

THE UNIVERSITY OF CHICAGO

OPTICAL CONTROL AND SENSING TECHNIQUES FOR THE
NITROGEN-VACANCY CENTER IN DIAMOND

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE PRITZKER SCHOOL OF MOLECULAR ENGINEERING
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
PAUL JERGER

CHICAGO, ILLINOIS
DECEMBER 2021

Copyright © 2021 by Paul Jerger

All Rights Reserved

For Emily

Table of Contents

LIST OF FIGURES	vi
LIST OF TABLES	viii
LIST OF PUBLICATIONS	ix
ACKNOWLEDGMENTS	x
ABSTRACT	xviii
1 INTRODUCTION	1
1.1 The Quantum Frontier	1
1.1.1 What is a quantum computer?	2
1.1.2 Why is it so hard to build a quantum computer?	7
1.1.3 Why are quantum computers exciting?	11
1.1.4 Are there other important technologies in QISE?	19
2 THE NITROGEN-VACANCY CENTER	27
2.1 Physics of the Nitrogen-Vacancy Center	27
2.1.1 Structure and Formation	27
2.1.2 Spin and Optical Properties	30
2.1.3 State Preparation, Manipulation, and Measurement	34
2.1.4 Interactions and Coherence Properties	36
2.1.5 Low-Temperature Properties	40
2.2 NV Center Applications	42
2.2.1 Quantum Simulation	43
2.2.2 Quantum Communication	45
2.2.3 Quantum Sensing	48
3 ALL-OPTICAL SINGLE-QUBIT CONTROL	55
3.1 Lambda Systems	56
3.1.1 STIRAP	57
3.2 Superadiabatic Transitionless Driving	59
3.2.1 Adiabatic and Superadiabatic Control	60
3.2.2 Experimental Setup	63
3.2.3 Superadiabatic Performance	66

3.2.4	Influence of Dissipation	72
3.2.5	Conclusion	76
3.3	All-Optical Holonomic Gates	76
3.3.1	Holonomic Gates	77
3.3.2	Experimental Implementation	80
3.3.3	Resonant Gates	83
3.3.4	Detuned Gates	89
3.3.5	Conclusion	93
4	SPATIOTEMPORAL PHOTOCURRENT MAPPING IN MOLYBDENUM DISULFIDE	95
4.1	Electric Current Sensing with NV Centers	96
4.2	Experimental Setup	101
4.3	ac Photocurrent Measurements	107
4.4	Photocurrent Rise Times	112
4.5	Photocurrent Vortex Mapping	116
4.6	Measurement Sensitivity	123
4.7	Conclusion	125
A	DETAILS OF SUPERADIABATIC PULSE SHAPES	126
B	DETAILS OF HOLONOMIC QUANTUM GATES	132
B.1	Holonomic Gate Parameters	132
B.2	Quantum Process Tomography	133
C	THERMAL MODELING OF MOLYBDENUM DISULFIDE LASER HEATING	135
	REFERENCES	138

List of Figures

1.1	The hierarchy of quantum-relevant complexity classes.	14
2.1	Diamond lattice with NV center.	28
2.2	NV center ground state level structure.	31
2.3	NV center photoemission spectrum.	32
3.1	A low-temperature Λ system of the NV center.	57
3.2	Concept of superadiabatic evolution.	60
3.3	Superadiabatic STIRAP pulse shapes.	62
3.4	Experimental hardware setup.	64
3.5	Optical Rabi oscillations on a single NV center.	66
3.6	Improved STIRAP efficiency due to SATD and MOD-SATD.	67
3.7	Photoluminescence during the STIRAP protocol.	68
3.8	Speedup of SATD and MOD-SATD protocols.	69
3.9	Phase coherence of STIRAP and fractional STIRAP.	71
3.10	Robustness of SATD and MOD-SATD to deviations in pulse shape.	73
3.11	Experimental and theoretical data on the influence of dissipation on STIRAP.	74
3.12	Superadiabatic transfer efficiency with deliberate one-photon detuning.	75
3.13	Photoluminescence excitation sweep of a single NV center.	80
3.14	Evolution of the NV center Λ system during holonomic gates.	83
3.15	Measurement of square optical pulse rise and fall times.	84
3.16	Rabi oscillations between the NV center ground and excited states.	85
3.17	Resonant holonomic gates as θ is varied.	86
3.18	Resonant holonomic gates as ϕ is varied.	87
3.19	Quantum process tomography of resonant holonomic gates.	88
3.20	Detuned holonomic phase gates.	90
3.21	Fidelity of the detuned holonomic phase gates.	91
3.22	Detuned holonomic X gates.	92
3.23	Quantum process tomography of detuned X and Y gates.	93
4.1	NV sensing concept.	98
4.2	Time- and frequency-domain filter functions for sensing sequences.	99
4.3	Atomic force microscope images of MoS ₂ on a diamond substrate.	103
4.4	Experimental schematic for measuring photocurrents in MoS ₂ with NV centers.	104
4.5	Optical characterization measurements of MoS ₂ and NV centers.	105
4.6	Coherence properties of near-surface NV center ensembles.	106

