2. Typical beam equivalent pressures of each source during first growth campaign. Note that the first set of flux measurements has a larger percentage drop while the shutter is open, likely due to As build-up around the source.

The flux profile will look significantly different between the first and second times the shutter opens. This is probably due to As buildup around the source that is mostly gone after the first time you open the shutter - see figure B.2. Beam fluxes should not be taken every day. Taking the fluxes takes around 45-60 minutes which wastes a considerable amount of material over the years, and it also coats the beam flux gauge with a lot of metal (we already had problems with the gauge partially dying, presumably from an electrical connection being shorted by Ga/Al). Beam fluxes should be taken once a week or if a cell is being used at a significantly different growth rate than when the flux was taken (e.g. if you use a Ga cell at 0.5 ML/s and its flux was last taken at 1.0 ML/s). If you do not take beam fluxes, just use the

percentage drops from the previous time the fluxes were taken to calculate the target RHEED rate and note in the summary sheet that you assumed a given flux drop.

B.3 RHEED Warm-up

Once you've finished taking the fluxes, start warming up the RHEED wafer. Since the thermocouple on the substrate manipulator died an untimely death, we have to run the heater in constant power mode. Temperature ramps are thus accomplished by using a code snippet. In the recipe editor, open "Ramp Substrate Power" in the "Frequently Used Recipes" folder. The first line is the target power output. Set this to 80 (this will get the wafer to ∼600C). The ramp time should be 900 (this is in units of seconds). Save the recipe, load it in the stepper, and start it. A word of caution is necessary at this point. Code snippets allow you to do things that can break the machine. It is therefore critically important that you pay attention to what you are changing in code snippets and that when you run a code snippet you keep an eye on what the machine is doing. It is easy to write a code snippet that will take a heater filament from 0% to 100% power in zero time. This could of course break the filament and ruin everyone's day (or several months). So whenever you start the "Ramp Substrate Power" recipe, check the actual output of the CAR heater. The idle level (∼150C) is 16%. So if you are ramping to 80%,the first output level should be something like 16.7%. If you start the recipe and check the output and see that it is suddenly 100%, there will be issues.

Start rotating the CAR index to growth (0) and open the As valve to the value used for the previous growth. Set the As valve position in the status menu of Molly. Do not change the speed (ramp rate). When the valve gets all the way open make sure the valve driver is not moving. If it is bouncing back and forth (within 0.1 mil of the target), set the valve controller to "manual", manually open or close the valve a tiny amount (∼ 0.1 mil), and set the controller back to "remote". The controller must be in "remote" control for the end-of-recipe script to run correctly.

Figure B.3. Check the power output % as soon as you start ramping the CAR heater.

Make sure all the Ga/Al shutters are closed and that the main shutter is up (Overview screen in Molly) so you don't damage the RHEED wafer. Once the CAR index reaches 0, start the wafer rotating so that it will heat uniformly.

Once the CAR index reaches 0, make sure the scribe marks on the index rotary feedthrough (ROMO) are lined up (see figure B.6). It is possible for the stepper motor to slip on the ROMO. If this happens, the controller will think it is in the growth position, but the wafer will not be completely vertical. To help prevent this, every time you check the position of the scribe marks, make sure the ROMO thumb-screw is backed out all the way (i.e. all the way counter-clockwise) so that it can't add any resistance to the stepper motor-ROMO coupling

If the stepper motor does slip on the ROMO, make sure the controller thinks it is at the 0 position, loosen the screws on the knurled connector and the gold screw, align the ROMO so the scribe marks are lined up, and tighten the screws back down. The screws need to be tight to keep things from slipping, but be careful to not strip any of the screws. Before doing anything else, double check:

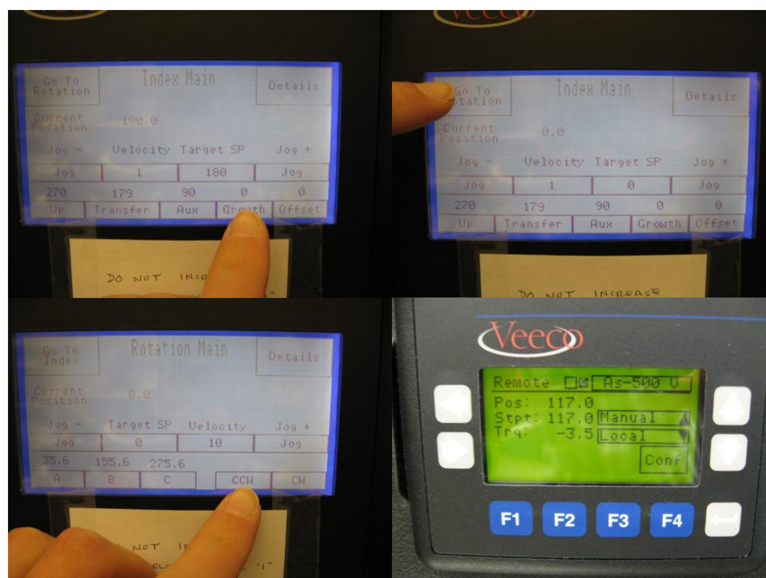

Figure B.4. View of CAR control screen and As valve driver controller.

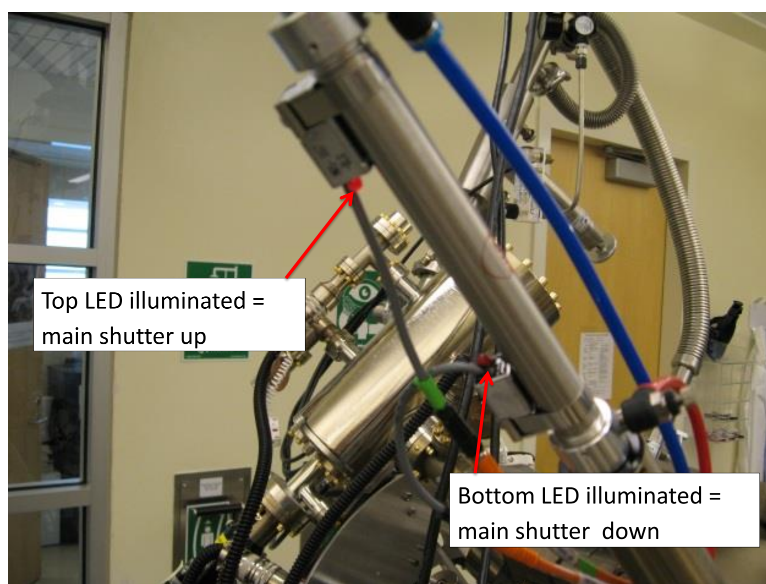

Figure B.5. Main shutter with LED position indicator lights.

1. As valve open

2. Main shutter up

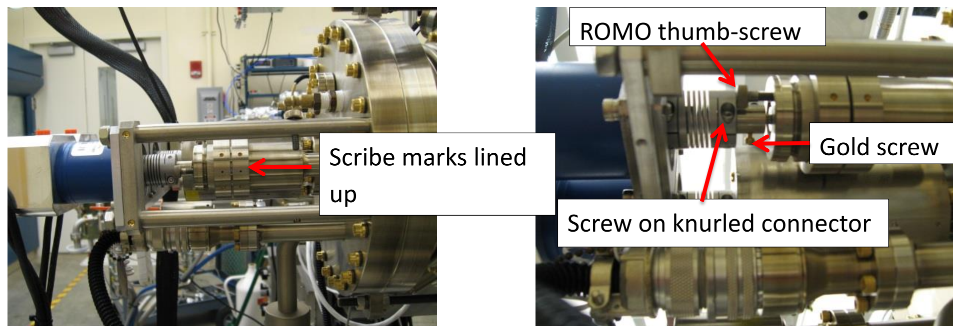

Figure B.6. View of the scribe marks on the CAR index position.

3. CAR in growth position
4. Wafer rotating 10 RPM CCW
5. CAR heater ramping steadily to 80%
6. Viewport shutters all closed

Now would be a good time to enter the flux measurements in the recipe spreadsheet and print it off so you have easy access to your target growth rates (the spreadsheet will calculate this automatically if you input the starting and ending fluxes for each source you will be using).

B.4 RHEED Measurements

Next, hook up the RHEED camera and load the Labview program ("RHEED Specular Measurement v1.1" located in \My Documents\Labview VIs). As an aside, another useful program (for taking pictures of RHEED) is "Grab and Annotate Pic" in the same folder. Once the Labview code is loading (it takes a while for the computer to settle down after it starts loading), start warming up the RHEED supply. First, turn on the beam blanking on the controller shown in figure B.7 (the red light will turn on when the beam is blanked). This is a deflection voltage inside the electron

Figure B.7. RHEED controller and power supply.

Figure B.8. Comparison of the 2× reconstruction (left) and 4× reconstruction (right) as seen on the RHEED screen.

gun that will keep the electron beam from exiting the RHEED gun. Then go back and forth between increasing the voltage and current in steps of 1.0 kV and 0.1 A, respectively. Stop at 15.0 kV and 1.4 A.

Once the CAR heater has reached 80% output (should be 600C on the pyrometer at steady state) you are ready to start the RHEED measurement. For the sake of consistency, always do the growth rate measurements on the "2×" reconstruction. The first time after I load a new RHEED wafer, I always set the 2× reconstruction to one of the three presets on the CAR rotation screen. However, due to the 120 degree

symmetry of the Ta block, you will never know in advance whether preset A, B, or C will line up the RHEED wafer so that you will see the 2× reconstruction. The easiest thing to do is look for the 4× reconstruction because it is very distinct (see figure B.8). There are many rotation positions that will look like the 2× reconstruction, so the only way to be sure is to first find the 4× reconstruction. Do this by sending the rotation position to 90 degrees above or below one of the presets. Eventually, you will find the 4× reconstruction 90 degrees off of A, B, or C. Once you have found the 4× reconstruction, send the rotation to the corresponding preset position (i.e. the 2× reconstruction).

Sometimes the vibrations in the machine will cause the RHEED pattern to shake. If this happens, change the rotation position (say 180 degrees) and then go back to the 2× reconstruction. You may have to do this a few times before the pattern stays still. If the diffraction pattern is dancing around, you are guaranteed to have really bad looking oscillations (and hence an uncertain growth rate). The electron beam should be kept blanked whenever you are not actually measuring the diffraction pattern. This keeps the wafer from getting charged up which will cause the diffraction pattern to drift and make it impossible to measure the growth rate. This is also why we now use the beam current at 1.4A and not 1.45A.

Next, set the shutters you will be using to manual mode so you can control them with the PDA. For each port (i.e. effusion cell), there are 4 boxes: Open, Close, Remote, and Manual. The open/close state in parentheses in the remote and manual boxes lists the state that the shutter will go to when you switch to that control mode. For example, suppose the shutter is currently closed and in remote mode, but you already tried opening it with the PDA. The Remote (Close) button will be depressed, and the Manual (Open) button will be unlatched. If you press Manual (Open), the shutter will be opened and the control will be in manual mode. The Open and Close boxes do essentially the same thing as the PDA when the shutter is in manual mode.

Now start the Labview code. Once you start the code you should be able to see the diffraction pattern in real time on the computer screen. If necessary, adjust the

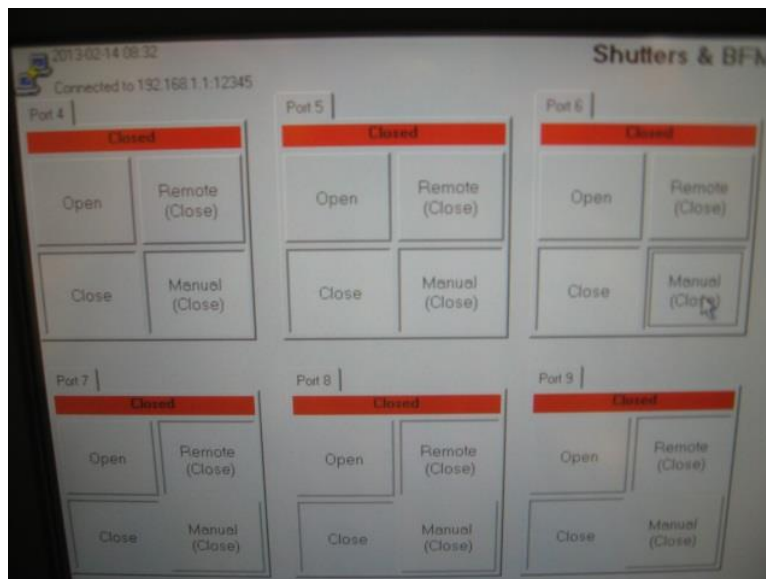

Figure B.9. Shutter control screen on the e-rack touch-screen computer.

X knob on the RHEED controller. DO NOT ADJUST ANY OTHER KNOBS ON THE RHEED CONTROLLER. The grid, focus, and rocking knobs all take a lot of time to tune. The Y knob should not be adjusted because this will put the electron beam on a different part of the wafer which will result in a different growth rate being measured. This is one of those instances in which we need to keep systematic errors from becoming random. The growth rates measured from RHEED are most certainly (slightly) wrong, and probably wrong by different factors for different cells since the flux profile across the electron beam is different for each cell. For growing high mobility 2DEGs (as opposed to, say, quantum cascade lasers), this is not the end of the world, but we do need things to stay consistent. If our 30 nm quantum wells are actually 29 nm when grown with Ga2 that is fine, but they need to stay 29 nm every time we grow them and not vary from 29 to 31 to 28 to 30 nm, etc.

For the first set of oscillations of the day you will need to define a new rectangle within which the program will average the pixel intensity. Once you have the diffraction pattern looking good, click “Start Data Acquisition” and press “Define

New Rectangle" in the pop-up window. Make a square over the most circular portion of the specular beam (the most circular portion is the part that will consistently exist over the course of the oscillations). When prompted to save the data, use the format of MM-DD-YY-CellAs-oscillation#. In other words, the file name should look something like 02-14-13-Ga1As-1. Once the file is saved the program will start taking data. Open the shutter and take 7 full periods of oscillations for GaAs or 4 full periods for AlAs if you are growing GaAs $\sim$ 1 ML/s and AlAs $\sim$ 0.33 ML/s. If you deviate significantly from this in growth rate, keep the shutter open for a comparable amount of time. Again, the number of oscillations is part of keeping systematic from becoming random errors. For instance, increasing the number of Ga2 oscillations beyond 7 will result in a faster measured growth rate by a couple percent, so just stick with 7.

Always measure GaAs oscillations first. If you measure GaAs oscillations right after AlAs oscillations this seems to result in a faster growth rate (again just by a percent or two). For a given cell, measure the growth rate twice. If they are the same within a percent or two, good. If not, measure a third time to see which one of the first two measurements was a fluke and then adjust the cell temperature accordingly (if needed). This seems to be especially important for the Al cells since the Al can suddenly move around and change the growth rate significantly. This always seems to happen when you have already been struggling to get the growth rate set for a long time and finally have the growth rate where you want it. Then you measure the growth rate a second time (i.e. what you think will be the last time), and then you end up having to change the cell temperature 3 or 4 degrees to get the growth rate back where you want it.

Once you have finished the GaAs oscillations for the current iteration and changed its temperature (if necessary), measure the AlAs growth rate. After each set of AlAs oscillations, open the Ga shutter for about 5 seconds to smooth out the RHEED sample. Once you are done with this iteration of oscillations, make sure the Ga shutter is open 5-10 seconds so that the surface will be ready for GaAs oscillations in the next iteration of measurements. Once you adjust the cell temperature, wait 10

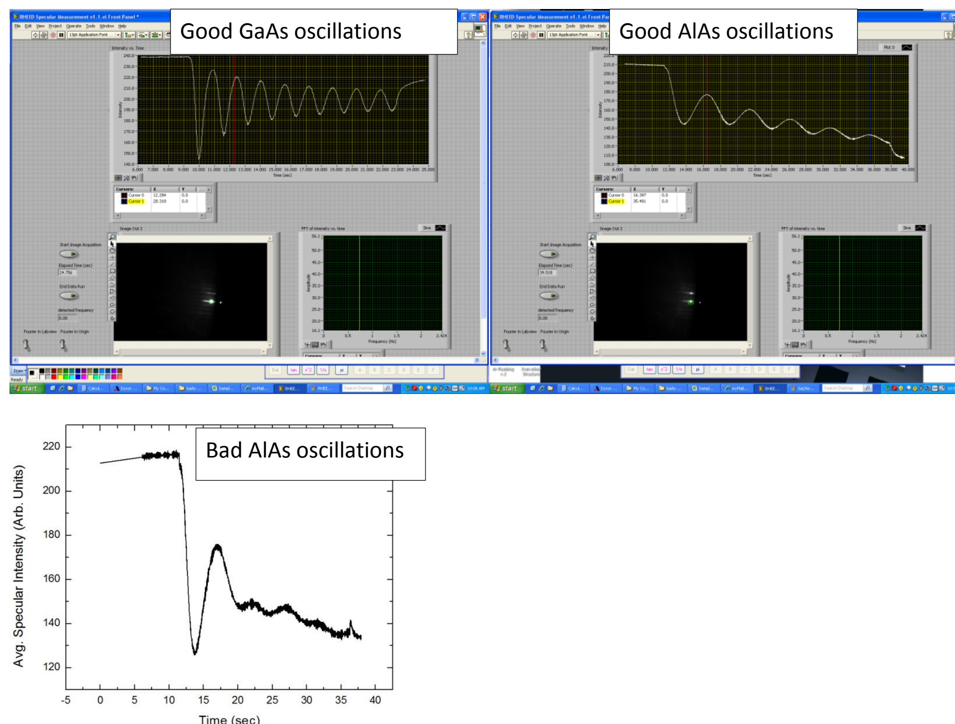

Figure B.10. Examples of good and bad GaAs and AlAs RHEED oscillations.

minutes before using that cell for measurements again. Also, if you measured a bunch of AlAs oscillations (followed by the GaAs smoothing), make sure the Ga shutter is closed at least 5 minutes before you measure its growth rate again since it takes ~ 5 minutes for the shutter transient to build back up.

Usually you will only be adjusting the Ga cells ~ 1 degree and the Al cells 1 to 3 degrees from the temperatures that previously gave the desired growth rate. If you have to adjust the temperatures more than this, you should make sure you are making adjustments based on good oscillations. Figure B.10 shows what the oscillations should look like. If they look really bad (like the bottom plot), do not expect the growth rate you extract to be accurate.

If the oscillations are looking really bad, there are a few things you can do. First, you can try moving the CAR rotation position to a few different values and then back

to the 2× reconstruction. I think this sometimes just allows the Ta block to shift around in the substrate manipulator a little bit. If the block is not perfectly vertical or if the RHEED wafer is not perfectly flat on the block, the oscillations will look kind of funny and the measured rate will be way off. If that doesn't fix things, the next thing is to try smoothing out the wafer. Heat it up to 635C, send the CAR to continuous rotation (10 RPM CCW), and grow GaAs (once the wafer has heated up to 635C) for 5-15 minutes, and then cool the wafer back to 600C and try again. Of course, once you stop growing GaAs, make sure to let the shutter stay closed for > 5 minutes before measuring the GaAs growth rate again. If none of this works, it may be time for a new RHEED wafer; talk to others in the group before loading a new wafer, though. Of course, if the oscillations are bad, make a note of which oscillations are bad. This will allow troubleshooting in the future. If the electron density comes out way off, your notes can serve as an indication that the reason for the bad density was a bad growth rate.

B.5 Loading New RHEED Wafers

When a new RHEED wafer is loaded, you need to give it some TLC before trying to use it for growth rate measurements. For starters, you need to get the oxide all the way off to keep the wafer from getting hazy. The last oxide species seem to come off above $\sim$ 610C on the pyrometer. So once the wafer is above 610C (in an As flux of course), let it sit for 10 minutes. Then you will need to smooth the surface out since the oxide steals Ga from the surface, creating small pits. It seems to smooth out the fastest when you use short GaAs growth periods followed by a pause. You can watch the wafer smooth out (i.e. watch the RHEED pattern get streakier) as you are doing this. I usually do this by hand for 5 minutes or so and then load a recipe that does our standard smoothing sequence (10nm GaAs, 20 sec pause, repeat 50 times). You may also need to throw in some AlAs or AlGaAs layers followed by the GaAs and smoothing pause layers to get the oscillations looking nice and symmetric for both

GaAs and AlAs. Again, if the oscillations don't look like the pictures shown above, I would be somewhat skeptical of the growth rates you are extracting. The growth rate is very likely wrong if you see any beating in either the GaAs or AlAs oscillations.

Loading new RHEED wafers doesn't happen very often, though, so at this point let's assume you have all your growth rates set. Next, cool the wafer back down (16% target power, 900 sec ramp time in the ramp substrate power recipe), and be sure to keep the substrate rotating during cooling. Now would be a good time to start writing the recipe in Molly. If you are just updating a previous recipe (e.g. slightly increasing the doping or changing the Al concentration), be sure to "save as" a new file so you dont overwrite the original recipe. Save recipes according to the naming convention MM-DD-YY-growth# (e.g." 02-14-13-1.cmd"). Including leading zeroes in the month and day makes it easier to find old recipes.

One word of caution when writing recipes. The code snippets used to ramp the substrate power for the doping steps and the dopant source ramp snippets require caution when using. As mentioned previously, code snippets allow you to do potentially harmful things to the machine, so be careful you haven't forgotten any semi-colons or added any extra zeroes after target currents, ramp rates, etc. If you ever need to write more advanced codes, I would suggest using the copy of Molly on the RGA computer. This way any infinite loops, etc. will only crash the RGA computer and will have no effect on the MBE. Rather than controlling hardware, you can just have the output of each function be sent to an echo command that will print the value to the screen. Once you are convinced you have made a robust code that will not do something unpredictable to the machine if you enter nonsensical input parameters, just copy and paste the code to the Molly computer and replace the echoes with the commands that actually control the hardware.

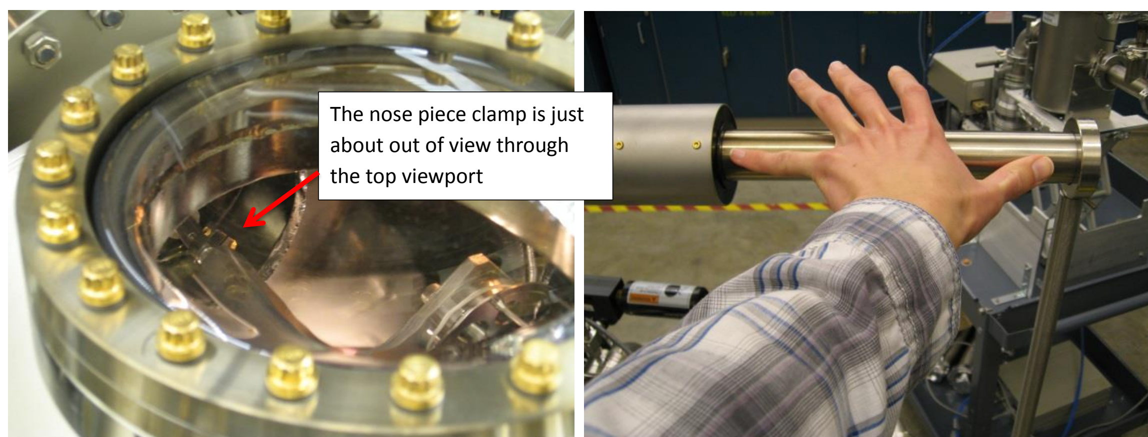

Figure B.11. Slide the nose piece towards the growth chamber 9-10 inches before opening the gate valve.

B.6 Pre-Growth Transfer

Now comes the trickiest part of the whole growth – transferring in the growth chamber. Once the CAR heater has gotten down to 16%, stop the rotation and send the index to 180. Once it gets there, the beam flux gauge will be reading the As flux. The target starting As flux should be 8.5 to 8.7E-6 Torr. The goal really is to get the flux close to 8.5E-6 Torr at the end of the growth. Unfortunately, the As flux tends to drop over the course of the run, sometimes it drops worse than others, so just look back through the notebook at how much it has been dropping recently and set the As valve accordingly. Try not to take too long when you do this because the As will eventually destroy the beam flux gauge filament. Once you have the As flux set, close the main shutter, send the CAR index to 150 and then to 179 (transfer position). Once the index reaches 179, make sure the rotation position is at 0. This will minimize the parallax when you are trying to line the pins on the block up with the slots in the CAR, and it seems to give consistent alignment between the nosepiece and the CAR.

At this point you want to get everything ready so that you can minimize the time the gate valve between the buffer chamber and growth chamber is open. Eventually

Figure B.12. The buffer chamber pressure must be sufficiently low, and the LL gate valve must be shut before the gate valve to the growth chamber can be opened.

the sealing surface in this valve will get coated with enough As that it will no longer seal, so we want to try to prolong that as long as possible. That being said, when the gate valve is open, work efficiently and not frantically. The last thing we need is someone tripping onto the MBE or dropping a block in the growth chamber. First, move the transfer arm about 9-10 inches towards the growth chamber as shown in figure B.11. This lets some of the pressure burst from moving the transfer arm get pumped away in the buffer chamber before the gate valve to the growth chamber is opened. Do not slide the transfer arm in too far; you don't want to whack the gate valve.

Next, start cooling the wafer you will be using for the growth. Send the heated station to 100C at 10C/min. Finally, double check that the following conditions are met so that the gate valve is safe to open: buffer pressure in 10E-11 Torr range, main shutter down, LL gate valve closed, CAR in transfer position.

Now open the gate valve all the way as shown in figure B.13. Then slide the transfer arm in as far as it can go while still leaving room to open the main flap to the growth chamber without hitting the nose piece. Open the main flap and slide

Figure B.13. The main shutter should also be down before the gate valve is opened (left). The gate valve is fully open when the second line on the arm is visible (right).

the transfer arm in all the way to the CAR. Depress the As shield and rotate the nosepiece counterclockwise (from your view) until the pins of the block are locked all the way in the slots on the nosepiece. Sometimes the block will get stuck a little bit when you try to rotate it. If this happens, don't panic. Just rotate the nosepiece back to the starting point (but do not slide the transfer arm back at all) and then try rotating the nosepiece counterclockwise again. The second try almost always will get the block off. Two words of caution are in order here. First, don't press the As shield back too hard. The magnet can slip off of the ferromagnetic chunk in the back of the transfer arm which will then allow the nosepiece to do whatever it wants (i.e. this is a good way to drop a block). Second, once you start to take the block off the CAR, you have to finish the job. It is very tricky to get the As shield far enough back to lock the block back in to the V's in the CAR if you have already started rotating the block off.

Once you have the block in the nosepiece, slide the transfer arm back and close the main flap. Leave the RHEED block on the "HS" station in the buffer. Now increase

the ramp rate of the heated station to 40 degrees/min. This will minimize the power spike that happens when you take the warm block off the heated station. Now take the block for the growth into the growth chamber and load it on the CAR. When you have the block on the CAR, efficiently (not frantically) walk to the other side of the machine and check that the pins on that side are locked in the V's. If the block is fully engaged, slide the transfer arm ALL THE WAY back and close the main flap. As a general rule, whenever the transfer arm is not in active use it should be slid all the way back so it cannot get caught in a gate valve or the buffer carousel. Start to close the gate valve, and double check that the transfer arm is all the way back. The gate valve is closed when the plastic sleeve is up to the top line in the metal and the indicator is in the closed position. Now send the CAR index to 150 then 180, open the main shutter, and check the As flux again. If it is has drifted, adjust the valve accordingly and record its position in the notebook. Follow the same procedure as before to warm up the wafer. Send the CAR to 85% output power in 900 seconds, send the CAR index to 0, make sure the main shutter is up, check the scribe marks on the CAR index ROMO when the CAR index reaches 0, and start substrate rotation (10 RPM CCW) when the CAR index gets to zero.

B.7 Wafer Warm-Up

Check the wafer's morphology by examining the RHEED pattern once the wafer is warming up, is in the growth position, and is rotating. Typically it is "faint and hazy" as shown in figure B.15. This is due to the remaining oxides (Ga oxides) that have not yet desorbed. A wafer that has not been outgassed at all should be even hazier and have an amorphous ring in the first Laue zone. The fact that you can see some diffraction streaks means the oxide layer is relatively thin. Record the description of the RHEED pattern along with the time the wafer spent at 350C and 200C in the notebook. Cool down the RHEED at the same rate you warmed it up and remember to turn off the beam blanking. Double check now that all the viewport shutters are

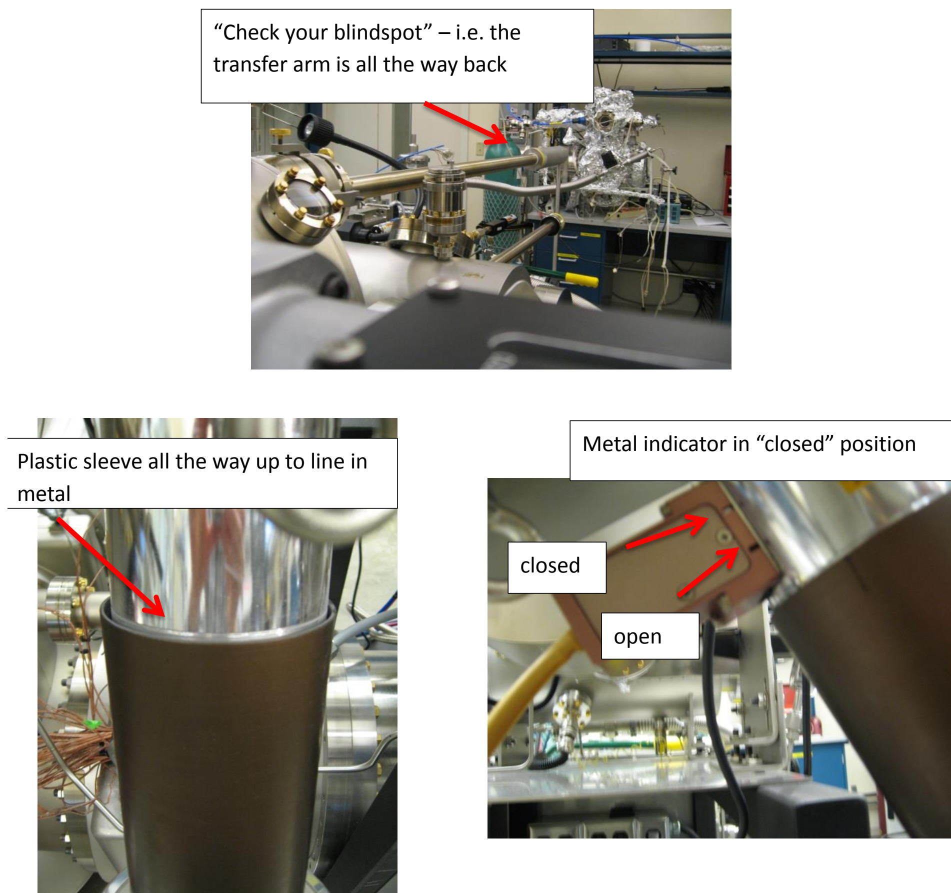

Figure B.14. Double check that the transfer rod is all the way back before closing the gate valve. The gate valve is closed when the plastic sleeve reaches the top mark and the indicator is fully in the "closed" position.

closed, the main shutter is up, and the wafer is rotating. Now would also be a good time to remember to update the wafer position sheet on the side of the electronics rack, put the shutters back in remote mode, and load another wafer onto the heated station (do not start outgassing a new wafer if the outgas will not finish before you transfer the first growth out of the growth chamber).

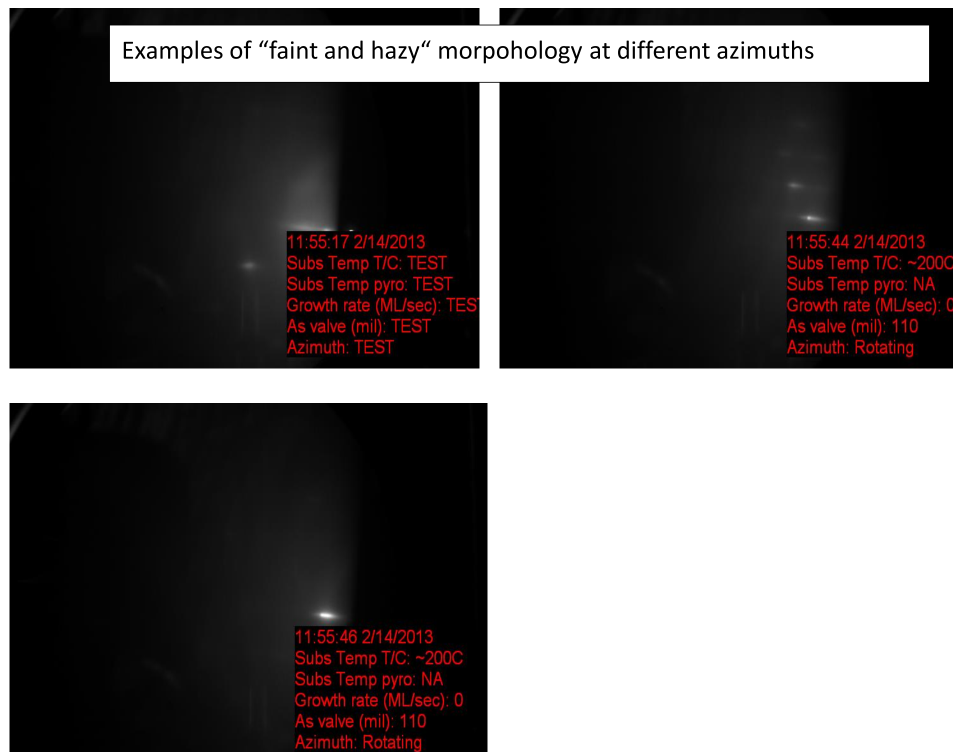

Figure B.15. Examples of "faint and hazy" morphology immediately after the wafer is loaded into the growth chamber.

B.8 Growth Startup

Now your time will be getting a bit tight since you don't want to leave the wafer hot for too long (this will cause cross-hatching morphology). The most important thing to do while the wafer is heating up is to double and triple check that the Molly recipe matches the Excel recipe and that both of them are indeed the structure you intend to grow. Growing the wrong structure wastes a lot of people's time (not to mention source material), and that makes the boss cranky. Life is never better with a cranky boss, so do your part to keep him not cranky. At this point you should double check the following items: the shutter times are correct, the correct shutters are being opened, the dopant source is being sent to the correct current, there is a "layer" used to close the Ga/Al shutters prior to temperature ramps for delta doping

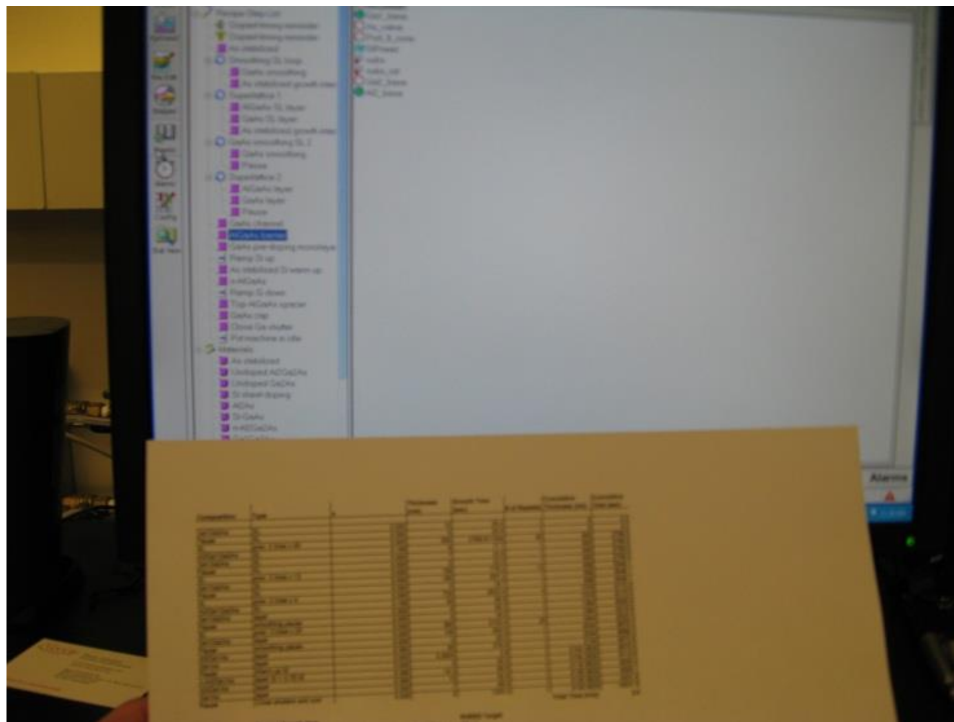

Figure B.16. Don't forget to double check the recipe in Molly line-by-line to make sure it matches what you intend to grow in the Excel sheet.

steps, the correct end of recipe script is being called (e.g. end of day vs. between run shut down), the "sleep cold"/"sleep hot" times in the dopant ramp codes are correct for the block you are using, and the total time matches between the Excel sheet and the Molly recipe.

Watch how the pyrometer responds during the warm-up. It should smoothly increase as the power is ramped up. If it does, write this down in the notebook. If the pyrometer increases rapidly, then slowly drops a degree or two, etc. write this down. If the pyrometer is oscillating during the warm-up it either means that there is an inconsistent oxide on the wafer or that the Ga wetting on the back of the wafer is not good. Either way, this will generally lead to rough morphology. Once the pyrometer reads above 610C, start a timer. The growth should start 10 minutes after the wafer reaches 610C. Do the final adjustments to get the temperature $\sim$ 633C;

the wafer will heat up $\sim$ 2C once the sources start shining on it. If the recipe has temperature ramp steps in it, make sure to update the "hot percent" values to match the power output needed to get the wafer to growth temperature.

Once the recipe is finalized, load it into the stepper (be sure you are loading the recipe from the correct year, e.g., 2-14-13.1 and not 2-14-12.1). Once the 10 minutes is up, start the recipe and press resume through the dopant warning message (i.e. the only safeguard against ramping the dopant sources too fast). Write down the time you started the recipe, and as a safeguard, write down the exact name of the recipe you see running in the stepper. We can only go back later and double check the recipe saved in Molly, not the recipe that you actually ran, so this is important if you want to be able to say, "No boss, I did everything correct with the recipe, it is just the design or characterization that gave us a bad answer."

B.9 Wrap-Up of Growth Setup and Checks During Growth

Write down the pyrometer temperature, power output (in Watts and percent), the layer during which you measured the temperature, the reactor pressure, and the beam flux pressure. Double check that all the viewport shutters are closed, the gate valve is closed, the correct source shutters are opening and closing, and that there is another block sitting on the heated station so it doesn't have to work too hard to stay at 100C. This probably also will be the start of your lunch break. If you time lunch right, you can usually get back and record the temperature at the end of the GaAs smoothing superlattice which is the most accurate representation of the ultimate growth temperature. Once the GaAs/AlGaAs superlattice starts, the pyrometer-measured temperature can drift around due to optical effects. Even if you don't get back before the superlattice starts, you should still record the pressures/layer at each time you check on the machine. It is also a good idea to check in on the machine during the doping steps to make sure the I-V characteristics of the doping filament are normal and to record the temperature of the wafer during doping (this has been

changing over time due to a bug in Molly's ramping the CAR output down slower than it did originally).

B.10 End of Growth

When the growth ends, the recipe will call either the end of day shutdown script or the between run shutdown script. The between run shutdown will cool down the substrate, stop the rotation, send the index to 180, wait for $\sim$ 10 seconds (to record the As beam flux in Epitrend), close the main shutter, rotate the index to 150, and rotate the index back to 179 (transfer position). The end of day script does all this plus it will cool the sources to their idle setpoints and close down the As valve. Be careful that the recipe has actually finished all the steps; Molly will usually say the recipe has ended before this script finishes and then start back up at what looks like the start of the growth. Ignore all this and do not try to start/stop/resume anything. Molly is just a very buggy program, but so far it has not done anything that can harm the hardware.

When the shutdown script really is finished (i.e. the CAR index is at 179), repeat the transfer procedure as before. If you are starting a second growth, you should have gotten that wafer up to 200C during the first growth (in addition to the 3.5 hour 350C outgas that would have been done earlier). Move the first growth from the growth chamber to the buffer (remember which carousel station you transfer it to) and load the new wafer into the growth chamber. If you are not doing a second growth, load the RHEED block onto the CAR. There should always be a block on the CAR. The block protects the hot filaments from getting eaten up by the As. Close the gate valve (remember to check the transfer arm). Check the ending As beam flux in Epitrend. Remember to multiply the pressure in Epitrend by 0.9507 to get the pressure that matches the front panel display of the ion gauge controller.

If you are doing another growth, repeat the previous procedure. If you are done for the day, run the "Outgas Cracker" recipe in the "Frequently Used Recipes" folder.

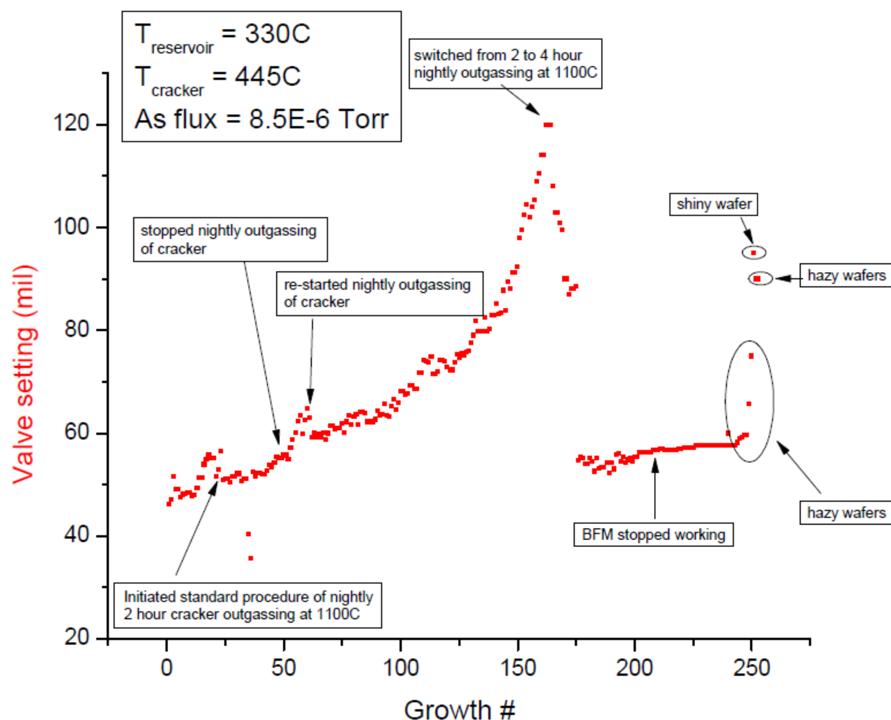

Figure B.17. As valve setting during part of the first growth campaign.

This will ramp the dopant filaments up to 1A (their idle state) and heat the cracking zone up. This clears out As that has built up in the conductance tube of the cracker during the growth. This is really important. The As valve setting necessary to get sufficient As flux will increase quite rapidly if the outgassing routine is not performed. Figure B.17 is a plot showing the As valve setting as a function of growth number before and after we increased the time the cracker is hot during its outgassing routine.

If you are done with everything for the day, go through the "End of day checklist" sheet to make sure everything is safe to leave. Be sure to write down the phase separator level from the LN2 controller in the galley; this needs to be tracked to see if the valve that feeds the phase separator is continuing to degrade over time. If there is an "Emergency Fill Condition" alarm, be sure to tell everyone immediately.