4.7	ac photocurrent measurement protocol.	108
4.8	Detection of ac photocurrents and dependence on optical power.	109
4.9	Comparison of photocurrent measurements between $ -1 \rangle$ and $ +1 \rangle$ branches. . .	111
4.10	Phase shift of photocurrent signal as a function of location.	113
4.11	Example signal shapes with different relationships between maximum and average amplitude.	114
4.12	Amplitude correction β calibrated against phase shift θ_{opt}	115
4.13	Spatial mapping of a photocurrent vortex in MoS ₂	117
4.14	Temperature and magnetic field dependence of the observed MoS ₂ photocurrent.	118
4.15	Temperature simulations in monolayer MoS ₂	120
4.16	Spatially-resolved photocurrent rise time measurements.	122
4.17	Smallest resolved photocurrent.	123

List of Tables

B.1	Experimental parameters for implementing holonomic gates.	132
C.1	Simulation parameters for 2D heat equation model.	136
C.2	Simulation parameters for laser heating source.	136

List of Publications

This thesis represents the motivations, results, and conclusions from the following works:

- [1] Accelerated quantum control using superadiabatic dynamics in a solid-state lambda system. B. B. Zhou, A. Baksic, H. Ribeiro, C. G. Yale, F. J. Heremans, **P. C. Jerger**, A. Auer, G. Burkard, A. A. Clerk, and D. D. Awschalom. *Nature Phys.* **13**, 330-334 (2017). <https://doi.org/10.1038/nphys3967>

- [2] Holonomic Quantum Control by Coherent Optical Excitation in Diamond. B. B. Zhou, **P. C. Jerger**, V. O. Shkolnikov, F. J. Heremans, G. Burkard, and D. D. Awschalom. *Phys. Rev. Lett.* **119**, 140503 (2017).
<https://doi.org/10.1103/PhysRevLett.119.140503>

- [3] Spatiotemporal Mapping of a Photocurrent Vortex in Monolayer MoS₂ Using Diamond Quantum Sensors. B. B. Zhou,* **P. C. Jerger**,* K.-H. Lee, M. Fukami, F. Mujid, J. Park, and D. D. Awschalom. *Phys. Rev. X* **10**, 011003 (2020).
<https://doi.org/10.1103/PhysRevX.10.011003>

Acknowledgments

It is impossible for me to adequately express my gratitude to everyone who has been a direct or indirect influence on my graduate studies, or even to name all of them here, lest these acknowledgments begin to resemble overly long film credits without any hope of a bonus ending afterward. My experience has been everything I could have hoped graduate school would be, although it took surprising turns many times, and I feel so fortunate to have been able to study here at the University of Chicago, in the heart of world-class research facilities and surrounded by friends, family, and supportive mentors. The people around me over the last six years have been responsible for making my time in grad school enjoyable, fruitful, and worthwhile, and I feel so privileged to have had such a positive experience. It has been a long road to graduation, and many of my mentors have been supporting me from the very first day.

My advisor, David Awschalom, has done so much to make his group and the PME a great place for grad students, and I'll forever be grateful for the opportunity to join the group. I still marvel at how well the timing worked out, as I certainly didn't imagine when I joined in 2015 that the PME would grow so quickly. David, I had good feelings about the group from my very first visit, as it was immediately apparent that both scientific creativity and supportive, lighthearted personalities abounded. I am humbled to have been invited into the fold, and have found that my initial impression has only been reinforced over time. I'm thoroughly impressed by your knack for finding students and postdocs that mesh really well with the group. Thank you for clearing the road to make it easy on your students, for

supporting all of our goofy scientific ideas, and for continually providing us with opportunities to collaborate scientifically and grow professionally.

Nobody has had a bigger impact on my development as a scientist than Brian Zhou, our former postdoc. He took me under his wing from day one and gave me the guided tour of how first-class scientific research is done. Brian taught me everything from how to solder to how to cope with catastrophic equipment failure (even when caused by your clumsy junior grad student...), by example as much as by instruction. Brian, I've always admired your well-rounded research skills – you're pragmatic and deft with hardware engineering, thorough in understanding the physics of your experiments, creative in explaining complex topics, and an excellent writer. It's a great combination for a professor to have, and I hope your hard-won research successes so far snowball into more great discoveries for you and your crack “NV Seal” team. I'm sorry you got the Bob Buckley treatment during the Jones Lab move, but I hope you'll be happy to know the SureShots are still prominently featured in LL215 in your honor. Thanks for keeping things fun with fantasy sports, and for coaching me on my first APS presentation. I'm hoping to grab Bob Cha one more time before my time here ends, and it always reminds me of you.

Joe Heremans is every grad student's best friend, as a repository of group lore, experimental knowledge, scientific contacts, and unfailing friendly advice and support. I've called on him more times than I can count – although maybe now we could scan the Slack records to find out – and he's always been there to answer questions and provide direction when I'm flailing about. Joe, it's been great to find a kindred spirit in note-taking and organization, and always fun to hear about Belgian racing snails and all the other quirky things in the Heremans family. I hope Horatio II can one day rise to the level of his father before him and bump a Nobel Laureate in the head. Thank you for rolling up your sleeves whenever I come looking for help even though you've got a billion things going on, for providing so much help with CNM fabrication, and for running with my crazy CVD ideas.

The pandemic of the last two years has had numerous detrimental effects on research, but

one of the biggest ones for me has been losing out on a lot of time holed up in LL215 with Chris Anderson. Chris is everything you could want in a mentor and colleague: encouraging, knowledgeable, willing to take any amount of time to teach a new concept, and doing all the thankless jobs behind the scenes. It was rare for me to find a cool paper that Chris hadn't already seen, and there was nothing I liked better in lab than bouncing some goofy new research ideas around together, even if (especially if?) there was no chance I'd ever get to try them. Chris, it's no accident that you got the first grad student mentorship award; I heard you were the overwhelming, obvious choice among that first round of nominations. I'll be looking out for you to become a professor – keep dreaming big science dreams! It's been a blast swapping background jams and memes. I'll put one more tally mark up on the whiteboard on my last day.

It has been a delight to share LL215 and NV center experiments with Masaya Fukami. I don't expect to ever know as much physics as he does right now, and it's been a sort of cheat code having an experimentalist with his understanding of theory around to double-check my reasoning. Once, I asked him if a theory paper looked like it was missing a factor of two in an equation, and in response he virtually rederived the entire paper for me. But Masaya's deep physics knowledge isn't nearly as important as his humility. He's quick to assume that he's wrong (even when he rarely is) and gentle in offering corrections when others are mistaken – believe me, I have tested that quality in him. Masaya, I've thoroughly enjoyed our conversations over the years, and I'm forever grateful for your guidance during the conference in Japan. It was so memorable!

Though we unfortunately never shared lab space in the same way, I am grateful to have joined the group with Alex Bourassa and Erzsebet Vincent in addition to Masaya. Especially at a time when there were few other students in the department, it was comforting to navigate courses, conferences, and lab moves with three fellow students. I've looked up to all of you over the years, and it will be bittersweet to all go our separate ways. Erzsi, I admire your willingness to dive into the unknown on the optical gating measurements, as I've taken the

opposite tack with very familiar NV center physics. Alex, your enthusiasm for science was catchy and really contributed to a passionately creative group atmosphere.

Speaking of group culture, it's been a joy to interact with so many iconic group members that overlapped with me. Chris Yale helped bring me up to speed on the diamond experiments and made me feel comfortable as a naive first-year. Paolo Andrich brought a spark to the group that just made everybody feel good – not to mention, he was always generous with his time. David Christle was tech-savvy and meme-savvier, and might be surprised how often I fondly recall some of our good times (like the internet speed test saga). Gary Wolfowicz was always a trusted source of sound opinions and could be counted on to cut through the hype to get to the important scientific core of a matter. I was also fortunate enough to see Sam Bayliss come and go. While it wasn't until the end of your tenure that we worked on anything together, I was chuffed to bits to have you around. Like Chris, I think you'll make an amazing professor and your students will be lucky to have you as a mentor. Andrew Yeats, Charlie de las Casas, Will Koehl, and Joost van Bree all added their own good humor, keen scientific insights, and sage advice to the mix, as well, and played a big role in setting the precedents for us younger students to follow.