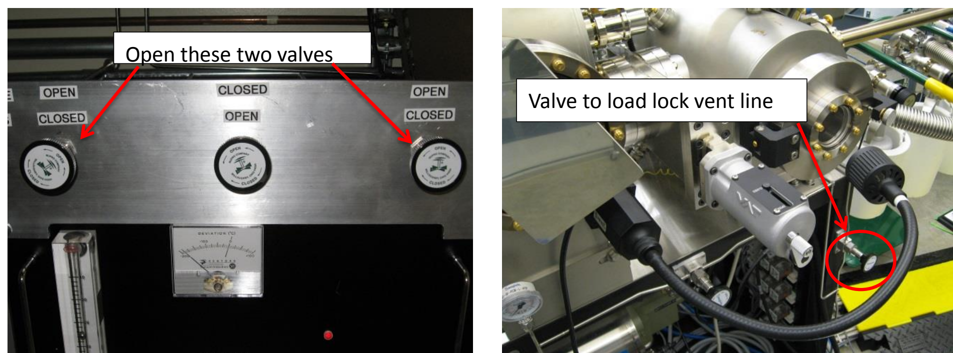

Figure B.18. Valves to open when warming up the gettering furnace.

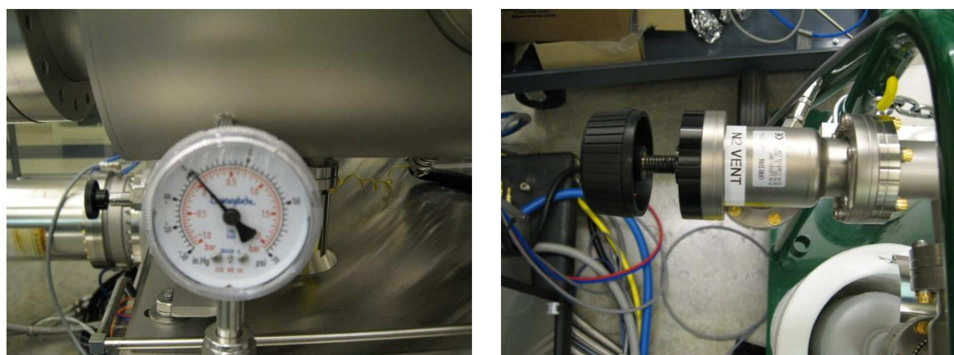

Figure B.19. LL vent line mechanical gauge and “N2 Vent” valve.

B.11 Unloading Wafers

First make sure all the wafers you want to unload are in the load lock. Next get the gettering furnace warming up. Flip the power switch on and open the two valves to operate it in furnace mode. The gettering furnace will overheat if there is no gas flow through it when it is on, so also open the valve to the load lock vent line called out in figure B.18. Watch the pressure on the load lock vent line mechanical gauge shown in figure B.19, and open the “N2 Vent” valve on the pump cart when the gauge reaches atmospheric pressure. Put the Ga beaker on the hot plate with a gloved hand to start melting the Ga (this takes a while). Wipe down the load lock ROMO and

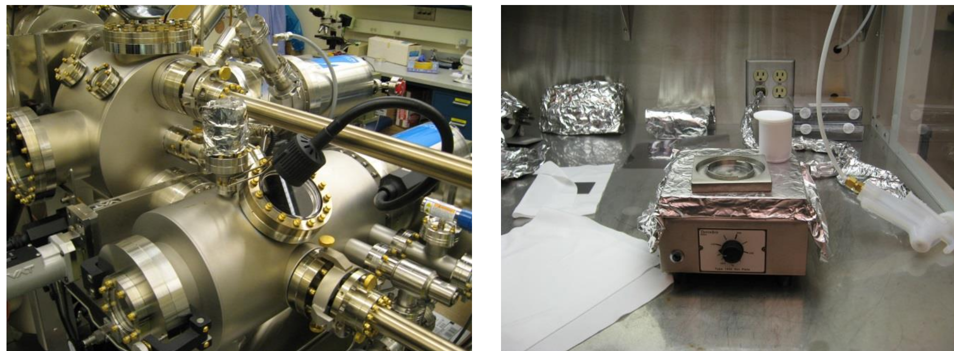

Figure B.20. Prepare the LL ROMO and Ga beaker for re-loading the LL.

door with some methanol and wrap a clean piece of foil around the ROMO as shown in figure B.20 so that the surfaces you touch during the unloading will be as clean as possible.

Get out as many new substrates as you will need, take them out of the wrappers, and leave them in the clean hood. Get more clean wipes in the hood if needed. Use only the TX1009 wipes (i.e. the green bag). Do not use the TX1109 wipes in the blue bag since they leave fuzz on the tweezers and sharp edges. Send the load lock heater lamp ("Intro Bake" control zone in Molly Status screen) down to 10C at 5C/min. If you don't send it down below room temperature it will send 100% power to the lamp when you open the gate valve back up, and this probably shortens the life of the lamp. Once this is ramping down, make sure the gate valve to the buffer chamber is closed and turn off the load lock ion gauge (see figure B.21). Next, close the gate valve to the load lock cryo pump. Keep an eye on the buffer ion gauge to make sure it doesn't spike up when you close the load lock cryo gate valve (i.e. to make sure the gate valve between the load lock and buffer really is sealed).

Now start filling the nitrogen dewars for the sorption pumps. BEFORE filling with LN2, check that the rubber pressure relief corks are all plugged in so you don't try to pump the whole room into the sorb. Forgetting this step will likely earn you a good deal of ribbing from the boss. It usually takes about four trips to get the

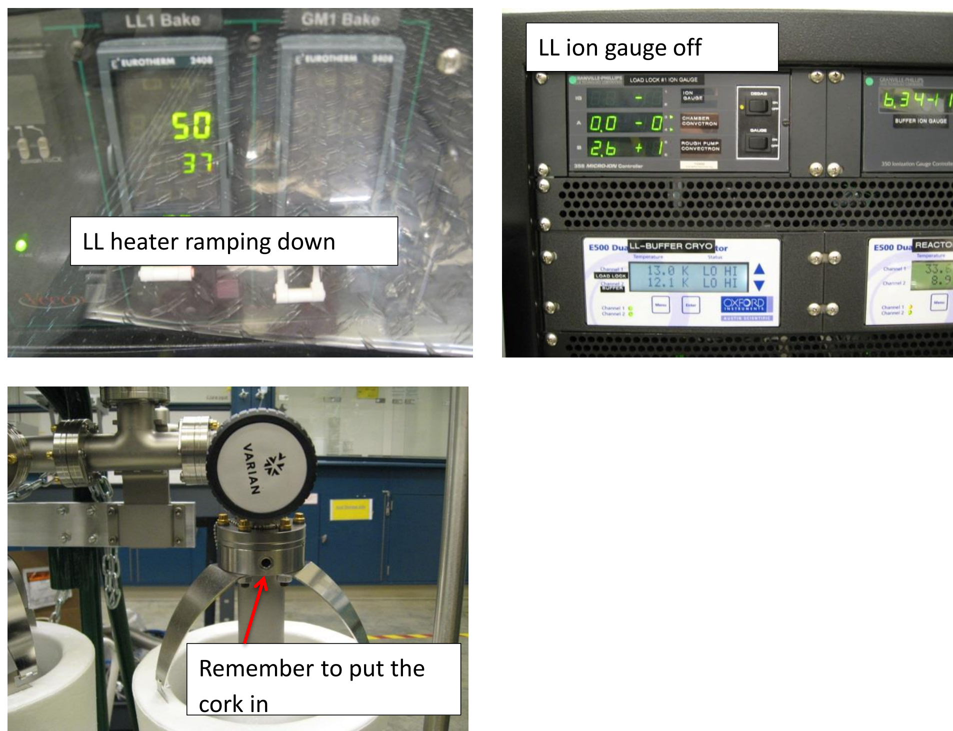

Figure B.21. Ramp down the LL heater, turn off the LL ion gauge, and check that the sorption pumps are plugged before cooling down the sorption pumps.

dewars all the way full and cold. Try not to splash a bunch of LN2 on the electrical or waters lines on the floor next to the sorbs; plastic and rubber generally don't like cold shocks very much. Once the dewars are full and the gettering furnace is warm (the orange light will be flashing rather than just staying on continuously) you are ready to unload. Put on the "clean jacket" and a pair of the purple gloves and blow yourself off with one of the nitrogen guns. Then double check the following items:

- LL ion gauge off > 5 min
- LL heater off > 5 min
- Gate valves to LL cryo and to buffer closed
- Gettering furnace hot

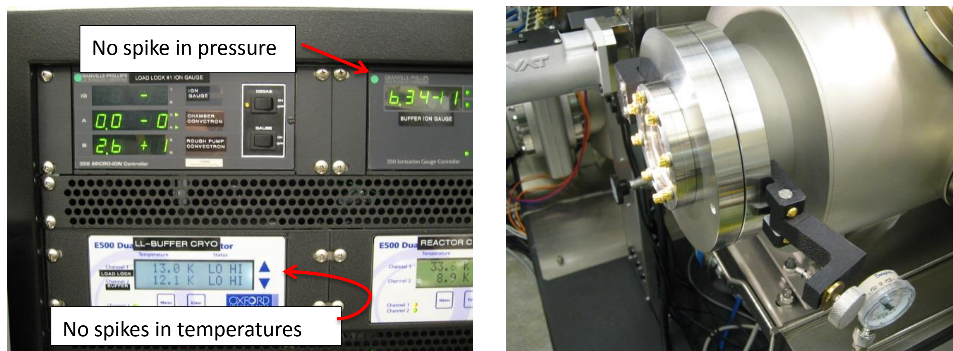

Figure B.22. Watch for evidence of leaking gate valves when the LL is first vented. When the LL reaches atmospheric pressure, open the thumb screw to relieve the pressure.

- Argon flowing
- Wafers you want to unload in LL
- New substrates ready to go in clean hood
- Ga beaker warm, plenty of clean wipes in hood
- Foil on ROMO, lights in position and turned on
- Big gate valve from buffer to MBE closed

Now open the all-metal valve to vent the load lock and simultaneously close the "N2 Vent" valve on the pump cart so that no air can get sucked into the load lock. As you are opening the all-metal valve, watch the buffer chamber pressure and the temperatures of the load lock and buffer cryo pumps as shown in figure B.22. Any spikes in any of these temperatures or pressures means one of the gate valves isn't sealed. The buffer chamber pressure will always show a small blip up from the vibrations of opening the all-metal vent valve, but this will only be a few percent. If the pressure suddenly doubles there is a problem. Assuming there are no problems, watch the pressure of the mechanical gauge. Once it gets to atmosphere, unscrew

the thumbscrew on the load lock door and let the pressure from the argon push the door open (you may have to help it a little bit to get the Viton unsealed). Keep the door as closed as possible while you are mounting/de-mounting wafers to minimize the amount of air that gets in the load lock.

Now put on a pair of the “Accutech Ultraclean” latex gloves over your purple nitrile gloves. The purple gloves leave fingerprints on clean sheets of tantalum while the latex gloves do not, so the latex gloves are evidently quite a bit cleaner. Blow the hand-held nosepiece out with the nitrogen gun before removing the first wafer from the load lock. When you unload the wafers, always start with LL carousel position 1, then 2, then 3 so that you don’t mix up the wafers. Gently place the block on the stainless steel block on the hot plate. Double check which wafer was on which carousel position and lay out the old wafer trays in the same order along with the new substrates that will be replacing them. The unloading process is probably the easiest time to mix up the wafers, so make sure you have a standard sequence in your mind so that you keep track of the wafers.

Wipe the wafer tweezers off on a clean wipe and blow them off with the nitrogen gun. Slide the wafer off of the block with the tweezers. This sometimes requires a lot of force, particularly if the block is new. Be patient. If you can’t slide the wafer off from one direction, try another direction. You can also try going around the edge of the wafer and gingerly prying the edge up a little bit. You have to be really careful to not crack the wafer while doing this. As soon as you see the surface of the wafer start to flex a little bit, stop. Since the wafer sits on a the pedestal in the center of the block, the edge of the wafer sits a little higher than the edge of the block, so keep the wafer tweezers at a bit of an angle so that you don’t slip and make a big scratch across the surface of the wafer. Put the wafer face-down in its tray once you get it off the block. Wipe the Ga off the wafer tweezers and blow them off.

Next, wipe off the Teflon spatula with a clean wipe and blow it off. Add some Ga to the block to replace the Ga left on the wafer you just removed. I typically add two drops on the corner of the spatula as shown in figure B.23. There should always be

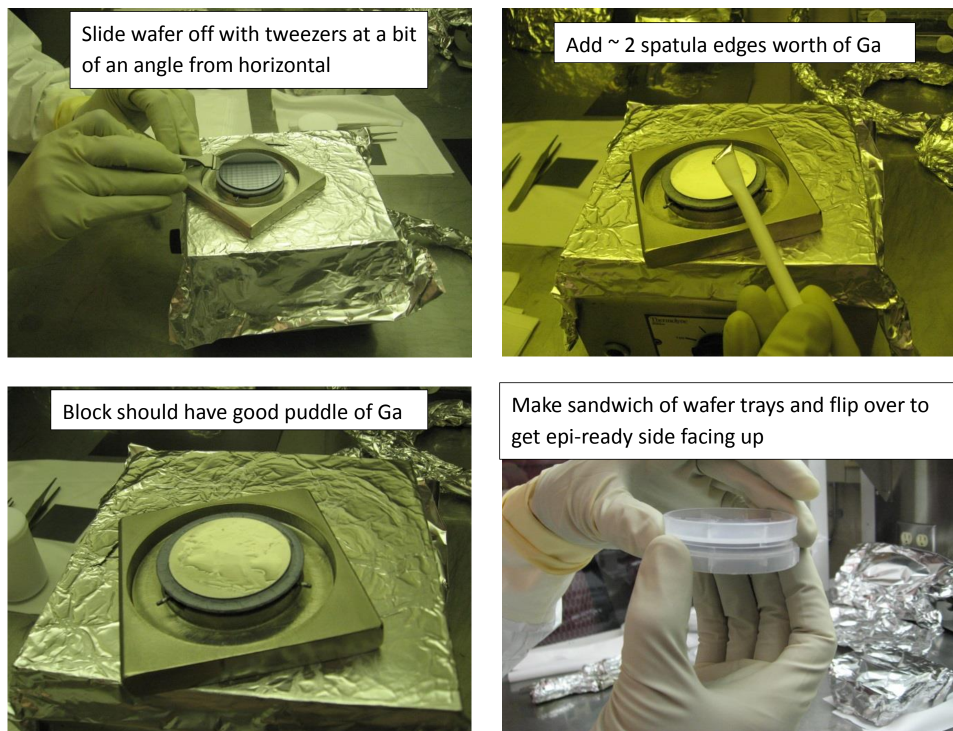

Figure B.23. Procedure for mounting a new wafer on a block.

a good puddle of Ga on the block. If there is not enough Ga, the As from the back side of the wafer will sublimate during the growth and create splotchy dry patches on the block. Once this happens it is very difficult to get good Ga wetting on these dry patches, more As will continue to sublimate off the wafers on subsequent growths, and you can eventually wind up with cold spots on the wafer due to the poor thermal link between the block and wafer.

Next, open up the new wafer tray, remove the spider with tweezers (try not to touch the wafer), and check the surface for any dust. The epi-ready side is face down in the tray, so you need to flip it over before putting it on the block. Using tweezers is a bit dicey since there is always a chance the wafer can slip and get scratched by the tweezers. My approach is to put a second wafer tray face down on top of the new wafer tray and then flip the sandwich over. Now inspect the epi-ready side for any dust and use the wafer tweezers to pick the wafer up and set it on the block.

Push the wafer around on the block a couple times to get the back side coated with Ga. If you slide the wafer about halfway off of the block you can peek at the back side to see how good the coverage is. Again, try to keep the wafer tweezers at a little bit of an angle when you are pushing the wafer around so that you don't slip and smear Ga on the surface. If you do get Ga on the surface, get a new substrate and set the dirty one aside for future use in etch testing, etc. Make sure the wafer is as centered as possible on the block, but also try to minimize the amount you push the wafer around. Every time you touch the wafer you potentially add dirt to it or nick the edge.

Once you are satisfied with the Ga coverage and wafer placement, blow out the inside of the nose piece (try not to blow anything on to the wafer), blow off the wafer, and inspect it for any visible dust. Load the new wafer into the load lock. Make one last check for dust. The best way to do this is to move the ROMO back and forth and look at the surface for little white specks. You will partially blind yourself with the reflection from the light, but I guess that is just part of our contribution to science. Once you have all the new wafers loaded in to the LL and are satisfied with their cleanliness, tighten down the thumbscrew on the LL door and immediately close the Ar valve to the LL.

Next, turn off the gettering furnace and close the two valves on top of the furnace. Now open the valve to sorption pump stage 1. As it pumps down the LL you will probably be able to tighten down the thumbscrew a little more. Tighten it down finger tight, but do not over-tighten it. The screw itself is stainless steel, but the body it sits in is brass. The genius that designed this fixture did not think about how soft brass is. You, however, must have been deemed pretty smart if you are being entrusted with the MBE, so you will know that the threads will strip and jam if you overtighten the thumb screw. I've already had to clean these threads out once, and I'm not sure if the brass body will survive getting jammed again. Before you forget, throw any wipes/gloves with Ga/As/GaAs/etc. in the Ga waste container shown in figure B.25. Once the LL pressure stops falling (should get to 2-3E-2 Torr),

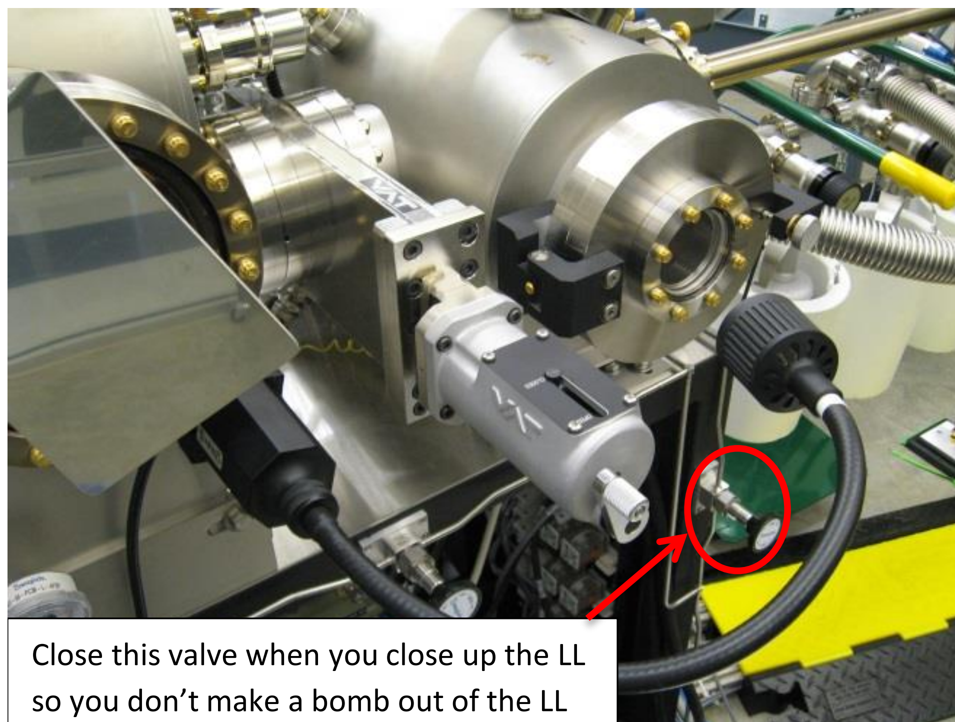

Figure B.24. Do not over-pressurize the LL.

first close the valve to sorb stage 1 and then open the valve to sorb stage 2. This second stage should get the pressure down below 1 mTorr as measured by the LL ion gauge. Once the pressure has stopped falling and is below 1 mTorr, close the all-metal valve finger tight with the small handle and then open the gate valve to the cryo pump. Immediately get the torque wrench out to tighten the all-metal valve down all the way. The LL all-metal valve is currently closed at 31 Ft-lbs. Remember to always check that the torque wrench is set to 31 Ft-lbs (e.g. make sure no one else borrowed it for the other MBE and set it to a different torque). The copper seats in the all-metal valves get deformed when you increase the sealing torque, so if you mistakenly tighten the valve down to say 45 Ft-lbs, you will always have to seal it at 45 Ft-lbs in the future. Once the sealing torque reaches $\sim$ 200 Ft-lbs (if you can even apply 200 Ft-lbs of torque without ripping the flange off the machine) the valve has to be refurbished.

Figure B.25. Ga waste container.

Now that the gate valve is open and the all-metal valve is sealed off, send the LL lamp to 50C at 1C/min and start the "Wait for LL pump down then bakeout" recipe in the LL tab of the stepper. Now close the valve to sorb stage 2 and double check that all the valves on the sorbs and Ar lines are closed. The machine should be in idle at this point. Run through the end of day checklist again to make sure you didn't forget anything. Be sure to read and record the phase separator level in the galley.

B.12 Wafer Inspection

The machine should now be safe to leave, but unfortunately you are not done yet. Now you need to examine the wafers you unloaded. First, look at the Ga on the back of the wafer and write the condition on the summary sheet for that wafer. For the foreseeable future, you will just have to write "Very good Ga coverage with Ga-free

ring around the edge." The blocks with the pedestal are a huge improvement over the original blocks with a completely flat surface. The pedestal blocks consistently wet the wafer very well since there is no crud building up around the edge of the block to prevent the wafer from making good contact with the Ga. The blocks may end up in a downward spiral towards having dry patches, however, if you are stingy with the Ga . If there is not enough Ga, the wafer can sublimate and make a dry patch on the block. This dry patch is then not wetted very well with Ga the next time you load a wafer, this wafer sublimates some more, and so on. If that happens you will likely see patches of shiny GaAs on the back of the wafer when you unload it. If there is a patch on the wafer without Ga, sketch out that patch on the summary sheet.

Next, flip the wafer over (you can usually do this without tweezers) and set it in the clean hood (where the lighting is the best). There are a number of defects you need to keep an eye on. The first is the infamous island defect that is apparently unique to our lab. This defect happens when the Ga freezes under the wafer and punches out chunks of the wafer. This plagued almost all our growths for the first ~ 4 months we ran the machine until we figured out that the cryo pump in the LL radiatively cools the inside of the LL to something like 13C which is more than enough to solidify the Ga. The solution is to keep the LL lamp at 50C all the time. Note that this can also (rarely) happen to the wafers after they are out of the machine. One time I had a half piece of wafer face down in its wafer tray with a quarter sitting on its back (i.e. the Ga-covered sides were touching). I touched the Ga with some cold metal tweezers which flash-froze the Ga on the back of the two pieces and created an island defect on one of my pieces of wafer. If there are any island defects on the wafer, sketch out where they are on the summary sheet.

The second, more ubiquitous, defect you will see is "Ga-spitting" or "Ga-grit". This will look like dust on the surface of the wafer. Just write down the degree to which you see the spitting. This is pretty subjective, but I usually use terms like "Minimal Ga spitting" or "significant Ga spitting" if I can see a lot of spitting outside

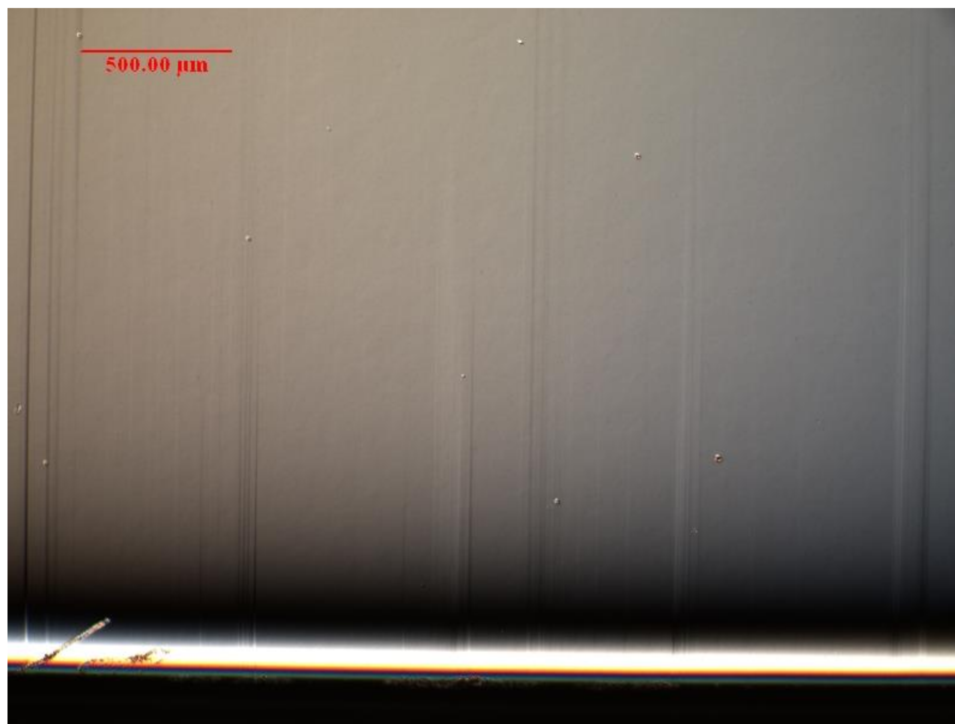

Figure B.26. Nomarski phase-contrast optical micrograph of slip lines at the edge of the wafer.

the LL. Wafers grown with Ga1 usually have "significant Ga spitting" while wafers grown with Ga2 usually have "Minimal Ga spitting."

Another typical defect is slip lines. These are caused by thermal gradients between the center and edge of the wafer and typically don't show up until the wafer gets heated all the way to growth temperature. The new custom blocks are thicker than the original blocks, so there is better thermal uniformity and fewer slip lines than the old blocks. Figure B.26 shows a picture (using Nomarski contrast) of slip lines at the edge of a wafer grown on block H (the first custom block).

There are varying degrees of slip, though the slip has been minimized a lot since we got the pedestal blocks. Here are my guidelines for classifying the amount of slip:

- No Slip: no slip lines visible to the naked eye

- Minimal slip: probably 3-4 slip lines visible to naked eye around the flats and directly opposite each flat
- Some slip: Several slip lines at the flats, slip lines extend more than 2-3 mm
- Significant slip: Many slip lines at each flat/across from the flats, some slip lines extend more than 1 cm

I have found the best way to see the slip is to move your head back and forth while looking at the wafer at a glancing angle. You should see a bit of a bend in reflections on the wafer surface at the slip lines. This will probably take some practice to find the best way to find the slip lines. Try to compare your sketch of the slip lines with what you see in the Nomarski (more about this later).

A similar defect is what I term “distortion”. This used to happen when too much crud got built up on the edge of the block and the wafer was not sitting perfectly flat on the block. It still happens sometimes for the first one or two growths on a new block (or right after a block is etched) and there is not good Ga coverage. It just looks like slip lines in the center of the wafer that do not extend to the edge of the wafer. Evidently these defects are caused by serious mechanical stress on the wafer during growth. To find these, look at the reflection of a straight line and scan your head around. If you see the straight line bend somewhere in the center of the wafer, this is distortion. Write it down and try to sketch out where the distortion line is. These lines do not typically follow the major crystal axes. If there are no distortion lines, write “no distortion”.

Next, there is “interference”. This is just optical interference due to film non-uniformity from the center to the edge of the wafer. It is more pronounced in wafers grown with more Al. For most x = 24% wafers I say they have “minimal interference”. For the higher x wafers, if they have very visible interference rings, I call it “some interference.” For wafers grown with large amounts of AlAs the wafer may be pretty green in some places, so I call that level of interference “significant interference.”

We have also had problems with "haze" at various points in the past. The haze is pretty easy to see; it looks like a white/hazy region of wafer. Sometimes it is easier to see it in the LL with the bright lamp shining on it and then is easy to miss outside of the LL. Haze can be caused by a few different things. First, too little As flux will result in a hazy wafer. This isn't usually a problem for us since we don't change the As flux from run to run. If too little As is the cause of the haze, the haze will be circularly symmetric since the wafer is rotated during growth. Remember that it is possible to have the wafer at growth temperature with the As valve completely closed for short periods of time (you can do this and look for a phase change in the surface reconstruction to calibrate the pyrometer). So for insufficient As to be the cause of haze, the wafer has to be starved for As for an extended period of time, which to my line of thinking requires the haze to be circularly symmetric if the wafer is rotated. Second, insufficient outgassing at growth temperature can leave some oxide on the surface of the wafer which will result in hazy regions. If the pyrometer reading is not increasing smoothly when the wafer is first warmed up in the growth chamber, this may be due to a non-uniform oxide. This is almost certainly the case if the pyrometer stabilizes after the wafer has been above 610C for a few minutes. If this happens, you really have to make sure the wafer is $>$ 610C for 10 minutes to get all the oxide off. This will probably show up as random spots of haze on the wafer. Occasionally you will see small slivers right on the edge of the wafer that look white-ish. My assumption is that this is just a region that never got hot enough to get rid of the oxide. I generally just ignore tiny slivers of haze like this since they are right on the edge of the wafer. Third, we have had at least one batch of wafers from Wafertech that were not properly cleaned prior to shipment. The haze usually showed up in a crescent shape when this was the case. Figure B.27 is a picture (again from the Nomarski) of what the hazy region looked like on the dirty Wafertech wafers. I do not currently have any pictures of other haze under the microscope; the other sources of haze may cause a different looking morphology. All the previously mentioned defects should be noted/sketched on the "wafer surface" portion of the summary sheet. All

Figure B.27. Nomarski phase contrast micrograph of a hazy wafer. The field of view is ~ 2.5 mm wide.

these defects are visible to the naked eye, though they may take some practice to spot on a regular basis. Once you have done this initial inspection it is time to look at the wafers with the Nomarski microscope.

B.13 Nomarski Exam

The microscope with the Nomarski contrast capability in the electrical characterization lab has turned out to be an extremely useful piece of equipment in diagnosing problems with the growths. The Nomarski contrast (also called "DIC") allows you to get rid of a lot of glare on the surface and see features that you would otherwise miss in bright field or dark field modes. Getting the microscope to give good images on the camera is a little tricky, though, so I've outlined all my settings here. First, enable the tool in Coral. It is silly that we get charged $\$20-30$ even if we just look at

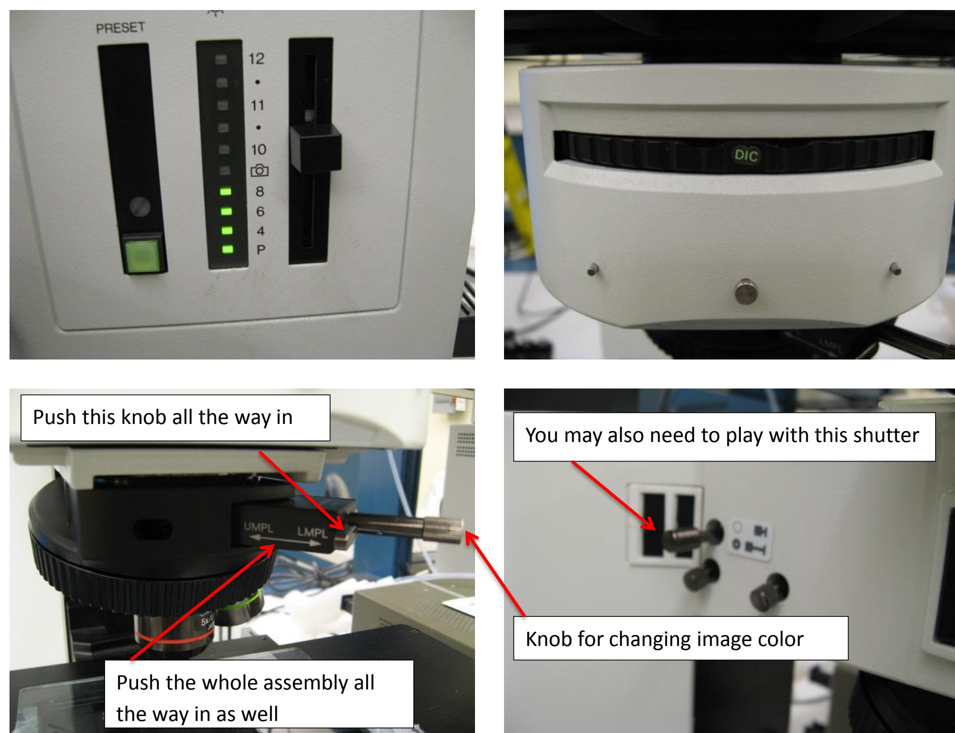

Figure B.28. Microscope setup.

the wafer for 2 minutes, but the rules are the rules. Next, turn the lamp power to 8 (remember no gloves on the microscope or computer), put the central rotation piece to “DIC”, and make sure the Nomarski slide is slid in all the way in to the optical path.

Normally you don’t have to adjust the shutter on the left side of the microscope, but recently someone has started messing with it. If you can’t see any light on your wafer when the microscope is in DIC mode you probably need to open this shutter up more. Once you have the microscope set up, turn on the camera and open the “Spot Advanced” software. If it says, “No camera found”, make sure the camera is on, close the software and re-open the software; it sometimes takes a while for the computer to realize the camera is there. Once the software is running, start the live image, and select “JDW DIC” from the image setup menu (also can be selected at the bottom

right corner of the screen). This has the right camera settings to get good images. If someone else messed with my settings, the numbers to use in the image setup are:

- Exposure: 88 msec
- Gain: 1
- Full Chip Imaging
- No color correct
- Bit depth: 8
- Red: 1.136
- Green: 1.000
- Blue: 1.400

Adjust the focus and the knob on the Nomarski filter until you get a good image. Figure B.29 is a picture of a wafer with good morphology (Wafertech wafer grown on block H) with significant Ga spitting (used Ga1 and Ga2).

The morphology can be pretty hard to see sometimes because its visibility is very dependent on the exact focus/contrast. So as a test, make sure you can see the slip lines at the edge of the wafer as shown in figure B.26. Figure B.29 has a surface that looks free of major features but may have a little bit of what I would call "orange peel" texture. Other growers have also seen orange peel, but to my knowledge it is not a major problem for nanostructure fabrication or FQHE quality. If you can't see any texture at all, I would guess you just haven't got the image settings correct. Again, this will take some practice to know what to look for and how to tweak everything. Once you get the image settings correct, scan over the surface of the wafer a few times. Don't worry about noting the position of Ga spitting defects (even the big ones); they are just the price we pay for using high mobility effusion cells (i.e. cold lipped cells). The main thing you need to look for is cross hatching. This was another defect that

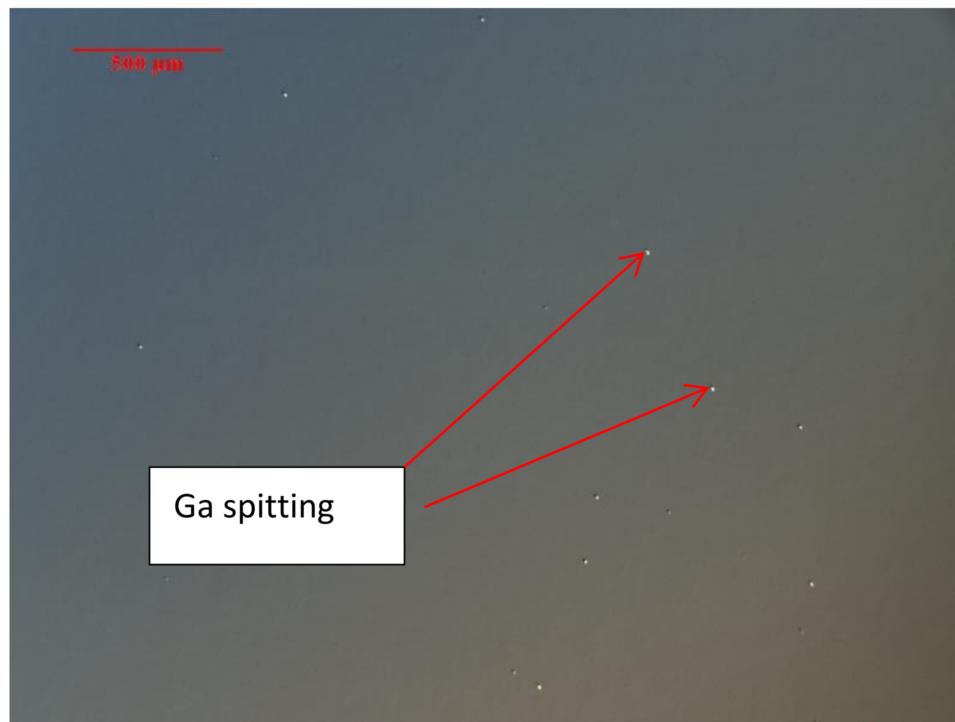

Figure B.29. Nomarski phase contrast micrograph showing a wafer with good morphology.

made grad school less than fun for many months. Figure B.30 shows an image (again with the Nomarski on the same scale as the previous picture) taken before I was good with the microscope/camera settings (hence why the image is blue). This image came from the center of the wafer. This kind of morphology is obviously bad news for anyone who wants to make nanostructures of any kind on the surface. That being said, it does not appear to affect the transport. I am pretty sure that our best ν=5/2 sample so far (that had $\sim$ 600 mK energy gap) had this kind of morphology. In fact, we didn't even know all our wafers had this morphology until Lisa Tracy pointed it out to us.

Cross hatching is apparently due to roughening of the surface during growth. To get rid of the cross hatching we changed a few things. First, we reduced the thickness of the initial smoothing layers after the oxide is desorbed from 100 nm + 100 sec pause

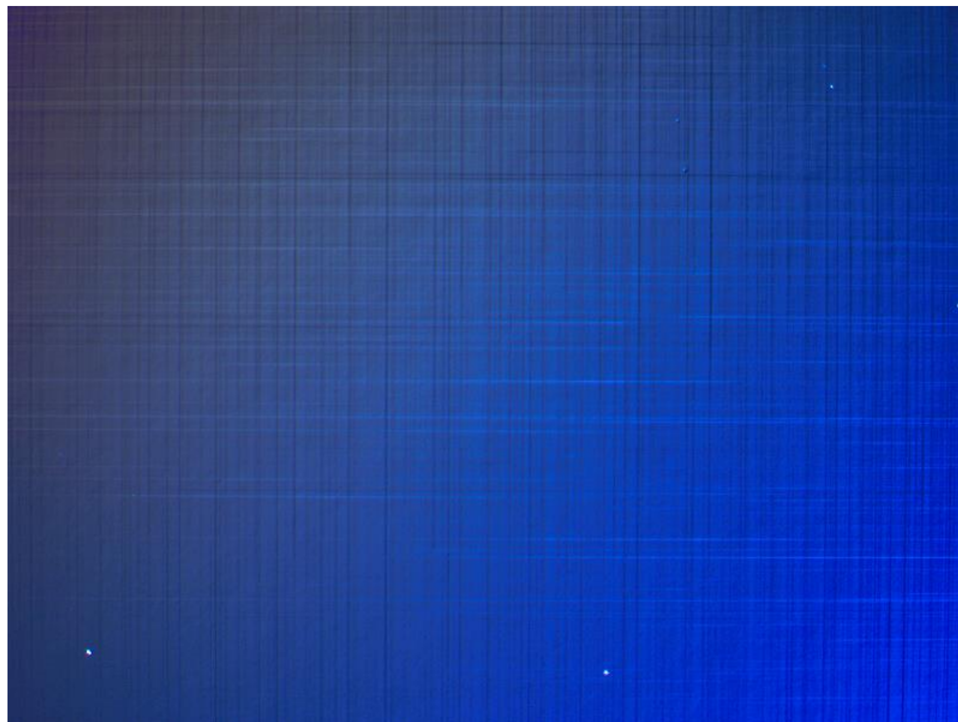

Figure B.30. Nomarski phase contrast micrograph of a wafer with "cross hatch" morphology. The field of view is $\sim$ 2.5 mm wide.

$\times$ 5 repeats to 10 nm + 20 sec pause $\times$ 50 repeats. Just from watching how a RHEED wafer smooths out it is obvious that thin layers with lots of pauses do a better job at getting the RHEED pattern streaky than a bunch of thick layers do. Second, we reduced the time the wafer sits hot at growth temperature prior to the growth from 30 minutes to 10 minutes. Third, the morphology is somewhat dependent on the wafer manufacturer. The Wafertech wafers give smoother morphology than the AXT wafers do (the Wafertech wafers start out smoother and have a thinner oxide). With our current growth procedures the AXT wafers don't typically have fully developed cross-hatching, but they are still noticeably rougher than the WT wafers. Finally, the pedestal block design again seems to have improved the morphology relative to the original blocks with flat surfaces.

Once you are done scanning the wafer with the Nomarski, sketch out any cross hatching regions on the "Nomarski exam" portion of the summary sheet. Try to get this map accurate. If there is cross hatching, there may be parts of the wafer that are still smooth, and it would be good to know this when we decide which piece of a wafer to send to a collaborator. If there is no cross hatching or roughness, just write "good morphology everywhere." Now that you are finished with the wafer exam, put the camera settings back to "factory defaults" so no one else edits my settings, close out of everything on the computer, turn off the camera and lamp, and log out of coral. Finally, make a photocopy of the summary sheet to send over to the physics building with the wafer.

Lastly, make sure you have copied all the structure sheet, summary sheet, and the wafer exam details into the database. If someone else has it open, e-mail or text them and complain about how they forgot to close it again. Finally, send an e-mail out to Mike plus anyone else who needs to be cc'd giving an update of the state of the machine, anything unusual in the growth, and when/how the wafers will be getting over to the physics building to be characterized (communication between everyone is key to keeping wafers from getting lost or sitting uncharacterized for long periods of time).

Congratulations, you've finally completed a day's work in the MBE lab! Now go home and get some sleep.

C. Wafer Characterization Standard Procedure

C.1 Sample Preparation

C.1.1 Initial Bookkeeping

- Look at the sketch of the wafer morphology on the sample summary sheet and choose a region with good morphology as the source of your samples

- Wear gloves while handling all clean tools

- Cleave a 4 mm wide strip out of the wafer (try to take a chip from as close to the center as possible). If the wafer has already been characterized (i.e. if the strip has been sitting in air > 24 hours) you need to have four fresh edges for your chips or the contact resistances may be very high. If the existing strip is wide enough, you can cleave off a tiny sliver (~ 500 μm) from each pre-existing edge, or you can cleave out a new strip. In all cases try to conserve material as much as possible. If you can, just cleave a strip out of a quarter wafer rather than a half wafer to preserve more material for potentially shipping to collaborators.

- Be careful to wipe the Ga off of the tweezers each time you grab the strip to minimize the amount of Ga on the surface of your square

- Make sure to not make your samples too big (e.g. 6 mm is too big). If the sample is too large, it will waste a lot of material and there will be larger sample-to-sample variation in density and mobility

- Draw a map of where the samples came from on the summary sheet

- Naming convention example: "1-31-11.1-a"' (".1" indicates first wafer grown on 1-31-11, "a" indicates which square the sample is – MAKE SURE TO INDICATE SQUARE LOCATION ON SAMPLE MAP)

- Store wafer top side facing down in wafer tray, leave small pieces sitting on larger half-wafer pieces (i.e. Ga-covered sides touching each other). Try to secure all loose pieces with the plastic spider, avoid getting Ga on any material that came from close to the wafer center.

C.1.2 Applying Contacts

- Scratch the square at each point you will make a contact – use a single scratch $\sim$ 0.5 mm long

- Record the alloy used for the contacts on the back of the summary sheet (use the $\sim$ 6.9% Sn InSn mix for standard n-type contacts)

- MAKE SURE YOU ARE USING THE CORRECT SOLDER TIP (mixing the wrong tip with the wrong alloy will contaminate the alloys and ruin the contacts)

- Turn off soldering iron when done and scrape tip with its associated razor blade

C.1.3 Mixing Up Contact Alloy

This section is a bit of an aside in case the contact alloy stops working or you run out. The InSn alloy rarely goes bad, but the InZn alloy has a tendency to suddenly stop working after a month or two.