I owe a big debt of gratitude to the trailblazing grad students who joined the first class of the then-IME: Peter Mintun, Kevin Miao, and Berk Kovos. Though technically PSD students, I count Alex Crook and Sam Whiteley among them, and they all helped establish the path for us in the next cohort. A year older and wiser, I often asked them for advice and tried to follow their examples. They set the tenor for the group after the move to the ERC, and made every day in lab a fun one. Peter, your tearable puns will be remembered, unfortunately. Kevin, I was glad you joined us first-years in Japan to see the Sukaitsuri and get selfies with deer. Berk, your strong opinions always made for entertaining lunch conversation to make the day interesting. One day I will eventually try a branzino and see for myself why they're all the rage. And as I promised, here's a shoutout to Costco for supplying us grad students with quality goods and foodstuffs at low prices! Alex, I was

constantly grateful for your intellectual humility, which set a great example. The pervasive goofiness of the lab, which you had a lot to do with, was a big reason why I wanted to join the group in the first place. Sam, it was great to go to Konstanz with you – that conference was fun for a lot of reasons, including our train-ride and late-night plaza science discussions, spitballing ideas and enjoying the sights. Although inevitable, it has been sad to see you all start graduating and moving on, as you’ve been the core of the group ever since I joined. We had a great run of a few years where we made up at least half of the PME softball team, and while we didn’t win much, it was a such a fun part of every summer.

Of the younger students, I’ve worked most closely with Jonathan Karsch and Ben Soloway, and I’ve really learned a lot through our interactions. They are both dedicated, careful scientists who ask the right questions and continue to come up with creative experimental ideas. With the way the experiments have been running lately, I’m confident that the future of diamond is bright! Jonathan, you’ve totally spoiled us with your baked goods; we don’t deserve such luxury. Keep up the great work with your outreach activities, as you’ve done an excellent job in all the science communication activities I’ve seen you in. Ben, thanks for your hard work learning a very complicated setup and pushing the experiments forward. I’m impressed with the progress you’ve been able to make on multiple fronts. You’re miles ahead of where I was in my research four years ago, and your efforts will definitely pay off in time.

I’d also like to acknowledge the entire Awschalom group, as a great group environment is only an emergent property of the efforts of its constituents. Our group has had a bit of a reputation for seeking our own company at department socials, but I take that to be a sign of just how well we all get along. Thanks to our postdocs Leah, Yeghishe, and Mouktik for setting good examples, being generous with your mentoring, and corralling us into doing what’s best for us. It’s been amazing to see what the younger students have brought to the group as the years have passed. Pratiti, Elena, Grant, Yizhi, Joseph, Cyrus, and Jacob, you all bring unique skills and ideas to the group and are so willingly collaborative. It

shocks me to see how prepared and insightful each new student is, and it bodes extremely well for the coming years. To the first-years, Swathi, Marquis, and Will, while we may not overlap much, you're in a great research group and will be in excellent hands. Don't worry, you'll figure it out as you go – that's all the rest of us are doing (nobody has a plan!). I will be waiting to hear of each new development in the group, and based on the last few years, they'll come fast and furious.

With the PME starting out so small, it was particularly gratifying to find myself in the company of an amiable and motivated cohort of grad students. In that first year, we all took thermodynamics together, explored Chicago, and started some of the friendships which spanned our respective graduate studies. It was great to have counterparts in the one other IME experimental quantum group at the time (the Cleland group), in Joel, Hung-Shen, and Youpeng. But in true IME interdisciplinary style, it was just as nice to get to know people in very different fields, including Peyman, Elizabeth, Cody, Chitavi, Taylor, Josh, Roulan, Ryan, Xiao, Ruben, and Marcello. Thank you all for your support, and it has been amazing to see people graduating as time goes on. At the department's current size, there are so many more grad students, both within and outside of the quantum area, that I have enjoyed meeting and learning with. I can't possibly begin to enumerate them all!

My thanks as well to our Argonne collaborators. I have had the pleasure of working with Nazar, Sean, and Michael, and these ongoing collaborations bring fresh waves of new ideas, opportunities, and banter to the group. You have our thanks for navigating the commute and the bureaucracy to help push new scientific ideas forward, and I expect these channels to strengthen (and maybe even entangle) in the future.