- Clean a glass petri dish (no plastic!) with some IPA and a clean wipe (try to avoid leaving a bunch of fuzz in the petri dish)

- For mixing InSn, weigh out a couple pieces of In along with one piece of Sn. Shoot for $\sim$ 7% Sn by mass

- Mix up the In and Sn with the appropriate soldering iron tip in the petri dish (use a new tip if the last batch of InSn got contaminated). Mixing $\sim$ 3-4 minutes is usually sufficient

- For InZn, take a couple pieces of In plus one piece of Zn and mix together in the petri dish. The solder tip should be at a low enough temperature to melt the In but not easily melt the Zn. Mix the Zn around in the puddle of In until the puddle starts to get a little sticky, then discard the remaining Zn pellet. It seems that it only takes a few atoms of Zn per contact to make them work. Adding too much Zn can cause strange behavior like the Hall trace being non-linear. Zn has a very high vapor pressure, so our theory is that if there is too much, it can evaporate during the annealing and form a weakly conducting film across the sample that starts to short out the Hall voltage. Too much Zn in the annealer is also a concern in terms of contaminating the n-type InSn contacts.

- Once you've tested that your new alloy works, label it (include the date) and add this name to the database so that future samples can be tagged as having been made with this particular batch of alloy.

- If the previous batch was contaminated and your new batch works, throw out the old petri dish (in the sharps container) and the old soldering iron tip (in the metals waste bin)

C.1.4 Annealing Contacts

- Purge the tube for 10 minutes with forming gas after you load your samples

- Use flow rate 5 sccm (most of the time this does not need to be adjusted)

- Use variac = 42% (the variac should always be left on and at this percent output)

- Make sure the thermocouple is inserted all the way into the tube

- After purging, insert the tube into the furnace and start the timer for 15 minutes
- At the end of the 15 minutes pull the tube out of the furnace and keep the gas flowing until the samples cool ($\sim$ 10 minutes)

C.1.5 Mounting Samples on Header

- File down the forks on the header a little bit to clean off any oxide so the In will stick better
- Solder wires to your sample using the same solder tip that you used for applying the contacts. If the wire breaks, please re-thread it
- Add some In to the forks on the header
- Glue your sample into the header with a dab of rubber cement
- Solder the wires to the header
- Label the sample boxes for each sample (including chips that you don't immediately mount on a header)

C.1.6 Room Temperature Checks

- Check that all the contacts are continuous using a 2-terminal measurement with the lock-in (the high mobility "standard structures" should have $1.5 - 5$ kΩ 2-terminal resistances to ground at room temperature)
- Check the the LED is working (use 2 mA current)

C.2 4K Characterization

C.2.1 Cooldown and Measurement

- Check that there is enough helium in the dewar (20L is the bare minimum)

- After loading your sample and the LED, screw the magnet back on and make sure its connections are continuous
- Mount the probe on the dewar and start lowering it slowly. Lower at a rate of $\sim$ 1 inch/minute
- Monitor the resistance of the magnet with a multimeter while you are lowering. Once the resistance gets to $\sim 500\ \Omega$ stop and let it cool for $\sim$ 5 minutes
- Continue lowering the probe until the magnet resistance bottoms out (usually $\sim 1\ \Omega$)
- Now switch the multimeter to the RuO resistor and continue lowering until the resistance saturates $\sim 1057\ \Omega$
- Measure and record the 2-terminal resistance values and the resistivity
- Hook the magnet leads up to the magnet supply and setup to measure the Hall data
- Save the data using the existing organizational scheme − each sample gets its own folder in the 4K data directory (don't add sub-folders inside of the folder for other samples from the same wafer).
- If illumination is necessary, ground all the contacts and illuminate with 2mA for 2 minutes (remember to take the magnet to zero first)

C.2.2 Warm-Up

- Pull the thinner part of the probe up all the way and wait a few minutes
- Pull the bigger part of the probe up $\sim$ 3-4 inches at a time, wait $\sim$ 5 min in between each step
- Once the probe is all the way up, wait until the magnet resistance $> 1000\ \Omega$ before removing the probe from the dewar

- Double check that the vent valve on the dewar is open before you leave

C.2.3 Wrap-Up

- Update the database on the group drive with the results of your characterizatoin. Fill out all the fields on the 4K input form. Do not put text in the numeric fields. If you couldn't measure something, leave that field blank and state why couldn't measure it in the comment field
- Return equipment (e.g. heat gun, curve tracer, etc.) to its home, and store the samples in the prep room in chronological order
- Update the group with the results of your measurement

D. Room Temperature Hall Effect System Standard Operating Procedure

This appendix will give a short overview of preparing and measuring samples with the MMR room temperature Hall effect system. This tool is used for measuring bulk doping concentrations to calibrate the MBE dopant sources and also to quantify background impurity concentrations when the material is really dirty (for instance at the beginning of a growth campaign prior to outgassing the sources). The lowest density it can reliably measure is in the mid- 10^{14} cm^{-3} range.

D.1 Sample Preparation

Begin by cleaving out a square piece of the wafer $\sim 6 \times 6$ mm^2. Apply contacts as described in Appendix C, but only apply contacts to the corners (the MMR Hall effect system can only handle contacts at the corners). If you are trying to measure the background impurity concentration, you should try either pure In or InZn contact alloy. Note that when the background concentration drops into the mid- to low- 10^{14} cm^{-3} range, the density the MMR spits out may be lower for pure In contacts than for InZn contacts. After applying the contacts, anneal them as described in Appendix C.

After the contacts have been annealed, wire them up with un-insulated copper wire (there should be a spool on the shelf above the microscope in the sample prep room). The normal gold bond wire used to wire up samples for low temperature measurements is not strong enough to surive the mounting process on the Hall effect system. Since the copper wire is of unknown cleanliness, use a “dirty” soldering iron tip (e.g. the one labeled “NiAuGe soldering”) and its associated In for solder. You should give yourself ~ 0.4 inches of wire to ensure that you can wire the sample up

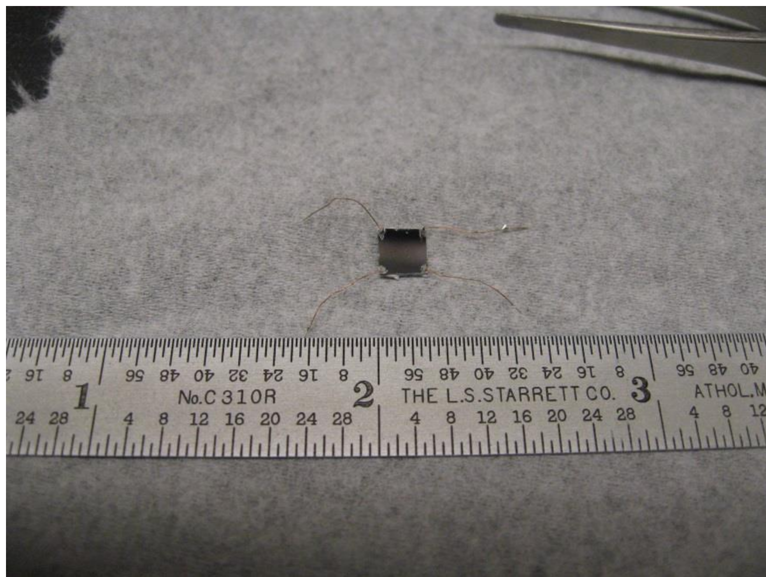

Figure D.1. A bulk doped sample after being measured. The scale shown is inches.

to the MMR system. Figure D.1 shows a finished sample. Since the copper wire is thicker and stiffer than the gold bond wire, try to avoid putting too much stress on the solder joint (e.g. if you need to bend the wire, hold the wire close to the bond with one pair of tweezers and bend it with a second pair of tweezers). Since all the contacts have to work, make yourself 2-3 squares before you head over to Birck to measure them so you don't waste the whole day going back and forth between the two labs. Once your samples are ready, pack them up in a sample box with a kim-wipe to keep them from rattling around and head over to Birck.

D.2 Sample Measurement

Your training on the MMR system should have already prepared you for using the system, but here is a reminder in case your training was cursory or in case the

Figure D.2. View of the inside of the MMR Hall effect system sample chamber.

instruction sheet[1] in the lab goes missing (I've also put a copy of the original instruction sheet on the group drive in the "/GaAs growth data/Standard Operating Procedures" folder). First, open up the sample chamber (remember to rest the sample end on the high-tech "sample stage"). Set your sample on the end of the suppport, and *very* lightly tamp it down. The sample sits on the end of a very fragile (and expensive) cantilever, so if you push too hard you will break the support (this has happened before). The original procedure I was shown involved holding a wood q-tip horizontally from the fuzzy end and lightly pressing down on the sample with the other end. The sample just needs to be flat, so you don't need to over-do it. Once the sample is in place, solder your wires on to the four solder pads shown in figure D.2. The soldering iron is not in very great shape, so you may need to clean some of the oxide off with sandpaper followed by a good wipe with IPA. Make sure you can wet the tip with some solder *before* you try melting the solder balls on the MMR

[1]The original operating procedure upon which this section is based was prepared by Jeremy Schroeder.

pads. If the soldering tip is not in good shape you will melt the plastic in the MMR before you melt the solder (another expensive mistake). Once your sample is wired up and flat on the support, screw the lid back on the chamber. Don't over-tighten the screws; you don't need to do any temperature studies, so it doesn't matter if the chamber is hermetically sealed.

Slide the chamber into the magnet housing and tighten down the thumbscrew underneath the chamber. Make sure the H-50, MPS-50, and K-20 controllers are all off and then plug in the ribbon cable along with the four triaxial cables. Be careful not to bend any of the pins on the ribbon cable connection.

Turn on the H-50 and MPS-50 controllers and start up the Hall effect software. Open the communications setup window and make sure the following boxes are checked: Splitter, Com1, MPS-50 present. Next open the H-50 manual control window and type SC0.00977 and press "send command". This will set the Hall sensor sensitivity constant to 0.00977 V/kG.

Next open up the "Experiment Setup" window. First you should check the contacts as the original standard operating procedure says. Open up the Van der Pauw tab, select "single point", and click "options". Set the "Max Voltage" field to 2.3V, click the "Fixed Setting" radio button, set "Coefficient %" to 100, and click "Start Measurement". This will apply a voltage in both directions between each pair of contacts. If the contacts are working, the I/V values should be approximately the same for each polarity for a given contact pair. If you see a variation of more than a few percent between the two polarities, you should try re-soldering the offending contact to the MMR wiring harness. If you do this process a couple times without improving the contact, move on to your next sample. You should have near 100% yield for doping concentrations $> 10^{17}$ cm^{-3}, but you may have finicky contacts when trying to measure background impurity concentrations. The maximum current the MMR can supply is 20 mA, so if one of your contact pairs gave a current higher than that, decrease the "Max Voltage" value and re-run the measurement. Note the lowest current value and round it down to the nearest two significant digits. This will

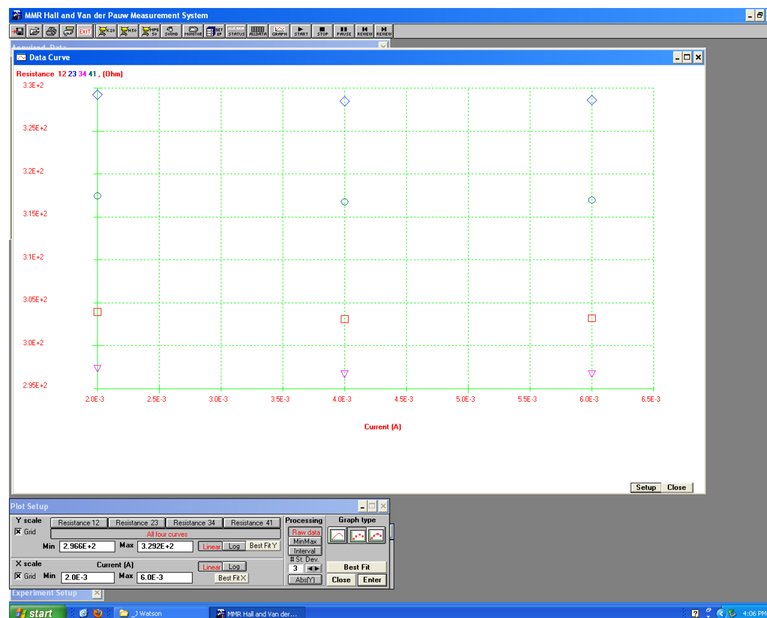

Figure D.3. Graph of data from the linearity check measurement.

be the maximum current you should specify in the Linearity Check and Hall Effect measurement sections.

Return to the “Experiment Setup” window and click on the “Linearity Check” tab. Set the mode to “Curve”, select the “Current” option for the experimental variable, click the “Linear” radio button, set the “Finish” field to the current value you wrote down from the Van der Pauw measurement step, and set the “Start” and “Step” values to convenient values so you get 4-5 steps from the initial to final current values. In the “Advanced” window, set the number of repetitions to 1 and ignore everything else. Click “Set” in the Experimental Setup window and press “Start” on the top menu bar. This will measure the 2-terminal resistance of each contact pair for the values of the current you specified. Once the measurement is done, click the graph button on the top menu bar, click the “All four curves” option, select “Raw data”, select the scatter plot graph type (i.e. the midddle option), and click “Best Fit” and “Enter” to display the data on a graph. You should see something like that shown in figure D.3. The resistance values for each contact pair should not vary by more than

a few percent over the measured current range (i.e. the contact should be Ohmic). If the resistance of any of the contacts does vary a lot, you can try re-soldering the offending contact or discard the sample. If re-soldering doesn't fix the problem, you can still continue on to the Hall measurement step, but you shouldn't put much stock in the answer the MMR spits out. Save your linearity data in case you want to come back later and check things (note that the software is old, so it will complain if the file name is long or has illegal characters like spaces).

Next open up the "Hall" tab in the Experiment Setup window. Select "Curve" mode, but set your current starting and ending values to both be equal to the working current you wrote down in the Van der Pauw measurement step. In the advanced setup window set the number of measurements to 1 or 2. The original operating procedure suggests doing several measurements, but this is only because the people who first used this tool had no idea how to make Ohmic contacts and consequently got wildly varying Hall measurement results (including changes in the sign of the Hall coefficient). If your contacts are working, the standard deviation of the density values it measures will be several orders of magnitude smaller than the average reading. Remember to set the thickness of your sample appropriately, make sure the "Field" field is set to 3300 G, click "Set", and press "Start" on the top menu bar. The system will then proceeed to measure the resistivity at zero field and the Hall voltage at positive and negative 3300 G for different contact configurations. At the end of the measurement, it will spit out a screen with all the results; all we really care about is the density and the sign of the carriers. Save your data and move on to the next sample. The boss was initially somewhat skeptical of this machine since it doesn't actually show the Hall data anywhere. If you want to give him some more convincing data, go back to the Experiment Setup window, click the "Field" tab, set the start/finish/end values to take data at several field points up to 3300 G, and start the measurement. Save the resulting data to a USB key as a CSV file and take it back to your office to extract the data. Open the file in Excel and try to decipher the meaning of all the sub-scripts in the header sections. You should be able to figure

out which data fields correpond to the Hall voltage at each field value. Once you find these values, you can make a plot of R_{xy} vs. B and extract the density. The density you extract this way should match the density the MMR spits out, and the plot should help reassure everyone that the MMR blackbox is actually working.

Once your measurements are all done, open the "H-50 Manual Control" window and type FI0.0 and press "Send Command". Shut down the software and turn off the H-50 and MPS-50 controllers. Remove your sample, remember to turn off the soldering iron, put everything away, and clean up your mess.

E. Helium Transfer from Liquefier Standard Operating Procedure

E.1 Pre-Transfer Bookkeeping

- Check on the Labview gas meter page that no more than one other person is transferring. The compressor in the sub-basement can only handle $\sim$ 18 CFM, so the sum of all the helium being boiled off in all the labs must be less than this to avoid blowing a hole in the bag in the attic
- Hook your storage dewar up to the helium recovery line
- Check that there is enough liquid in the liquefier to do your transfer. There always needs to be at least $\sim$ 150L in the liquefier
- If the liquefier was just started up, make sure it is stable before you transfer. It is best to just ask Keith if it is ok to transfer, but if you can't find him the screen on the laptop should look something like that shown in figure E.1
- Measure the liquid level in your storage dewar with the thumper dipstick (the brass quick-connect fitting should be in the top right drawer of the desk)
- Turn on the liquefier level meter if it is not already on (switch is on the back) and put it in continuous mode (denoted by the *) by pressing < MENU >, < ENTER>, < MENU >
- Fill out the Google Docs spreadsheet. The "LHe" column refers to the starting amount of liquid in your storage dewar (in L), "Dewar cm" refers to the level of helium (in cm) in the liquefier, "meter (scf)" refers to the reading on the gas meter

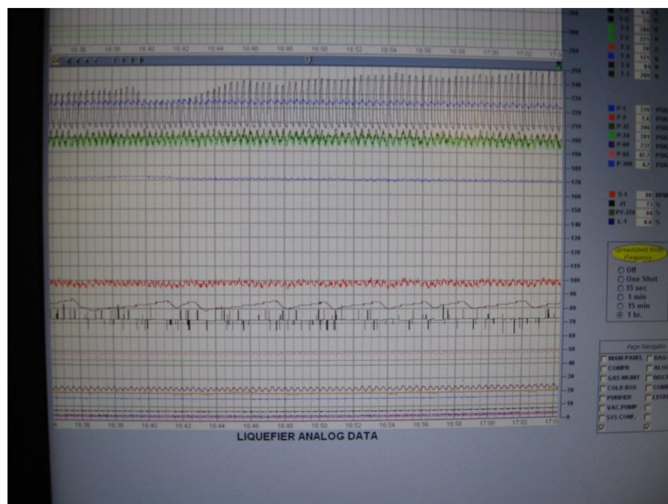

Figure E.1. View of the liquefier laptop when the liquefier is running smoothly.

- Fill out Keith's Excel spreadsheet with the date, time, "P impure" (from the Ashcroft gauge shown in figure E.2), and liquefier dewar liquid level (in cm)

E.2 Starting the Transfer

- Put the appropriately sized quick-connect fitting on top of your storage dewar to mate with the transfer tube
- Close the recovery valve on the big dewar (see figure E.3) so that you can pressurize the dewar
- Open the valve on top of the liquefier and slide the tube down a few inches
- Immediately open the valve on the transfer tube to let the gas start flowing out to purge the tube
- Once the tube is sufficiently purged ($<$ 1 minute if the liquefier is pressurized, $\sim$ 1 minute if it is not pressurized) insert the transfer tube into your storage dewar

Figure E.2. Pressure gauge behind the computer showing the “P impure” reading.

- Slowly lower the transfer tube into the liquefier and your storage dewar. It should take $\sim$ 3-4 minutes to get the tube all the way down.

- If the tube starts to get stuck (it usually does since it is kinked a bit), gently push away from you (towards the tool bench) while simultaneously pushing down. It usually doesn’t require a lot of strength, just a little finesse. If you’re having a really hard time stop and get some help (preferably Keith). Towards this end, plan to do your transfers during regular business hours until you’ve been transferring regularly for a couple months so that people will be around in case you need help.

- Once you get the transfer tube down, make sure neither o-ring is leaking (i.e. no hissing). If there is a leak, stop the transfer (see the end of this operating

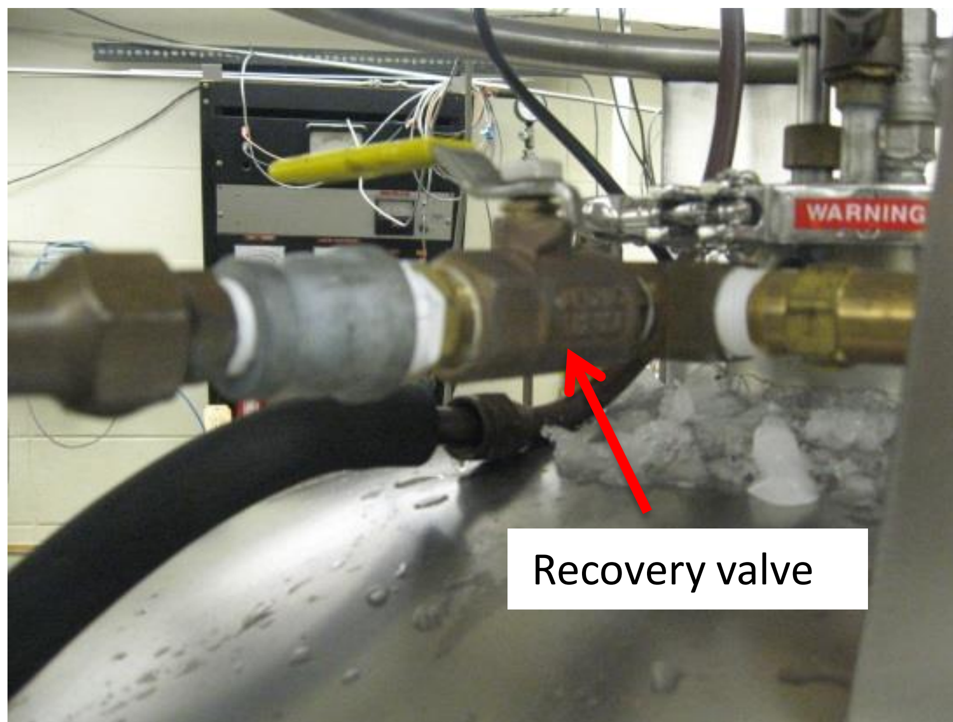

Figure E.3. Recovery valve on the liquefier dewar.

procedure) and replace the o-ring (there should be some spares in the top left desk drawer). If you see that we are low on this particular size of o-ring, let Keith know.

- If the liquefier is running, make sure that the pressure stays below 6 psi. If it gets higher than this, the liquefier will shut down

- If the liquefier is running, do NOT pressurize the dewar with any gas. The liquefier will keep itself pressurized.

- If the liquefier is not running, hook up the rubber helium hose from the gas cylinder to the pressure-building valve (after purging the hose for ~ 10 seconds).

Figure E.4. Liquefier dewar pressure. Do not let this pressure exceed 6 psi if the liquefier is running.

- If the liquefier is not running, pressurize the liquefier dewar to $\sim$ 3.5 psi. This should be high enough to complete the transfer, though you may need to check if you are filling a completely empty dewar.

- If the liquefier is running, your storage dewar should start filling by the time you get the transfer tube all the way down. If the liquefier is not running, it will probably take $\sim$ 5-10 minutes to get liquid transferring (evidenced by the sudden large drop in the boil-off rate from your storage dewar). Once liquid starts flowing, it should take $\sim$ 1 hour to fill an empty (but still cold) dewar. Do not try to fill a warm dewar. If the dewar is warm, you need to talk to Keith about how to pre-cool the dewar. Attempting to transfer directly into a room temperature dewar would waste an unacceptable amount of helium.

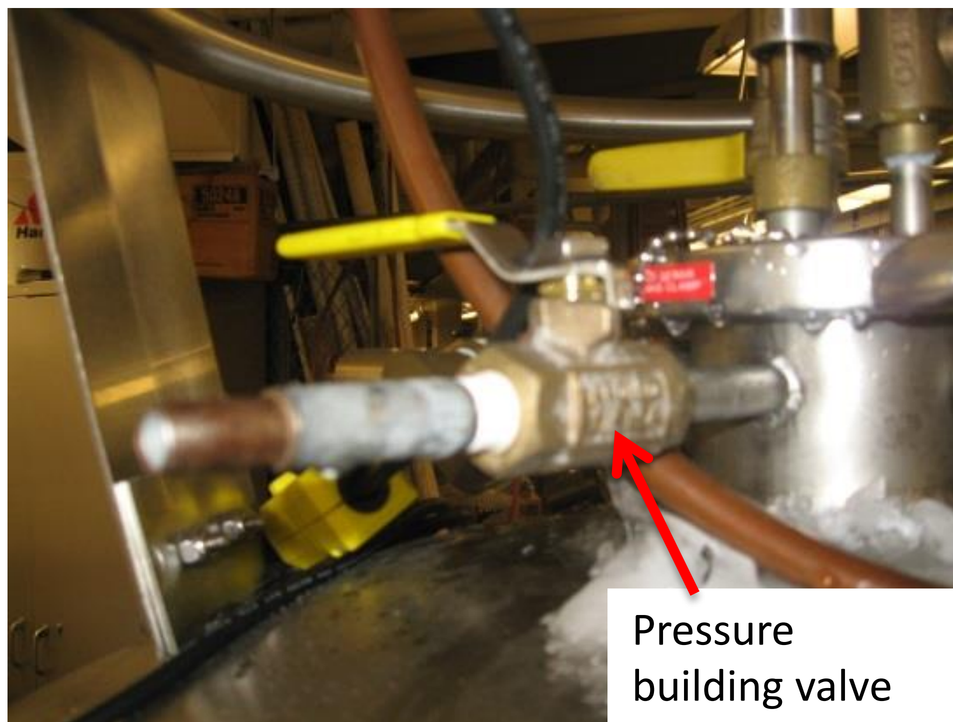

Figure E.5. Pressure building port.

- Once the recovery line starts to get cold close to the gas meter (usually ~ 30 minutes into the transfer), turn on the heat gun below the recovery meter to keep the meter from freezing up (see figure E.6. If the meter freezes, the rubber seals inside will crack and we will start losing a lot of helium. This happened once in the past (fortunately no one in our group was involved). Suffice it to say, you don't want to be the next guy to freeze up the meter. With this in mind, always double check that the heat gun is actually putting out hot air. The heat gun filament burns out once every few months, and people don't always alert others to this situation.

- Monitor the gas recovery rate during the transfer. Once the liquid starts transferring, it will drop down as low as 1 CFM but then gradually increase over the course of the transfer, reaching ~ 10 CFM by the end. If you see a boil-off rate

Figure E.6. Heat gun used to keep the gas meter from freezing.

consistently above 11-12 CFM, though, this means you have probably over-filled the dewar.

- As the transfer is proceeding, you can enter the liquefier dewar level and the gas meter reading into the Google Docs spreadsheet to estimate the liquid level in your storage dewar. If the liquefier is running during your transfer, the spreadsheet estimate will be $\sim$ 10-12L low at the end of a 100 L transfer. This is of course because the liquefier was making liquid during your transfer.

E.3 Ending the Transfer

- Close the valve on top of the transfer tube.
- Put on pair of gloves. You only have to burn yourself once with helium before remembering this step becomes easy.
- Start pulling the transfer tube out of the liquefier and your storage dewar. Raise it about 6 inches at a time and don't waste time. If you take too long, the quick-connects will freeze up more, and it will be more difficult to pull the transfer tube out. When pulling the tube out of the storage dewar, make sure to hold on to the top half of the brass quick-connect so that you don't pull it out of the storage dewar (pulling the quick connect apart is a good way to give yourself frostbite).
- Once you get the liquefier end of the transfer tube about 2/3 of the way out, pull the tube all the way out of the storage dewar. Try to do this slowly and after each time it moves a little bit try (gently!) closing the valve on the top of your storage dewar. Be careful not to smash the transfer tube in the valve. Once you get the tube out, hang it on the hook. During all this, try to avoid bending the liquefier end of the transfer tube any more than it already is.
- Pull the liquefier end of the transfer tube all the way up. If it gets stuck, push away from yourself (towards the tool bench) gently while simultaneously pulling up with one hand and pulling on the rope with your other hand.
- Once the transfer tube is all the way up, close the valve on top of the big dewar
- Slightly open the liquefier's recovery valve if the liquefier is not running. Open it just enough to hear some hissing ($\sim$ 1 CFM on the Labview page if nothing else is hooked up to the meter). You will probably have to turn off the heat gun temporarily to hear this. If the liquefier is running, do NOT open the recovery valve. The tube between the liquefier proper and the big dewar carries liquid

to the dewar and gas back to the liquefier. Opening the recovery line while the liquefier is running will cause problems for the liquefier.

- Measure the liquid level in your storage dewar with the thumper. It should be pretty easy to find the top of the liquid level, but don't waste time with the dipstick far down in the dewar as the nitrile glove membrane will freeze up. If you can't feel the change in vibration frequency, just pull the thumper out, let it thaw out, and try again. Getting this measurement right is important for all the helium accounting, so don't guess.
- Adjust the starting gas meter number in the Google Docs spreadsheet to get the final storage dewar reading to match what you actually measured.
- Update Keith's excel spreadsheet like you did when you started. Make sure to mark down how much helium you actually added to your storage dewar.
- Put the level meter back in sample/hold mode by pressing < MENU >, < ENTER>, < MENU > again. Turn the meter off if the liquefier is not running.
- Thaw out the rubber hose on your storage dewar.
- Close both valves on the regulator if you had to pressurize the big dewar.
- Take your dewar back to the lab, hook it up to the recovery line, and just slightly open the recovery valve so that the pressure is released slowly as the dewar settles down.

E.4 Final Checks

After you take your dewar back to the lab, go back to the liquefier and double check the following:

- Valve on top of the dewar closed
- Recovery valve opened (ONLY IF THE LIQUEFIER IS NOT RUNNING)

- Valves closed on gas cylinder
- Pressure building valve on big dewar closed
- Level meter in sample/hold mode if liquefier is running, shut off if liquefier is not running
- Heat gun is off
- Both spreadsheets updated and correct
- If you unhooked any dewars from the recovery line, make sure they are hooked back up with their recovery valve(s) open

F. ^{3}He Fridge Standard Operating Procedure

F.1 Sample Loading

Assuming that the probe is all the way up and warm, use the following procedure for loading your new sample on to the probe.

- Make sure that the gate valve is shut (see figure F.1a).
- Close the Speedi-valve to the probe (see figure F.1c) and disconnect the probe pump-out line.
- Double check that the gate valve is shut.
- Open the lower Nupro valve (see figure F.1d) to the sample space.
- Break the KF-50 seal *above* the gate valve while keeping one hand on the load lock to steady it.
- Use the winch to pull the probe up an inch or two as shown in figure F.2a.
- Remove the green clamp (figure F.2b) and pull the load lock up a couple inches to reveal the sample holder.
- Change the sample (pin 1 is marked with the red paint). Be careful not to touch the wires underneath the teflon tape shown in figure F.2c.
- Check that all four sets of set screws, the aluminum stand-offs, and both sets of brass nuts shown in figure F.2d are all tight. If they are loose, tighten them while being very careful to not touch any of the wiring.

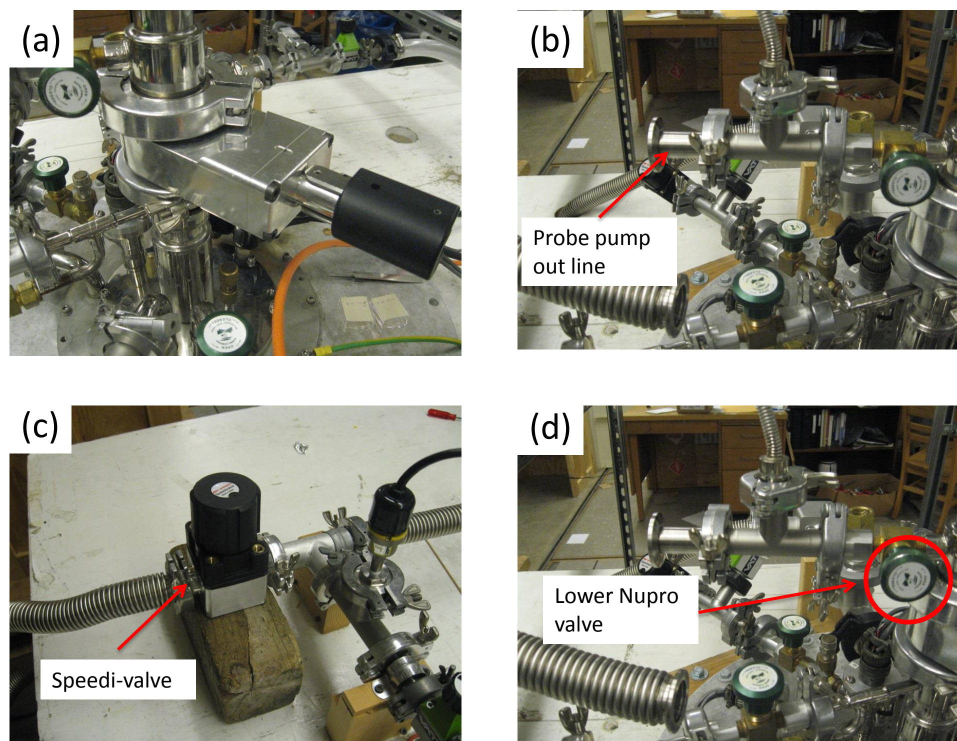

Figure F.1. (a) Gate valve. (b) Probe pump-out line. (c) Speedi-valve. (d) Lower Nupro valve (circled in red).

- Check that the heat-shrink tubing around the thermometer wires are not sticking out past the edge of the sample mount. If they are sticking out, very gently push them back so that they don't get caught anywhere inside the fridge.
- Slide the load lock back down all the way and put the green clamp back on the probe.
- Lower the probe back down to the KF-50 flange with the winch. Be careful to steady the probe with one hand while the other hand operates the winch.
- Clamp the probe back on top of the gate valve with the KF-50 clamp.
- Hook the probe pump-out line back up.

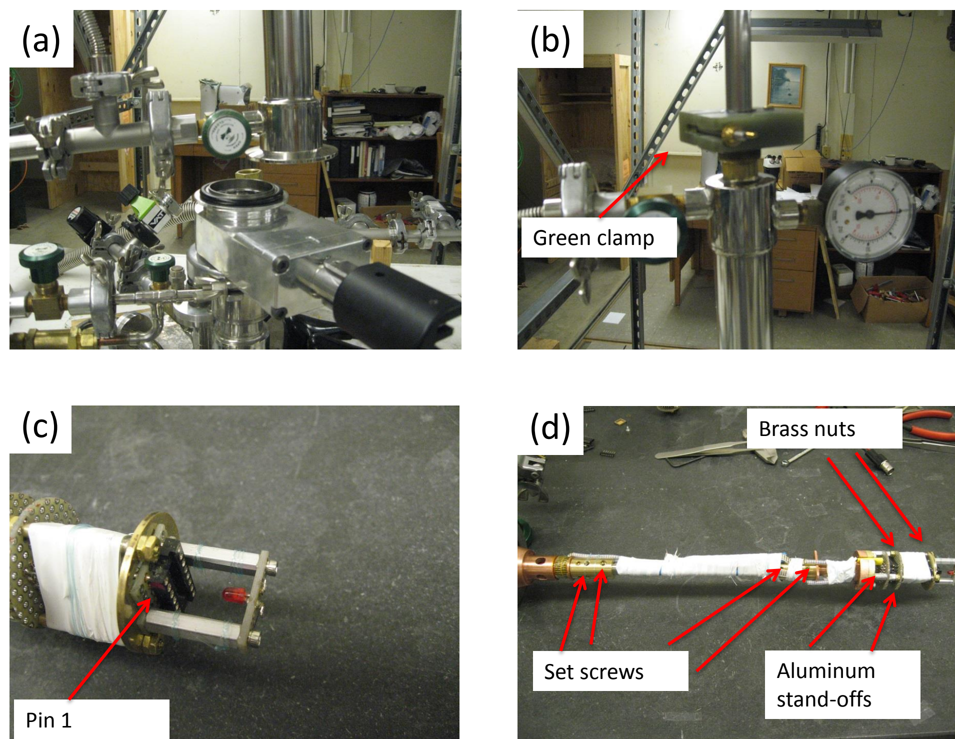

Figure F.2. (a) Probe suspended above the gate valve by the winch. (b) Probe clamp. (c) Sample mount. (d) Tail. Make sure all set screws, stand-offs, and nuts are secure.

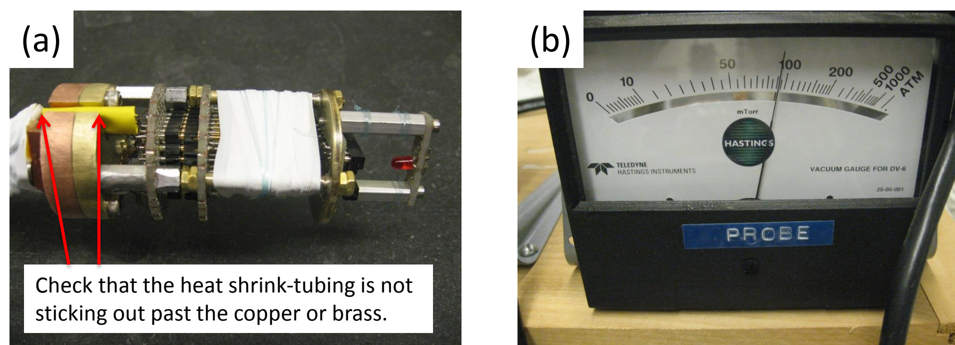

Figure F.3. (a) Check that the heat-shrink tubing around the thermometer leads is not sticking out past the copper or brass. (b) Wait for the probe pressure to fall to $\sim$ 20 mTorr before proceeding.

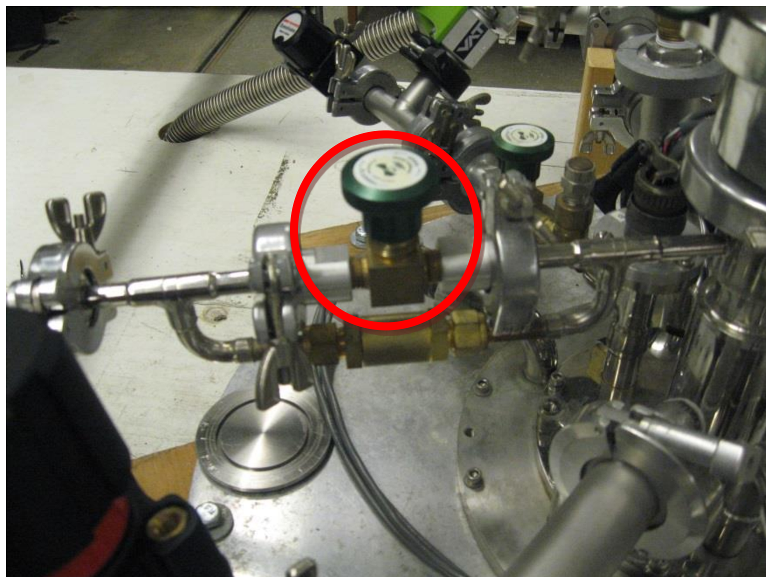

Figure F.4. Close the valve between the dump and the fridge.

- Open the Speedi-valve to pump out the probe. *Do not open the gate valve while pumping out the load lock.*
- Wait for the probe pressure to read ~ 20 mTorr. Check that all your contacts work while you are waiting.
- Once the probe is pumped out, close the lower Nupro valve to isolate the load lock from the pump.

F.2 Sample Cool Down

- *Double check that the lower Nupro valve is closed and that the upper Nupro valve (i.e. the sliding seal valve) is open.*
- Ground the sample.

- Open the gate valve and watch that neither the probe pressure mechanical gauge nor the probe pressure electronic gauge show any pressure spikes. If they do, close the gate valve back up and ask for help.

- Open the needle valve to get the 1K pot pressure $\sim$ 7 Torr.

- Close down the valve between the ^{3}He insert and the dump (see figure F.4) so that the ^{3}He all stays inside the fridge while you are doing the condensing.

- Set the charcoal setpoint to 15 K and set the heater output range to 625 mW.

- Lower the probe $\sim$ 6 inches.

- The probe must always stay lubricated with vacuum grease. If you cannot feel any grease on the probe, add a thin layer of grease to the dry section of the probe. Wear a glove while you do this and remove the greasy glove before you spread the grease to everything else in the lab.

- Watch that the sliding seal pressure (measured by the "probe pressure" electronic gauge) does not increase at all while you lower the probe.

- Fill out the cool-down log sheet as you go.

- Once the charcoal reaches $\sim$ 14 K, set the output range to 6.2 W and set the setpoint to 20K. Keep slowly increasing the setpoint until the charcoal is at 30 K.

- Keep lowering the probe a couple inches at a time. Aim for getting the probe all the way down over the course of about 1 hour.

- Try to keep the 1K pot at "T-UNDR" with the lowest possible 1K pot pressure. The 1K pot pressure will increase as there is an increasing heat load on it, but if you started at 7 Torr with no heat load, it should not be necessary to open the needle valve if you increase the charcoal temperature slowly enough.

- While you are still lowering the probe, the ^{3}He pot RuO should reach > 17.2 kOhm. The maximum resistance it will reach will depend on the helium level in the dewar (refer to previous log sheets for what numbers to expect).
- Once the probe is all the way down, tighten the c-clamps on top of the probe to compress the thermal anchor into the 1K pot.
- Turn off the charcoal heater.
- The sample should cool to 300 mK in $\sim$ 20 minutes. Check the 2-terminal resistance of your contacts while the samples finish cooling.
- Record any unusual behavior (e.g. scraping, rises in probe pressure, difficulty in sliding the probe through the sliding seal, etc.) in the log book and e-mail this information to everyone else who works with the fridge.
- At the end of your measurements, record in the log book how long the sample was at 300 mK. Normally you should be able to finish your measurements before the ^{3}He all boils off, so you should record something like "Hold time > 5 hours."

F.3 Sample Warm Up

- Ground the sample.
- Turn off the magnet switch heater.
- Make sure the charcoal is at 4 K and the 1K pot reads "T-UNDR".
- Remove the c-clamps from the probe and start raising the probe up. You should be able to raise the probe $\sim$ 30 inches the first time.
- Take the slack out of the rope with the winch, but do not try to pull the probe up out of the fridge with the winch.
- Wait 5-10 minutes for the probe/sliding seal to warm up.

- Pull the probe up another few inches; stop when the probe pressure starts to rise above $\sim$ 30 mTorr, the probe starts to frost, or the sound of the probe pump changes.
- Never force the probe up if it feels stuck. If the probe does feel stuck, lower it a little bit and then try pulling it back up. If this does not fix the problem, find someone to help you diagnose the problem.
- Once the probe is all the way up, close the gate valve. Make sure the probe is really all the way up and did not slide down through the green clamp while you were getting down from the table.
- Tighten down the needle valve to seal off the flow of helium into the 1K pot.
- Open the valve between the dump and the fridge
- **Before leaving, double check the following:**
    1. **Needle valve tightened down.**
    2. **Switch heater off.**
    3. **Gate valve closed.**
    4. **Valve from fridge to dump open.**
    5. **Sample database updated and closed.**

F.4 Changing the Sample Mount

The new probe design allows us to measure large samples mounted on DIP headers as well as wire-bonded devices mounted on LCC chip carriers[1]. Due to the difficulty in wiring the probe and sample mounts, great care should be taken when changing out the sample mounts.

[1]Spectrum Semiconductor part number LCC03201, www.spectrum-semi.com

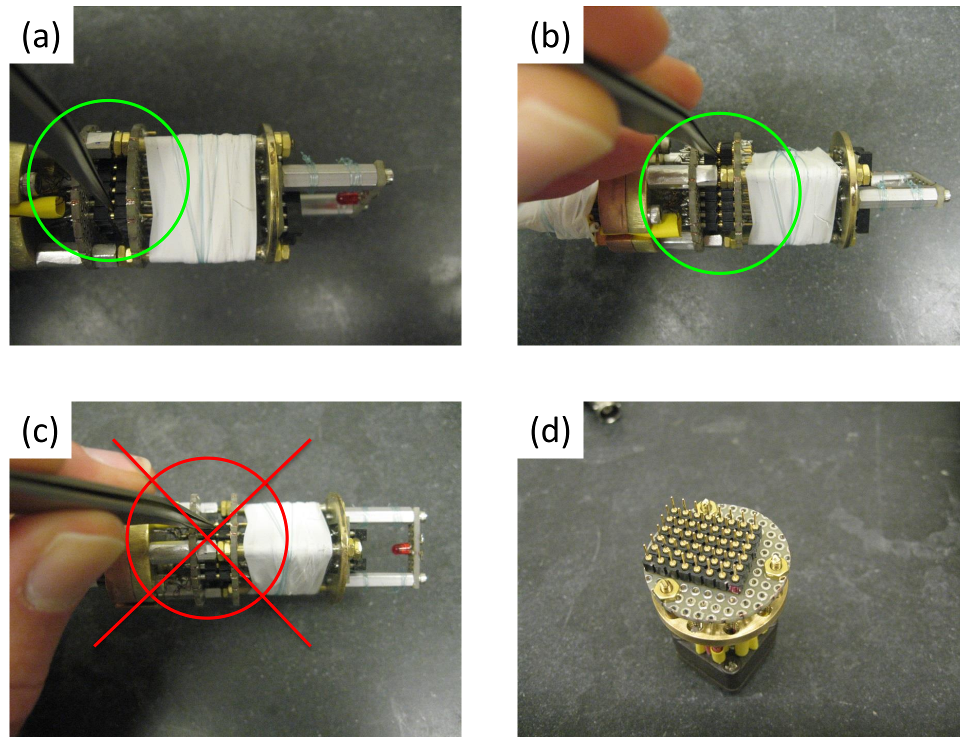

Figure F.5. (a) Start by prying the connectors apart with the tweezers sitting across multiple rows of connectors. (b) Move to the other end of the sample mount and pry the connectors apart a little more. (c) Do not try to pry the connectors apart with the tweezers parallel to rows of connectors. (d) Check that none of the pins are bent before trying to install the sample mount on the probe.

- Start gently prying the black connectors apart. Insert a sharp pair of tweezers across the rows of connectors and pry the connectors apart just a little bit as shown in figure F.5a.

- Next, pry the connectors apart at the other end of the column of connectors as shown in figure F.5b.

- Next, continue prying the connectors apart on the other side of the sample mount.

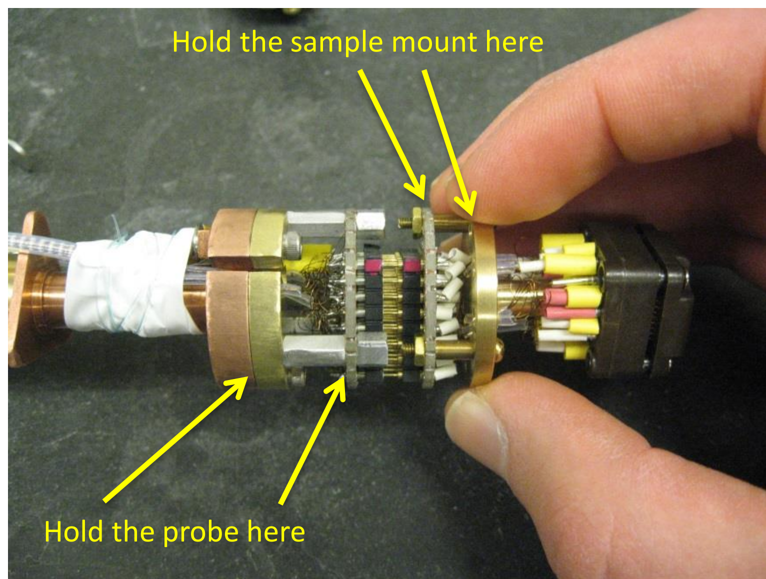

Figure F.6. Support the sample mount and probe by holding either the circuit boards or metal disks. Do not touch the wires.

- Do no insert the tweeezers parallel to each row of connectors as shown in figure F.5c. This will put unnecessary stress on the connectors.
- Keep slowly working the connectors apart with the tweezers until the sample mount is completely disconnected. Do not try to pull the sample mount off by hand.
- Check that all the pins on the sample mount that you plan to install are straight as shown in figure F.5d.
- Line up the red marks on the mating connectors (the red mark denotes pin 1).
- Very carefully push the pins on the sample mount into the sockets on the probe. Hold on to the sample mount and probe by the circuit boards or copper/brass disks as shown in figure F.6. Be extremely careful to not grab the wiring or slip and bend any of the pins.

- Inspect the probe to make sure nothing is loose or sticking out past the edge of the circuit boards or metal disks.

G. Device Fabrication Standard Operating Procedure

This appendix will cover the recipes I've used for processing devices in the clean room for a number of different projects. Some points will be unique to the in-situ back-gated devices, but for the most part the process flow would be the same for fabricating other types of devices. I will also try to include what I've learned along the way about more general ideas like mask design and process development in the hope that the reader may not have as many wasted process runs as I did early on in my time in the clean room.

Sumit Mondal and I were the first members of the Manfra group to start processing in the Birck clean room, so we learned a lot about processing from a number of theses from other groups. I would particularly recommend the theses of Jeff Miller[164] and Doug McClure[165] from the Marcus lab at Harvard as these are the most detailed, but I would also recommend the reader look through other theses [166, 167, 111, 168, 169, 170, 171, 172, 173, 174, 175, 176] to get a sense of what other groups do. Reference [174] in particular has a lot of detail on the microscopic details of Ohmic contacts which was helpful in understanding some of our rather bizarre initial results (more on this later). There is, of course, a lot of variation among all these theses in the specific details of the recipes, so it is useful to see what recipe steps can be tweaked and still give good results.

G.1 Initial Preparations

At the risk of being pedantic, start out by choosing which wafer you are going to process. If you are already well into a project, this won't really be an issue as you will know which wafers are working well for your project and which ones aren't; but if you are, for example, doing your first processing run in the clean room you will need

to find a less-than-stellar wafer to practice with. Make sure you look at the transport that was taken in the ^{3}He system to be sure that the transport is homogeneous and isotropic. Some of the mediocre doping well samples that we've grown over the years show anisotropic transport. In other words, at 300 mK R_{xx} measured in one direction may not get above 20 Ω even in the insulating states while the resistance at $\nu = 5/2$ may be 150 Ω in the perpendicular direction. If the transport was not recorded in both directions for a particular wafer, don't bother trying to process it. Your processing may be fine, but if the transport comes out looking anisotropic you won't know if this was due to damage during processing or was intrinsic to the wafer.

Once you've decided on a wafer to use, make a note in the sample database in the "Collaborator Records" section about what material you are using (e.g. "Top R 1/4") and what your goal is. Fold up half a small kimwipe and stick this on top of your wafer in a plastic sample box to keep it from rattling around when you take it over to Birck. When you take the sample into the gowning room, label the box with the label maker (no masking tape allowed in the clean room), and replace your dusty kimwipe padding with a strip of one of the clean room wipes. Ira keeps a pair of scissors on the top shelf of the rack with all the new clean room garments; if you borrow his scissors, make sure to put them back when you are done.

G.2 Tools and Tool Preparation

You should assume that all beakers, jars, tweezers, petri dishes, bottles for photoresist, etc. that you bring in to the clean room are dirty, so you will need to clean them before you start using them with your samples. At a minimum, you should clean everything with solvents in the sonicator. If you are cleaning a large beaker, just fill the beaker itself with the solvent and stick it in the sonicator. If you have smaller beakers, tweezers, or other small tools, use a larger (already clean) beaker to hold the solvent. Sonicate first with toluene to remove grease, sonicate with acetone to remove the toluene, sonicate with IPA or methanol to remove the acetone, and

finally sonicate in DI water to remove the remaining solvent residues. If you are cleaning a beaker on its own (i.e. not inside a larger beaker), wipe the outside with a solvent soaked wipe after each sonicating step since the outside of the beaker will also be dirty.

If the tool will survive a piranha etch, this will do the best job cleaning things up. Glass beakers, teflon tools, and stainless steel tweezers are all candidates for piranha cleaning. If you do this, though, you should first read up on proper safety when using piranha. Be careful to get rid of any organic contaminants and solvents before sticking something in a piranha bath. Never attempt to clean plastic in piranha. Be sure to rinse the tool(s) thoroughly (at least 5 minutes) in running DI after etching to make sure there is no acid residue that could interact with your processing chemicals later.

Perhaps not surprisingly, it is no small feat to get a small chip (sometimes as small as 4×4 mm) through several rounds of lithography, etching, metalization, etc. without scratching the surface of the chip, dropping a chip with wet photoresist face down, leaving it in a chemical for too long because you couldn't grab it easily, etc. Each mistake like this can potentially degrade the quality of your device, so you should try to make each processing step as dummy-proof as possible. Use shallow, wide beakers for wet etches, developers, etc. If you drop your chip, it will be a lot easier to fish it out of a beaker that is shallow and wide than one that is tall and narrow. The last thing you want to do is ruin your last good piece of wafer because you over-etched it by mistake.

In addition, don't be cheap when choosing your tweezers. The "economy" tweezers may claim to be a great value, but the tips will not close as uniformly and as a result you will drop a lot more samples than if you get the "high precision" model. I personally like the plastic "carbo-fib" tweezers from Techni-Tool[1] since the soft tips are less likely to scratch the chip and they are relatively high precision, though note

[1] www.techni-tool.com catalog #758TW0304

that these tweezers are not suitable for HCl or photoresist developer. For processes involving HF acid, get a pair of teflon-coated stainless steel tweezers[2].

Furthermore, you need a reliable way to carry your samples around and protect them from light exposure, rogue solvent squirt bottles, other people bumping your storage box, etc. My personal approach is to take a 4 inch petri dish[3] and cover the lid with foil to keep the light out. The samples don't slide around too much on the glass, it is easy to pick the samples up and set them down on glass (as compared to crinkly foil), and you can clean your petri dish with solvents each time you clean your samples to keep the dish itself clean.

Most of your other supplies can be obtained from Ira in the Birck stockroom. For your photoresist bottle, use the wide-mouth container with the hard plastic lid and not the narrow-mouth jar with the built-in dropper. The rubber on the dropper can dry and crack, potentially leaving rubber particles in your photoresist. Whenever you run out of photoresist, get a new bottle from Ira. It's not worth the risk of having old, crusty photoresist contaminate your fresh resist. Whenever you fill up your small photoresist bottle, label it with the expiration date so you know when to replace the resist in case it expires before you use it all up. Each time you use the resist, wipe the mouth of the bottle with a clean wipe to try to prevent resist from building up on the lip of the jar; this build-up could turn into crusty flakes that will mess up your lithography.

G.3 Ga Removal

The first step in processing a sample is to clean the gallium off the back of the wafer. The staff don't want gallium getting on any of their equipment (particularly the evaporators and the RTA), so this step is important. Surprisingly, this step is not mentioned in any of the theses I listed in the beginning section of this appendix. Evidently, staff at other universities are not as paranoid about the presence of gallium

[2]www.tdiinternational.com catalog #TDI-2A-SATCE5

[3]These can be purchased from the chemisty stockroom

in their equipment, or the students just never mentioned to the staff that the samples had gallium on the backside. Regardless of the reason, we had to develop this process on our own.

Start by blowing your sample off. Everything in the physics building (including our sample prep room) is really dusty, so try to get any large dust bunnies off before you do anything else. Next you need to coat your chip with resist to protect the epilayers while you clean the gallium off. Enable the spinner in Coral and then start setting up. First, line the black bowl with foil to make clean-up easier at the end. Lay out two clean wipes in the hood (one for setting your stuff on, the other for cleaning resist off your tweezers). Assume all wipes left in the hood by previous users are dirty. You never know what chemical residues may be left on a wipe, so just get yourself some new ones straight from the bag. You can use the ones other people left behind for clean-up at the end. Go to the stock room (it is the door immediately to your right when you walk out of the air shower from the gowning room) and grab yourself a pipette, a bulb, and a glass slide. Set the pipette and bulb on one of your clean wipes and don't let the tip touch the dirty surface of the hood anywhere. Wipe the glass slide off with a clean wipe and some IPA or methanol, go set it on one of the hot plates in the lithography bay, and set the hot plate to 100C. Try to use the same hot plate whenever possible as each of the three hot plates has a slightly different offset from its setpoint. Get a small piece of foil and make a little ridge on it perpendicular to the length of the pipette you set on it so that the tip of the pipette is not touching the foil. Find a chuck that has a pedestal on it smaller than your sample and mount it on the spinner. Set the spinner to ramp up to 4000 RPM over 4s, sit there for 40s, and then spin down over 4s. Get a junk piece of GaAs and try spinning it. I would recommend having a stash of junk chips of various sizes ranging from 4×4 mm up to a quarter wafer so that you have test pieces roughly the same size as any of the samples you may potentially spin. The spinners get abused a lot, so you always need to test that the vacuum is good enough to hold your sample on before you fling your real sample into the jagged foil. Make sure the junk chip stays on for the entire 40s,

and then try mounting it a little off center, try abruptly stopping the spinner, etc. to make sure the vacuum is really good.

If you get the vacuum interlock error, try pushing the chuck on to the spindle more. If this doesn't work, take the chuck off the spinner and check for flakes of the white plastic insert on the o-ring inside the chuck. If there are flakes, try to clean them out with a q-tip. If the vacuum is ok but the door interlock prevents the spinner from starting, fold up a small piece of foil and set it on top of the lid interlock switch to cause the lid to push it down a little more. If this fixes it, report the problem in Coral so that the staff can adjust the sensitivity of the switch so you don't have to do the foil trick in the future.

Once you are convinced that the vacuum is good, do a test spin on your real sample. The last thing you want is to get your sample coated with a bunch of photoresist and then not be able to start the spinner because the vacuum is not good between the back of your sample and the chuck. You can cheat a little and stop the spinner once your sample starts spinning, but if you do this make sure you stop it in the first second or so of spinning so that it doesn't get up to full speed. If the vacuum is not really good, the sudden stop can sometimes fling your chip off the chuck. Once the spinner appears to be working, cover the surface of your chip with AZ1518 resist (remember to blow the bubbles out of the resist somewhere other than on your sample) and start the spinner. If you have additional chips, spin the resist on them now.

Bake your samples for 2 minutes on the hot plate at 100C. Start the timer as soon as you set the first one on the glass slide and space out subsequent samples $\sim$ 10s. When the 2 minutes are up, pick the samples up in the same order with the same $\sim$ 10s pause in between each one to ensure they all get baked for the same duration. Don't set your samples straight on the hot plate since the hot plates tends to have a lot of photoresist, epoxy, and other mystery goop baked on to them.

Once your samples are done baking, turn off the hot plate and let it cool down to $\sim$ 50C, or if another hot plate is free set it to 50C. Set a clean wipe on the 50C hot

Figure G.1. Gallium-covered backside of a wafer.

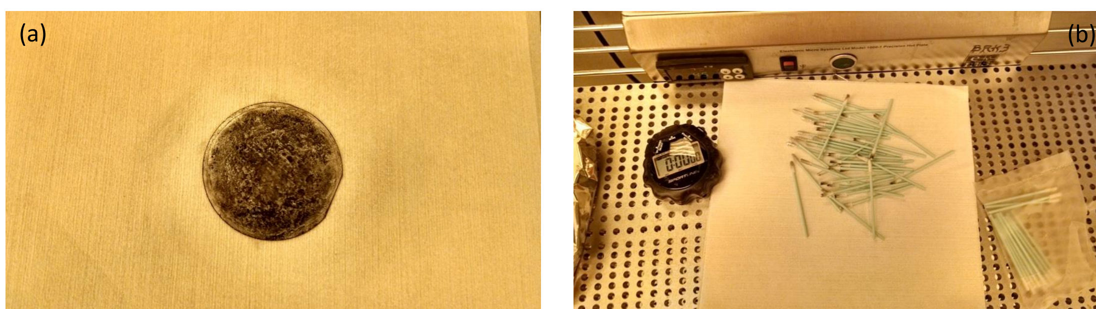

Figure G.2. (a) Backside of wafer after gallium has been wiped off. (b) Approximate number of q-tips necessary to wipe all the gallium off a full wafer.

plate and set your sample face down on the wipe as shown in figure G.1; this will keep the gallium from freezing. Get a bunch of q-tips and wipe the gallium off the back. You should be able to wipe 99% of the gallium off this way (see figure G.2). Be very careful, though, to not slide your sample onto a part of the wipe with gallium on it. In addition, try to slide your sample around as little as possible to avoid damaging the protective resist layer.

Once you have wiped all the gallium off, spin and bake a second layer of AZ1518 resist onto your sample as before. This will fill in any holes you may have made in the resist while wiping the gallium off. Clean up the spinner and disable it in Coral. Next, etch your sample resist-side-down in full strength HCl for 3 minutes to remove

any remaining gallium residue. You can use stainless steel tweezers to pick up your samples, but don't leave the tweezers in the acid for very long as the HCl will attack the steel and give off a yellow/green cloud of junk. Rinse your tweezers as soon as you take them out of the acid to prevent them from corroding while your sample etches. I use a dedicated beaker for the HCl etch to avoid any potential cross-contamination with other process steps. Have a rinse beaker filled and ready so that when the 3 minutes are up you can transfer your sample straight to the rinse. Let it sit in the rinse water for $\sim$ 30s and then rinse it with running water for a few seconds and blow it dry. Once all your samples are etched, rinse your beakers and tweezers thoroughly, blow dry, and clean up your mess.

G.4 Cleave and Initial Clean

Now you need to decide how to cleave up your samples for your process run. I would recommend that you process several samples through the Ohmic deposition and annealing since the mesa etch and ohmic deposition takes a significant amount of time but should have a relatively high yield. After the Ohmic deposition, you can either cleave out individual chips to anneal at separate temperatures, or you can anneal them all at the same temperature and then process them one at a time through subsequent steps (like e-beam patterning of depletion gates).

At this point you will need to consider your remaining processing steps before you cleave up your chips. Figure G.3 shows a sketch of the wire-bondable, 16-pin DIP chip carriers that we use[4]. The orange lines represent the chip carrier bond pads and the green rectangles represent the chip. I have found that designing my masks such that my chips are 4.2 $\times$ 5.5 mm after cleaving generally results in my chips fitting in the cavity of the chip carrier. From the drawing, though, it is obvious that this doesn't leave much room for error when cleaving, so it is best to make the scratch to define the cleave line under a microscope or magnifying glass. If you process a

[4]Purchased from Spectrum Semiconductor Materials, Inc. www.spectrum-semi.com catalog #CSB01648, manufacturer drawing IDK16F1-390GAL.

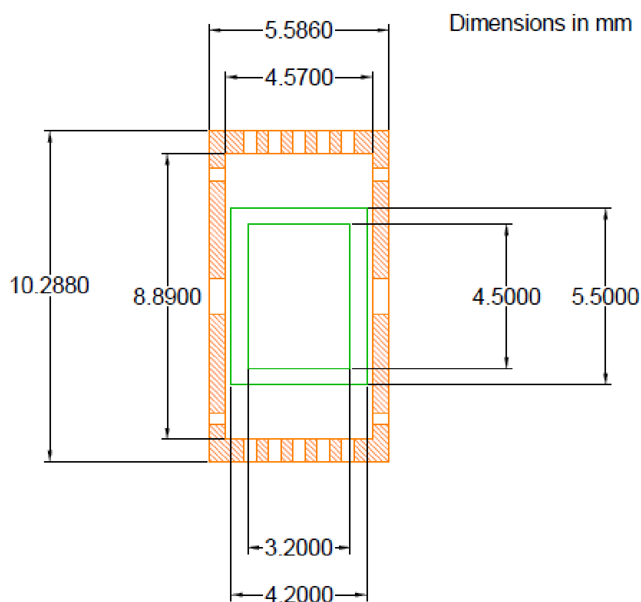

Figure G.3. Sketch of wire-bondable chip carrier and chip. The orange lines represent the chip carrier with the cross-hatched region representing the gold bond pad. The green rectangles represent the chip. The larger rectangle is the border of the chip and the smaller rectangle represents the area on the chip that will be free of significant photoresist edge-bead. All dimensions are in mm.

chip this small at any point during your fabrication, you will need to leave a ~ 500 μm border around the edge of the chip so that the photoresist edge bead does not interfere with your lithography. If you know for sure that you will not be processing any individual chips you could make your chips smaller and then cleave each chip out when you are completely done, but this will just depend on your process. The main idea here is that you need to have a clear plan at this point for how you will process your chips so that the cleaving step doesn't make it impossible to finish the fabrication.

Once you have thought through your processing plan, cleave out your chips. I keep a chip in my clean room storage box that is known to fit in the chip carriers,

and I use this chip as a template for cleaving. Make a single scratch with the scribe (kept in the Manfra group storage box), grab the chip on each side of the scratch mark with your tweezers, and cleave the wafer. Some people prefer to set the wafer face down on a wipe and then push on the backside of the wafer with their tweezers, but I don't find this method to be precise enough for my purposes. Once you cleave out your pieces, draw a picture in your notebook of where each chip came from on the wafer, what its orientation is with respect to the major flats, and try to come up with a way to tell each chip apart if you can. If the chips are all identical, you will just have to be careful for now to not mix them up. Once you start processing them further, the unique goobers on each chip will help you identify each chip if you mix them up accidentally.

Now you need to clean the photoresist off your chip. This particular cleaning step seems to have the largest impact on how clean your chip will end up looking after subsequent process steps, so you need to be particularly aggressive in cleaning at this point. Set up the acetone airbrush and have a beaker filled with acetone ready along with a dirty beaker. Use a squirt bottle to rinse the resist off into the dirty beaker; the goal here is to rinse all the junk off so that if you got any gallium on the resist while you were wiping it off the back of the chip, it doesn't have a chance to settle on the wafer surface. Start spraying the chip down with the acetone airbrush immediately after you rinse off the resist so that the acetone can't dry on the surface. The acetone airbrush will remove a lot of the more stubborn junk. Finally, stick the sample in the clean acetone beaker when you are done spraying it down with the airbrush. Once you have all your chips sprayed down with the airbrush, squirt them down with some methanol or IPA (don't let the acetone dry on the surface) and blow them dry with the nitrogen gun. Inspect your samples under the micrscope and use the Nomarski phase contrast filters as these will allow you to see a lot more debris/damage on your chip than a standard bright field or dark field image will. If everything looks reasonable, start the standard sonicated solvent clean.

Figure G.4. Sonicator setup used to clean the samples.

Sonicate your samples for 5 minutes each in toluene, then acetone, then methanol or IPA, and then DI water. The toluene will remove organic contaminants, the acetone will remove the toluene and any residual photoresist, the methanol/IPA will remove the acetone, and the water will remove the methanol/IPA. You need to be aware of a few things when using the sonicators. First, the sonicators themselves are really dirty. If the water in the bath looks really bad, dump it out, rinse out the sonicator, and fill it back up with clean water. Even if you do this, though, you should try to avoid getting any water from the bath in your beakers. The water tends to splash more if the tank is not filled up to the fill line. Filling the tank all the way to the fill line, though, means that your beaker will probably sink since the tray insert is so deep. To get around this I wad up some aluminum foil to prop up the tray so that the tray is only $\sim 1/2$ inch below the water level. You should also cover your beakers with foil, and I like to also always have two of the shallow 3 inch beakers in the sonicator at all times (see figure G.4). Having two beakers covered with foil in the sonicator will eliminate all chances of water splashing into your beakers. Water can still creep up the side of the beaker, though, if your two beakers touch each other or touch the walls of the tray, so try to avoid this.

While sonicating your samples, you also need to be careful to not allow your chips to flip over or to jump on top of each other to prevent your chips from getting damaged. One way to avoid this is to only have one chip in a beaker at a time, but

Figure G.5. Homemade teflon inserts used to sonicate multiple samples in parallel.

this will make each cleaning step take a really long time if you are processing multiple chips in parallel. To get around this, I made some teflon inserts for the beakers with slots in them for the chips (see figure G.5). This keeps the chips separated, aids in keeping track of the chips, and speeds up the cleaning steps. If you use these inserts or make similar ones, do not touch the inserts with your gloves as you will introduce a lot more junk to your solvent bath and make your samples dirtier than when you started. Only handle the inserts with clean tweezers.

Once the 5 minutes are up for a given sonication, squirt the chip down with methanol to remove the previous solvent and transfer the chip to the next bath. The only exception to this is after the methanol/IPA sonication. In this case, squirt the chip down with methanol or IPA, blow it dry, and then transfer it into the water bath. After the DI sonication, squirt the chip down with some running water and then blow it dry.

When you think the chips are all clean, take a picture of each one with the Nomarski filter and save them on your network drive. Figure G.6 shows a sample imaged with the Nomarski filter after the gallium removal and initial cleaning step. Dump the solvents from your beakers in the non-halogenated solven waste container, rinse out the inside of the beaker with a little methanol or IPA, wipe the outside of the beaker down with a methanol- or IPA-soaked wipe, and blow out the beaker with

Figure G.6. Surface of wafer after gallium removal and sonicated cleaning. Field of view is ∼ 2.5 mm wide.

nitrogen. Don't touch the inside of the beaker with the wipe or your gloves. Clean up your mess in the hood and move on to your next round of lithography.

G.5 Etching

Your first lithography step will most likely be an etch step, either for etching a mesa or vias to a buried gate. The AZ1518 resist is free and is robust enough for all the etching that I have ever done (as deep as $\sim 2\ \mu$m). You can use thinner S1805 resist if you are really concerned about the roughness on the edge of your mesa. The thinner S1805 resist ($\sim$ 500 nm as opposed to $\sim 1.8\ \mu$m thick AZ1518) allows the mask to get closer to the wafer surface and as a result gives sharper etch features as shown in figure G.7. Set up the spinner as described in the previous section. If you are using the AZ1518 resist, the spinning and baking parameters are the same as described before except that you only need to spin and bake a single layer of resist.

If you will use the S1805 resist, however, you should spin the resist at 5000 RPM for 40s and bake for 5 minutes at 80C. For historical reasons, I also do a 2 minute dehydration bake at 80C immediately prior to spinning the resist on the chip. I used to also do a 2 minute 100C dehydration bake before spinning the AZ1518, but I never

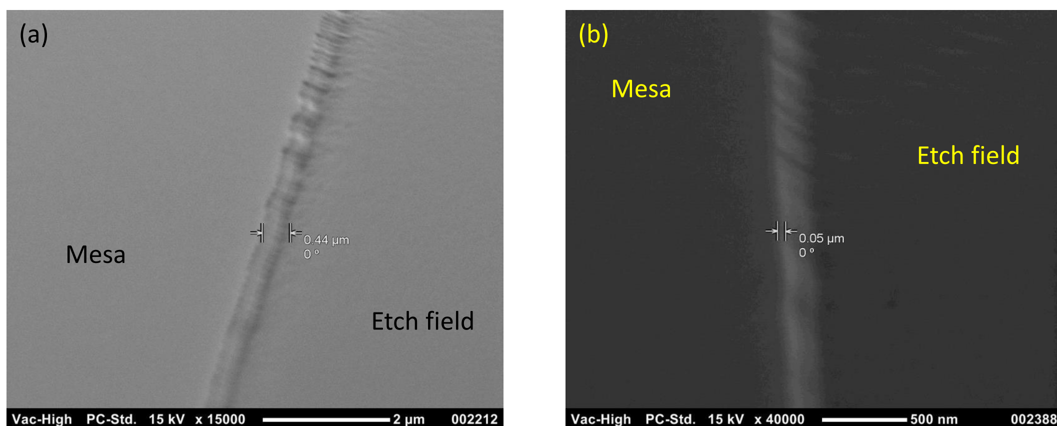

Figure G.7. Shallow-angle SEM view of mesa edge after etching using (a) AZ1518 resist and (b) S1805 resist. The thinner S1805 allows features to be defined more sharply and results in less edge roughness after etching. Plan-view SEM images (not shown) show that typical roughness in the plane of the original wafer surface is $\sim$ 150 nm for features defined with AZ1518 resist while the roughness drops to $\sim$ 50 nm for features defined with S1805 resist.

saw any noticeable difference in my devices, so I dropped this step. As a rule of thumb, the less you do to the chip the better. Each time you handle the chip you can potentially drop it, each time you heat it on a hot plate you bake junk onto the surface, and so on, so don't add steps if you can avoid them.

After the softbake, clean up your mess, disable the spinner in Coral, and enable the mask aligner. Both MJB3 aligners will work, but I have learned to avoid MJB3_1 because it is so hard to see anything through the optics. Start out by setting your exposure time and do a test exposure without a sample or mask in the aligner to check that the power supply is putting out 10 mW/cm^2 and that the timer is working correctly.

The self-leveling feature of the aligner will not really work with a small sample because the sample is not big enough to apply enough torque to the chuck to level it out, so you need to improvise. Cut out a strip of the white clean room tape and stick it on top of the chuck, making sure to cover up all the vacuum holes. Poke a

hole in the tape over the central vacuum hole so that your sample will stay stuck to the chuck and not stick to the mask. The tape is thick and soft enough to allow your chip to level out when it is brought into contact with the mask. This will cause the feature size of your resist to match that of the mask much more closely than if you didn't use the tape.

Next, load your mask (remember to put the chrome/iron oxide side down so it is in contact with the resist) and your sample. Turn the z-height knob clockwise several times so that when you bring the chuck into the "contact" position you don't smash your sample into the mask. The "separation" lever doesn't work very well for small samples with a lot of edge bead, so just leave the "separation" lever in the "contact" position and control the height of the chuck with the knob. Adjust the alignment of the chuck to align the edges of your sample with the mask. You can make this easier by designing in some large, straight features in your mask design. Figure G.8 shows the via and mesa designs from one of my masks. The mesa layer (black lines) has a bright field polarity while the via layer (purple lines) has a dark field polarity. The green rectangles are the same rectangles shown in figure G.3 which represent the chip and the edge bead. The long, skinny purple rectangles are used to help align the chip to the mask and center the region for the devices on the chip. All subsequent mask layers also have these long, skinny rectangles to aid in the gross alignment of the chip to the mask.

Once you have your sample aligned to the mask, bring the chip into contact with the mask by rotating the z-height adjustment knob counter-clockwise until you see the resist squash into the mask (this will probably happen primarily at the corners of the chip). Make sure that the "soft exposure" button is depressed and start the exposure. Expose 20s for AZ1518 resist but only expose for 6s with the S1805 resist.

After exposing all your chips, take your mask off the aligner, turn off the light for the microscope, turn off the aligner, and disable the tool in Coral. Since your mask now has photoresist smashed onto it, you need to clean it. Leave it soaking in a beaker of acetone for a couple hours while you continue your processing. Make

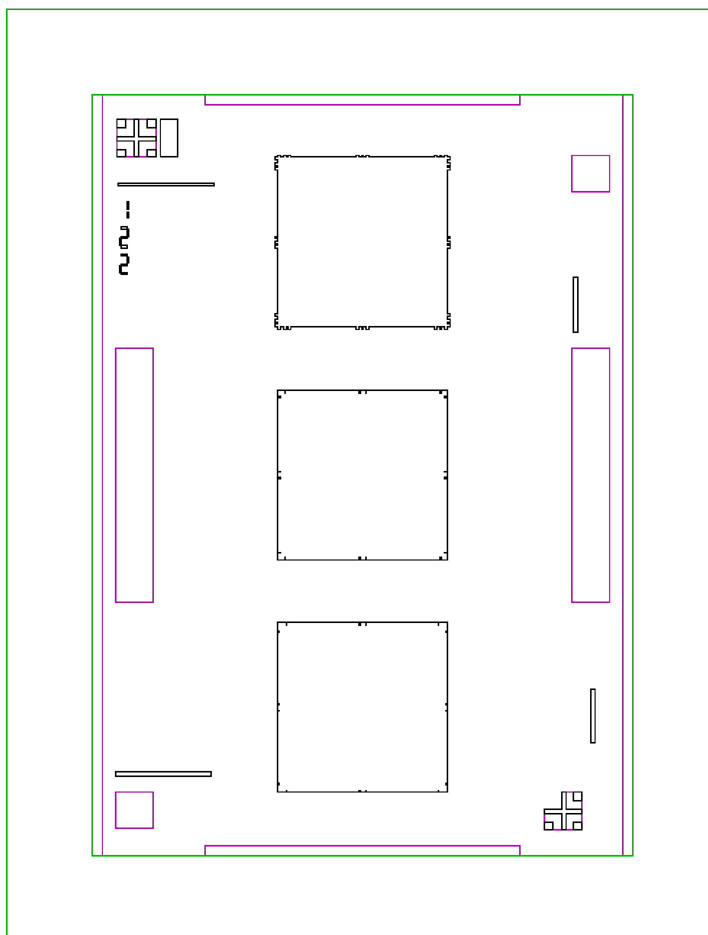

Figure G.8. Mask design showing via and mesa layers with large features for aligning the sample to the mask.

sure to cover the beaker with foil and push it to the back of the hood so no one else splashes any other chemicals into it. After the mask has soaked for a while, squirt it down with acetone followed by IPA, dump the beaker into the waste container, and then fill it back up with IPA and soak your mask in IPA for a couple minutes to make sure all the acetone is gone. Then squirt the mask down again with IPA and blow it dry quickly to try to avoid leaving any condensed water on any of the critical features. Always remember to set the mask chrome/iron oxide side up in the beaker so you don't scratch the chrome/iron oxide on the bottom of the beaker.

You should continue with the lithography while your mask is soaking in the acetone. The next step is developing the pattern. Use MF-26A developer (free and stocked in the cleanroom) for the AZ1518 resist and MF-319A for the S1805 resist (we have to buy this developer ourselves, so don't waste it). Set up two dedicated developer beakers in one of the lithography hoods. Do not use your developer beakers for other chemicals (including other developers); metal-ion containing developers and metal-free developers in particular should not be used in the same beakers as cross contamination in the part-per-million level can affect the development. Fill each beaker with the appropriate developer and also fill your rinse beaker with fresh water. You need to use two developer baths because the majority of the resist will come off in the first bath, but you will need the cleaner second bath to finish removing the resist to get the etch field really clean. Use the water to quench the developing process. For AZ1518/MF-26A, develop 20s in the first bath and 10s in the second bath. For S1805/MF-319A, develop 10s in the first bath and 6s in the second bath. Regardless of the developer, wave your sample around in the rinse water for $\sim$ 15s and then spray it down with some running water. If your developer bath starts to get pink, refill that bath with fresh developer. If you need a really clean etch field, you should probably change the developer out after ever 2-3 samples even if the developer still looks clear. After developing, checking the lithography under the microscope. Remember to use the UV filter so that if the developing is not finished you don't expose the resist any further with the bright white light from the microscope. If the developing does not look complete, develop another $\sim$ 10s for the AZ1518/MF-26A or another 5s for the S1805/MF-319A lithography and then check the lithography again. Once you are happy with the development, save a picture of each chip, rinse out your rinse beaker, and soak the samples in fresh DI water for 5 minutes. At the end of the soak, squirt the samples down with running water and blow them dry.

Next, you need to "de-scum" your samples to remove the last few nm of resist from the surface. Use the Branson asher to clean them with an Ar/O_2 plasma (note that the sample sits in a Faraday cage so there should not be many energetic ions

striking and potentially damaging the sample surface). Pump the chamber down to $\sim$ 110-130mTorr and then turn on the gas flow. The flow meters on this tool are pretty rough, so the units I list here should be considered arbitrary. Set the Ar flow rate to 120, and set the O_2 flow rate such that the top ball in the flow meter is $\sim$ 12 and the bottom ball is $\sim$ 5-6. It is difficult to adjust these flow rates very precisely, but it doesn't seem to matter much. The process pressure should be $\sim$ 1.3 Torr. Use 100W to generate the plasma (remember to check that there is no reflected power), and "de-scum" for 90s. At the end of the cleaning, shut off the RF power, turn off the gas flow, close the valve to the pump, and purge the chamber up to $\sim$ 200 Torr. Then close the purge valve and pump the chamber back down to $\sim$ 200 mTorr. Do this pump/purge process twice to make sure you pumped all the oxygen radicals out before you open up the chamber.

The AZ1518 resist is ready for etching at this point, but the S1805 resist should be baked another 5 minutes at 100C on the hot plate. I used to do additional baking of the AZ1518 as well, but once again I never saw this make much difference. I have never tried etching without doing this additional bake of the S1805, but I suspect it also would probably survive a short etch.

Once you finish baking the resist (if necessary) get setup in an acid hood to do your etching. Get yourself plenty of wipes so that you can dry your gloves off without having to go back and rifle through the bag; if you have acid on your gloves you want to keep that in the hood and not drip it all over the floor. Whenever you are working in the hood, treat any liquid you see as an acid. HF acid gets used a lot in all the hoods, so you should always be careful. Don't ever touch your face, goggles, or face shield with wet gloves, and always wipe up all liquid on the bench before you start and before you leave. Wipe your acid apron and gloves off before you take them off, and if someone else left liquid all over the apron or gloves, go get a different pair. Your white cleanroom gloves will not offer you much protection if you handle an acid-covered apron or glove. When you finish measuring out acids, wipe any drips off the bottle and rinse out the acid-soaked wipe.

We use a dilute phosophoric acid piranha, but other groups also substitute sulfuric acid or ammonium hydroxide for the phosphoric acid. For the appropriate ratio of acid or base to peroxide the etch will be fairly isotropic with the sidewall of your mesa sloping gradually. For larger ratios of peroxide to acid, the etch will become anisotropic, and this anisotropy can be used to determine the crystallographic orientation [170]. For etching a standard mesa, however, you want an isotropic etch so that your ohmics and gates can easily climb up over the edge of the mesa. Our etchant recipe was provided by Lisa Tracy and was/is used in the Eisenstein group at Caltech [111, 173, 172]. It consists of 50:5:1 water:phosphoric acid:hydrogen peroxide (30%). Measure out 500 mL of water with a volumetric flask (make sure it is calibrated "TD" - "To Dispense" and not "TC" - "To Contain"). Add 50 mL of acid with a volumetric pipette or a small graduated cylinder, and add 10 mL of peroxide with a volumetric pipette. Mix up the etchant thoroughly with the pipette and then rinse the pipette out thoroughly. Since a 500 mL beaker is rather large, you should pour some of this etchant into a smaller, shallow beaker so you can pick up your samples more easily at the end of the etch.

At one point, I was trying to track down the source of some funny looking transport and examined a couple of other etchants. Figure G.9 shows cross-sectional SEM views of mesas etched with different etchants. The ammonium hydroxide etchant recipe was provided to us by Bob Willett [177], and the phosphoric-based piranha etch shown in figure G.9b was taken from [176]. In all cases, the cross-sectional view in the orthogonal direction looks similar to the views shown in figure G.9. Since all three etchants gave similar looking edge profiles, we never changed our etch recipe.

Etch your test piece for 90s and then quench the etch in your rinse beaker for $\sim$ 30s; this should result in an etch depth $\sim$ 150 nm. If you need to etch farther in your real samples, scale the etching time of your test piece so that your test etch depth is comparable to your target etch depth in your real devices. Rinse with flowing water and blow the sample dry. Now take the test piece to a solvent hood and strip the resist for a minute or so in acetone. Rinse the chip with acetone followed by methanol

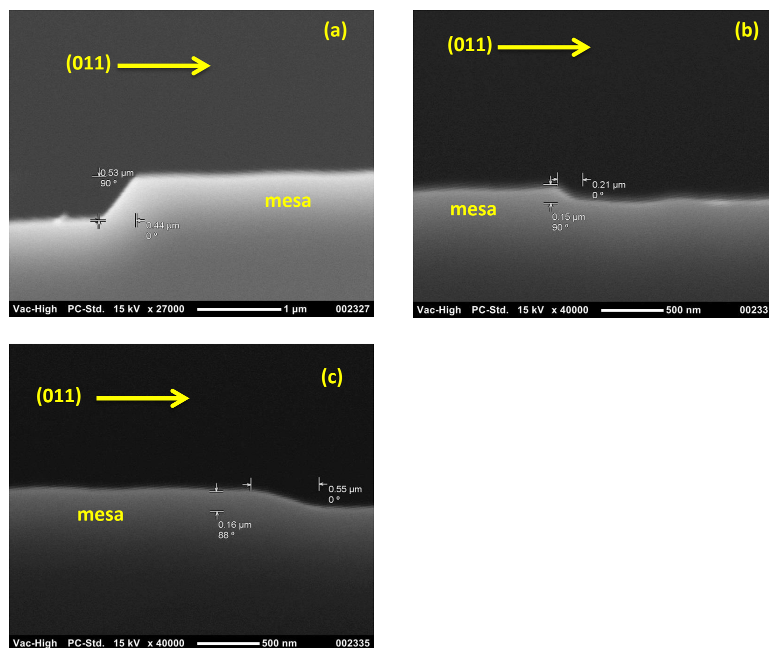

Figure G.9. Cross-sectional SEM views of different mesa etches. (a) Mesa etched with ammonium hydroxide piranha from [177]. (b) Mesa etched with phosphoric piranha from [176]. (c) Mesa etched with our standard phosphic piranha.

or IPA and blow dry. Enable the Bruker optical profileometer and measure the etch depth on your chip in a couple places. The only thing you need to change from the default measurement setup is to set the scan type to VXI and decrease the back-scan from 25 to 5 μm. Use the 10$\times$ objective. Try to get the chip as level as possible before you make the measurement; if you have to do a lot of leveling on the data, you may introduce more error.

Normally, you should use a blank piece of GaAs as your etch test piece to conserve your heterostructures. This will introduce a bit of a change in the etch rate, though. For etching our "standard structures" (doping well structures with x = 0.24 AlGaAs and a 200nm deep 2DEG) assume the etch rate will increase $\sim$ 15% from the test piece to your real chips. Once you have your corrected etch rate calculated, etch your real pieces all in one shot. For most wafers you only need to etch past the top

doping layer to electrically isolate the 2DEG. If you have a really low density 2DEG (i.e. really large dopant setbacks) you may need to etch all the way to the quantum well. Some of the students in the Marcus group at Copenhagen found in some of our low density wafers that adjacent mesas were only insulated from each other by a few kΩ when the etch stopped right below the top doping layer [178].

Once you etch your devices, strip the resist and measure the step height like you did with the test piece. Clean up your mess in the hood and disable the profileometer. Another approach to hitting your etch target is to measure the resist thickness prior to etching and then etch your real device in steps, measuring the height of the resist plus the etch depth each time [164, 165]. However, I have found that both methods result in similar accuracy in reaching the desired etch depth but that etching and measuring in steps tends to take longer. If you try this method, remember that you will have to use one of the stylus profileometers since the optical profileometer will not measure the depth of the resist plus etch step correctly.

Clean your samples again with the sonicated 3 solvent plus DI clean as before. If your first lithography step was etching vias to a buried gate, you can leave the samples overnight without any problems. If instead you etched a mesa, you should continue on to deposit the Ohmics so that the AlGaAs in the sidewall of the mesa doesn't oxidize too much. We and some of our collaborators [179] have found that the contact resistances are much lower when the Ohmics are deposited the same day the mesa is etched. If for some reason you have to leave an etched mesa overnight, cover the sample up with photoresist and softbake the resist to try to minimize the potential for oxidization.

G.6 Ohmic Contacts

The Ohmic contacts will be patterned by a lift-off process which means the photoresist must be thick enough to create a break in the metal. The S1805 resist is reportedly too thin for lift-off [176], so you will need to use the AZ1518 resist. Spin,

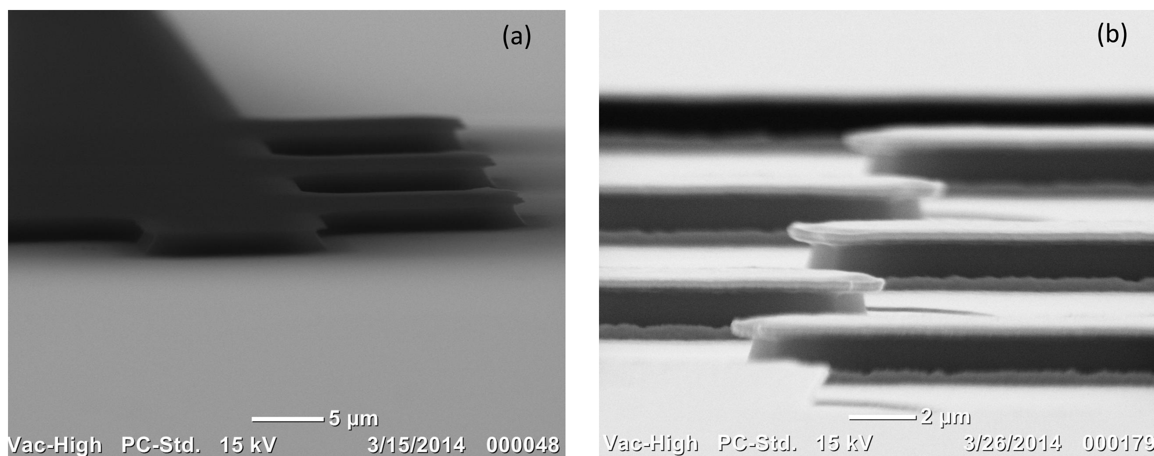

Figure G.10. Shallow angle SEM images of resist profile for lift-off. (a) AZ1518 resist after patterning and hardening with chlorobenzene. (b) AZ1518 resist after metalization but before lift-off. Light regions are metal and dark regions are resist.

bake, and expose the resist as previously described. Don't forget to clean your mask again. In terms of mask design, you may want to consider using transparent iron oxide for your mask rather than chrome. Since the Ohmic layer will have a dark field polarity you will not be able to see much of your chip through the mask unless you have a design with large Ohmics.