Of course, it is not only students and postdocs who contribute to the graduate student journey; the University faculty are all accomplished, and draw excellent collaborators from around the world. I thank Aash Clerk and Peter Maurer for years of collaborative brainstorming, and for their roles on my dissertation committee. In addition, I gained valuable insights and recommendations from discussions with Guido Burkard and David Cahill during

our interactions. Finally, I thank Giulia Galli and Andrew Cleland for joining my candidacy committee. There are too many faculty to name, both at the University of Chicago and elsewhere, that have aided my progress with inspiring talks, insightful comments, and by sharing teaching resources.

The PME has only just celebrated its 10-year anniversary, which means I have been a part of it for more than half its existence. Even at a remove, I have seen that it takes a tremendous amount of coordination and institutional resources to found and grow a department like this, and to create and sustain all of its programs. I have benefited greatly from the tireless efforts of many administrators and faculty as they established the department, built up research infrastructure, and fine-tuned the policies and procedures which allowed my doctoral research to occur. Much of it happened behind the scenes, unbeknownst to me, but I would like to mention at least the administrators whose impact is most obvious to me. Rovana Popoff and Novia Pagone led the Dean of Students office during my initial years, when the IME could offer very few courses and the degree program details were in flux year to year. They had to juggle many priorities, but managed to recruit students and keep us moving toward graduation. David Taylor and Lisa Abston-Leftridge have done well in taking over as the program has crystallized. Peter Duda and Sally Wolcott oversaw the construction and operation of the cleanroom, and Anna Mukhortova trained me (and countless others) in using a number of the instruments. Paul Julson and Jennifer Robles have managed the growing pains and potential emergencies of the ERC, along with the ever-changing lab requirements of the rapidly-grown PME faculty ranks. Craig Hamill has provided expert IT support, especially over the past two years when pandemic lockdowns suddenly changed our entire relationship to technology. Maurice, Marco, Josh, and Oscar have kept the loading dock running smoothly, despite the unusual shipping demands of an equipment-heavy research group. Laura Rico-Beck has done a fantastic job of organizing unique science communication programs to aid the graduate students in developing communication skills, which are such an integral part of our studies. It has been a delight to be involved in just a few of them,

and entirely to my benefit, as well. Last but certainly not least, Mary Pat McCullough has taken care of countless vexing issues, ranging from travel plans to shipping mixups and much more. It's been reassuring to know that there's a reliable safety net to turn to when new problems arise.

Despite the contents of this dissertation, grad school is about much more than research. I would like to recognize a few non-research groups that enriched me personally and professionally alongside my studies. First, the Science Policy Group (SPG) and its members helped me discover a professional interest I didn't even know I had when I started grad school. I found it so fulfilling to be a part of our grant projects, conference trips, and shared professional development. In particular, my thanks to Tim Steeves for taking the initiative to found SPG, grapple with the logistics of student organizations, and articulating a vision for the organization that really connected with me. SPG activities have been some of my favorite projects over the last few years. Second, Calvert House and the Catholic students of the University have been incredibly welcoming. Finding lifelong friends in this community has been a welcome source of stability, fellowship, and support especially throughout the pandemic.

I would not be at the University of Chicago pursuing a doctorate if it weren't for my parents. I am truly fortunate that they have supported my education at every turn, encouraging me to pursue my intellectual interests and opening doors for me to do so. Thank you, Mom and Dad, for everything.

My wife, Emily, has supported me every step of this long road to graduation. This has meant dealing with unpredictable hours, uncertain timelines, and too many Chicago winters. Emily, I can't really imagine what grad school would have been like without you here, and I'm glad I don't have to. Your love and support have guided me all the way here.

Abstract

Nitrogen-vacancy (NV) centers in diamond are excellent model quantum systems, due to a combination of long coherence times, high-fidelity control, and inherent stability. Their atomic scale, solid-state environment, and intrinsic spin-photon interface make them well-suited for a range of quantum communication and quantum sensing applications. The success of these applications depends on developing an array of techniques to initialize, manipulate, and measure the NV center's quantum state through optical control fields. To that end, this dissertation demonstrates a set of novel techniques for NV center qubit control and sensing which leverage optical interactions. The first class of protocols executes fast, high-fidelity ground state spin operations using an intermediate excited state. By shaping optical control pulses on nanosecond timescales, accelerated adiabatic state transfers and nonadiabatic holonomic single-qubit gates are performed. These techniques are designed to combine error resistance with short duration in pursuit of fault-tolerant qubit operations. In addition, optical control is integrated with established ac magnetometry protocols to enable detection of photocurrents in monolayer molybdenum disulfide. Synchronizing pulsed photoexcitation with spin-echo-based sensing sequences for a near-surface ensemble of NV centers creates a sensitive, local detector of photocurrent density. This capability is used to spatially and temporally map a micron-scale photocurrent vortex. These demonstrations expand the range of optical control methods for single NV centers and provide new methods for probing current distributions in 2D materials and thin films.