After exposing the resist, you need to harden the resist with chlorobenzene before developing. The chlorobenzene will create a hard layer on the surface of the resist that develops slowly. Once the developer eats through this hardened region, it will eat the resist underneath much more quickly and undercut the hardened layer. This will prevent the metal you deposit in the evaporator from creeping up the sides of the resist and forming a continuous film.

Soak your samples for 20 minutes in chlorobenzene in a dedicated beaker. Blow your samples dry and then soak them in DI water for 5 minutes to remove the solvent residue. Remember to dispose of the chlorobenzene in the halogenated solvent container. Blow the beaker out and let it and any solvent-soaked wipes air out in

the hood for a while for the fumes to dissipate. After rinsing your samples, blow them dry and then develop in MF-26A for 70s in the first bath and 20s in the second bath. Check the development under the microscope; the developing time required to fully develop the resist tends to vary more than when developing non-hardened resist. Once you are satisfied with the development, rinse in fresh DI water for 5 minutes and blow dry.

De-scum your samples as before except only expose the samples to the plasma for 15s. You have been working hard to put the Ohmics down the same day to minimize oxidization, so it doen't make sense to stick them in an oxygen plasma for an extended period. Early on I tried a couple times to do control devices that didn't receive the de-scum, but for one reason or another the whole round of devices was bad, so I never got a clear test of whether or not this de-scum actually improves the contact resistance or transport quality. I ultimately decided to stick with it since this step resulted in good devices in [164, 165] and since our own group found that resist residue could be removed by this de-scum process.

Since you likely oxidized your contact region during the de-scum, you need to remove that oxide before depositing the metal for your contacts. Get setup in a hood to do a short HCl etch, but don't do the etch yet. First enable and vent the CHA evaporator. Blow out as many metal flakes as you can. Be especially careful to blow out the region under the hearth around the filament so you don't short out the filament. Carefully inspect each of the sources. Set the source control on the electronics rack to manual and move the knob to which source you want to look at (don't forget to set it back to "Automatic" when you are done). The Ge may look a little blue; this is not uncommon and is ok. Make sure the Ni crucible is not cracked, and watch for black spots on any of the source material. If any of the sources have these little black spots, track down Kenny or Dave and let them know the source material is covered in graphite. If the electron beam is not adjusted correctly and hits the edge of the graphite crucible, the metal will get coated with graphite. If they do replace a source for you, pump the chamber down and do an evaporation without

your samples to clean up the source material. Also check the opening in the hearth for large flakes of metal; if there are any flakes, pick them out with a freshly cleaned pair of tweezers. You don't want to have any other metals fall into a hot crucible during the evaporation. Once you are satisfied that the evaporator is ready to go, lower the bell jar to keep the vacuum surfaces from getting too soaked with humidiity.

Get the Manfra group sample holder out of the group box. I made this holder to allow us to mount small samples; the clips are spaced more closely and the springs are not as stiff, so it is easier to clip your samples on without scratching them. Remember to not touch anything outside the evaporator with your clean gloves; the only "clean" thing in the clean room is the air, so don't get a bunch of goobers on your gloves if you want to get a clean evaporation.

Now go back to the acid hood and do a quick HCl etch and rinse on your samples. Etch for 20s in HCl and rinse $\sim$ 30s in DI water. Give your samples a quick squirt with running water and then blow them dry. As soon as all your samples have been etched, get them mounted and loaded into the evaporator so they don't re-oxidize too much. I typically get the evaporator roughing down $\sim$ 5 minutes after the last sample comes out of the acid.

Rough the chamber down to 40 mTorr, close the roughing valve, and open the gate valve to the cryo pump. If the chamber will not rough down for some reason, close the roughing valve and notify the staff. Never leave the roughing valve open for an extended period of time, and never allow the roughing valve and gate valve to be open at the same time. Both of these actions will cause oil to back-stream from the mechanical pump into the chamber, contaminating the source material. Back-streaming results in a lot of down time for this tool, and our group should not be a contributor to this failure mechanism.

Once the chamber pumps down below Kenny's approved evaporation pressure (this should take 20-30 minutes), de-gas the sources. Just start the deposition recipe as normal but abort the deposition right before the "deposition" step and do not open the shutter. Once the first source cools move on to the next source you will use. After

de-gassing all the sources, let the chamber pump down for another hour. If possible, I would recommend trying to time your work so that you get to this de-gassing step before 4:00 pm when the staff leave. This way if something goes wrong with the evaporator (e.g. filament burns out, turret gets stuck, source material is dirty, etc.) there is some chance the staff can fix it. If you start later than this you will have to leave your samples overnight if the evaporator breaks. Problems with the Ohmic deposition have ruined a lot of processing runs for me, so do whatever you can to make sure you can get a clean evaporation the same day you etch the mesa.

Once the evaporator has pumped down, deposit your Ohmics. There is a lot of lore regarding metal stacks and different wafers, but from what I have seen the metal stack doesn't make a whole lot of difference most of the time. The common thread is that you want an initial layer of Ni to help the metal wet the surface and improve diffusion of the Au and Ge into the crystal during the annealing. The Au and Ge thicknesses should have a ratio of 2:1. This is almost universal practice so don't change this ratio. The 2:1 ratio also turns out to give a mass ratio of 88:12 which is the same as AuGe eutectic which is often used in thermal evaporators due to its low melting temperature $\sim$ 360C. A Ni cap on top of the Au and Ge is optional, but it does seem to result in smoother contacts. For the in-situ back-gated wafers I have been using a metal stack of 8/80/160/36 nm Ni/Ge/Au/Ni since this was already known to work for these types of wafers [169]. The initial Ni layer is deposited at a rate of 1.5 Å/s while the other layers are all deposited at 2 Å/s. Our group has also used the metal stacks from [111, 172], and these metal stacks (which do not use a Ni cap) also give low contact resistances to a variety of wafers.

One peculiar problem we encountered early on with the Ohmics turned out to be due to the mask design. Evidently, it is easier for electrons to tunnel from the 2DEG into the contact along the (011) direction than it is along the $(01\bar{1})$ direction. What this means in practice is that if your Ohmic contacts are rectangular pads on the mesa your Hall bar will wind up with a working source and drain but dead voltage probes if the body of the Hall bar is oriented along (011) and just the opposite if the Hall

bar is oriented along $(01\bar{1})$. The work-around to this problem is to design saw-tooth or square-wave "scallops" into the Ohmic layer on top of the mesa so that electrons have a tunneling path into the contact along (011) for all your Ohmics. Figure G.11

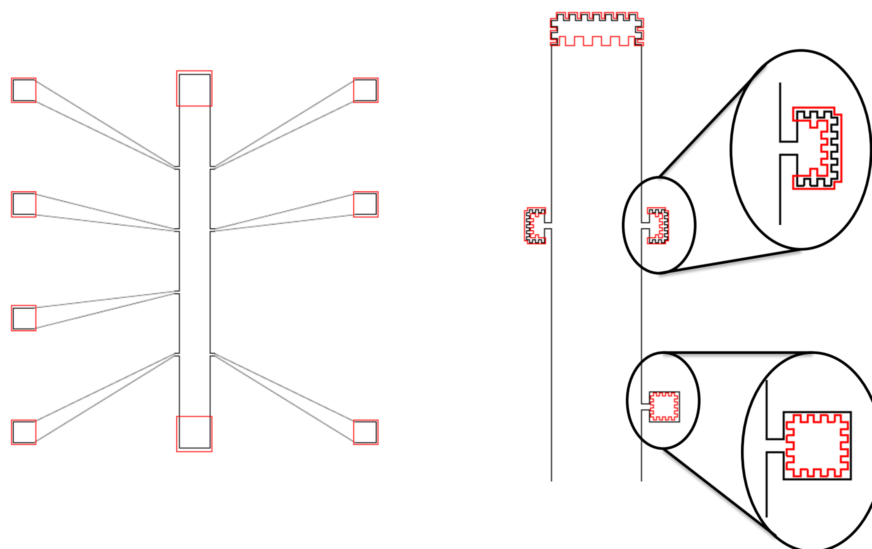

Figure G.11. Drawings of different mesa and Ohmic designs. The black lines represent the mesa while the red lines represent the Ohmics. The Ohmics on the Hall bar on the left have well defined directions in which the electrons in the 2DEG would have to tunnel into each contact. The Hall bar on the right has Ohmics with "scallops" which allow the electrons to tunnel into the contacts along both crystallographic direction for all the contacts. The contact region of the voltage probe is shown enlarged for clarity. The bottom contact on the right has no overlap with the edge of the mesa. Such "interior" contacts were consistently insulating at low temperature.

shows two different designs for a Hall bar; the black lines outline the mesa while the red lines outline the annealed metal for the Ohmic contacts. In the device on the left the electrons must tunnel along one direction to enter the source and drain and must tunnel along the perpendicular direction to enter the voltage probes. This results

in a large contact resistance anisotropy between the source/drain and the voltage probes. The device shown on the right has square-wave “scallops” on the Ohmics. In this case electrons can tunnel along both crystallographic directions to enter both the source/drain and voltage probes with the result that the contact resistances are similar for the source/drain and voltage probes regardless of device orientation. Also shown in the device on the right is an “interior” contact in which the annealed metal does not overlap the edge of the mesa. These contacts were consistently found to be insulating at low temperature, suggesting that diffusion of the metal through the sidewall of the mesa is also important for contacting the 2DEG. Similar behavior has also been reported in reference [153]. As a result, we also added the square wave scallops to the edge of the mesa in the contact region to increase the perimeter of the mesa in the contact region. This anisotropy in contact resistance has also been seen by other groups. It is specifically mentioned in [175, 174], and similar scalloping is used by other groups as well [164, 165, 171].

Once you finish the evaporation, start soaking your chips in acetone to initiate lift off. Use a dedicated beaker since it will get coated with metal flakes over time. Soak for $\sim$ 5 minutes and then squirt the metal off with some acetone. Do not let the acetone dry before you get all the metal off. If any of the metal lands on the surface of your sample and the acetone dries, it will be difficult to remove the metal. Once you get the metal squirted off, spray the chip down really well with the acetone airbrush and then transfer the chip to a clean beaker of acetone. Since your next step is to bake the samples in the RTA, you need to get the devices really clean. Any resist left on the surface now will get hard-baked on to the surface and never come off. Soak your samples in clean, hot acetone for 30 minutes. Set the hot plate to $\sim$ 70 $^{\circ}$C and cover the beaker with a watch glass or round bottomed flask filled with water to condense the acetone that boils off. You should see convection lines in the acetone, but the acetone should not actually start boiling. After soaking in acetone, spray the chips down with methanol, soak in methanol for 5 minutes, rinse with methanol, blow dry, soak in DI water for 5 minutes, rinse, and blow them dry.

Next enable the Jipelec RTA and do a test run to 400 °C to warm up the chamber. The RTA is designed to hold a 6 inch wafer, so each student has his/her own 6 inch Si wafer to act as a chuck on which to set his/her samples. In order to get the control and reproducibility in annealing temperature necessary for the in-situ back-gated project, I found it was necessary to use a special wafer with a thermocouple bonded to the top side of the wafer. It turns out that the top side of the wafer (the side facing the lamps) is $\sim$ 40 °C hotter than the backside of the wafer as measured by the built-in thermocouple that is pressed up against the back of the wafer. So if you use the backside thermocouple to control the temperature, be aware of this offset when trying to map out contact resistance vs. annealing temperature. The annealing time does not seem to matter a whole lot, but the temperature has a strong impact on the contact resistance. Shoot for a sample temperature of $\sim$ 375C for 1 minute in forming gas (4% hydrogen, 96% nitrogen) for the in-situ back-gated devices; if there is no buried gate present, anneal at $\sim$ 420C as this will give you lower contact resistances. Do not use the standard recipes that Dan has written for the RTA because the initial purge time is not long enough; the chamber must be purged for 10 minutes to avoid oxidizing the contacts. If you use a thermocouple wafer, do an annealing run without your samples to check that everything is working normally. Once you are done with the annealing, inspect your samples under the microscope and save a picture of each device. The contacts for the in-situ back-gated samples must be annealed at as low a temperature as possible to minimize gate leakage, so the Ohmics should not look much different before and after the annealing. Contacts annealed at higher temperature will show minimal roughening if a Ni cap is present but will show significant roughening and non-uniform color if there is no Ni cap.

If you managed to get through the mesa etch, evaporation, lift off, and annealing, congratulations. Go home and get some sleep. The time-sensitive portion of the processing is now done. Once you have the Ohmics annealed, you can set chips aside and process them one at a time if, for instance, you need to study the impact of different nanostructure designs.

G.7 Dielectrics

I have included this section on the deposition and patterning of dielectrics based on the initial work that I have done on FET-based devices for studying charge noise. These techniques should also prove useful for work with quantum Hall experiments since depositing a dielectric underneath the depletion gates will prevent the gates from leaking after illumination [59].

G.7.1 Deposition

I have used two different deposition systems and three different dielectrics. The Axic PECVD system can be used to deposit Si_xN_y, and the Fiji ALD system can be used for Al_2O_3 and HfO_2.

The recipe for the Axic PECVD system follows the recipe outlined in Sumit Mondal's thesis [180]. The chamber cleaning and the pauses in between deposition steps were found to be very important to eliminate pin-holes in the dielectric which caused signficant leakage. Below is a brief summary of the recipe:

- Manually clean the chamber to remove all dielectric flakes.
- Run plasma cleaning recipe: 15 min, 200 W, 250 C, 600 mTorr, 20 sccm CF_4, 100 sccm O_2.
- Condition the chamber with a nitride layer: 5 min, 150 W, 300 C, 600 mTorr, 100 sccm NH_3, electrode at 3 inches.
- Load your sample, pump down, and deposit the dielectric: 150 W, 120 sccm SiH_3, 100 sccm NH_3, 300 C, electrode at 3 inches, 600 mTorr. Deposit for 2 minutes total with plasma on/off as follows: 45 sec on, 30 sec off, 45 sec on, 30 sec off, 30 sec on.

This 2 minute deposition should give a film thickness $\sim$ 60 nm, but the deposition rate is not especially repeatable, so if you need a tight control on the thickness you

should first run a test piece, measure its thickness with the Filmetrics, and then adjust your deposition time for your real sample accordingly. I would also suggest keeping a chip with a known thickness of Si_xN_y in your box to check that the Filmetrics is giving you an accurate reading; I have had somewhat inconsistent readings with this particular tool in the past.

The Fiji ALD system is a fully automated "black box" system which makes it very easy to use but also rather inflexible (e.g. recipes can only be edited by the staff). To run the system, use the following procedure:

- Check that the previous user's "sytem idle" recipe is complete and that the chamber is sufficiently cool (probably anything < 150 C is okay).

- Vent the chamber

- Load your sample(s) onto the chuck. If your samples are small (e.g. $\sim 4 \times 4$ mm), try to load them in the center with their edges touching. The pump down of the chamber is very abrupt, and I'm told that if the samples are too small, they can get blown off the chuck.

- Pump down the chamber and let it sit for 5 minutes after reaching the base pressure before starting the recipe. I am not sure what (if any) impact this waiting period has on the quality of the dielectric-semiconductor interface, but since your sample will start out coated with a bunch of water from the air, I suspect this may have some impact on the interface quality. There are some reports in the literature that say that starting with a couple pulses of trimethyl-aluminum (TMA) can help remove the native oxide and improve the interface quality. The Fiji starts with a water pulse, so for the default recipe it may not make much difference how long the chamber is pumped down before the deposition starts, but if you want to try starting with some TMA pulses, this pump down time would likely be an important variable.

- Load the intended recipe. So far I have only used the "Thermal_HfO2_110C" recipe for depositing HfO_2 and the "Thermal_Al2O3_100C" recipe for depositing Al_2O_3.
- Set the number of layer repeats.
- Start the recipe. There is nothing you can control at this point, so just watch that the pressure spikes alternate like they should and then go do something else while it runs.

The deposition rate is unfortunately not as repeatable as it should be for the ALD process. You can assume deposition rates of 0.09 nm/cycle for Al_2O_3 and 0.15 nm/cycle for HfO_2. Particularly for the HfO_2, though, this deposition rate differs significantly from what is reported in the literature, but I have not found anyone else at Purdue who knows what is normal for the Fiji machine (evidently most people who use the Fiji don't actually measure their resulting film thicknesses). Because of this variation in deposition rate, you should make sure to simultaneously deposit the dielectric on a larger test chip coated with Au so that you can measure the thickness with the Filmetrics and also chop it up and use it for etch test pieces prior to patterning the dielectric on your devices.

In an initial test I found that for metal-insulator-metal capacitors, the ALD dielectrics give leakage-free capacitors at 4 K for areas as large as 100×400 μm. To be more precise, for a 28 nm thick Al_2O_3 film with Ti/Au pads for both capacitor plates the I-V was linear to 5 V with a resistance > 1 TΩ. The resistance did not scale with area which indicates that the small leakage current that was measured was likely through the wire insulation in the measurement setup. At room temperature, however, the ALD films were very leaky (measured in the dark). The same Al_2O_3 capacitors mentioned previously gave 1 nA of leakage for voltages as low as 0.87 V for an area of 2500 μm^2. The HfO_2 films (56 nm thick) gave similar results both at low temperature and room temperature. The PECVD Si_xN_y recipe is also known

to give leakage-free dielectrics at low temperatures (see references [180, 181]); to my knowledge the Si_xN_y films have not been tested for leakage at room temperature.

G.7.2 Dielectric Patterning

There are two options for patterning the dielectric: lift-off or etching. The lift-off process has been shown to work for low-temperature-grown ALD films [182, 183, 184, 166]. However, this lift-off procedure is not compatible with the high temperature of the PECVD chamber, and it may also be difficult to lift-off small, isolated features. In other words, I suspect it works better for bright-field polarity masks than for dark-field polarity mask designs. For this reason, I chose to pursue patterning the dielectric by wet etching with buffered oxide etch (BOE).

All three dielectrics can be etched with BOE. The approximate etch rates for full strength BOE are 0.09 nm/s, 2.8 nm/s, and > 4 nm/s for HfO_2, Al_2O_3, and Si_xN_y, respectively. The etch rates do still vary, however, so you should once again etch a couple test pieces for varying times to make sure you know the current etch rate before you try patterning your device. As the etch rates for Al_2O_3 and Si_xN_y are quite high, the feature defintion tends to not be particularly sharp for small (<5 μm) features. I suspect that the resist does not adhere well to the dielectric. However, I have not yet done a systematic study of the impact of hard baking the resist or using an adhesion promoter like HMDS.

The BOE will stop abruptly on Au as well as annealed NiAuGe, so there is no danger in damaging Ohmic contacts or gates with the wet etch. BOE will, however, etch the oxides on the surface of GaAs, so whatever metal layer you use as an etch stop should be several microns larger than the etch window you are defining in order to connect, for instance, an Ohmic contact to a bondpad.

If you want to make your own metal-insulator-metal capacitors to study leakage in different films, the TiAu of the bottom plate of the capacitor can be defined by etching in a two-step process. I suggest etching rather than lift-off to define the metal

pads because liftoff can often leave a ridge of metal at the edge of the pad which could be difficult to cover with the dielectric. The Au can be etched with aqua regia. The etch rate is extremely high, so just wave the chip around in the etchant until you see the pattern appear on the surface of the wafer. After rinsing thoroughly, the Ti adhesion layer can be removed with a quick etch in BOE (30 seconds should be sufficient to remove 20 nm Ti). The etch field will look quite rough, but the pads should be electrically isolated well enough to measure the leakage through the subsequently deposited dielectric.

G.8 Electron Beam Lithography

G.8.1 Sample Preparation

Electron beam lithography proceeds similarly to optical lithography except that you will use an e-beam resist (PMMA 950 A2 in our case) and expose the resist with an electron beam rather than a UV lamp. Start out with a clean sample. If you just annealed your sample last night, you don't need to do any additional cleaning since you did an extensive cleaning just prior to the annealing. If you've had your sample sitting in a petri dish or plastic box for a week or two, though, you should re-clean it. Be careful about sonicating the sample, though. You should check your test sample (i.e. the one on the blank GaAs wafer and not a real heterostructure) first to see if the sonication will cause the metal to peel off before cleaning your real sample. If you are concerned that the metal might come off, just soak your sample in each solvent and/or use the acetone airbrush or hot acetone to do more aggressive cleaning.

Next, get the spinner set up. Set it to spin at 4000 RPM for 45s. Pre-bake your sample at 100C for 2 minutes and try to minimize the time that moisture can coat the surface before spinning the resist. Once you've spun all your chips, bake them at 180C for 10 minutes. Remember to let the hot plate warm up and stabilize for a while before baking the PMMA.

G.8.2 Raith Setup

Clean up your mess, disable the spinner, and take your sample out to the Raith. I would recommend using something other than foil to carry your sample so you don't drop it or flip the sample over and scratch the resist. Enable the Raith in Coral, pull the sample holder out of the load lock with clean gloves, and set it on a fresh clean wipe. Don't use the wipes other people have left out on the bench, this is poor vacuum hygiene. Load your samples onto the sample holder with plastic tweezers so you don't scratch the sample holder. Place the holder in the load lock and shut the load lock door.

Open up the e_LiNE software on the left computer screen (henceforth referred to as [L]) and log in using your username/password from your initial training. Click on the "Navigator Loadlock" icon and select "via Loadlock", "Load sample", "Ok". Follow the prompts that the software gives you to pump down the load lock and transfer the sample holder into the main chamber. Click "Yes" when prompted to reset the u/v alignment. Enter a name for your sample (this doesn't really matter).

Next, turn on the electron gun (EHT) on the lower right of the right computer screen (henceforth referred to as [R]). Set the accelerating voltage to 20 kV and select the 10 μm aperture. The accelerating voltage and aperture will, of course, need to be selected based on your feature size; for now, let's assume we are just writing a quantum point contact with a minimum feature size $\sim$ 200 nm. On [L], open the "stage control" window and drive the stage to the Faraday cup. On [R], select the "Detectors" tab and select "Signal A = InLens" for accelerating voltages of 20 kV and under. Turn off the beam blanker on [L]. Set the working distance on [R] to 18 mm and focus on the Faraday cup. Once you have it in focus, zoom in to $\sim$ 20 kX. On [L], select the "Exposure" tab and click "Measure". This will measure your beam current. Check your previous notes to make sure the system is operating as you expect.

Next, you need to set the working distance for writing your pattern. Everyone at Birck uses a working distance of 10 mm, though I have never heard an explanation of why except that it is known to work. If you end up having to push the resolution limits of the machine, you may want to look into what impact the working distance willl have. To set this distance, you first need to focus on something on the surface of your chip. On [L], go to the stage control tab and drive to the clip that is holding your sample. Once the stage has stopped, unblank the beam and zoom in to the clip so that you don't start exposing your sample. Carefully drive towards your sample and find one of the corners of the chip. Find something to focus on and zoom in a bit to get a rough focus and then blank the beam. On [L], click the "Stage Control" button, select the "Drive" tab, enter your desired working distance (10 mm) in the "W" field, make sure the "Absolute" and "mm" radio buttons are selected, and press "Start". This will drive the stage up to the 10 mm position. Since you were focused on your sample when it was at an 18mm working distance, though, you now need to re-focus the electron beam. On [R], set the working distance to be 10 mm. Unblank the beam to check that your sample is still in focus.

Now you need to do the hard work of getting the beam focused well enough to write your pattern. This is accomplished by burning so-called contamination dots into the resist to check the spot size and shape of the electron beam. First, find some debris on the surface of the resist and focus as well as you can. Next, you need to align the aperture. First click on the "Reduced Raster" icon on [R] to view a smaller area with the SEM so that your refresh rate is faster. In the "Apertures" tab on [R], select "Aperture Align" and turn on the focus wobble. The speck you have focused on will now move in and out of focus and also probably dance around on the screen. Adjust the aperture alignment until the speck stays still on the screen and only moves in and out of focus. Once you think you have the aperture aligned, turn off the focus wobble, select the "Stigmation" option on [R] and make fine adjustments to the stigmation to try to get slightly better focus. In order to really adjust the stigmation, though, you will have to examine your contamination dots.

Move away from your dust speck to a clean area of resist, zoom in to 70 kX, and center click on the “Spot short/Spot long” button” on [R] to start burning the dot. Since the beam is likely not well focused, you will need to start out with a longer exposure to get a visible dot. Start out with $\sim$ 60 seconds. At the end of the 60 seconds, center click on the “Spot short/Spot long” button again to stop the exposure and return to normal SEM view. If you can’t see your dot, adjust the focus (remember to make sure you have the focus and not the aperture align or stigmation selected on [R]), and try burning another dot slightly away from your first dot. Keep repeating this procedure until you can actually see something that looks like a oval-shaped white blur.

Once one of your dots is visible, focus on it and then burn another dot. Keep repeating until you have a well focused dot. If the dot is not a perfect circle, tweak the stigmation slightly and burn another dot. Iterate in this fashion until you have a nice round doughnut shaped dot. Now you should be able to start backing off on the exposure time for each dot. After burning each dot, try to improve the focus. As the exposure time decreases and your focus improves, the size of the dot should decrease. Always aim for getting a 20 nm dot (you can measure it with the SEM tools on [R]), but don’t work harder than you have to. If your smallest feature is 200 nm, a 30 nm dot will still work just fine.

G.8.3 Sample Alignment

Now you need to get your sample coordinates set and align the write field so that your pattern will come out where you want it on your chip. Due to the details of how the Raith figures out the coordinate transformations, the order of these steps is important. First, you need to do the 3-point alignment in order to tell the Raith where to drive the stage in order to expose your pattern. The tricky part here is that you need to find where you are on your chip without looking at anything on one of your mesas to avoid exposing the resist. It will be much easier to orient yourself on

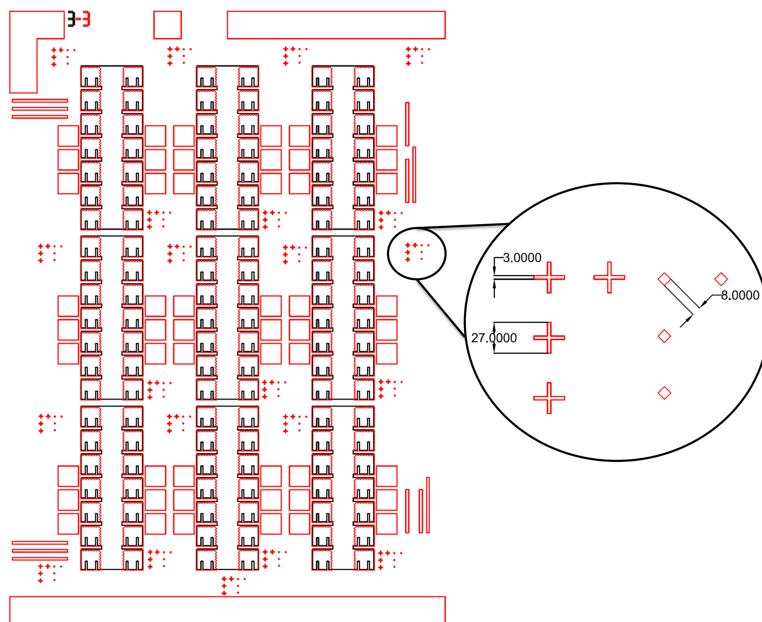

Figure G.12. Overlay of mesa (black) and Ohmic (red) designs. The large red rectangles at the top and bottom of the chip serve as initial orientation markers during the initial e-beam alignment. The Hall bars shown in the design are $\sim$ 1 mm long. The smaller L-shaped arrays of alignment marks are used for the actual 3-point alignment procedure to create the coordinate transformation between the mask and stage coordinates. In addition, the fact that these features are large enough to be seen with the naked eye enables one to mount the chip on the sample mount in a pre-determined orientation. The inset shows a magnified view of the alignent markers used for the 3-point alignment. These markers are arranged so that the array lacks rotational symmetry; this lack of symmetry makes errors during the alignment less likely. Dimensions shown are in microns.

the chip if you have some large features at the edge of the chip. Ideally, these features should not be symmetric so that as soon as you find one of them you know exactly where you are on the chip. Figure G.12 shows an overlay of the mesa and Ohmic layers of one of our mask designs which feature these orienting features at the edge of the chip. Once you have a sense of what part of the chip you are looking at, carefully drive the stage to the nearest set of alignment markers.

Once you find a set of markers, center the cross hairs on one of the markers and blank the beam. Open up your GDS file in edit mode on [L] and drag flags #1, 2, and 3 to the positions on your design that corresponds to the center of the alignment marks you will use. For right now, just use three points that are in a single alignment mark cluster. Open up the "Adjust UVW" window on [L] , select the "3-Points Adjustment" tab, and press the eye dropper to capture the current stage coordinates. Check the box to lock the UV/XY coordinate pair. Next, shift the stage to match the coordinates at your 2nd and 3rd markers. Press "Adjust" to tell the software to calculate the coordinate transformation.

Now that you have a rough alignment done, you can drive anywhere on your sample by holding CTRL + R Click on your GDS file. In order to get a more accurate coordinate transformation, re-do your 3-point alignment using alignment marks spaced as far apart as possible to minimize the impact of errors during the alignment. Drive around to various features on your chip using the CTRL + R click method to verify that your alignment is good enough. Since our goal in the group right now is primarily to looks at QPCs or larger devices for studying quantum Hall physics, the alignment requirements are not that stringent. If at some point in the future it is necessary to get really tight alignment between multiple e-beam layers, it would probably be necessary to write the alignment marks on the chip with e-beam prior to the mesa etching. As of this point, though, as long as your alignment is good to within $\sim$ 2-3 μm when you drive all the way across the chip, you should be fine.

Now that you have your stage-mask coordinate transformation set, you need to get the stage-beam transformation set by doing a write field alignment. The write field alignment essentially consists of the computer moving the stage slightly and then deflecting the electron beam (i.e. center of the SEM view) to where it thinks the stage moved. The user tells the computer where the stage actually moved, and the computer is thereby able to correct its error.

To perform the write field alignment, drive back to your speck of dust that you used for your initial focusing and center the cross hairs on the sharpest protrusion of

the speck. Do not adjust your focus, aperture, or stigmation. If necessary, adjust the contrast and brightness to get a good image. Open a new positionlist from the "File" menu on [L]. In the "Microscope Control" window, select the desired write field size. In the "Scan Manager" window, select "Align Write Field Procedure", then select "Manual". Drag the "100 μm WF - Manual ALWF 25 μm marks" line into the positionlist, right click in the positionlist window, and press "scan". Each time the SEM view pops up, hold CTRL and drag the cross hairs back to the point on which you originally centered the cross hairs. Once you get the cross hairs adjusted, press proceed. This window will pop up three times. At the end of the procedure, accept the corrections. Repeat this alignment procedure with the "100 μm WF - Manual ALWF 5 μm marks" and "100 μm WF - Manual ALWF 1 μm marks" procedures. This should give you a decent alignment of the write field. Poor write field alignment will result in stitching errors at the write field boundaries in your pattern.

G.8.4 Exposing and Developing

Now that everything is aligned, you are ready to start the exposure. Drag your GDS file into a new positionlist window, R click, and select "properties". First, select the layer you wish to expose. Make sure you only select your e-beam pattern if you GDS file also contains the Ohmic layer. Next, click the button next to the working area coordinates to update the working area coordinates. The working area is the portion of your file that will actually get exposed, and its boundaries determine where the Raith will put each write field (more on this later). Click the icon next to the UV coordinates field to update the field based on your design. These fields correspond to the center of the lower left-most write field in your working area.

Next, you need to calculate the exposure times. Open the "Exposure Parameter" window and set the step size and dose (for now just set the step size to 20 nm). The dose will depend on your design (more on this later). Click on the calculator icon to set the dwell time. Go back to the properties window and click on "Times" to

calculate the total exposure time. Assuming this is what you expect, click "Ok" to go back to the positionlist. Right click on the positionlist window and select "scan" to start the exposure. While the pattern is being exposed, keep an eye on the UV coordinates in the bottom right corner of [L] to make sure they are roughly where you expect the pattern to be written.

Repeat the 3-point alignment and write field alignment for each chip that you have loaded. Once you have finished your writing, turn off the EHT on [R]. Open the navigator window on [L], click on "unload", and follow the prompts. Disable the system in Coral and take your sample back to the cleanroom.

Develop the samples for 30 seconds in MIBK:IPA 1:3 (this is pre-mixed). Quench the developing with IPA (make sure you don't dilute your developer bath in the process) and blow dry. Examine your device under the microscope to make sure the pattern was exposed where you intended. Assuming it did, load the sample into the evaporator.

As with photoresist, you should not leave the resist on the surface for extended periods if you want to have repeatable results. As such, make sure you have an evaporator reserved so that you can do the evaporation the same day as the exposure. Assuming you are using the CHA, pump down the chamber for 1 hour and then evaporate 5 nm Ti and 20 nm Au at rates of 1.5 and 2.0 Å/s, respectively.

After the evaporation, soak the samples in acetone for a few hours and finish with $\sim$ 5 seconds in the sonicator to clean up the edges. Rinse with methanol or IPA and blow dry. Take pictures of the device at this point (assuming it still looks ok). Next, do another aggressive solvent clean to make sure all the PMMA is really gone before you move on to your next lithography step. Soak in toluene for 5 minutes, hot acetone for 10 minutes or more, and methanol for 5 minutes.

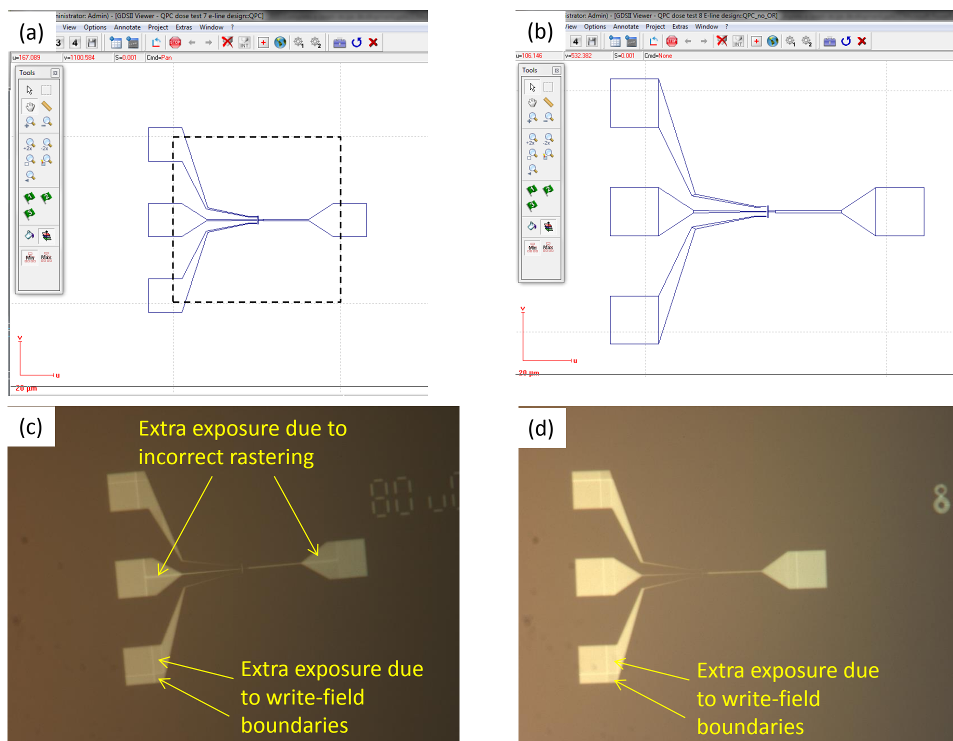

Figure G.13. Example of the influence of e-beam pattern design on exposure uniformity. (a) Pattern in which each arm of the multi-QPC device is "OR-ed" into a single large polygon. The black dashed box denotes the 100 μm write field used for the exposure. (b) The same device broken up into multiple polygons. (c) Optical micrograph of resist after exposing and developing using the design in (a). The device shows brighter lines at the boundaries of the write fields (where the design was exposed twice) as well as through the central arms of the device. (d) Optical micrograph of resist after exposing and developing using the design in (b). No excess exposure is seen in the central arms of the QPC structure.

G.8.5 Design Tips

I cannot claim to be a real expert at e-beam writing, but I have picked up a few useful points while learning to use the Raith. The most important detail is that the GDS pattern should be broken up into multiple polygons to get consistent exposure doses for your entire pattern. Figure G.13 illustrates why this is important. In figure G.13a, each of the four arms of the device were "OR-ed" into individual, large

polygons since this was our standard procedure for generating optical masks. Figure G.13b, on the other hand, shows the same design but with each arm of the device broken up into smaller, simpler polygons. Panels (c) and (d) show the PMMA after exposing and developing. These patterns were some of the under-exposed regions of a dose test. The important point, however, is that because these features were intentionally under-exposed, inconsistencies in the pattern exposure become evident. There are, not surprisingly, bright lines in the exposed regions of both figure G.13c and G.13d that correspond to the boundaries between write fields. However, figure G.13c also shows an additional bright line through the central arms of the QPC structure. This additional bright line was evidently due to the Raith making a poor decision on how to raster the central features which, as a result, received a larger exposure dose. One word of caution, however, is in order regarding breaking the design up into polygons. The rastering can also sometimes leave small gaps between polygons, so be sure to overlap your polygons a little bit to ensure your features are continuous.

As a result of this design-dependent dosing, a dose test must be performed for all new pattern designs, even if the design is a small perturbation from a previous design. Figure G.14 shows an example of a successful dose test using the multiple polygon design from figure G.13b. As the dose is increased, progressively smaller and smaller features survive the lift-off procedure until the minimal clearing dose is reached. It is important to note that the dose test must always be followed by a metalization and lift-off since it is impossible to tell from optical or SEM imaging whether the PMMA was completely cleared out of the smallest features. An SEM examination of the metallization following the dose test is important for examining the smallest features, as shown in figure G.14. However, processing lore states that actual devices should not be examined with the SEM following lift-off as this could potentially embed charge in the active region of the device.

Another important design consideration is the definition of the working area. As figure G.13 shows, features at the boundary of each write field will get a higher dose than features away from the boundary. As a result, it is important to keep fine

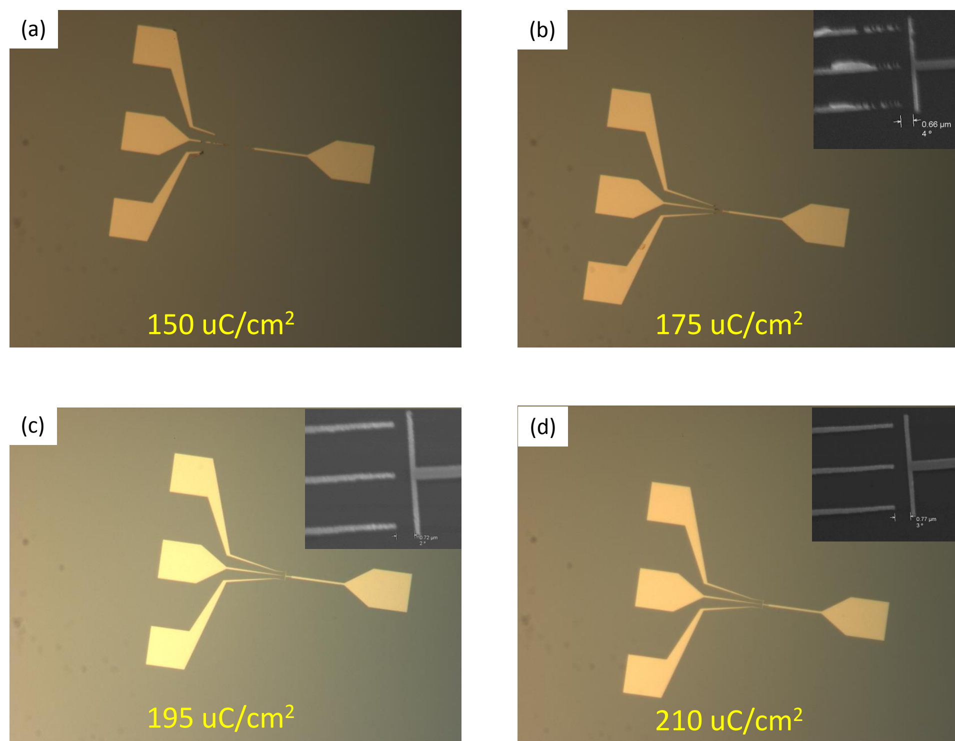

Figure G.14. Example of a successful dose test. Main panels show optical micrographs after metal lift-off while the insets show SEM micrographs of the center of the device. As the dose is increased from (a) to (d), smaller and smaller features should survive the lift-off procedure.

features away from write field boundaries. In addition, poor write-field alignment can also cause rotational and/or translational discontinuities at the edge of the write fields. This is not a big issue for large features, but it could easily ruin small ones. The position of the write fields is determined by the working area. While you are designing your device, set the working area such that it will be tiled with an integer number of write fields with one of the write fields centered on your smallest features. Name this working area something obvious like “Use me” so that you don’t make any mistakes in telling the Raith which working area to use.

Finally, avoid curves in your design if possible and instead draw quasi-curves with multiple polygons. The Raith can write curves without too much trouble, but curves will not translate correctly if the design is created in AutoCAD and then imported

as a DXF file into e_LiNE. You could do all your designs in e_LiNE to get around this issue, but I personally think AutoCAD allows much more control in designing your features (the object snap functionality is particularly nice). In addition, you can download AutoCAD for free as a student and use it whenever you want, whereas using e_LiNE requires you to check out the license key from the Appenzeller group office which is often times not convenient.

G.9 Optical Gates and Bond Pads

Once your sample has been thoroughly cleaned of all resist residue, do the lithography the same as you did for your Ohmic contacts. Don't worry about doing a de-scum or oxide removal etch; you just need to make electrical contact between the Ohmics/bond pads and/or the e-beam and optical gates, so the invisible residue of photoresist is not a big deal. Pump the CHA down for about an hour after loading your samples. If the CHA is not available, the Lesker is another good option for Ti/Au. The Varian can also do TiAu, though it will take significantly longer to pump down and will heat your sample significantly during the evaporation. The Airco is faster but has had problems in the past with Cr contamination in the Au which results in a bunch of black spots in the deposited metal layer. Evaporate 20/150 nm Ti/Au at a rate of 2/2 Å/s. Thinner metal layers may work for wire bonding, but decreasing the thickness is kind of risky. I know from personal experience that 5/75 nm Ti/Au is not sufficient for wire bonding, but I have not tried anything in between these two values.

Once the evaporation is done, do lift-off in acetone, spray the chip down with the acetone airbrush, and soak in methanol for a few minutes to remove the acetone residue. Don't worry about doing a DI water soak. Take pictures of your completed devices, make photocopies of your notes from this round of processing, pack your samples up in plastic boxes with clean-wipe padding, and take them over to the physics building for measurement.

G.10 Sample Mounting and Wire-Bonding

Now that you are back in the dusty physics building, you need to be careful with how you handle your samples if you want to have an easy time with the wire bonder. Clean your tweezers with some IPA or methanol, but be warned that they are likely still not very clean. If you touch any bond pad with your tweezers, assume that you will have a hard time bonding to that pad. This is particularly important if you need to cleave chips out of a larger strip. If you do this, try to grab the strip with the tweezers as close to the edge as possible.

For initial testing you should use the commercial ceramic chip carriers[5]. These are easy to bond to, and their poor thermal conduction makes no difference when the sample is submerged in liquid helium in the dipper probes and the ^{3}He system where you will do your initial characterization. Handle the chip carriers with clean gloves and try to avoid touching the bonding surfaces as much as possible. Snip off the shorting bars on the legs with wire cutters. Mount your chips in the carriers with rubber cement; try to avoid getting any of the stringy cement on the bond pads. Since the bonds can be somewhat fragile, cut out a square of the black foam in the prep room to stuff into the bottom of your small plastic sample box. This will keep the chip carrier from rattling around and also give a (weakly) conductive path to all the pins on the chip carrier in case your device is ESD sensitive. Once your samples are all mounted, take them upstairs to the wire bonder.

Always start the bonds on the chip carrier since this will give a stronger bond than the TiAu on your sample will. Our bonder (a wedge bonder) works by applying pressure and an ultrasonic vibration to the wire to scrub it across the bond pad, expose clean Au atoms on the pad and wire, and cause the two pieces to form a cold weld. The wire is fed through the back of the wedge so that the end of the wire is right under the wedge where the pressure and vibration is applied. During the bond process, the wire gets deformed into a long, narrow foot. As a result, the bond is very

[5]16-pin CERDIP chip carrier from Spectrum Semiconductor Materials, Inc. www.spectrum-semi.com. Part number CSB01648, Mfg. Dwg. IDK16F1-390GAL.

weak if the wire is pulled perpendicular to the length of the bond. This means that you need to plan out your bond path and pull the wire straight back from the initial bond site. If you pull the wire too much to the side, the first bond will likely pop off. Remember to write down your bond parameters in the log book when you are done and turn off the microscope light.

The bonding settings that have worked well for me are as follows:

- Bond 1 power: 2.5
- Bond 1 time: 2.5
- Bond 1 force: 3.3
- Bond 2 power: 1.8
- Bond 2 time: 3.9
- Bond 2 force: 2.0
- Tail: 2.2
- Pull: 0.6
- Stage temperature: Room temperature
- Tool temperature: no heat
- Bond wire size: 1 mil

If your sample and the wire are both clean you should have a success rate $> 90\%$. The most likely cause of difficulty in bonding is a dirty sample. If the bonder is not breaking the wire after the second bond, examine the clamp. Sometimes the wire starts to slide around in the clamp, and this causes problems with breaking the wire after the second bond and/or advancing the wire before the first bond. If you are still having trouble, make sure the sample is at the correct height; the bond surface should be 3 inches above the base. If the bond surface is not at the correct height, the wedge will contact the bond pad at an angle and not be able to form the bond.

G.11 Condensed Checklist for In-situ Back-gated Devices

G.11.1 Ga Removal

- Spin AZ1518 40s 4000 RPM, 1000 RPM/s ramps up/down. No pre-bake necessary.
- Bake 2 min 100C
- Wipe off Ga
- Spin and bake second layer AZ1518
- Etch 3 min HCl, full strength
- Rinse thoroughly, blow dry
- Cleave as necessary. Make sketches of chips if possible to keep them in order.
- Rinse resist off with acetone into dirty beaker, spray down with acetone gun, soak in clean acetone
- Rinse with methanol, soak in methanol few minutes
- Rinse with methanol, blow dry
- Check sample under microscope
- Sonicated triple solvent clean: 5 minutes each toluene, acetone, methanol, DI. Rinse with methanol after each step except after DI rinse. Blow dry after methanol step. Rinse with running water after DI sonication and blow dry.
- Take pictures

G.11.2 Via Etch

- Spin AZ1518 40s 4000 RPM (no pre-bake)
- Bake 2 min 100C
- Expose 20s 10mW/cm^2 on MJB3_2
- Develop 20 + 10s MF-26A, rinse 30s
- Check development with UV filter. Take pictures if development complete
- Soak 5 min DI, rinse in running water, blow dry
- De-scum 90s Branson asher. 130 mTorr base pressure, 120:5.5/12 Ar:O_2, 1.3 Torr process pressure, 100 W power (check reflected power)
- Do not hard bake resist
- Etch test piece 50:5:1 water:phosphoric:peroxide. Assume 1.7 nm/s etch rate and aim for close to desired etch depth in real devices.
- Strip resist on test piece in acetone, rinse with methanol, blow dry
- Measure etch depth with Bruker optical profileometer. Use VXI scan type, 5 μm backscan, 10$\times$ objective, default parameters
- Assume 15% increase in etch rate from GaAs to standard high mobility structure
- Etch real samples, aim for 160 nm above gate layer (will etch the rest of the way during the mesa etch)
- Strip resist, measure actual etch depth for each device
- Sonicated triple solvent + DI clean
- Take pictures
- Leave overnight

G.11.3 Mesa Etch

- Pre-bake 2 minutes 80C
- Spin S1805 40s 5000 RPM
- Bake 5 minutes 80C
- Expose 6s 10 mW/cm^2
- Develop 10 + 6s MF-319
- Rinse 30s, blow dry
- Examine developing. Take pictures if complete.
- Soak 5 minutes in DI, rinse with running water, blow dry.
- De-scum 90s Branson asher. 130 mTorr base pressure, 120:5.5/12 Ar:O_2, 1.3 Torr process pressure, 100 W power (check reflected power)
- Bake 5 minutes 100C
- Etch test piece 90s 50:5:1 water:phosphoric:peroxide
- Strip resist, measure etch depth, calculate rate for real devices assuming 15% increase in rate
- Etch devices. Aim for 160 nm etch depth for standard doping well structures with 110 nm deep doping layer and 200 nm deep 2DEG.
- Strip resist, measure etch depth
- Sonicated triple solvent + DI clean
- Take pictures
- Minimize time mesa sidewalls exposed to air. It is possible to have sidewall exposed for as little as 3-4 hours prior to Ohmic deposition

G.11.4 Ohmics

- Spin AZ1518 40s 4000 RPM (no pre-bake)
- Bake 2 minutes 100C
- Expose 20s 10 mW/cm^2
- Harden 20 minutes in chlorobenzene
- Blow dry, soak 5 minutes in DI water, blow dry
- Develop 70 + 20s MF-26A, rinse 30s, blow dry
- Examine developing, if necessary develop longer in 10s intervals
- Take pictures when developing complete
- Rinse 5 minutes in DI water, blow dry
- De-scum 15s Branson asher. 150 mTorr base pressure, 120:5.5/12 Ar:O_2, 1.3 Torr process pressure, 100 W power (check reflected power)
- Vent evaporator. Blow out metal flakes. Check source cleanliness. Lower bell jar and leave chamber purging. Get sample holder ready
- Etch devices 20s HCl, rinse $\sim$ 30s, blow dry.
- Load samples into evaporator and pump out chamber.
- De-gas metals once pressure low enough. Do NOT open shutter.
- Pump down 1 hour
- Evaporate 8/80/160/36nm Ni/Ge/Au/Ni 1.5/2/2/2 Å/s.
- Lift-off few minutes in acetone, squirt down with acetone. Do not let metal flakes settle on surface of device. Spray down with acetone airbrush and transfer to clean acetone.

- Soak in 70C acetone 30 minutes
- Squirt down with methanol, soak in methanol 5 minutes. Spray down with methanol and blow dry. Soak 5 minutes in DI water, squirt down with running water, blow dry.
- Take pictures
- Warm up Jipelec RTA with test run (no devices)
- Anneal 1 minute in forming gas at target temperature (375C for in-situ back-gated devices). 1 minute ramp up, ramp down as fast as possible. Purge chamber 10 minutes prior to annealing.
- Take pictures
- Leave overnight

G.11.5 E-beam Gates

- Clean devices if they have been sitting for an extended period
- Pre-bake 2 minutes 100C
- Spin PMMA 950 A2 45s 4000 RPM
- Bake 10 minutes 180C
- Expose in Raith
- Develop 30s MIBK:IPA 1:3, squirt down with IPA, blow dry
- Check development, take pictures
- Evaporate 5/20 nm Ti/Au 1.5/2 Å/s
- Lift-off few hours in acetone, finish with 5s sonication

- Take pictures
- Soak 5 minutes in toluene, > 10 minutes hot acetone, 5 minutes methanol, 5 min DI water

G.11.6 Optical Gates and Bond Pads

- Spin AZ1518 40s 4000 RPM (no pre-bake)
- Bake 2 min 100C
- Expose 20s 10 mW/cm^2
- Harden 20 minutes in chlorobenzene
- Blow dry, soak 5 minutes in DI water, blow dry
- Develop 70 + 20s MF-26A, rinse 30s, blow dry
- Examine developing, if necessary develop longer in 10s intervals
- Take pictures when developing complete
- Rinse 5 minutes in DI water, blow dry
- No de-scum or de-oxidization etch
- Evaporate 20/150 nm Ti/Au 2/2 Å/s
- Lift-off few minutes in acetone, spray down with acetone airbrush, squirt down with methanol, soak few minutes in methanol, squirt down with methanol, blow dry.
- Take pictures
- Photocopy notes, pack up samples

G.11.7 Mounting and Wire Bonding

- Cleave as necessary
- Keep samples as clean as possible
- Glue into chip carriers with rubber cement
- Wire up in bonder
  - Bond 1 power: 2.5
  - Bond 1 time: 2.5
  - Bond 1 force: 3.3
  - Bond 2 power: 1.8
  - Bond 2 time: 3.9
  - Bond 2 force: 2.0
  - Tail: 2.2
  - Pull: 0.6
  - Stage temperature: Room temperature
  - Tool temperature: no heat
  - Bond wire size: 1 mil
- Fill out log sheet and turn off microscope light
- Measure device, publish paper, graduate, get paid the big bucks

H. Kelvinox Dilution Fridge Standard Operating Procedure

H.1 Introduction to the System

The Kelvinox fridge was originally put together with the goal of being able to measure the fractional quantum Hall gaps of different wafers with a large throughput to act as feedback on heterostructure design. As a result, we designed a tail and associated headers that can hold up to four 4 mm Van der Pauw squares. Each sample has its own red LED pointed at it so that if multiple samples are examined after

illumination, they should all have experienced identical illumination conditions. This is important since the illumination has a very large effect on the gaps and transport quality. In addition, the samples are mounted on homemade headers that have a copper strip that can be screwed into the copper tail so that there is a continuous metal connection between the sample and the mixing chamber. Finally, the Labview code was written in such a way that all four samples could potentially be measured simultaneously with the data saved to separate data files with corresponding log files automatically generated by the code to help keep track of all the measurement parameters. In principle, then, if everything is working (a big if) and if the samples are roughly the same density, the gap at say $\nu = 5/2$ could be measured from four different wafers in the space of a week or so.

The downside to having such ambitious goals when setting up a fridge is that the system has very tight design parameters. Namely, there is a very small gap between the samples and the radiation shields meaning that the possibility of a thermal short between the samples and the ($\sim$ 50 mK) radiation shield is always present. In addition, the headers don't have much space for mounting samples, and to get a good thermal link between the sample and the fridge it is necessary to use a conductive adhesive (silver paint) to mount the samples which introduces the extra complication of potentially shorting contacts to the fridge. This is not too much of a problem for large Van der Pauw squares with soldered contacts, but it makes mounting processed samples (which require wire bonding) very challenging at times.

Finally, this system was put together from parts of a couple fridges and is therefore very "homebrew" (which is I suppose fitting for a piece of equipment which runs "mash" through a "still"). In other words, this standard operating procedure should be taken as a guideline for what normally works and how the system normally behaves, but it should not be followed blindly without understanding how each part of the fridge works and what can potentially go wrong.

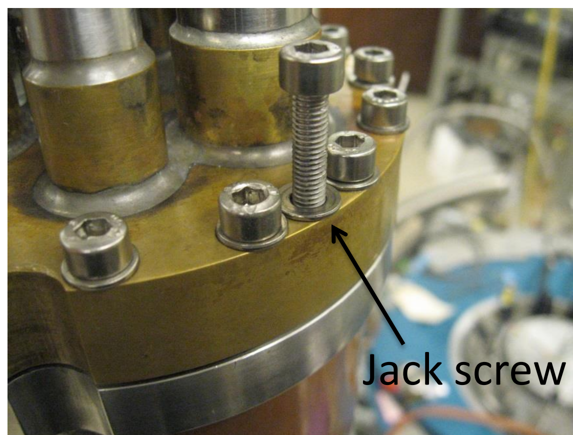

Figure H.1. Top of the IVC with jack screw called out. The jack screws are used to break the indium seal.

H.2 Preparing Samples

Let's assume that you are starting with the fridge at room temperature but with the sliding seal, inner vacuum can (IVC), and radiation shield in place and the copper headers still in the fridge. The first step is to get the headers off of the fridge. First, remove the sliding seal (you may need to get someone else to help hold the sliding seal while you unscrew the last couple screws). Next, make sure the fridge is high enough above the floor that you will be able to slide the IVC off. Remove all but two of the screws from the IVC and start threading screws into the jack-screw threads shown in figure H.1. Loosen your remaining screws holding the IVC on a little bit but do not completely un-thread them. Then go back and forth tightening the jack screws until you break the indium seal; the IVC should be resting on the two remaining screws at this point. Now hold the IVC with one hand (again you may want a buddy to make this easier) while removing the last two screws with your other hand. Once the screws are out, lower the IVC off the fridge very carefully; pay special attention not to catch the IVC on the serpentine pre-cooling line and be extremely careful not to scratch any of the silver heat sinks with the IVC (if you do you will wind up having to re-wire the fridge). Once you get the IVC off, set it on the bench and put something

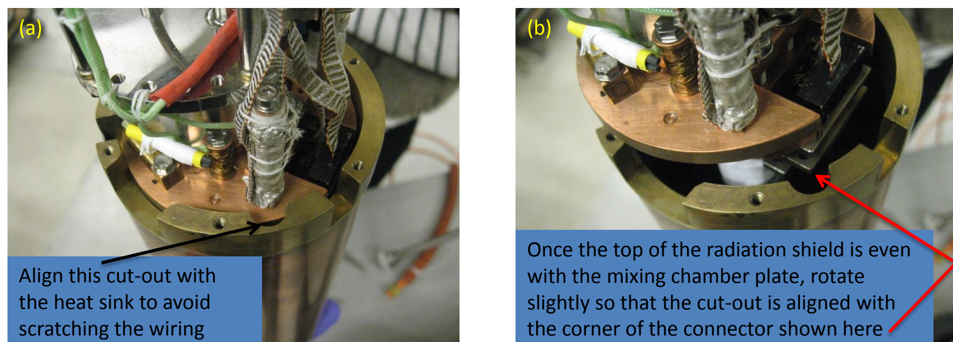

Figure H.2. Procedure for removing the radiation shield. (a) Align the cut-out in the radiation shield flange with the heat sink to avoid scratching the wires on the heat sink. (b) Once the flange is safely past the heat sink, rotate it slightly to align the cut-out with the corner of the connector.

heavy in front of it so that it can't roll off the table. Peel the indium off of the IVC or fridge (wherever it stuck) and dispose of it in the zip-loc bag in the green cabinet (we will eventually recycle this and hopefully recover some of the cost of the metal). Next, put on some clean gloves (get a new pair if you touched the chain hoist with your old gloves) and unscrew the radiation shield and slowly lower it off the fridge. Getting the top flange past the mixing chamber is a bit tricky. First, as shown in figure H.2a, you have to align one of the cut-outs in the flange with the silver heat sink closest to the edge of the mixing chamber to get the flange all the way down to the mixing chamber plate level. Once the flange clears the heat sink, rotate the flange about 15° to get the same cut-out to align with the corner of the micro-d connector on the tail as shown in figure H.2b. Do NOT force the shield off. If you have the shield aligned properly you should not feel resistance. If you do feel resistance, stop and think about what you are doing. This is another point at which you could scratch some of the wiring and short a bunch of wires to each other. Once you get past the micro-d connectors, just keep sliding the shield off slowly and be careful to not whack the LEDs when you get the skinny part of the radiation shield past the end of the tail. Once the shield is off, set it on the bench and again be careful to not let it roll

Figure H.3. Copper headers mounted on the tail. Care should be taken when removing the radiation shield to not bump the LEDs with the shield.

off the table. At this point it would also be a good idea to make sure that the LEDs on the end of the tail are above the level of the fridge stand so that if you bump the fridge the LEDs don't get bent.

Figure H.3 shows the copper headers mounted on the fridge. To remove them first unscrew the stainless steel screw (note that this is the only screw on the whole fridge that is not a metric thread) and then gently start prying the headers off with a blunt pair of tweezers (do not use sharp tweezers as these will scratch the copper). While removing the headers there are a few things to keep in mind. First, be careful what you touch on the tail. I have found the best place to hold the tail to steady it is the part of the teflon tape with the copper oxide rub marks on it. This point on each side of the tail is free of any wires that you might inadvertently shift around. Try to avoid touching the RuO_2 thermometer or any of the other wiring. Second, try to pry the header out evenly so that you don't bend the pins on the header or put any unnecessary stress on the epoxy bond that glues the copper strip to the plastic header. Finally, do your best not to scratch up the copper surfaces that mate between

the header and the tail since these surfaces are what link your sample to the mixing chamber.

Once the headers are off the fridge, take them into the prep room and remove the old samples. You should be able to just push the sample off the header with a pair of tweezers to break the silver paint bond. Do not use the clean tweezers in the sample preparation supply drawer to avoid contaminating any samples and the tube furnace with silver paint; instead, use a dirty pair of tweezers from the dil fridge lab. To avoid bending the pins on the header I like to put the header in the socket glued to a piece of copper that we usually use for wiring up samples. Once the samples are off, put them back in the appropriate sample boxes for long term storage. Next, you need to clean the headers in order to ensure that you get a solid thermal connection with your new samples. The silver paint is soluble in acetone, so put a little acetone on a q-tip and wipe the paint off. Try to minimize how much acetone you get on the plastic since the plastic is not very resistant to the acetone. Once all the paint is off, put some acetone on a fresh q-tip and go over it all once more to make sure it is really clean. Once all the paint is gone it is a good idea to remove the copper oxide and acetone residue with some very fine grit sandpaper[1]. Figure H.4 shows some technology I developed to make cleaning the headers and other small parts easier. Be sure to get the oxide off the top and bottom of the copper strip. Once you have the copper looking nice and shiny, wipe it down with some methanol or IPA on a q-tip to remove the sandpaper grit. The last step before you mount your samples is to make sure you have good indium blobs on all the solder forks on the header. When you do this, be sure to use the “dirty” indium and associated soldering iron tip. These are currently labeled “NiAuGe soldering” since they are also used for making solder connections to large evaporated NiAuGe ohmic contacts.

Now that the headers are clean, it is time to mount your samples. If you are going to mount two samples on a header I would recommend that you first put the

[1]I typically use 2000 or 3000 grit sandpaper. Very fine sandpaper like this is often used for autobody work and is available at most auto parts stores

Figure H.4. Homemade sandpaper tools for cleaning copper oxide off of small pieces. 2000 and 3000 grit sandpaper is glued to wooden sticks with quick set epoxy to aid in cleaning small copper pieces such as the headers.

screw in the header so that you can position your samples appropriately. You should also plan out now which sample(s) will go on each header. The headers are not interchangeable on the tail (i.e. the header with the yellow paint marking pin 1 must go on the socket with the yellow paint marking pin 1 in order to get the stainless steel screw threaded). Take your samples (presumably already measured in the ^{3}He system) off of their original headers or chip carriers and try to keep track of the sample orientation so you can use the same contact sets from the ^{3}He system for taking your data. To get the best thermal contact you will need to remove the rubber cement from the back of the sample. Hold the sample by the edges with a pair of tweezers and use your other hand to roll the rubber cement up into a little ball with a pair of tweezers. The cement should stick to itself better than anything else so you should

be able to roll it all into a little ball and then grab it with the tweezers. Once you remove the rubber cement, get the silver paint from the green cabinet in the dil fridge lab[2]. Stir up the silver paint a bit with a wooden stick and then shake the bottle vigorously to get the silver uniformly distributed in the solvent. Put one drop of the silver paint on the header immediately followed by your sample and push your sample down with the clean end of a wooden stick. Remember to put the cap back on the silver paint right away since the solvent evaporates quite quickly.

This gluing procedure can be a bit tricky. If you get to much silver paint on the header, it will creep up the side of the chip and short your ohmics together or short the back gate to the fridge if the sample has a back gate. For right now, just inspect the chip for paint creeping up the side. If you think there is a short, remove the chip, wash it thoroughly in acetone and then methanol or IPA and re-clean the header. Once you have the samples mounted, solder the contacts to the header like you normally would if your sample is a large Van der Pauw square. If you have processed samples which require wirebonding, the next step is a little more challenging since you need a bond wire with a wirebond on one end and a solder connection on the other.

First, move the header to the sockets glued to an aluminum block. These sockets are glued on with silver-powder infused GE Varnish which acts to short all the leads of the sockets together in case you have an ESD-sensitive device. The two adjacent rows of sockets come in handy for wiring up processed samples on the copper headers. Put your header in the single socket row and then put clean chip carriers in the adjacent double socket row and take everything upstairs to the wirebonder. The goal here is to use the chip carrier as the site of your first bond and then drag the wire over the solder fork that you want to use before bonding to your sample (see figure H.5). If the bond pops off the chip carrier, you can sometimes reposition things so that the wire dangling from the wedge of the wirebonder falls on top of the solder fork. Make as many bonds like this as you can and then take the aluminum block with your sample back downstairs to the prep room. If you planned your bond path well, the wire

[2]Pelco 187 silver paint from Ted Pella, Inc., catalog number 16045

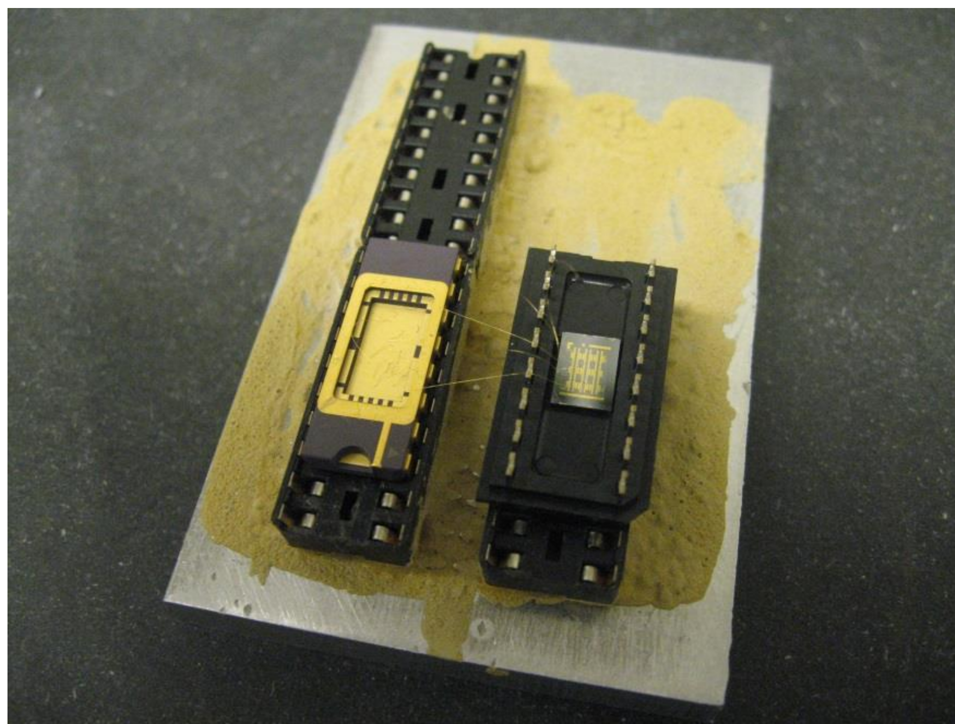

Figure H.5. Illustration of how to wire up processed samples on the dilution fridge headers. Start the bond on the chip carrier (left) and drag the wire over the solder fork before bonding to the sample (right). Note that a standard plastic header without the copper insert is shown here because the copper headers were both in the fridge during the preparation of this manuscript.

should be bonded to your sample with the mid-section of the wire touching one of your indium blobs. In this case, melt the indium blob with the soldering iron before ripping the bond off of the chip carrier. If the wire is not touching the indium, try to gently break the bond on the chip carrier without breaking the bond on your sample (i.e. try to pull away from the chip carrier and towards your sample). Do your best to have the soldering iron ready and solder the wire to the header. Because we have a wedge bonder, the bonds are pretty weak if the bond footprint and the wire are not collinear, so try to move the wire as little as possible. If it pops off, just repeat the process as before. Because this bonding process is so difficult, I would recommend

making your bondpads no smaller than 150 × 150 μm so that you have enough real estate to make several bonding attempts before the bondpad is completely destroyed. If you can put annealed metal underneath the bondpads, this will also help since the annealed metal sticks to the sample better than plain TiAu does. If you go a few hours without getting all the wires bonded, take a break, watch a stupid cat video on YouTube, get some fresh air, pump some iron, take an aspirin, etc. to burn off some steam before you go back to finish up the wiring. Once the samples are all wired up, double check under the microscope that all the wires are continuous and not shorted to anything. Before your samples can be mounted on the fridge there are a few other checks to do, so just set your samples aside for now.

H.3 Preparing the Fridge

It is a good idea to check that all the solder joints on the fridge survived the thermal cycling to room temperature before you mount your samples on the fridge. First, check the resistance from the switch panel to the sockets on the tail. Hook up the D-SUB cable from the switch panel to the appropriate switch box on the top of the fridge (they are labeled #1 and #2) and make sure you have the switch box in the "A" position so that the switch panel in the rack is connected to the sockets. Then float all the switches on the switch panel and check the resistance to each pin on the socket with a hand-held multimeter. Look back through the notebook to make sure that the values you are measuring match the previously recorded values. If the resistance of one connection increased significantly since the last cooldown, this may be an indication of a bad solder joint. If there is a bad connection, you should try to fix it. Keep in mind that the room temperature connections are much less likely to have problems than the joints that get thermally cycled. So far most, if not all, of the problem joints that I have found have been on the room-temperature side of the blue connectors on the 1K plate; so when in doubt, start by re-soldering this connection first.

Once you are confident that all the wires are continuous, ground all the switches on the panel (except the one you have the multimeter hooked up to) and check for shorts. Be careful not to hook the multimeter up to any of the LEDs when you have the multimeter set to measure resistance because this may drive an unsafe current throught the LED. To test the LEDs, put the multimeter in diode testing mode and use the cheat sheet taped to the switch panel to determine the correct polarity for each LED. Again, compare the values you measure with their historical values and record your measurements in the front of the log book.

Next, move filter box #1 on top of the fridge to Fischer connector #1. This will connect pins 1-24 on the switch panel to the thermometers. Check the resistance of all the combinations for each thermometer and record the readings in the front of the log book. The log sheet has some helpful reminders about the wiring of the thermometers. Once again, you should only see a couple Ohm variation from one cooldown to the next. If you see a larger variation than this, a solder joint may need to be fixed. Once you are done checking the thermometers, plug filter box #1 back into Fischer connector #2.

Assuming all your wires are continuous, there are no shorts, and your thermometers and heaters all match their historical values, it is time to mount the headers on the fridge. Put on a clean pair of tight-fitting gloves (you will need good dexterity to avoid breaking any wires) and get your samples and tweezers within arm's reach of the fridge. Hold the header upside down with the tweezers and align it with the socket. Once the pins are lined up with the receptacles in the sockets, try to pinch the header and socket sandwhich with your fingers and remove the tweezers. Then try to go back and forth between each end of the header, pushing the header into the socket a little bit at a time with the tweezers. Try to avoid pushing on the header inside the solder forks (i.e around your bond wires). Also, make sure to work the header into the socket evenly so you don't break the epoxy joint on the header. Once you get the header $\sim 2/3$ of the way on the socket you can try threading the screw into the tail. Be very careful to not cross-thread the threads. If you feel resistance,

stop and back the screw out. Copper is very soft and these threads are very small, so it will not take much to ruin them. If the pins on the header are pushed into the socket enough, you can usually get them the rest of the way by tightening down the screw. Go back around with the tweezers and make sure the header is pushed all the way into the socket and then tighten the screw a little more. Don't over-tighten the screws to avoid damaging the threads.

Once you get both headers on the fridge, re-install the radiation shield. Line up one of the cut-outs with the corner of the micro-d connector to get the flange up to the mixing chamber plate and than rotate $\sim 15°$ to align this cut-out with the heat sink closest to the edge of the mixing chamber plate. Once the radiation shield is all the way on, get all the screws started before tightening any of them down all the way. Once they are all started, go around tightening them a couple times. The threads in the radiation shield are brass so they are a bit more durable than the copper threads in the tail, but you should still be careful to not over-tighten them since the stainless steel screws are certainly capable of stripping the brass threads.

Once the radiation shield is on, test the 2-terminal resistance to ground of all your contacts (hopefully you remembered to write down what connections go to which contacts on the samples). As an aside, if you have any samples with an in-situ backgate, you should check that the gate is not shorted to the copper strip before loading the header on the fridge. Tie a gold wire around the hole in the copper strip and solder it to an unused pin on the header. This will allow you to check for shorts with the 4K dip stick. The wire connection to the copper strip is necessary because the socket in the dip stick is all plastic and thus the copper strip is insulated from the body of the probe and ground. Try to go slowly with cooling and warming the sample; the silver paint can sometimes lose its bond after being thermally cycled. Assuming you already did this, take the cap off the bottom of the radiation shield, lie down directly underneath the fridge and look for any contact between the LEDs, headers, teflon tape, etc. and the radiation shield. If there are not any thermal shorts, screw

the cap back in and tighten it gently with a socket; both threads are copper so they can get jammed very easily.

If your contacts resistances are all reasonable, then it is now time to install the IVC and start pumping. First, clean the 4K flange on the fridge and the flange on the IVC with some IPA or methanol and a clean wipe. If any indium is stuck to the 4K flange, try to scrape it off with a wood stick. Do not use metal tweezers to scrape the 4K flange; the flange is made of soft brass and could be scratched by stainless steel tweezers. Next, put four or five screws in the flange; you will use these to raise the IVC the last little bit. Stand the IVC up on the bench and get a clean wipe, vacuum grease, and IPA or methanol ready within arm's reach. Pinch off a length of indium long enough to go all the way around the flange and overlap the ends of the indium around the "x" scratched into the side of the flange (this will ensure that the tails of the indium don't get bumped by the serpentine pre-cooling line. Wipe the indium wire off a couple times with some IPA or methanol and then coat the wire with a thin layer of vacuum grease. You only need the vacuum grease for sealing really small cavities in the flanges, so don't use too much. Lay the wire around the flange and bend the tails down so they don't get caught on anything. Lastly, inspect the wire for any fuzz that might compromise the seal. If everything looks good, start sliding the IVC on to the fridge. You will probably feel it hit the bottom of the radiation shield when the skinny part of the IVC reaches the bottom of the radiation shield. Just try to wiggle it around gently until the two shields get lined up. Just like when you took the IVC off, be careful to not catch the IVC on the pre-cooling line or the heat sinks. Once you get the flange up far enough to touch the screws, start threading the screws into the IVC can and try to avoid touching the top flange with the indium wire. Once you get four or five screws started, you can leave the IVC hanging. Make one last check for any debris on the indium wire and then go around tightening each screw a little bit at a time; the goal here is to uniformly bring the wire into contact with the top flange. Once the wire is in contact with the flange, insert the rest of the screws into the flange and tighten them finger-tight. Next go around in a star pattern

tightening $\sim 1/4$ turn each time to compress the wire evenly. After doing this a few times, you can start tightening around in a circle until everything is snug. Lie down under the fridge again and shake the bottom of the IVC a little bit; check that the radiation shield wiggles around freely in all directions to check that it is not touching the IVC anywhere.

Now that the IVC is on, the guts of the fridge are protected so you can lower the hoist a bit without as much risk of damaging something with the dewar stand. Fire up the leak detector and hook it up to the IVC pump-out port. Don't try to pump out the IVC with the turbo pump; the helium background is too high for the turbo pump to make any progress. Once the diffusion pump on the leak detector is heated up, start pumping the IVC out with the rotary vane pump in the leak detector. This will probably take 10-15 minutes, so every 5 minutes or so rotate the valve on the leak detector back to position #2 to pump on the backside of the diffusion pump. This may not actually be necessary, but I get nervous about leaving the diffusion pump hot for long periods of time with no backing pump. Once the gauge on the leak detector gets down to 2 mBar, go ahead and open the throttle valve to the diffusion pump. This is above the pressure that the manual says to rough the chamber down to before opening the throttle valve, but the fridge has a huge surface area and a lot of very tortuous pumping paths, so 2 mBar is about as low as the rotary vane pump will get the IVC. The diffusion pump is pretty efficient at pumping helium, so once it is pumping on the IVC you should only have to wait a few minutes before you can turn on the ion gauge. Let it pump until it shows a leak rate $< 1 \times 10^{-8}$ atm cc/s. Leak check the IVC with a liberal amount of helium, and don't forget to check the window on the bottom of the IVC since it also has an indium seal. Let the leak detector pump on the IVC for an hour or two (if things went smoothly this may be a good time for a late dinner).

Once you are convinced that the IVC is leak tight, seal it off and shut down the leak detector. Hook the turbo pump up to the condenser and start pumping down the dilution unit. Be very careful to minimize the strain on the condenser port; the

tube is very thin and there is a long lever arm. You may need to add an elbow to the condenser port to keep the bellows from putting too much torque on the system. Pump down the bellows before opening the condenser valve. Once the turbo pump is up to speed and you are convinced there are no leaks in your connections (the pump should get down into the low 10^{-4} Torr range within a minute or two of reaching its top speed), turn off the ion gauge, valve off the turbo pump with the valve on top of the pump, and open the condenser valve. Record the pressure that the hose comes up to according to the convectron gauge. This should be a reasonable indication of how much junk made it past the traps and into the fridge (or how much mixture you left in the fridge and have now pumped into the room). If this looks ok (i.e. a few Torr), slowly open the valve on top of the turbo pump. Be careful to do this slowly so you don't slam the pump with a large gas load. If there is still an appreciable helium partial pressure in the system, the pressure will not really move. If this happens, seal the condenser off, shut down the pump, let the turbo spin all the way down, vent the hose, pump the hose back down, and then open the condenser valve back up. Adding the air to the line should help flush the helium out of the pump. If the pressure still gets stuck at 10's of mTorr, that is a problem since it indicates there was a lot of mixture left in the fridge. Seal everything off and talk to the big man before proceeding. Assuming that whoever used the fridge before you was careful, though, and got all the mixture out of the fridge, this won't happen and the pressure should start falling pretty quickly. Once you see that the pressure is below 1 mTorr and is steadily falling, turn off the ion gauge and leave the system pumping like this overnight. If loading the samples went reasonably well, this will probably be the end of the day anyway so you will be ready for a break.

When you come back the next morning, the condenser pressure should be at the base pressure of the pump ($\sim 2 \times 10^{-5}$ Torr). Assuming this is the case, seal off the condenser and shut down the turbo pump. Let the pump spin down for ~ 5 minutes before you move it so that you don't crash the fan blades into the housing of the turbo pump. Next, you need to add some exchange gas to the IVC to get it to cool

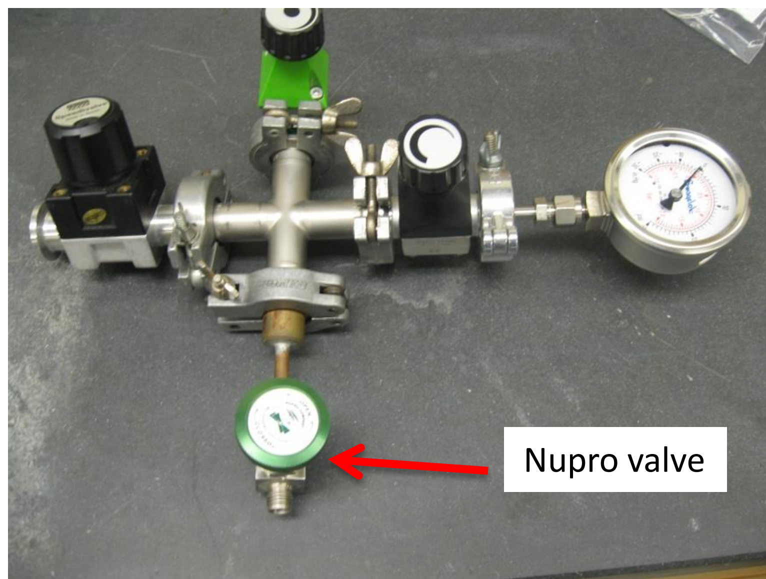

Figure H.6. Manifold used for adding exchange gas to the IVC. Fill the 4-way cross to ∼ - 24 inHg with helium and then add this to the IVC to act as exchange gas.

efficiently to 4K. The goal here is to add a couple hundred mTorr of ^{4}He to the IVC. I put together a little manifold, shown in figure H.6, that makes this easier. It is just an NW-16 4-way cross, but it does the trick. Hook this manifold up to the turbo pump using one of the o-rings from the "helium saturated o-rings" bag (keep these o-rings separate since they are no good for anything you need to leak check). The manifold is kind of heavy, so try to rest it on the frame of the pump cart. Hook up the helium line to the manifold and flow some helium through it to purge the air, then seal off the green Nupro valve on the manifold. Hook the manifold up to the IVC port and pump it down with the turbo pump. Let it pump on the hose for a few minutes and convince yourself that there are no leaks. Then close the two NW-16 valves on the manifold to isolate the 4-way cross from the pump and the bellows and turn off the pump. Slowly open the Nupro valve to fill the cross with helium up to ~ -24 inHg. If you overshoot, you can bleed some pressure off into the pump assuming it has spun down all the way. Once you hit your pressure target, seal off the Nupro

valve and open the valve to the bellows and then the IVC. This will fill the IVC with some exchange gas so that the guts of the fridge cool along with the rest of the fridge to 4K. Over time I've slowly been cutting back on how much helium I add to the cross to make it easier to pump out the exchange gas at 4K (more on this later).

Once you've added the exchange gas, seal off the IVC, put a blank on the IVC port as a safety (you should also have blanks on the condenser and still ports), and move the pump out of the way. Now put the sliding seal back on the fridge (you may need to find a helper for this). Don't over-tighten the screws on the sliding seal; you are sealing a rubber o-ring, not an indium wire, and you just need to seal ~ 1 psi of helium, so don't go bananas and squash the o-ring. Check the homemade red rubber gasket on the bottom of the sliding seal for debris and cracks. This piece of rubber is needed because the plastic o-ring that mates with the sliding seal has a small dent in it that leaks a fair bit of helium. If the gasket is cracked, you can make another one with an Exacto knife and some thin gasket rubber (I think I bought the rubber at Menard's or maybe Ace Hardware).

The last thing to prepare on the fridge itself is the 1K pot. You need to get it filled with dry helium before cooling it down so that you don't freeze the needle valve shut or block the inlet or outlet with ice. Hook the turbo pump up to the 1K pot line and pump it down (make sure the needle valve is all the way shut). Ultimately, you need to pump out the 1K pot (to a few mTorr) and fill it back up with helium 3-5 times, but you have to be careful with the turbo pump. It doesn't pump helium very well, and hitting a thin fan blade spinning at 20,000 RPM with a bunch of gas is a pretty risky move. So to prolong the life of the pump, valve off the pump from the manifold (you should still have the 4-way cross hooked up after adding the exchange gas to the IVC), and let the pump spin down. In the mean time, fill the 1K pot up to atmospheric pressure with helium. Once the pump has spun down, start it back up and immediately open the valve to the 1K pot. This will force the diaphragm pump to do most of the pumping on the helium and allow the turbo pump to spin up gradually under the gas load. Once you get the pressure down to -30 inHg on the mechanical

gauge, valve off the pump again and shut it down. Repeat the pump/purge process 3-5 times. Leave the 1K pot pressurized at 1-2 psi the last time you fill it up with helium. This will ensure that whatever is leaking through the needle valve is helium leaking out of the fridge and not air leaking into the 1K pot.

H.4 Cooling the Fridge to 4K with a Cold Magnet

The fridge itself is now ready to go. At this point what you do next will depend on whether or not the magnet is already cold. Since I already told you to pump and flush the 1K pot, let's assume the magnet is still sitting at 4K from the last cooldown. Hook the sliding seal recovery port up to the recovery line (try to match the sharpie marks on the the bellows and the valve) and raise the fridge with the hoist enough to slide it into position over the dewar. You will have to remove the D-SUB cables if you have not already done so. Before you do that, make sure that the switch boxes on top of the fridge are in position "B" or "C" so that your samples are grounded to the fridge. The initial lowering of the fridge into the dewar to mate the sliding seal with the top of the magnet is not a very graceful process. You are opening a large hole in the dewar and shining a lot of radiation into the helium, so things will start boiling off pretty quickly. In addition, the IVC extends quite a ways below the end of the sliding seal, so you will also be adding a lot of warm metal to a region previously filled with cold helium gas. Before you can do this you need to plan ahead a bit. First, the liquid level in the dewar should be no higher than ~ 5 inches. If it is higher, you will boil off a lot of liquid, pressurize the dewar a lot, and hit the helium recovery system with a very large gas load. Second, you need to make sure no one else is transferring too much. You will be boiling off a lot of liquid for the first ~ 20 minutes, so you should wait if someone else is transferring out of the liquefier since transferring out of the liquefier boils off a lot of helium. It also wouldn't be a bad idea to go around and talk to the guys in the other labs and ask them to hold off transferring for the next 30 minutes so that the helium recovery compressor can safely handle the gas load.

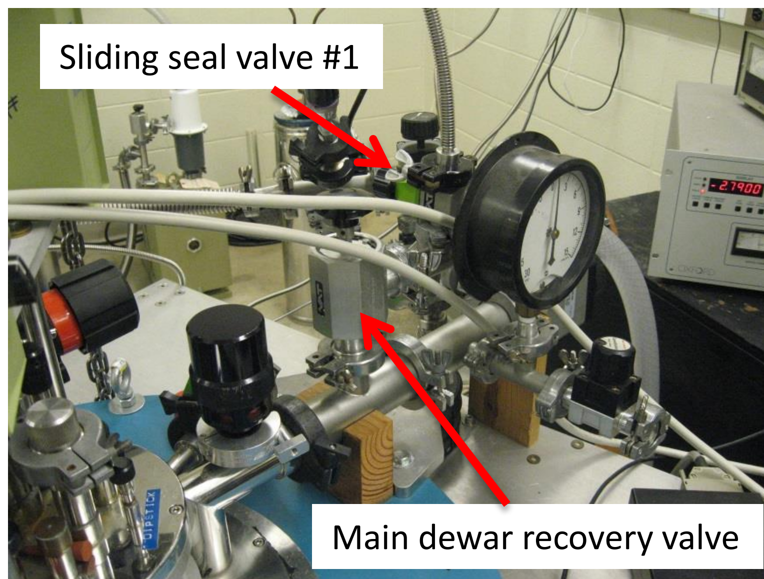

Figure H.7. Dewar exhaust port to helium recovery system.

Remember that the maximum gas flow rate the compressor can handle is 18 CFM, so the sum of the boil-off from all the labs shown on the gas meter webpage needs to stay below 18 CFM for the compressor to keep up.

Once the helium recovery system is ready, find a buddy to help you lower the fridge into the dewar. Get the fridge centered over the dewar, pull the cap out of the magnet, set the cap someplace where it won't roll off onto the floor, and start lowering the fridge as quickly as possible while your buddy guides it into the magnet opening. Stop lowering once the sliding seal is about 1 inch into the magnet. Open the valves on the sliding seal recovery port shown in figs. H.7 and H.8 and close the main dewar recovery valve (you may need to thaw this out with a heat gun) to force the boil off to leave via the port on the sliding seal. This will take advantage of all the cold He gas you are boiling off to help cool the fridge gradually before it reaches the liquid surface in the dewar. Do not close the main dewar recovery valve before you get the fridge in so you don't create a bomb. Even with the main recovery valve open, the dewar pressure will likely rise to $\sim$ 3-4 psi.