Chapter 1

Introduction

1.1 The Quantum Frontier

Quantum information science and engineering (QISE) has come a long way from its theoretical origins to its current state as a field with a growing set of applications, technologies, and initialisms. Its inspiration is often attributed to Richard Feynman, who in the early 1980s envisioned the ability to simulate an interesting quantum system with not just classical digital computers, but other, well-controlled quantum systems [4]. In the wake of the initial conception, theoretical work on what might be possible with such a “quantum computer” proceeded apace, while experimental scientists endeavored to achieve the requisite level of control of light and matter to engineer the corresponding hardware platforms. Today, it is possible to look backward and marvel at the recent advances in QISE which would hardly have been imaginable forty years ago. An undergraduate student in physics can now assemble fully controllable quantum systems in the laboratory which would have been the envy of every physicist of Feynman’s time. Major corporations, venture capitalists, and government R&D arms are buying into the prospects of transformative technological advances and racing to build the biggest and most capable quantum devices. At the same time, it is easy for a skeptic to poke holes in the ever-expanding hype balloon. Few near-term practical applica-

tions of truly quantum technologies seem possible, let alone worthwhile, despite the decades of development and investment. Large obstacles still exist to scaling up the proof-of-principle demonstrations, and it will be necessary to scale many orders of magnitude before quantum computers are expected to start churning out useful calculations. QISE therefore seems to be at an inflection point, where the coming years might hold breakthroughs, snowballing investment, and the transition to an established industry, or instead a desert of useful applications, investor fatigue, and QISE’s relegation to a fringe pursuit. As the field slowly charts its course, the potential utility of a functional quantum computer and the continual march of progress have made it an exciting place for academic research.

To contextualize the recent advances in QISE, it will be useful to review the basic questions of the field. In short:

- What is a quantum computer, and how is it different from a classical computer?
- Why is it so hard to build a quantum computer?
- Why are quantum computers exciting?
- Are there other important technologies in QISE?

1.1.1 What is a quantum computer?

As implied by the similarity between the phrases “classical computing” and “quantum computing,” quantum computers are often easier to understand in relation to the design and capabilities of the computers we use every day. Several key aspects of classical computers and information theory have direct analogues in quantum computer design, and in an abstract sense the purpose of a quantum computer is much the same: to take in digital or digitized information, perform some calculations on it, and output digital answers that humans can make use of. The impetus to build a quantum computer is founded on the strong belief that quantum computers are “better” than classical computers at this information-processing

goal, and a comparison of the two will make it clearer where classical and quantum devices diverge, and where they still overlap.

The fundamental unit of modern information processing is the bit. Conceptually, a bit is anything that can take on two different states to represent values in a calculation, usually labeled 0 and 1 for simplicity. Historically, bits have taken on many different forms. Early computers used holes punched into paper cards as bits, as each location on the card could have a hole or not, indicating a 0 or 1. The intervening years have seen a continuous effort to design bit systems which are smaller, faster, and easier to rewrite than the punch card, with remarkable success. Modern bits come in many types, including the common magnetic memory and flash memory types. In magnetic memory, a grid of bit locations is defined on a magnetic disc, and a small magnetic read/write head travels around the disc as it spins, like a vinyl record player. At each bit location the disc's magnetism can be up or down, and the read/write head can tell the difference, using the local magnetic orientation to represent 0 and 1. In flash memory, an array of microscopic electrical circuits is patterned on the surface of a semiconductor material, usually silicon. Each of these tiny circuits is a bit, since each one is accompanied by an electrical "gate" which can either allow or prevent electrical current from flowing through the circuit. Just like with other types of bits, these two states are the 0 and 1 which form the basic element of computing. The drastic design differences between the different varieties of bits show that there is no natural or preferred way to build a computer. The choice of what kinds of bits to use depends on many factors, including how many bits can be put on a device, how quickly the bits can be read and written to, and how much wear and tear the device must handle. These key metrics, in turn, are the end result of carefully choosing the component materials, optimizing the device layout, and perfecting fabrication techniques. These considerations are equally relevant for quantum bits, and their development is following a similar path.