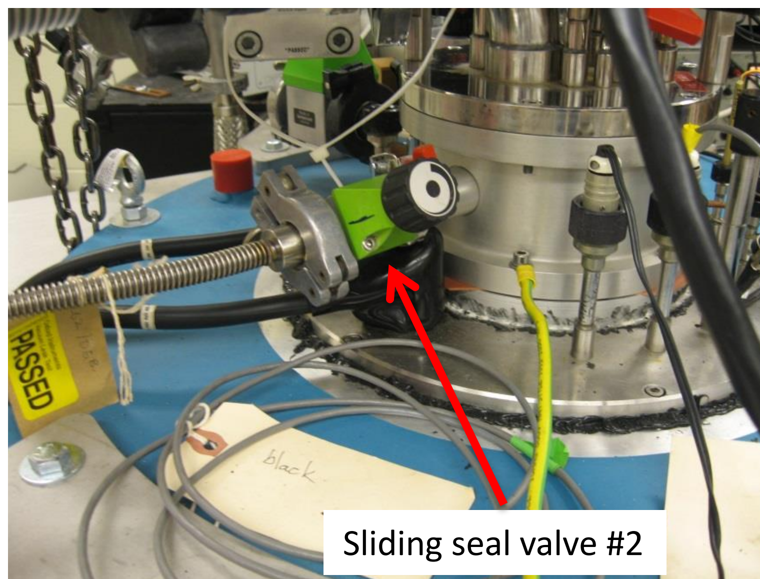

Figure H.8. Sliding seal exhaust valve #2.

Keep lowering the fridge slowly and periodically check the liquid level in the dewar and the boil-off rate on the gas meter webpage. Try to keep the boil off rate $\sim$ 4 CFM. This means you will probably lower the fridge $\sim$ 1 inch every 5 minutes. Be vigilant to wipe the condensation and frost off the sliding seal as well as you can. If you start to build up too much frost, you either won't be able to get the fridge lowered farther or (if you can keep lowering it) you may damage the plastic o-ring that makes the connection with the G-10 of the sliding seal. Once you get the fridge low enough, connect the resistance bridge pre-amp to Fischer connector #1. Turn the bridge on, set the channel to "0", range to "2 K", excitation to 30 μV, input to "MEAS", and leave everything else unchanged. Start up the Labview VI "Monitor Bridge Temperatures", select whichever Agilent multimeter you have the bridge output connected to, select the temperature output option, select the sorb thermometer, and start the code. You will have to adjust the "x" parameter to give you the correct reading. This is just a scale factor to get the order of magnitude right for the resistance the bridge is reading. Adjust this up or down in factors of

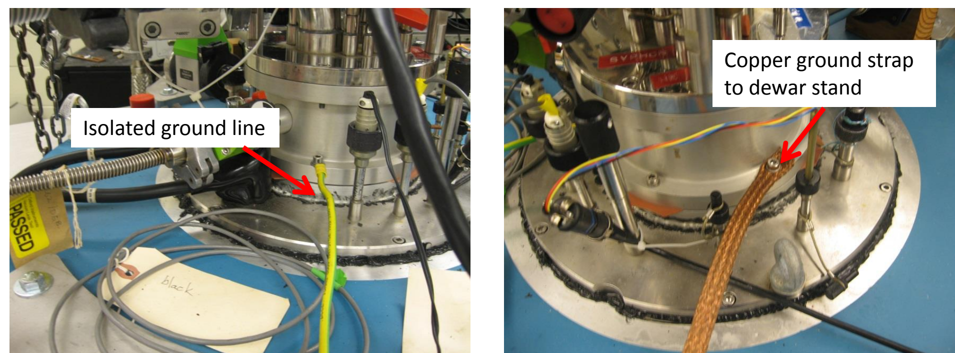

Figure H.9. Fridge ground lines. Both these ground lines must be securely connected to the fridge to keep the resistance bridge from showing overload at base temperature.

10 until you get the right answer (it should show something like 200 K at this point since the calibration for the sorb thermometer is basically flat from 200 K to room temperature).

Once you get the fridge all the way down, screw it down into the magnet and open the main dewar recovery valve back up. When screwing the fridge down be sure to connect the copper ground strap between the fridge and the fridge stand and the green and yellow ground line to connect the fridge to the building isolated ground as shown in figure H.9. Try to get all these screws tight. There is an as-yet unfound ground loop or noise source in the system that causes the bridge to show an overload signal when Labview is talking to the multimeters if these screws are not all tight and if the copper ground strap is not well connected.

H.5 Helium Transfer

By this point, the liquid level in the dewar will probably be getting pretty low so it would be a good idea to transfer helium. If you are running the dil fridge, you are probably already pretty familiar with transferring into the ^{3}He system, but you should pay attention to the details in this section anyway because it is a little trickier

transferring into the dil fridge. Since you will always need to transfer shortly after loading the fridge, be sure to plan things out in advance and have your storage dewar full before you load the fridge. Assuming that you have plenty of liquid in the storage dewar (you will probably use $\sim$ 60L to fill the dewar all the way if the liquid level is around 2 inches), check the gas meter webpage (wait to transfer if more than one other person is transferring). Go get the taller step ladder from the ^{3}He system, move the storage dewar into position in front of the fridge, and start lowering the transfer tube into the storage dewar (make sure the extension on the transfer tube is threaded in tightly). If the storage dewar is mostly full, you should not need to pressurize it to get the liquid flow started. As soon as fog starts coming out the end of the transfer tube, pull out the cork on the transfer port and insert the transfer tube a few inches. You will probably hear the gas meter start spinning faster and the dewar pressure may rise a little bit, but this should slow down in a minute or so. Slowly lower the transfer tube until the 90° elbow in the transfer tube is about even with the top of the green switch boxes on the fridge. Right now it doesn't matter so much, but when the magnet is energized you do not want to put the transfer tube down all the way. If the tube is pushed all the way into the cone on the magnet, everything coming out of the transfer tube will wind up at the bottom of the magnet. If there are any warm gas bubbles coming along with the liquid, there is a small chance this could cause a magnet quench. This is not very likely, but given that this magnet can go as high as 15T without the lambda plate, a quench would probably vaporize all the liquid in the dewar (and remember that 1L of liquid helium turns into 26.6 ft^3 of gas at room temperature).

Once you get the transfer tube down into the fridge, start pressurizing the storage dewar. It will probably take 75-90 minutes to fill the fridge all the way to a liquid level of 16.7 inches. Close the pressure building valve on the storage dewar when the boil off rate gets up to $\sim$ 6 CFM, and open it back up again when the boil off drops back down to $\sim$ 3.5 CFM. Try to plan things out so that you end the last pressurization of the storage dewar when the the liquid level in the fridge is $\sim$ 15.6

inches. This will let the fridge and dewar settle a little bit before you pull the transfer tube out. Don't get caught off guard; the last half an inch goes a lot faster than the rest of the transfer. When you see that the transfer is almost finished, valve off the helium regulator, take the rubber hose off the pressure-building valve, get the step ladder in position, and grab yourself a pair of leather gloves. When the liquid level hits 16.7 inches, crack the pressure building valve open to let some of the pressure off into the room, quickly loosen the quick-connect fitting on the fridge transfer port (be careful not to unscrew it all the way), and pull the transfer tube out in one smooth, quick motion. If you go too slow or pause at any point, you will probably freeze and rip the o-ring. In case this does happen I try to keep a spare o-ring right in front of the level meter so it is within arm's reach in case the o-ring on the fridge breaks. If you rip that one too, there is a lifetime supply of the appropriately sized o-rings in the plastic organizer drawer on the bench by the door (look for the drawer labelled "Oxford o-rings").

As soon as you get the transfer tube out of the fridge, plug the port with the cork. I like to have the cork in my left hand while I pull the tube out with my right arm to plug the hole as quickly as possible. This dewar gets pressurized a lot more during the transfer than the ^{3}He system does, so you will lose helium a lot more quickly if you are fumbling around with the cork. You also will probably burn your fingers if you aren't wearing gloves at this point. Once you get the cork in, tighten it down, open the pressure building valve on the storage dewar the rest of the way to minimize how much liquid you are continuing to pump into the transfer tube, climb up the step ladder and pull out the transfer tube. Be careful not to whack anything in the electronics rack with the transfter tube as you get down. Close up the valves on the storage dewar and put everything away.

H.6 Preparing to Condense

Now that the fridge is full, you can start your final preparations to condense and circulate. You could have done some of this before or during the transfer; most likely the order you do things will be decided by the liquid level in the fridge, when other people are transferring, etc. The first thing to do is to get the fridge hooked up to the pumps and gas handling system. Hook up the 1K pot line to the 1K pot pump using an o-ring with a plastic centering ring and an all-plastic clamp to keep the fridge electrically isolated from the pump. Do the same for the still and condenser lines but note that the still line isolation occurs where the NW50-NW40 adapter attaches to the gate valve on the fridge, so you can use a metal centering ring when you attach the NW50 hose to the still port. Make sure you line up the sharpie marks on the NW50 hose with the marks on the NW50-NW40 adapter so that you don't twist the connections when you pump out the lines and the hose shrinks up.

Start up the leak detector and hook it up to the pump-out port on the sand-bucket manifold shown in figure H.10. Make sure the valves to the helium trap and to the long still hose are shut (essentially all of the valves on the gas handling circuit should be shut at this point so you don't risk pumping out any mixture). Pump out the still and condenser hoses and leak check the connections you just made with the leak detector. You should be able to get the leak rate below 1×10^{-8} atm cc/s. If you have a background higher than this, let it pump for a while to make sure it falls below this level. Once you are convinced there are no leaks, seal off the sand-bucket manifold and remove the leak detector. If it is handy, hook up the turbo pump and let it pump on the lines that were exposed to air for a while; the longer you can pump these out the better in terms of removing water from the vacuum surfaces.

Once you have the leak detector free, hook it up to the IVC port and pump out your lines. In the meantime get ready to start putting some heat in the sorb. Hook one of the small green HP power supplies up to the sorb connection on the blue Pomona box shown in figure H.11. Turn on the power supply and monitor its output

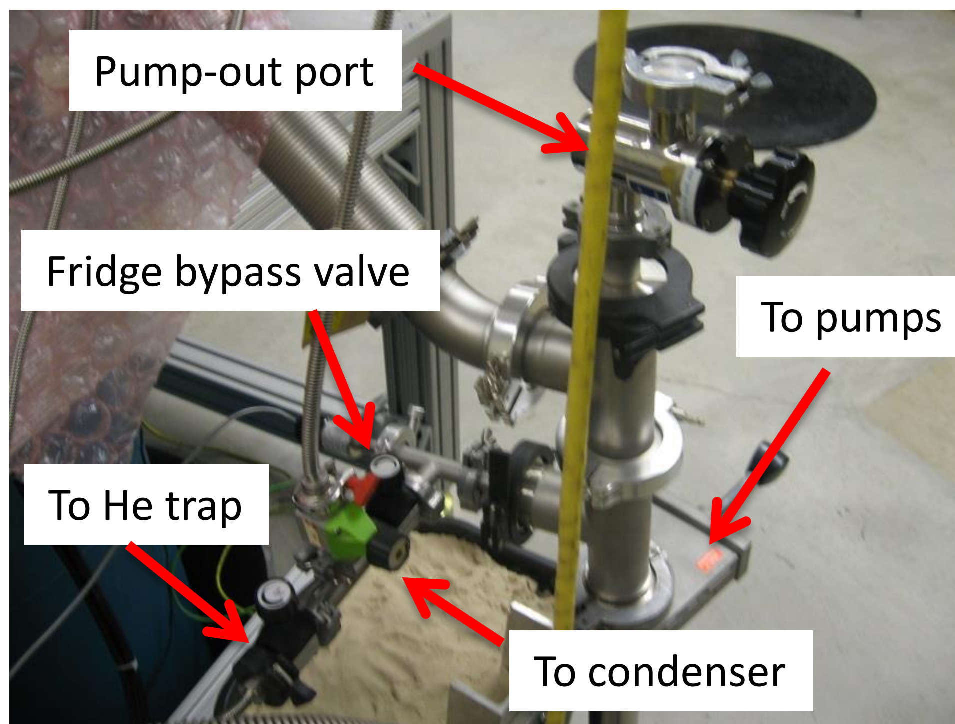

Figure H.10. Sand-bucket manifold.

with the Keithley multimeter (make sure the output is as close to zero as you can get it). Once the leak detector has the line pumped down start pumping on the IVC. As a precaution (in case of a leak in the indium seal) first close the throttle valve on the leak detector and then open the IVC valve. Assuming there is no leak, the pressure the leak detector sees should still be well below 0.1 mBar, so you can just open the throttle valve back up without using the rotary vane pump to rough it down. If you turn the ion gauge on, you should see that the pressure is below the minimum detectable level but the helium leak rate is probably over full scale. Shut off the ion gauge so you don't damage any of the electronics. Finally, switch the sorb from "NC" (not connected) to "C" (connected) and turn the HP supply up to 4.6V slowly. You should see the sorb reading in Labview rise to $\sim$ 40K. Try to keep the sorb at this temperature; you will have to decrease the voltage over time as you pump out more of the exchange gas and the sorb becomes less and less coupled to the 4K walls of

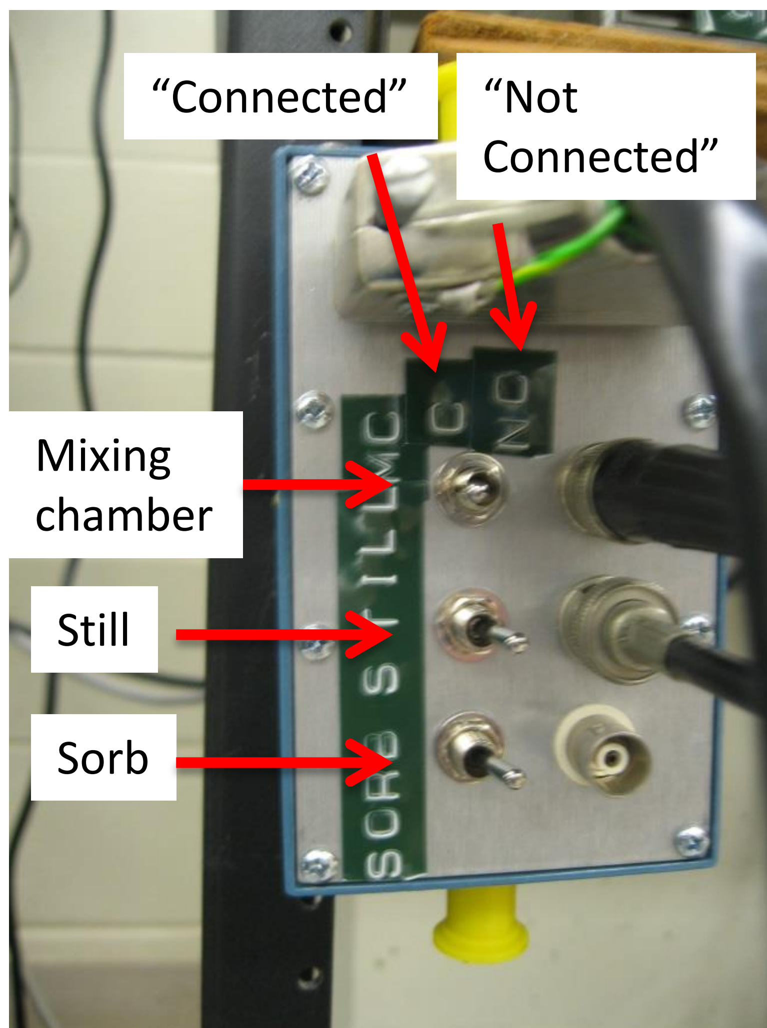

Figure H.11. Fridge heater connections.

the IVC. Pump like this for about an hour and make sure you keep the nitrogen trap in the leak detector full so you don't pump a bunch of oil into the fridge. Once you pump for an hour or so, turn off the heat to the sorb and seal off the IVC. As the sorb cools back to 4K the charcoal gauze should act to trap the rest of the helium left in the IVC so that the dilution unit is thermally isolated from the 4K bath. Make sure the sorb is all the way cold before you cool down the 1K pot; if you cool down the 1K pot too soon, you can create a superfluid film of ^{4}He on the dilution unit which will put a heat load on the mixing chamber and limit the base temperature.

If you haven't already stopped pumping on the still and condenser hoses with the turbo pump, seal everything off, shut down the turbo pump, and get it out of the way once it spins down. You next need to start "washing" the mixture, meaning that you will circulate it through the pumps and traps but bypass the fridge. Start by plugging in the Pfeiffer sealed pump. Make sure you plug it into the correct power cord because the phases in each power cord may be wired differently, so plugging the pump into the wrong cord could potentially cause it to run backwards. When you plug it in, keep an eye on the mechanical pressure gauge on the wall manifold. It should come up a little bit as it pumps the mixture that backstreamed through the pump while it was off, but it should not rise above -7.5 inHg. Once you've had the sealed pump on for a minute or so, plug in the roots blower. It is very important that the sealed pump always be running when the roots blower is on; if it is not and the outlet pressure of the roots pump gets above atmospheric pressure, the roots pump could be damaged.

Next, go around and start opening up the valves in the circulation path. Start with opening the fridge bypass valve shown in figure H.10 and then open the helium trap to the sandbucket manifold and watch what the still pressure gauge shows. Next, open valve 11 on the wall manifold shown in figure H.12 and watch the condenser gauge. Next, open valve 14 so that the wall manifold and the sandbucket manifold are connected. At this point the still and condenser gauges may be reading a few hundred mTorr from small leaks and outgassing, but this is not a big deal. As long as the mechanical condenser gauge is still showing -30 inHg, it is fine. Next, open the gate valve on the sand-bucket manifold to connect the He trap volume to the long still hose. The pressure should drop when you do this because the connections on the long hose seem to pretty good. Lastly, open valve 10 on the wall to connect the nitrogen trap to this whole volume. As long as you don't see any response on the mechanical condenser gauge there is not much gas in the lines. You need to check that this is the case before opening the lines to the pump; if there is a signifcant leak such that this large volume has an appreciable amount of gas, you could over-pressurize the back

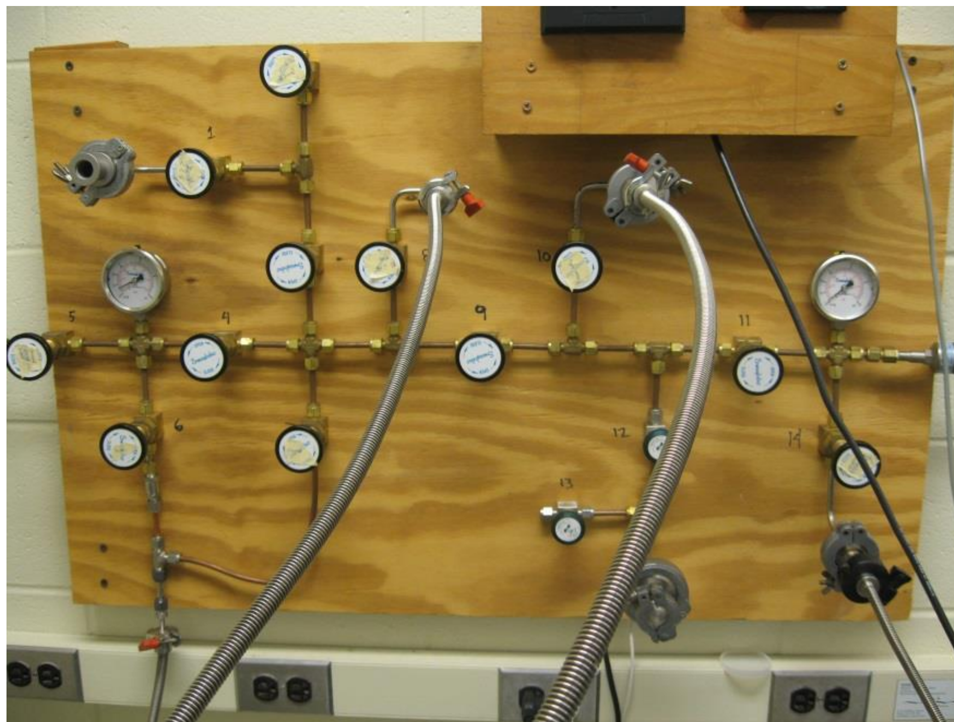

Figure H.12. Primary gas handling manifold.

side of the pumps. Once you are convinced everything is safe, open the gate valve with the red tape on the handle on top of the roots blower shown in figure H.13 to connect the pumps to the rest of the circuit. Do NOT open the bypass gate valve on top of the roots blower. This valve should never be opened while the roots blower is running since it would allow the inlet and outlet of the roots pump to equilibrate and damage the pump. This bypass valve was included when we originally put the fridge together so that we could circulate with just the sealed pump (it was not *a priori* obvious whether or not using the roots blower would raise or lower the base temperature of the fridge). Finally, make sure valve 7 is shut to keep the majority of the mixture sealed in the dump, open valve 4 all the way, and just crack valve 8 to start letting a little mixture circulate. Keep slowly opening valve 8 until the system equilibrates with the still gauge showing $\sim$ 300-400 mTorr.

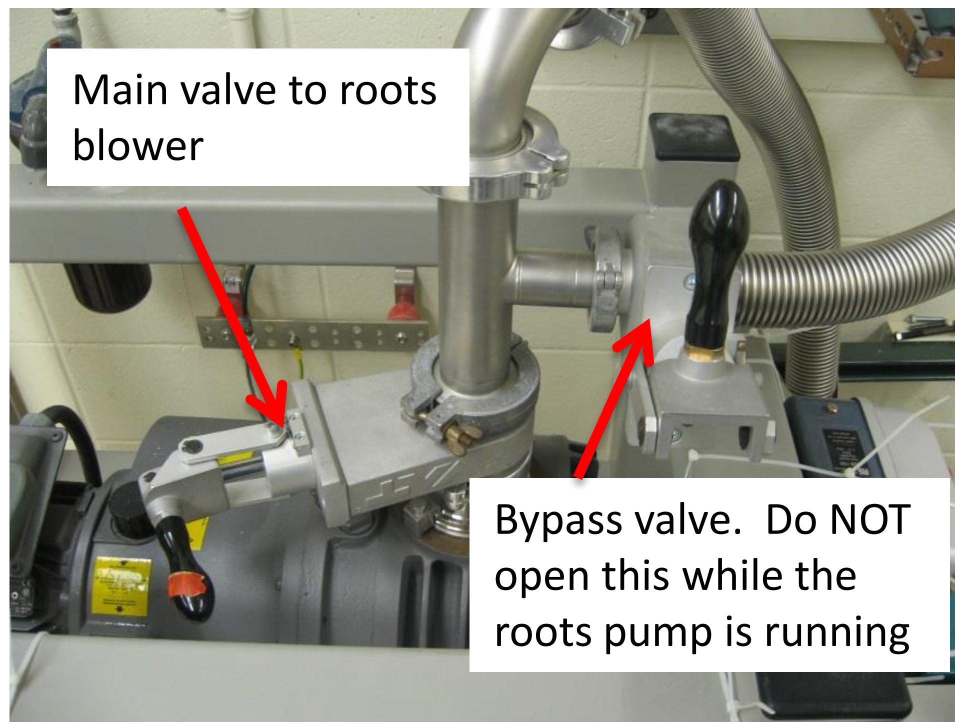

Figure H.13. Roots blower manifold.

After a few minutes when you are convinced that the circulation rate is stable, slowly start cooling down the nitrogen trap. Try to stabilize the hoses on the trap so that you don't put a bunch of strain on the connections on the trap itself or on the wall manifold. Once the trap is cold, let the circulation continue for 5 or 10 minutes. Once you are convinced the circulation rate is stable, cool down the helium trap. Since you are already circulating mixture you need to be extremely careful while you are moving the trap into position over the fridge so that you don't put a lot of stress on a connection and open up a leak. Pull the cork out of the port on top of the fridge and slide the trap in a few inches. The weld seam on the bottom of the trap is a bit too large, so it can be a little tough to get it started. Try to wiggle it around a little bit but be very careful to not bend the trap at all. Slowly lower the trap into the fridge over the course of a few minutes and keep an eye on the condenser gauges and

the still gauge to make sure there are no sudden jumps in the readings that could indicate a leak or plug.

Once you have both traps cold, you can go ahead and open up valve 7 to allow the mixture in the dump to join the circulation. The goal with this whole washing procedure is just to make sure that whatever junk may have leaked into the lines gets caught in the traps before you start condensing. You may need to adjust valve 8 on the wall a bit to keep the still pressure around 400 mTorr since the conductance of the traps changes a bit as they cool. Do not let the pressure get above this; at higher pressure (i.e. higher circulation rate) the outlet of the nitrogen trap gets very cold and this could cause the rubber o-ring to start to leak. If everything looks stable, let the mixture circulate like this for at least half an hour.

H.7 Sample Checks and Illumination

While you are letting the mixture wash, you can keep yourself busy by checking the 2-terminal resistance of all the Ohmic contacts on your samples. Remember to change the switch boxes on top of the fridge to the "A" position after you hook up the D-SUB cables. If you need to illuminate your samples, you can also do that now. If the Keithley source-meter is available I prefer to use it for powering the LEDs since it can limit the voltage being supplied (a useful safety in case you have an LED hooked up backwards). I have generally been using just LEDs 1, 2, and 3. LED #4 has its negative terminal shorted to ground. In principle, you could still use it as long as you keep this leg as the end of your chain of series LEDs, but I've been a little nervous about what it could potentially do to the samples. Set the current to 2mA. The voltage drop across LEDs 1, 2, and 3 in series should be 8.2V, so set the voltage limit on the supply to 8.3V. Make sure all your contacts are grounded and illuminate with the LEDs for 15 minutes. If you have a gated sample now would also be a good time to check that the gate is not shorted to anything.

H.8 Condensing and Circulating

Once you have determined that your samples are still working, start the condensing process. First, make sure the back of the 1K pot pump is open to the helium recovery lines and turn on the 1K pot pump. Open the valve on the fridge to the 1K pot. The 1K pot pressure reading should initially jump up a lot and then quickly fall to near zero since the needle valve is completely shut. Switch the bridge and your Labview temperature monitor code to read the 1K pot and start opening up the needle valve. The needle valve tends to stick a bit when you open it up the first time, but if it is really stuck don't force it. Once you open the needle valve and get the pressure up to a few Torr, the 1K pot should cool to 1.7K in less than 30 seconds.

The fridge is now ready to start condensing the mixture, so note this in the log sheet and notebook. First, close the fridge bypass valve on the sand-bucket manifold and close the gate valve on top of the roots blower. Next, open the condenser valve on the fridge and watch the 1K pot pressure and temperature reading in Labview. You may also need to adjust valve 8 on the wall. You should try to keep the condenser pressure $\sim$ -26 inHg and the 1K pot temperature $\sim$ 1.7K. Keep monitoring the dump pressure over time. You should be able to condense 95% of the mixture in 60-90 minutes. Also keep an eye on the mechanical gauge on the backside of the pumps. Once it gets to -20 inHg, close valve 4. The manual says to keep the outlet pressure > -22.5 inHg at all times, though it doesn't say what might happen if the pressure drops below that point. Keep condensing into the fridge; you will probably have to keep opening valve 8 as you go to keep the condenser pressure $\sim$ -26 inHg. Once the dump gets to -28 inHg, you have gotten as much mixture out as you can with the 1K pot (the 1K pot is evidently not cold enough to get the vapor pressure of the mix any lower than this).

Next, close the condenser valve on the fridge to seal most of the mixture in the fridge. Now you need to pump the dump out with the pumps. Open the bypass valve on the sand-bucket manifold and the gate valve on top of the roots blower.

You should see the mechanical gauge on the backside of the pumps start to come up slowly (it will probably get to somewhere between -15 and -10 inHg). Keep pumping on the dump like this until the still pressure gauge drops below 100 mTorr. Beyond that you aren't going to really gain much by continuing to pump. Close valve 7 to seal off the empty dump.

Now put the fridge back in the condensing configuration to put this excess gas into the fridge. Close the fridge bypass valve on the sand-bucket manifold, open the condenser valve, close the gate valve on top of the roots blower, and slowly open valve 4 on the wall to let the mixture from the backside of the pump condense (do not open valve 4 all the way). The pressure on the backside of the pump should drop pretty quickly since you have increased its volume by a lot.

Now you are ready to start trying to circulate, so quickly make a note of this in the notebook and logsheet. Open the still gate valve on the fridge; you will probably hear some gas rush through it as it fills the big hose. Make sure the gate valve on top of the roots blower is shut before you do this so you don't slam the pump with a large pressure. Now, very slowly start to open the gate valve on top of the roots blower. This gate valve opens slowly enough that you can throttle the flow a little which is important in this case. Until you get the still cooled down, the still pressure will be very high and you will over-pressurize the gas handling system and warm up the 1K pot if you pump too much gas out of the still.

As you are opening the gate valve on top of the pump, watch the mechanical gauge on the wall. Close the gate valve back up when the pressure gets up to atmospheric pressure. Go over to the computer and check the 1K pot and still temperatures. The 1K pot should be staying $\sim$ 1.7-1.8K, and the still should have cooled a bit below the 1K pot temperature, though it is probably rising again now that you closed the gate valve. Go back to the roots pump and, as long as the mechanical gauge on the backside of the pumps is $<$ -15 inHg, slowly open the gate valve again. Close it back up once the pressure gets up to atmospheric pressure. After you do this cycle a couple times, you should see the still pressure start to drop into a measurable range

(i.e. 400-500 mTorr) at which point the still temperature has probably dropped to $\sim$ 1K. Once the pressure on the backside of the pumps stays below -10 inHg with the gate valve open, you are done. Just leave the gate valve open and let things start to settle as you monitor the temperatures and pressures of the system. The cooldown to base varies a lot from cooldown to cooldown. Sometimes the mixing chamber will cool to 25mK in as little as 20 minutes, other times it will take 90 minutes. Anecdotally, it seems that the total time to condense and cool to base is roughly constant. If you do the condensing quickly, it takes longer to cool the rest of the way to base because the still temperature (and hence circulation rate) drops quickly; whereas if you take a long time to condense, it seems that the still stays warmer and the mixing chamber cools more quickly. Once the pressure on the backside of the pumps has dropped below -16 inHg, open valve 6 on the wall. This opens a path for the mixture back to the dump through a 10 psi check valve in the event of a plug. While things are settling down you may need to adjust valve 4 a little bit. A typical condenser pressure is $\sim$ -28 inHg which is too low for the sealed pump, so you need to create a bit of a pressure drop across valve 4. Try to set the pressure on the backside of the pumps to $\sim$ -20 inHg. Typical parameters for the fridge at base temperature are as follows:

- 1K pot $\sim$ 1.7 K
- Still $\sim$ 700-800 mK
- Cold plate $\sim$ 30 mK
- Mixing chamber $\sim$ 11 mK (sometimes the first day after cooling to base it only gets to $\sim$ 12-13 mK)
- “J1” (i.e. tail) $\sim$ 1 mK higher than mixing chamber
- 1K pot pressure $\sim$ 3-7 Torr
- Still pressure $\sim$ 90-100 mTorr
- Condenser pressure $\sim$ -28 inHg

- Back of pumps $\sim$ -18 inHg
- Still power at base temperature = 0

Once you get the mixing chamber close to base and everything looks stable, your experiment will dictate what you do next. If you haven't already done so for these samples, you should probably check the resistivity (as well as the gate leakage, if applicable) before you start sweeping the field. If you already have this information from a previous cool-down, you might as well sweep out to $\nu = 2$ to let the sample sit overnight so you can take a slow down sweep through the 2$^{\text{nd}}$ LL the next morning to see how things look. If this is the case, stop the Labview bridge monitor program, but before you do tell it to measure "Mike's RuO" (this is meant for channel 4 on the bridge). Once Labview records at least one data point, go ahead and tell the program to stop and create a folder to save your data (preferably in My Documents/Instrument Logs/Temperature Logs/...). Please give the folder a meaningful name like "2014-09-01 cool fridge from 300K to base" so that you or someone else can track down the information in the future. Labview will save the data from each thermometer in its own file (hence why you need to create a new folder), and the independent variable will be a time stamp (i.e. it includes the date and time). I created an import template and graph template in Origin to allow you to import this data easily and display it all on a graph with a time stamp axis (this makes it easier to compare the data in the graph with notes in the notebook or logsheet). If you haven't collected data for each thermometer, though, the program has problems and won't actually save any data (hence why you needed to give it a couple fake data points for the "Mike's RuO" thermometer). Once Labview has given up control of the GPIB bus, you should be ready to start your experiment.

H.9 Cooling Down the Magnet

If the magnet was also warm before you started, there is obviously some additional work you will need to do. Preparing the samples and loading them onto the tail is,

of course, the same as is the room temperature preparation of the fridge. The first difference in the procedure is loading the fridge into the dewar. Since there is no liquid in the dewar, you can load the fridge into the dewar quickly. Just remember to pull the plug out of the transfer port once you get the sliding seal mated with the top of the magnet to let the pressure out as you lower the fridge (the dewar should be valved off from the helium recovery system at this point so you don't lose all of the department's helium). Once you get the fridge down, screw it into the magnet and hook up the ground straps like I previously described.

Hook up the still and condenser lines, leak check them, and start washing the mixture through the nitrogen trap like I described before. Since the fridge may have been sitting for some time since it was last used, you need to be extra careful about checking for any large gas load in the gas lines as you start pumping on the circuit. There is a small leak somewhere in the system (possibly one of the pumps) that you should be aware of. As a result, you should plan to wash the mixture through the nitrogen trap overnight and then warm up and clean the trap out the next morning so that you don't start with a large gas load in the trap. To clean the trap, first close valve 8 to stop the circulation of mixture and pump the remaining mixture out of the trap. Wait until the condenser gauge drops to $<$ 15 mTorr to be sure that you got all the mixture out of the trap. You should also confirm that the dump pressure came up to -7.5 inHg. While this is pumping out, get the turbo pump and hook it up to the pump-out port on valve 12 on the wall (you may need to find someone to help you lift the pump over the still line). Be very careful with the pump cart around the dump; the connections on the dump are very flimsy, so if you run into them with the pump cart you could potentially lose a lot of mixture. Pump out the line and make sure the pressure drops enough to be sure you don't have a leak, but do not open valve 12. Turn off the ion gauge once you are convinced there are no leaks so that the filament can cool down.

Once the trap is completely empty of mixture, close all the valves on the gas handling system except valves 10 and 11 so that you can monitor the pressure in

the trap with the condenser mechanical gauge. Get a heat gun and some zip-ties within arm's reach and then pull the trap out of the dewar while being careful to minimize the strain on the connections. Zip-tie the trap to the handle on the dewar so it doesn't fall over if you have to walk away. Start gently heating the trap with the heat gun. Don't leave heat gun in the "hot" position for very long; just heat the filaments up and switch to "cold" mode so that you blow luke-warm air over the trap. You should be able to hold the heat gun an inch from your arm without burning yourself. Periodically turn the filaments back on for 10-20 seconds to keep the air warm. Try to thaw the whole trap out uniformly so that you don't develop too much thermal stress on the trap. While you are doing this you should also be keeping an eye on the mechanical gauge. Don't let the trap get above atmospheric pressure. If it does start to go positive, stop heating the trap, close the valve on the top of the turbo pump to protect the pump, and open valve 12 on the gas manifold to vent some of the gas into the vacuum hose. Then you can stop the pump, let the blades spin down, and then pump out the trap. Normally, though, this won't happen. As long as the trap pressure doesn't go positive, keep heating it until the whole trap is warm (but not hot) to the touch. Record the pressure it came up to in the front of the log book along with how long it was sitting since the last use. Pump out the trap with the turbo pump. There will be a large helium background, so you aren't going to accomplish much by leaving the pump running for more than 20 or 30 minutes. The condenser gauge will probably show that it is stuck somewhere around 2 Torr. Seal everything back up at this point so you don't put unnecessary strain on the pump.

Before you can start cooling down the dewar, you will need to pump and flush it with dry nitrogen or helium so that you don't clog anything up with ice. Hook up a gas line to the dewar purge port shown in figure H.14. Pump the dewar down with the 1K pot pump by opening the dewar pump-out valve. Open this slowly and stop when the pump starts howling. You can open the valve a little more once the noise dies down. Monitor the progress with the 1K pot gauge by opening the 1K pot manifold valve shown in figure H.15. You should be able to pump the dewar down

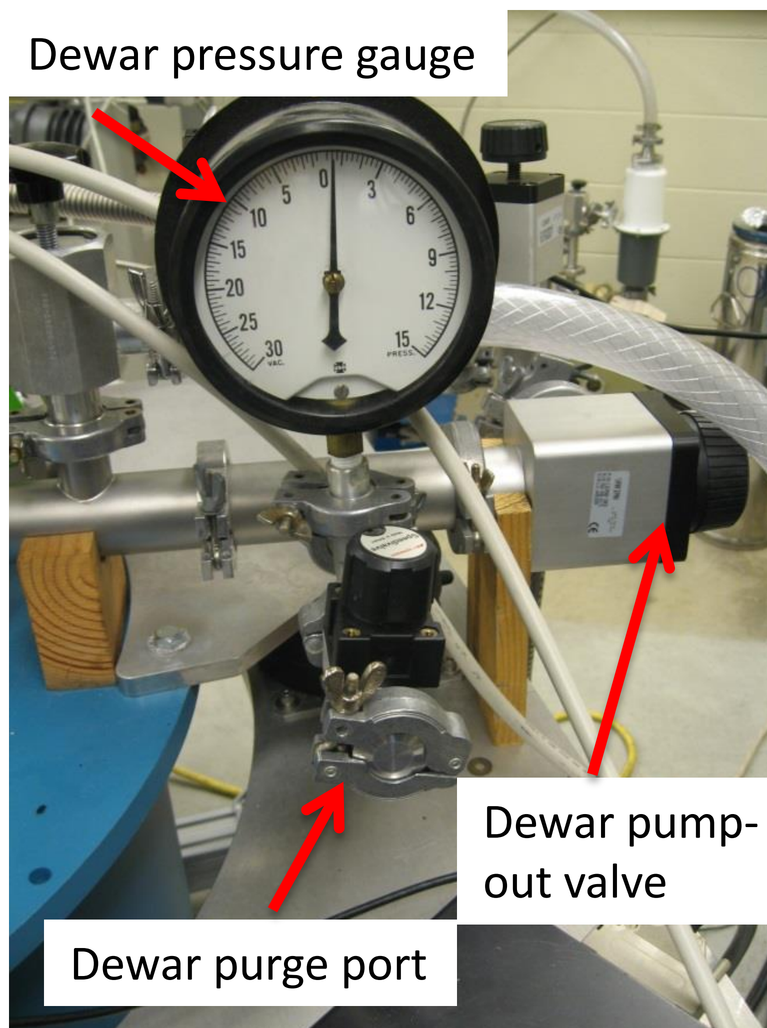

Figure H.14. Dewar purge manifold.

to < 1 Torr. If you can't get this low, look around for leaks, try tightening down the screws bolting the fridge to the magnet, etc. Once the pressure bottoms out, close off the dewar pump-out valve and fill the dewar back up with dry gas through the dewar purge port. Do this pump/purge cycle 3 times. When you fill the dewar up with gas the last time, leave it pressurized $\sim$ 1-2 psi.

Next, pump and purge the 1K pot (use helium and not nitrogen for purging). Since the 1K pot gauge is not sensitive around atmospheric pressure, you will have to hook up some kind of tee to watch its pressure (the NW16 4-way cross used for

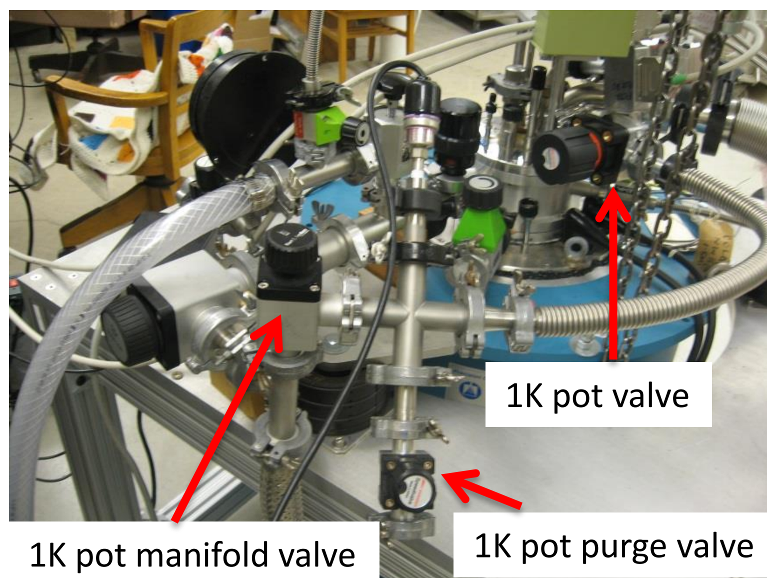

Figure H.15. 1K pot manifold.

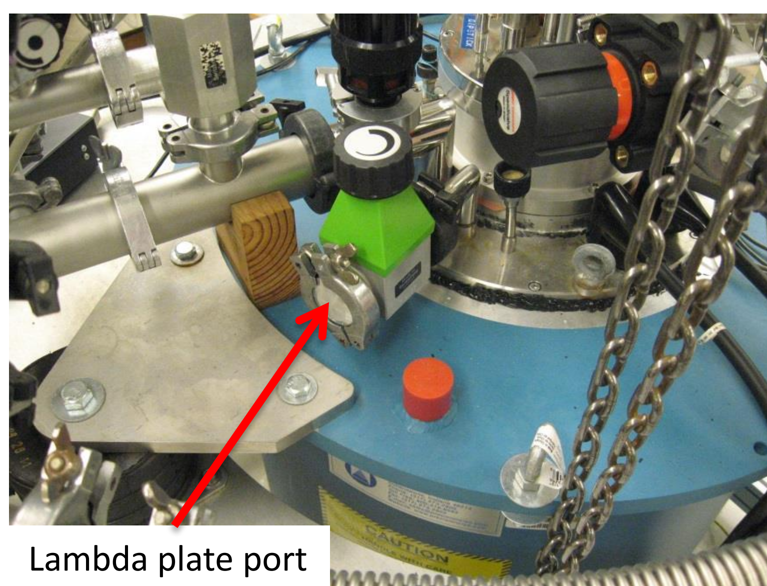

Figure H.16. Lambda plate port.

adding exchange gas to the IVC is a good option). Pump and flush the 1K pot 3-5 times and leave it pressurized $\sim$ 1-2 psi with helium when you are done. Close down

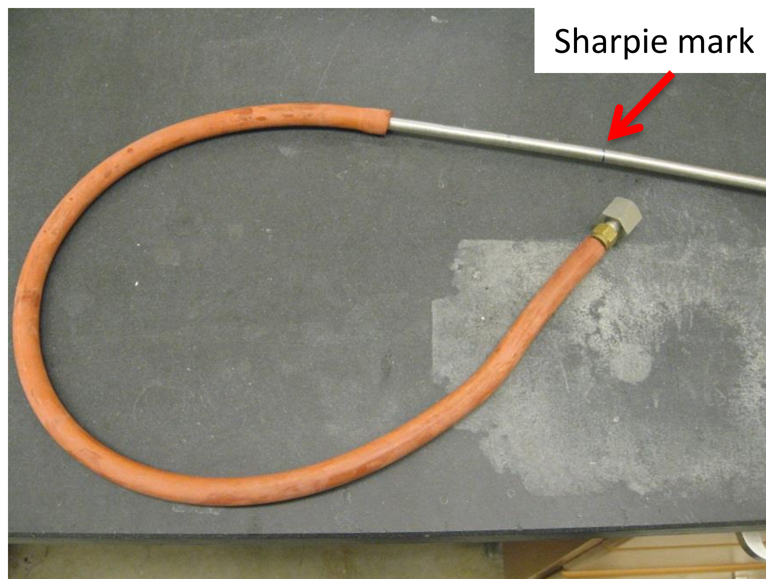

Figure H.17. High-tech LN_2-fridge coupling unit.

the 1K pot valve and then hook up the 1K pot manifold to the lambda plate port shown in figure H.16. You will probably need to grab a more flexible vacuum hose to do this. Make sure the lambda plate needle valve is shut all the way before you start pumping and purging. Pump and flush the lambda plate 3-5 times and leave it pressurized $\sim$ 1-2 psi with helium.