Quantum bits, or "qubits," are the building blocks of quantum computers in the same way bits are for classical computers. And much like bits, qubits are usually designed with

two states, $|0\rangle$ and $|1\rangle$, with the $|\ \rangle$ notation indicating that these are quantum states. Now, what distinguishes qubits from bits is their ability to be placed in *superposition*, such as in this example, where the qubit $|\psi\rangle$ is a superposition of $|0\rangle$ and $|1\rangle$:

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle.$$

The superposition property bestows quite a few strange behaviors on a qubit. The relation above indicates that the qubit's value is not either 0 or 1 alone, but partly 0 *and* partly 1. α and β are called the *amplitudes* of $|0\rangle$ and $|1\rangle$, respectively, and represent how much of $|0\rangle$ and $|1\rangle$ are in $|\psi\rangle$. If the qubit is examined (in other words, a measurement is performed on it), it will be found either in $|0\rangle$ or $|1\rangle$. Exactly which state it will end up in is a matter of chance and can't be predicted; if the state $|\psi\rangle$ is created and then measured many times, it will sometimes be $|0\rangle$ and sometimes $|1\rangle$, and only the probability of finding $|0\rangle$ or $|1\rangle$ can ever be calculated. However, the $|0\rangle$ and $|1\rangle$ basis states are special – a qubit in one of these states can be measured over and over again without further changing its value. It is only the superposition states which are so fragile that a single measurement can disturb them. The challenge of quantum computing is that these fragile superposition states are absolutely essential to outdo a classical computer at any problems.

The superposition behavior of a qubit is unfamiliar to the world of classical physics, and is hard to accurately describe. It is sometimes described as being “both $|0\rangle$ and $|1\rangle$ at the same time,” but the qubit is only ever found in one of those two states, never both. Likewise, a superposition is not something between the two states, as if there were “1/2” or “2/3” states made from a mix of $|0\rangle$ and $|1\rangle$. A spinning coin on a tabletop is a closer analogy – while it is spinning, is it heads, or tails? The coin isn't really either, and it's not an average of the two; both faces are “in play” until it stops spinning. Quickly flattening the coin on the table is akin to measuring a qubit, as it suddenly disrupts the delicate state and produces a single, stable outcome (heads or tails, $|0\rangle$ or $|1\rangle$) which can then be checked over and over

without any more disturbance.

Multiple qubits can be put in special joint superposition states known as *entangled* states. If two or more qubits are entangled together, their values are no longer independent, like classical bits; they become correlated. The truly strange part is that these correlations exist even though the qubits' values are still unpredictable. As another example, consider a pair of qubits, each of which has $|0\rangle$ and $|1\rangle$ states, leading to the possibilities $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. These two qubits can be entangled in a superposition state $|\phi\rangle$:

$$|\phi\rangle = \alpha |00\rangle + \beta |11\rangle.$$

As before, this relation indicates that these two qubits have some probability to be measured in $|00\rangle$ and some probability to be measured in $|11\rangle$. Notice that only two of the four possibilities can occur, as there is no possibility to measure the individual qubits in different states; $|01\rangle$ and $|10\rangle$ have no amplitudes in $|\phi\rangle$. Somehow, the qubits appear to coordinate their outcomes, and this can be true even when the qubits have no physical connection through which to coordinate. It sounds like a magician's trick: they spin two coins on a table which always gave the same result – heads-heads or tails-tails – but never opposite. The surrounding apparatus becomes suspect; is it a special table? Weighted coins? Sleight of hand? Scientists have been searching for the trick up nature's sleeve for nearly a hundred years, and have largely managed to confirm that superposition and entanglement, bizarre though they may seem, are indeed an accurate description of qubits and other quantum systems.

With the key differences between bits and qubits in mind, classical and quantum computers can be contrasted. Structurally, quantum computers will resemble classical computers, using arrays of qubits in place of bits. The bit-level operations of quantum computers are far more complex, however. Classical computers need only read one bit at a time, checking if it is 0 or 1, and perhaps writing 1 or 0 instead. Quantum computers also read out single qubits,

but the write operations could entail putting the qubit in $|0\rangle$, $|1\rangle$, or any superposition of its states. In addition, quantum computers need to operate on multiple qubits at once in order to create entanglement. The requirement of potentially connecting any two of the device's qubits together at each computing step creates some major design challenges, which will be covered in more depth in section 1.1.2.