Now that all the cavities in the fridge are either under vacuum or filled with dry helium, you can start pre-cooling the system wiith liquid nitrogen. You will need to borrow one of the big nitrogen dewars from Keith. Talk to him in advance and remember that he really needs all three dewars when he is running the liquefier, so try to plan your cool down to coincide with a day when the liquefier is shut down. Hook your nitrogen dewar up to the nitrogen-fridge coupling unit shown in figure H.17 (it should be on the bottom shelf in the storage room). Make sure the end of the stainless steel tube (should be propped up against the door jamb by the green cabinet in the dil fridge lab) with the sharpie mark is the end you have hooked up to the rubber hose. This marks how far you have to shove the tube into the fridge to

Figure H.18. Primary lab helium recovery valve.

get the tube all the way in to the cone on top of the magnet. Temporarily close the dewar recovery valve and remove the plastic recovery hose. You will be venting a lot of nitrogen out this port during the cool-down, and you don't want to fill the helium recovery system up with nitrogen. You should also make sure the primary lab helium recovery valve shown in figure H.18 is shut so that you don't lose any helium from the recovery system by disconnecting the fridge.

Once you have the recovery line disconnected, crack open the valve on the nitrogen dewar to start letting a little nitrogen purge the hose and your stainless steel tube. You don't want a large gas flow yet, so just open it a little. Pull the cork out of the dewar transfer port and slide the stainless steel tube in quickly so you don't get too much air into the dewar. As soon as you insert the tube into the fridge, open the dewar recovery valve so that you don't blow up the dewar. Get the stainless steel tube all the way inserted so that you feel it stop in the cone on the magnet and see that the sharpie mark is just barely visible above the plane of the dewar transfer port. Get some zip-ties and tie off the liquid nitrogen hose to help support the weight and

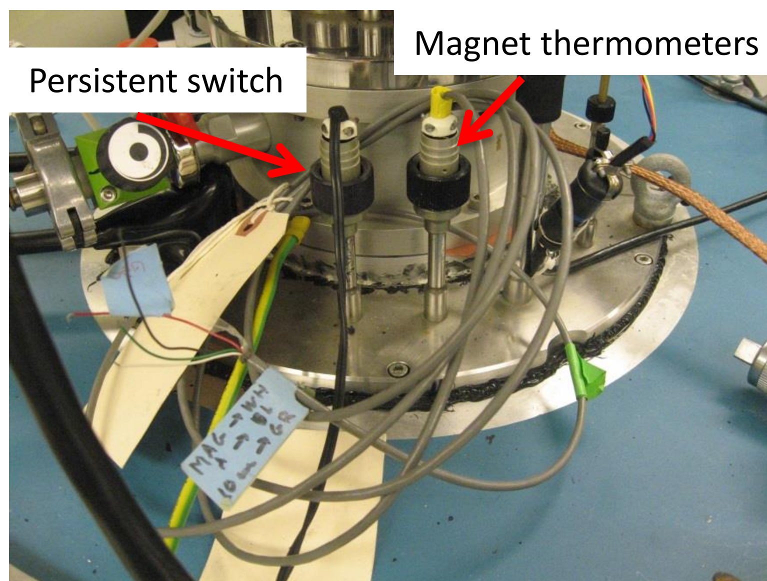

Figure H.19. Magnet thermometers and persistent switch.

then open the valve on the nitrogen dewar a bit more. You should feel a strong, cool breeze coming out the dewar exhaust port, but you shouldn't hear any howling or see any fog yet. According to Oxford, this particular magnet design can be damaged during the pre-cool in one of two ways. First, if you don't get the tube all the way into the cone on top of the magnet, you will shower the top of the magnet with liquid, and this thermal shock can cause damage. By getting the transfer tube into the cone, you instead spray the bottom of the dewar with liquid and the resulting vapor does the initial cooling of the magnet. Second, even with the transfer tube inserted all the way, you can also damage the magnet if you cool down too fast. We unfortunately don't have a good way to quantify this, so you will just have to compare your cooling rate with historical values. You can watch the magnet cool by hooking the Fluke multimeter up to measure the magnet thermometer leads shown in figure H.19. The red wire is a common ground, the white wire goes to the Allen-Bradley resistor on the bottom of the magnet, the black wire goes to the resistor on the lambda plate on the top of the magnet, and the green wire goes to a resistor 10 cm above the lambda plate

(approximately the height of the mixing chamber). Record the resistance readings for each resistor over time in the notebook and compare with previous cool-downs to gauge if you need to open or close the valve on the nitrogen dewar. Keep filling with nitrogen until the resistance of all three resistors has plateaued. This will probably take over two hours. Once a steady stream of fog is shooting out the dewar exhaust port, try to angle the exhaust up with some aluminum foil so you don't blast the magnet supply with cold, wet air for an extended period.

Once all three magnet thermometers and the sorb thermometer stabilize, you are ready to cool to 4K. Depending on the time you may decide to leave the nitrogen in the fridge overnight. If you do this (or if you need to leave the system for very long with no liquid flowing into the fridge), put a blank on the exhaust port but do not use an o-ring. This will cut down on the conductance enough so that no air will get into the dewar and freeze, but the nitrogen boil-off will still be able to escape. When you do decide to blow out the nitrogen, first crimp the rubber hose on the liquid nitrogen-fridge coupler and disconnect your stainless steel hose from the liquid nitrogen dewar. Go find a few small nitrogen dewars to collect the liquid you blow out of the fridge. Then hook up a helium gas line to the dewar purge port shown in figure H.14. When you are ready to start blowing out the liquid, close the main dewar recovery valve, un-crimp your rubber hose so the nitrogen can get out and start pressurizing the dewar with helium. Pressurize the dewar to ~ 2 psi to force out the liquid. If you don't have the stainless steel tube all the way down into the cone on the magnet, you won't be able to blow out the liquid, so you might want to check that this is really all the way down before you start trying to blow it out. You will know all the liquid is out when you stop seeing liquid trickle out the nitrogen hose and see the pressure on the dewar gauge suddenly drop to close to atmospheric pressure. When all the liquid is out, keep the dewar pressurized with helium, warm up the transfer port with a heat gun to thaw out the o-ring, and then pull out the stainless steel tube and cork the transfer port. Remember to valve off the dewar purge port at the same time so you don't keep pressurizing the dewar.

Next, you need to pump and purge the dewar since it still has a lot of nitrogen gas left in it. Do this the same as you did at room temperature, but remember that since the dewar itself is a lot colder, it will be holding a lot more gas and take longer to pump down. As long as you can get it down to ~ 1 Torr don't worry too much. If it gets stuck a lot higher than this, though, you may still have liquid in the bottom. If that is the case, you need to try blowing it out again. If there is any liquid nitrogen left in the bottom of the dewar, you will boil off all your helium trying to cool down and freeze the nitrogen (not to mention you may end up plugging up the 1K pot sipper or lambda plate sipper) when you start transferring helium. Once you get the dewar purged a few times, leave it pressurized with helium gas. You should now be ready to fill up the fridge with helium, so if you have a full storage dewar start transferring as you normally would, but start out slow. You will be converting all the liquid to gas for roughly the first 30 minutes, so make sure the recovery system is ready to handle a big gas load. You also, of course, need to make sure that the fridge is hooked up to the recovery line and that the main valve to the lab shown in figure H.18 is open. While you are getting setup, though, keep the recovery valve on the fridge itself closed so you don't suck air from your recovery lines into the fridge. As soon as you start transferring, make sure you open the recovery valve on the fridge so that the gas has somewhere to go.

The transfer will probably take ~ 2 hours. You can watch the magnet cool down with the resistors just like you did when you pre-cooled it with liquid nitrogen. The liquid level meter starts just above the "10 cm" resistor, so once the resistance of this resistor plateaus wait a few minutes and then try turning on the liquid level meter. Don't turn the meter on too early as you can damage it by turning it on if it is too warm. A full 100L storage dewar should be able to fill the fridge up to ~ 12 inches on the level meter. Once the transfer is done, continue cooling to base as previously described.

H.10 Warming Up the Insert

Once your experiment is done, you will need to get the fridge back to room temperature so you can load a new set of samples. The first step is to pull all the mixture back to the dump. Start up the bridge temperature monitoring program in Labview, make a note in the notebook and logbook, and then close valve 8 to cut off the flow of mixture to the fridge. You will need to get everything in the fridge up to 4K to get all the mixture out, so close down the 1K pot needle valve to dry it out as much as possible and then close the 1K pot valve to stop pumping on it so that it will warm up. You can also open the fridge bypass valve on the sand-bucket manifold so that you pump on both sides of the fridge. Then start adding some heat to the still and mixing chamber to accelerate the process. Your ultimate goal is to put 2.9V into the still and 3.3V into the mixing chamber, but work up to these values over the course of $\sim$ 30 minutes so you don't pressurize anything in the dilution unit. The pressure on the backside of the pumps will first rise quickly as you pump all the He3 out but then slow down after that. The still will also probably warm up suddenly when you get all the ^{3}He out. While you are pumping things out, keep an eye on the pressure gauge on the backside of the pumps. When this gets to $\sim$ -10 inHg, open valve 7 to let some gas into the dump, but don't let the pressure get too low to avoid damaging the sealed pump. After $\sim$ 30 minutes, though, you should have enough gas out that you can leave the valve open to the dump. Pump like this until the still pressure drops suddenly to 0 and the cold plate warms up to 4K. Once this happens, close the valve to the He trap on the sandbucket manifold so that you are only pumping on the fridge and no gas from the traps can sneak into the fridge. Pump on just the fridge for 10 minutes to be sure everything is out and then valve off the condenser and still and turn off the still and mixing chamber heaters. Open the valve to the traps back up and pump them out until the condenser gauge drops down to $\sim$ 10 mTorr. By this point the dump pressure should be at -7.5 inHg. If it's not, you lost some mixture or didn't pump it all out.

Once you are convinced the mixture is all back in the dump, clean the traps. I already described how to clean the nitrogen trap, and the helium trap cleaning is very similar. Valve off everything but valves 11 and 14 on the wall so that you can watch the trap pressure on the mechanical condenser gauge. Make sure the heat gun is within arm's reach and then pull the trap out of the fridge slowly in steps. Try to thaw the o-ring out as you go, and be very careful not to bend the trap or yank the vacuum hoses. Once you get the trap all the way out, put the plug in the port and thaw out the trap. Record the pressure in the trap and then pump it out with the turbo pump through the port on valve 12.

Next, disconnect the still and condenser lines. Make sure the valves on the fridge are shut, of course, and also make sure the gate valve on the sand-bucket manifold is shut as well as the valve on the sand-bucket manifold to the condenser. Remember to put blanks on both the still and condenser ports on the fridge in case the valves leak. Next, valve off the 1K pot manifold valve shown in figure H.15, disconnect the 1K pot line from the fridge, and put a blank on the 1K pot port on the fridge. Shut off the 1K pot pump and close the valve connecting the pump to the helium recovery line.

Since you've got all the mixture back to the dump, you can shut down the roots blower and the sealed pump now, too. Just make sure that the gate valve on top of the roots blower is shut so that the mixture that backstreams through the pump can't make it any farther than the gate valve. Shut down the roots blower first so the sealed pump can keep the outlet pressure of the roots blower low. Once you shut down both pumps, the mechanical gauge on the backside of the pumps will drop due to mixture flowing backwards through the pumps. Make sure that valves 4, 6, and 7 on the wall are shut so that most of the mixture stays in the dump.

Now that the gas lines are all taken care of, disconnect the D-SUB cables on the green switch boxes. Remember to first ground your samples to the fridge by putting the switch boxes in position "B" or "C". Next, unscrew the fridge from the magnet,

remove the ground lines, and hook the fridge up to the hoist. Finally, valve off and disconnet the sliding seal recovery line.

Start raising the fridge slowly. Remember that the hoist can easily lift a couple thousand pounds, so if the fridge doesn't start moving right away double check that you don't have anything still trying to hold it in place. Start by raising the fridge $\sim$ 6 inches and then go in steps of a couple inches after that. Try to wipe off as much frost as possible while you are pulling it out. The G-10 on the sliding seal is 27 inches long, so keep track of how close you are to having the fridge out. It will probably take about an hour and a half to get the fridge to the bottom of the sliding seal. When you have about an inch of G-10 left in the magnet, go find a friend to help steady the fridge while you pull it out. As soon as the sliding seal clears the magnet, a lot of helium is going to start pouring out, so raise the fridge quickly. Plug the hole in the magnet as soon as the bottom of the IVC is out of the magnet.

Once the fridge is out, you can either let it sit overnight for the guts to warm to room temperature, or you can poison the IVC with some helium. If you add some helium exchange gas to accelerate the process, it will probably take $\sim$ 2-3 hours for the inside of the fridge to be warm enough to take the IVC off without condensing a bunch of moisture on the dilution unit.

If you need to transfer into the dewar to keep the magnet cold while the fridge is out, be warned that the cap is not held in place very firmly so it will pop out of the magnet if you pressurize the dewar very much (like at the beginning of a transfer). There are a couple of lead bricks under the sink, and setting one of these (gently) on top of the cap is an easy way to keep the cap from jumping up. If you are not planning on cooling the insert back down, just leave the dewar alone with the valve open to the recovery. There is still quite a bit of helium in the tail region that the level meter can't measure, so you need to let this all boil off before you close up the recovery system. It will probably take over two weeks for the bottom of the dewar to get all the way to room temperature, so don't leave any ports open before that point. If you need to accelerate the process for some reason, you can try blowing some nitrogen gas

down the transfer port, but try not to thermally shock the magnet. Once the fridge is all the way warm, I would recommend shutting the main recovery valve to the lab. Every once in a while the department loses a good chunk of helium for one reason or another, and if this happens you want to be able to point to the closed valve as proof that it was nothing in your lab that leaked.

H.11 Condensed Checklist

H.11.1 Cooling Down Fridge (Magnet Cold)

- Check resistors, LEDs, and wiring for bad solder joints.
- Attach radiation shield.
- Load samples, check 2-terminal resistance for shorts/opens.
- Attach IVC and pump out/leak check/add exchange gas. Use leak detector for initial pump down. Use turbo pump for adding exchange gas.
- Pump out condenser with turbo pump.
- Pump/purge 1K pot 3-5 times, leave pressurized $\sim$ 1-2 psi with helium.
- Install sliding seal, hook up to recovery line.
- Double check samples are grounded to fridge.
- Load fridge into dewar. Make sure helium level $<$ 5 inches. Make sure helium recovery ready for big gas load.
- Force boil-off through sliding seal. Monitor sorb temperature with Labview.
- When fridge all the way down, screw into magnet and connect ground straps.
- Transfer helium

H.11.2 Condensing and Cooling to Base

- Hook up gas lines, pump out, and leak check
- Heat sorb and pump out exchange gas with leak detector for 1 hour
- Wash mixture through LN2 trap. Start circulating mixture before cooling the trap. Do NOT open roots blower bypass valve.
- Cool down He trap after washing configuration is stable
- Check 2-terminal resistances and gate leakage
- Illuminate (if necessary) 15 min, 2mA, 8.3V limit if using LEDs 1-3.
- Cool down 1K pot, close fridge bypass valve, open condenser valve to start condensing
- Monitor 1K pot temperature, condenser pressure (should be $\sim$ -26 inHg), and pressure on back side of pumps (don't let it get $<$ -22.5 inHg).
- Seal off fridge and pump out dump once dump pressure $\sim$ -28 inHg. Pump until the still gauge reads $<$ 100 mTorr.
- Start circulation. Remember to open valve on top of roots pump very slowly. Close valve back down when backside of pumps $\sim$ 1 atm.
- Fridge should cool to $\sim$ 25 mK in $<$ 2 hours and to $\sim$ 11 mK overnight.
- Open valve 6 to open path to dump through check valve.

H.11.3 Cooling Down the Magnet

- Load fridge, hook up and leak check gas lines, add exchange gas to IVC as before.
- Wash mixture through LN2 trap overnight.

- Clean the LN2 trap.
- Pump/purge dewar 3 times with 1K pot pump. Leave pressurized $\sim$ 1-2 psi. Should pump down to $\sim$ 500 mTorr.
- Pump/purge 1K pot and lambda plate, leave pressurized $\sim$ 1-2 psi with helium.
- Pre-cool magnet with LN2. Make sure lab helium recovery valve is shut. Remember to open dewar recovery valve to let nitrogen exhaust to the room.
- Watch magnet resistors to set the pace of the cool-down. Should take $\sim$ 2 hours.
- Blow out the nitrogen.
- Pump/purge dewar with helium 3 times. Remember pumping will go slower than at room temperature.
- Transfer helium. Make sure recovery system is hooked up and ready for big gas load.

H.11.4 Warming Up

- Pull mix back to the dump. Close down 1K pot, add heat to still and mixing chamber (2.9V to still, 3.3V to mixing chamber). Pump on both still and condenser lines. Remember to open valve to dump.
- Once fridge empty, pump on only fridge (valve off traps) for another 10 minutes to be safe.
- Seal off fridge, turn off heaters.
- Pump on traps until condenser gauge $\sim$ 10 mTorr.
- Dump should rise to -7.5 inHg.
- Clean the traps

- Valve everything off and disconnect gas lines.
- Shut down pumps (roots blower first).
- Ground samples and disconnect electrical lines.
- Unscrew fridge from magnet.
- Pull fridge out (be careful when you start that it is not caught on anything).
- G-10 on sliding seal is 27 inches long.
- Let fridge warm up overnight or 2-3 hours if IVC poisoned.
- Let magnet warm up over $\sim$ 2 weeks.

REFERENCES

REFERENCES

[1] P. W. Anderson, Science **177**, 393 (1972).

[2] N. W. Ashcroft and N. D. Mermin, *Solid State Physics*, Thomson Learning, 1976.

[3] J. H. Davies, *The Physics of Low-dimensional Semiconductors: An Introduction*, Cambridge University Press, 1997.

[4] K. v. Klitzing, G. Dorda, and M. Pepper, Phys. Rev. Lett. **45**, 494 (1980).

[5] S. D. Sarma and A. Pinczuk, editors, *Perspectives in Quantum Hall Effects*, Wiley Interscience, 1997.

[6] R. Landauer, Philos. Mag. **21**, 863 (1970).

[7] M. Büttiker, Phys. Rev. Lett. **57**, 1761 (1986).

[8] M. Büttiker, Phys. Rev. B **38**, 9375 (1988).

[9] S. Datta, *Electronic Transport in Mesoscopic Systems*, Cambridge University Press, 1997.

[10] H. L. Störmer, Rev. Mod. Phys. **71**, 875 (1999).

[11] D. C. Tsui, H. L. Stormer, and A. C. Gossard, Phys. Rev. Lett. **48**, 1559 (1982).

[12] R. B. Laughlin, Phys. Rev. Lett. **50**, 1395 (1983).

[13] J. K. Jain, Phys. Rev. Lett. **63**, 199 (1989).

[14] W. Kang, H. Stormer, L. Pfeiffer, K. Baldwin, and K. West, Phys. Rev. Lett. **71**, 3850 (1993).

[15] R. Willett, R. Ruel, K. West, and L. Pfeiffer, Phys. Rev. Lett. **71**, 3846 (1993).

[16] X. Wen, Phys. Rev. Lett. **64**, 2206 (1990).

[17] X. G. Wen, Phys. Rev. B **43**, 11025 (1991).

[18] R. Willett et al., Phys. Rev. Lett. **59**, 1776 (1987).

[19] M. P. Lilly, K. B. Cooper, J. P. Eisenstein, L. N. Pfeiffer, and K. W. West, Phys. Rev. Lett. **82**, 394 (1999).

[20] M. Lilly, K. Cooper, J. Eisenstein, L. Pfeiffer, and K. West, Phys. Rev. Lett. **83**, 824 (1999).

[21] J. P. Eisenstein, K. B. Cooper, L. N. Pfeiffer, and K. W. West, Phys. Rev. Lett. **88**, 076801 (2002).

[22] J. S. Xia et al., Phys. Rev. Lett. **93**, 176809 (2004).

[23] N. Deng, J. D. Watson, L. P. Rokhinson, M. J. Manfra, and G. A. Csáthy, Phys. Rev. B **86**, 201301 (2012).

[24] W. Pan et al., Phys. Rev. B **77**, 075307 (2008).

[25] A. Kumar, G. A. Csáthy, M. J. Manfra, L. N. Pfeiffer, and K. W. West, Phys. Rev. Lett. **105**, 246808 (2010).

[26] A. Kitaev, Ann. Phys. **303**, 2 (2003).

[27] S. Das Sarma, M. Freedman, and C. Nayak, Phys. Rev. Lett. **94**, 166802 (2005).

[28] G. Moore and N. Read, Nucl. Phys. B **360**, 362 (1991).

[29] R. H. Morf, Phys. Rev. Lett. **80**, 1505 (1998).

[30] E. H. Rezayi and F. D. M. Haldane, Phys. Rev. Lett. **84**, 4685 (2000).

[31] M. R. Peterson, T. Jolicoeur, and S. Das Sarma, Phys. Rev. Lett. **101**, 016807 (2008).

[32] M. Storni, R. H. Morf, and S. Das Sarma, Phys. Rev. Lett. **104**, 076803 (2010).

[33] W. Bishara and C. Nayak, Phys. Rev. B **77**, 165302 (2008).

[34] S.-S. Lee, S. Ryu, C. Nayak, and M. P. A. Fisher, Phys. Rev. Lett. **99**, 236807 (2007).

[35] X.-G. Wen, Phys. Rev. Lett. **70**, 355 (1993).

[36] P. Fendley, M. P. A. Fisher, and C. Nayak, Phys. Rev. Lett. **97**, 036801 (2006).

[37] I. P. Radu et al., Science **320**, 899 (2008).

[38] A. Feiguin, E. Rezayi, C. Nayak, and S. Das Sarma, Phys. Rev. Lett. **100**, 166803 (2008).

[39] C. de C. Chamon, D. E. Freed, S. A. Kivelson, S. L. Sondhi, and X. G. Wen, Phys. Rev. B **55**, 2331 (1997).

[40] E. Fradkin, C. Nayak, A. Tsvelik, and F. Wilczek, Nuclear Physics B **516**, 704 (1998).

[41] M. Dolev, M. Heiblum, V. Umansky, A. Stern, and D. Mahalu, Nature **452**, 829 (2008).

[42] M. Dolev et al., Phys. Rev. B **81**, 161303 (2010).

[43] X. Lin, C. Dillard, M. A. Kastner, L. N. Pfeiffer, and K. W. West, Phys. Rev. B **85**, 165321 (2012).

[44] S. Baer et al., Phys. Rev. B **90**, 075403 (2014).

[45] X.-G. Wen, Phys. Rev. B **44**, 5708 (1991).

[46] J. P. Eisenstein et al., Phys. Rev. Lett. **61**, 997 (1988).

[47] W. Pan et al., Phys. Rev. Lett. **83**, 820 (1999).

[48] T. D. Rhone et al., Phys. Rev. Lett. **106**, 196805 (2011).

[49] M. Stern et al., Phys. Rev. Lett. **108**, 066810 (2012).

[50] L. Tiemann, G. Gamez, N. Kumada, and K. Muraki, Science **335**, 828 (2012).

[51] J. Nuebler et al., Phys. Rev. B **81**, 035316 (2010).

[52] G. Yang and D. E. Feldman, Phys. Rev. B **88**, 085317 (2013).

[53] A. Bid et al., Nature **466**, 585 (2010).

[54] H. Inoue et al., Nature Communications **5**, 4067 (2014).

[55] V. Venkatachalam, S. Hart, L. Pfeiffer, K. West, and A. Yacoby, Nature Ph **8**, 676 (2012).

[56] M. Manfra, Ann. Rev. Cond. Matt. Phys. **5**, 347 (2014).

[57] N. Samkharadze et al., Physical Review B **84**, 12305(R) (2011).

[58] R. L. Willett, L. Pfeiffer, and K. W. West, Proceedings of the National Academy of Sciences of the United States of America **106**, 8853 (2009).

[59] R. L. Willett, L. N. Pfeiffer, and K. W. West, Phys. Rev. B **82**, 205301 (2010).

[60] S. An et al., Arxiv **1112.3400v** (2011).

[61] D. T. McClure, W. Chang, C. M. Marcus, L. N. Pfeiffer, and K. W. West, Phys. Rev. Lett. **108**, 256804 (2012).

[62] Varian, inc. vacuum technologies online catalog, 2012.

[63] Brooks cryo-torr online catalog, 2014.

[64] J. F. O'Hanlon, *A User's Guide to Vacuum Technology*, Wiley Interscience, 1980.

[65] M. A. Herman and H. Sitter, *Molecular Beam Epitaxy*, Springer, 1996.

[66] V. Umansky, R. de Picciotto, and M. Heiblum, Appl. Phys. Lett. **71**, 683 (1997).

[67] L. Pfeiffer, K. W. West, H. L. Störmer, and K. W. Baldwin, Appl. Phys. Lett. **55**, 1888 (1989).

[68] L. Pfeiffer and K. W. West, Physica E **20**, 57 (2003).

[69] G. Gardner, J. Watson, S. Fallahi, and M. Manfra, To be published.

[70] Y. Asaoka, Journal of Crystal Growth **251**, 40 (2003), Proceedings of the Twelfth International Conference on Molecular Beam Epitaxy.

[71] G. W. Smith, A. J. Pidduck, C. R. Whitehouse, J. L. Glasper, and J. Spowart, Journal of Crystal Growth **127**, 966 (1993).

[72] P. M. Petroff, R. C. Miller, A. C. Gossard, and W. Wiegmann, Appl. Phys. Lett. **44**, 217 (1984).

[73] W. J. Schaff, L. F. Eastman, B. V. Rees, and B. Liles, J. Vac. Sci. Technol., B **2**, 265 (1984).

[74] T. J. Drummond et al., Applied Physics Letters **42**, 615 (1983).

[75] I. Vurgaftman, J. R. Meyer, and L. R. Ram-Mohan, J. Appl. Phys. **89**, 5815 (2001).

[76] H. C. C. Jr. and M. B. Panish, *Heterostructure Lasers*, Academic, New York, 1978.

[77] E. F. Schubert, *Doping in III-V Semiconductors*, Cambridge University Press, 1993.

[78] P. M. Mooney, J. Appl. Phys. **67**, R1 (1990).

[79] E. F. Schubert and K. Ploog, Phys. Rev. B **30**, 7021 (1984).

[80] D. J. Chadi and K. J. Chang, Phys. Rev. Lett. **61**, 873 (1988).

[81] D. J. Chadi and K. J. Chang, Phys. Rev. B **39**, 10063 (1989).

[82] S. D. Sarma, E. H. Hwang, S. Kodiyalam, L. N. Pfeiffer, and K. W. West, Arxiv **1412.8479v1** (2014).

[83] T. Ando, A. B. Fowler, and F. Stern, Rev. Mod. Phys. **54**, 437 (1982).

[84] K.-J. Friedland, R. Hey, H. Kostial, R. Klann, and K. Ploog, Phys. Rev. Lett. **77**, 4616 (1996).

[85] V. Umansky et al., J. Cryst. Growth **311**, 1658 (2009).

[86] M. Samani et al., Phys. Rev. B **90**, 121405 (2014).

[87] J. B. Miller et al., Nat. Phys. **3**, 561 (2007).

[88] M. Pioro-Ladrière et al., Phys. Rev. B **72**, 115331 (2005).

[89] E. Buks, M. Heiblum, Y. Levinson, and H. Shtrikman, Semiconductor Science and Technology **9**, 2031 (1994).

[90] C. Buizert et al., Phys. Rev. Lett. **101**, 226603 (2008).

[91] D. H. Cobden et al., Phys. Rev. Lett. **69**, 502 (1992).

[92] O. E. Dial et al., Phys. Rev. Lett. **110**, 146804 (2013).

[93] M. J. Manfra et al., Phys. Rev. Lett. **99**, 236402 (2007).

[94] D. Brunner et al., Science **325**, 70 (2009).

[95] B. D. Gerardot et al., Nature **451**, 441 (2008).

[96] D. V. Bulaev and D. Loss, Phys. Rev. Lett. **95**, 076805 (2005).

[97] J. J. Heremans, M. B. Santos, K. Hirakawa, and M. Shayegan, J. Appl. Phys. **76**, 1980 (1994).

[98] R. J. Malik, R. N. Nottenberg, E. F. Schubert, J. F. Walker, and R. W. Ryan, Appl. Phys. Lett. **53**, 2661 (1988).

[99] J. Nagle, R. Malik, and D. Gershoni, J. Cryst. Growth **111**, 264 (1991).

[100] D. Reuter, A. D. Wieck, and A. Fischer, Rev. Sci. Instr. **70**, 3435 (1999).

[101] M. J. Manfra, L. N. Pfeiffer, K. W. West, R. de Picciotto, and K. W. Baldwin, Appl. Phys. Lett. **86**, 162106 (2005).

[102] J. D. Watson, S. Mondal, G. Gardner, G. A. Csáthy, and M. J. Manfra, Phys. Rev. B **85**, 165301 (2012).

[103] B. Grbić, R. Leturcq, K. Ensslin, D. Reuter, and A. D. Wieck, Appl. Phys. Lett. **87**, (2005).

[104] O. Klochan et al., Appl. Phys. Lett. **96**, (2010).

[105] L. A. Tracy, T. W. Hargett, and J. L. Reno, Appl. Phys. Lett. **104**, (2014).

[106] M. N. Khannanov et al., JETP Lett. **85**, 242 (2007).

[107] H. Zhu et al., Sol. St. Comm. **141**, 510 (2007).

[108] T. M. Lu et al., Appl. Phys. Lett. **92**, 012109 (2008).

[109] F. Pobell, *Matter and methods at low temperatures*, Springer, 2007.

[110] D. Kramer, Physics Today **64**, 20 (2011).

[111] I. B. Spielman, *Evidence for the Josephson effect in quantum Hall bilayers*, PhD thesis, California Institute of Technology, 2004.

[112] C. Gerl, S. Schmult, H.-P. Tranitz, C. Mitzkus, and W. Wegscheider, Appl. Phys. Lett. **86**, 252105 (2005).

[113] A. Kumar et al., Phys. Rev. B **83**, 201305 (2011).

[114] B. Grbić et al., Phys. Rev. Lett. **99**, 176803 (2007).

[115] S. P. Koduvayur et al., Phys. Rev. Lett. **106**, 016804 (2011).

[116] M. J. Manfra et al., Phys. Rev. Lett. **98**, 206804 (2007).

[117] Nextnano3 simulator ©1999-2008 walter schottky institute. http://www.nextnano.de/nextnano3/index.htm.

[118] E. H. Hwang and S. Das Sarma, Phys. Rev. B **77**, 235437 (2008).

[119] K. Lee, M. S. Shur, T. J. Drummond, and H. Morkoç, J. Appl. Phys. **54**, 6432 (1983).

[120] C. Jiang, D. C. Tsui, and G. Weimann, Appl. Phys. Lett. **53**, 1533 (1988).

[121] J. A. Seamons, D. R. Tibbetts, J. L. Reno, and M. P. Lilly, Appl. Phys. Lett. **90**, 052103 (2007).

[122] E. H. Hwang and S. Das Sarma, Phys. Rev. B **78**, 075430 (2008).

[123] S. Das Sarma and E. H. Hwang, Phys. Rev. Lett. **83**, 164 (1999).

[124] S. Das Sarma and E. H. Hwang, Phys. Rev. B **61**, R7838 (2000).

[125] J. P. Eisenstein, H. L. Störmer, V. Narayanamurti, A. C. Gossard, and W. Wiegmann, Phys. Rev. Lett. **53**, 2579 (1984).

[126] R. Winkler, *Spin-Orbit Coupling Effects in Two-Dimensional Electron and Hole Systems*, Springer, Berlin, 2003.

[127] B. Habib, M. Shayegan, and R. Winkler, Semicon. Sci. and Tech. **24**, 064002 (2009).

[128] J. D. Watson et al., Phys. Rev. B **83**, 241305 (2011).

[129] MJM thanks Ken West and Loren Pfeiffer from Princeton University for sharing their carbon source design upon which ours is based.

[130] G. Bastard, *Wave Mechanics Applied to Semiconductor Heterostructures*, Halsted Press, New York, 1988.

[131] F. Stern, Appl. Phys. Lett. **43**, 974 (1983).

[132] J. W. Harrison and J. R. Hauser, Phys. Rev. B **13**, 5347 (1976).

[133] D. Chattopadhyay, Phys. Rev. B **31**, 1145 (1985).

[134] W. Li, G. A. Csáthy, D. C. Tsui, L. N. Pfeiffer, and K. W. West, Appl. Phys. Lett. **83**, 2832 (2003).

[135] F. F. Fang and W. E. Howard, Phys. Rev. Lett. **16**, 797 (1966).

[136] T. Ando, J. Phys. Soc. Jap. **51**, 3900 (1982).

[137] J. C. H. Chen et al., Appl. Phys. Lett. **100**, 052101 (2012).

[138] H. Störmer, A. Gossard, and W. Wiegmann, Sol. St. Comm. **41**, 707 (1982).

[139] H. L. Störmer et al., Phys. Rev. Lett. **51**, 126 (1983).

[140] W. Walukiewicz, Phys. Rev. B **31**, 5557 (1985).

[141] E. H. Hwang and S. Das Sarma, Phys. Rev. B **67**, 115316 (2003).

[142] W. Bishara, P. Bonderson, C. Nayak, K. Shtengel, and J. K. Slingerland, Phys. Rev. B **80**, 155303 (2009).

[143] M. Levin, B. I. Halperin, and B. Rosenow, Phys. Rev. Lett. **99**, 236806 (2007).

[144] P. Fendley, A. W. W. Ludwig, and H. Saleur, Phys. Rev. B **52**, 8934 (1995).

[145] C. d. C. Chamon and X. G. Wen, Phys. Rev. B **49**, 8227 (1994).

[146] U. Meirav, M. Heiblum, and F. Stern, **52**, 1268 (1988).

[147] D. Ritchie et al., J. Cryst. Growth **111**, 300 (1991).

[148] A. R. Hamilton et al., Appl. Phys. Lett. **60**, 2782 (1992).

[149] Y. Hirayama, K. Muraki, and T. Saku, Appl. Phys. Lett. **72**, 1745 (1998).

[150] A. Kawaharazuka, T. Saku, Y. Hirayama, and Y. Horikoshi, J. Appl. Phys. **87**, 952 (2000).

[151] K. Muraki, N. Kumada, T. Saku, and Y. Hirayama, Jap. J. Appl. Phys. **39**, 2444 (2000).

[152] A. Kawaharazuka, T. Saku, C. A. Kikuchi, Y. Horikoshi, and Y. Hirayama, Phys. Rev. B **63**, 245309 (2001).

[153] A. Valeille, K. Muraki, and Y. Hirayama, Appl. Phys. Lett. **92**, 152106 (2008).

[154] N. Deng et al., Phys. Rev. Lett. **112**, 116804 (2014).

[155] M. Lenzlinger and E. H. Snow, J. Appl. Phys. **40**, 278 (1969).

[156] Z. A. Weinberg, J. Appl. Phys. **53**, 5052 (1982).

[157] J. Smoliner et al., Appl. Phys. Lett. **50**, 1727 (1987).

[158] T. W. Hickmott, P. M. Solomon, R. Fischer, and H. Morkoç, Appl. Phys. Lett. **44**, 90 (1984).

[159] K. L. Jensen, J. Vac. Sci. Technol., B **21**, 1528 (2003).

[160] S. Das Sarma and E. H. Hwang, Phys. Rev. B **90**, 035425 (2014).

[161] N. Deng et al., Phys. Rev. Lett. **108**, 086803 (2012).

[162] C. Zhang et al., Phys. Rev. B **85**, 241302 (2012).

[163] Y. Liu et al., Phys. Rev. Lett. **107**, 176805 (2011).

[164] J. B. Miller, *Electron transport in GaAs heterostructures at various magnetic field strengths*, PhD thesis, Harvard University, 2007.

[165] D. T. McClure, *Interferometer-based studies of quantum Hall phenomena*, PhD thesis, Harvard University, 2012.

[166] A. Kou, *Microscopic properties of the fractional quantum Hall effect*, PhD thesis, Harvard University, 2013.

[167] F. H. L. Koppens, *Coherence and control of a single electron spin in a quantum dot*, PhD thesis, Technische Universiteit Delft, 2007.

[168] J. R. Medford, *Spin qubits in double and triple quantum dots*, PhD thesis, Harvard University, 2013.

[169] J. Nübler, *Density dependence of the $\nu = 5/2$ fractional quantum Hall effect*, PhD thesis, Universität Tübingen, 2011.

[170] K. B. Cooper, *New phases of two-dimensional electrons in excited Landau levels*, PhD thesis, California Institute of Technology, 2003.

[171] L. Tiemann, *Phase coherence in the regime of bilayer exciton condensation*, PhD thesis, Universität Bielefeld, 2008.

[172] L. A. Tracy, *Studies of two dimensional electron systems via surface acoustic waves and nuclear magnetic resonance techniques*, PhD thesis, California Institute of Technology, 2007.

[173] M. J. Kellogg, *Evidence for excitonic superfluidity in a bilayer two-dimensional electron system*, PhD thesis, California Institute of Technology, 2004.

[174] O. Göktaş, *Small alloyed ohmic contacts to 2DES and submicron scale Corbino devices in strong magnetic fields: observation of a zero bias anomaly and single-electron charging*, PhD thesis, Universität Stuttgart, 2009.

[175] S. E. Foletti, *Manipulation and coherence of a two-electron logical spin qubit using GaAs double quantum dots*, PhD thesis, Harvard University, 2010.

[176] V. Venkatachalam, *Single electron probes of fractional quantum Hall states*, PhD thesis, Harvard University, 2012.

[177] R. L. Willett, L. N. Pfeiffer, and K. W. West, Appl. Phys. Lett. **89**, (2006).

[178] K. Rasmussen, Private communication.

[179] S. Harvey, Private communication.

[180] S. Mondal, *Two Dimensional Electron Systems for Solid State Quantum Computation*, PhD thesis, Purdue University, 2014.

[181] S. Mondal et al., Sol. St. Comm. **197**, 20 (2014).

[182] M. J. Biercuk, D. J. Monsma, C. M. Marcus, J. S. Becker, and R. G. Gordon, Appl. Phys. Lett. **83**, 2405 (2003).

[183] P. D. Ye, G. D. Wilk, E. E. Tois, and J. J. Wang, Appl. Phys. Lett. **87**, (2005).

[184] M. J. Biercuk, *Local gate control in carbon nanotube quantum devices*, PhD thesis, Harvard University, 2005.

VITA

VITA

John D. Watson

Department of Physics and Astronomy, Purdue University, 525 Northwestern Ave.

West Lafayette, IN 47907

Education

- Ph.D., Experimental Condensed Matter, Purdue University May 2015
  - Dissertation Adviser: Prof. Michael Manfra, Ph.D.
- M.S., Physics, Purdue University May 2010
  - Cumulative GPA 4.0/4.0
- B.S., Engineering Physics, Northwest Nazarene University May 2008
  - Cumulative GPA 3.97/4.0

Awards

- H.Y. Fan Award for Excellence in Condensed Matter Physics, Purdue, 2012
- Outstanding Student Poster, North American Molecular Beam Epitaxy Conference, 2011
- Purdue/Sandia National Lab Excellence in Science and Engineering Fellowship, 2009-2012
- Steven Wilson Outstanding Graduate, School of Health and Science, Northwest Nazarene University, 2008
- President's Scholar, Northwest Nazarene University, 2004-2008
- Top Freshman Physics Student, Northwest Nazarene University, 2005
- National Merit Scholar, 2004

Employment

- **Purdue Univeristy, West Lafayette, IN** Research Assitance, Department of Physics, 2012-2015 Graduate Fellow, Department of Physics, 2009-2012 Teaching Assistant, Department of Physics, 2008-2009
- **Northwest Nazarene University, Nampa, ID** Research Assistant, Department of Physics and Engineering, Summer 2008
- **Micron Technology, Inc. Boise, ID** Engineering Intern, Wafer Level Optics Group, Summer 2007
- **Boise State University, Boise, ID** Research Assistant, Department of Electrical Engineering, Summer 2006

Research Interests

- High-mobility growth of GaAs/AlGaAs heterostructures by MBE
- Fractional quantum Hall effect
- Micro/nano-scale semiconductor processing
- Cryogenics
- Quantum computing

Publications

1. J. D. Watson, S. Mondal, G. A. Csáthy, M. J. Manfra, E. H. Hwang, S. Das Sarma, L. N. Pfeiffer, K. W. West. "Scattering mechanisms in a high-mobility low-density carbon-doped (100) GaAs two-dimensional hole system." Phys. Rev. B **83**, 241305(R) (2011).
2. N. Samkharadze, J. D. Watson, G. Gardner, M. J. Manfra, L. N. Pfeiffer, K. W. West, and G. A. Csáthy. "Quantitative analysis of the disorder broadening and the intrinsic gap for the $\nu = 5/2$ fractional quantum Hall state." Phys. Rev. B **84**, 121305(R) (2011).

3. A. T. Hatke, M. A. Zudov, J. D. Watson, and M. J. Manfra. "Magnetoplasmon resonance in a two-dimensional electron system driven into a zero-resistance state." Phys. Rev. B **85** 121306(R) (2012).

4. J. D. Watson, S. Mondal, G. Gardner, G. A. Csáthy, and M. J. Manfra. "Exploration of the limits to mobility in two-dimensional hole systems in GaAs/AlGaAs quantum wells." Phys. Rev. B **85**, 165301 (2012).

5. N. Deng, J. D. Watson, L. P. Rokhinson, M. J. Manfra, and G. A. Csáthy. "Contrasting energy scales of the reentrant integer quantum Hall states." Phys. Rev. B **86**, 201301(R) (2012).

6. M. I. Hossain, Z. Ikonic, J. Watson, J. Shao, P. Harrison, M. J. Manfra, and O. Malis. "Strong heavy-to-light hole intersubband absorption in the valence band of carbon-doped GaAs/AlAs superlattices." J. Appl. Phys. **113**, 053103 (2013).

7. A. T. Hatke, M. A. Zudov, J. D. Watson, M. J. Manfra, L. N. Pfeiffer, and K. W. West. "Evidence for effective mass reduction in GaAs/AlGaAs quantum wells." Phys. Rev. B **87**, 161307(R) (2013).

8. G. C. Gardner, J. D. Watson, S. Mondal, N. Deng, G. A. Csáthy, and M. J. Manfra. "Growth and electrical characterization of $Al_{0.24}Ga_{0.76}As/Al_xGa_{1-x}As/Al_{0.24}Ga_{0.76}As$ modulation-doped quantum wells with extremely low x." Appl. Phys. Lett. **102**, 252103 (2013).

9. M. R. Delbecq, T. Nakajima, T. Otsuka, S. Amaha, J. D. Watson, M. J. Manfra, and S. Tarucha. "Full control of quadruple quantum dot circuit charge states in the single electron regime." Appl. Phys. Lett. **104**, 183111 (2014).

10. N. J. Goble, J. D. Watson, M. J. Manfra, and X. P. A. Gao. "Impact of short-range scattering on the metallic transport of strongly correlated two-dimensional holes in GaAs quantum wells." Phys. Rev. B **90**, 035310 (2014).

11. Q. Zhang, T. Arikawa, E. Kato, J. L. Reno, W. Pan, J. D. Watson, M. J. Manfra, M. A. Zudov, M. Tokman, M. Erukhimova, A. Belyanin, and J. Kono. "Superradiant decay of cyclotron resonance of two-dimensional electron gases." Phys. Rev. Lett. **113**, 047601 (2014).

12. S. Mondal, G. C. Gardner, J. D. Watson, S. Fallahi, A. Yacoby, and M. J. Manfra. "Field-effect-induced two-dimensional electron gas utilizing modulation-doped ohmic contacts." Sol. St. Comm. **197**, 20 (2014).

13. M. Samani, A. V. Rossokhaty, E. Sajadi, S. Lüscher, J. A. Folk, J. D. Watson, G. C. Gardner, and M. J. Manfra. "Low-temperature illumination and annealing of ultrahigh quality quantum wells." Phys. Rev. B **90**, 121405(R) (2014).