The discussion so far has focused on the abstract differences between classical and quantum, bits and qubits. In practice, qubits can take on many different physical forms. Researchers first identify systems that have qubit properties; in other words, they can exhibit superpositions and entanglement. There are quite a few possible qubit systems, including single electrons, atomic ions, microscopic superconducting circuits, and individual light particles (known as photons). This list makes clear that qubits tend to be much more delicate systems than bits – it is impossible to make qubits out of a punch card. In each case, different hardware must be built to protect and operate the qubits, and often additional hardware is needed to implement the complex entanglement operations. Several different types of qubit are under development at the leading quantum computing companies today, each with their own advantages and disadvantages. Superconducting circuits are popular due to their relative ease of fabrication and interconnection, as they leverage existing semiconductor fabrication processes and there are straightforward designs for configuring them into arrays. However, their relatively large size will prove challenging to scale to useful numbers of qubits. Similarly, single electrons trapped in semiconductors, known as quantum dots, can also be created using existing semiconductor fabrication techniques. Their smaller size is more favorable for scaling, but they are more difficult to link together into large arrays. A third platform, trapped atomic ions, takes an entirely new approach to designing a quantum computer, and completely eliminates the semiconductor base material. While this requires developing entirely new hardware and fabrication processes, the isolated atoms are ideal qubits with very high performance operations. However, as with quantum dots, linking them together into large arrays is challenging. It is not yet clear which quantum computing

platform will become the dominant paradigm, and it could even be a type of qubit not yet invented. Perhaps, like with different forms of classical bits, each will find an application niche that takes advantage of their respective strengths.

1.1.2 Why is it so hard to build a quantum computer?

The formative ideas of quantum computing have been around for forty years, and yet it is only in the last few years that real quantum devices have emerged with even a hope of performing any calculation which a classical computer couldn't do far more quickly. Even now, with money flowing into R&D from governments and venture capital firms, it is common to predict another twenty years or more until useful quantum computers are available. What is responsible for the slow development of quantum computing? In short, taking quantum computing from a vague idea of a novel type of computer to a functional device with real applications has required and will continue to require overcoming some towering experimental and theoretical obstacles.

It is easier to appreciate the experimental hurdles standing in the way of quantum computers. The hardware requirements of cutting-edge quantum computers entail networks of superconducting circuits cooled to a fraction of a degree above absolute zero; single electrons carefully isolated from their semiconductor surroundings and controlled with electrical gates only a few nanometers wide; or chains of atoms trapped together in a vacuum chamber and individually controlled with arrays of laser beams. Until recently, these architectures were entirely unattainable. The 'classical' supporting equipment of quantum computers is often overlooked, but most qubit systems could not function without the advances in cryogenic refrigerators, lasers and optical modulators, or control electronics which have emerged since the inception of the quantum computer. Even the pinnacle of scientific equipment today is not yet sufficient to run a sizeable quantum computer, nor is it known exactly what the final material and instrumental requirements will be. Each design iteration, as it pushes the limits of the existing technology, creates new questions and challenges to be answered by the

ensuing round of hardware innovation.

But these impressive feats of engineering raise a follow-up question: why are all these gadgets required to create a quantum computer? Why aren't the incredible capabilities of modern semiconductor foundries enough to create quantum processors as well as classical? After all, laptops and smartphones don't need to be cooled to near absolute zero or held in total vacuum to function. The answer stems from the intrinsic properties of qubits. Superposition and entanglement, those quirky properties only found in qubits, are delicate conditions. As previously described, any measurement of qubits in superposition results in the qubit states collapsing to a single, stable classical value like $|01\rangle$, which ends the quantum calculation. This is true not only for the deliberate measurements at the end of a calculation, but also for any uncontrolled, accidental measurements of the qubit. Even worse, the interaction of a qubit with its environment can be treated as just this type of uncontrolled measurement. Thus, unless carefully regulated, the noisy environment of a qubit will strongly disrupt any quantum calculations in progress, a process known as *decoherence*. The majority of experimental effort goes into isolating quantum systems from decoherence: superconducting circuits are vulnerable to thermal vibrations in the sample, and so they are placed in ultra-cold fridges which cool it until vibrations are tolerable. Trapped atomic ions will be lost if they collide with even a single molecule of atmosphere, and so they are held and tested with vacuum chambers that have 1/1,000,000,000,000th as much gas in them as ambient air. Despite the best efforts of scientists and engineers, it has so far proven impossible to remove all decoherence, and so current qubit platforms all come with a ticking clock – once put into superpositions, the qubits can only maintain their quantum information for a short time (ranging from roughly microseconds to seconds, depending on the platform) before decohering due to the environment and effectively becoming classical again.

The prospects for successfully running quantum computations appear even worse in light of a fundamental fact of quantum information theory known as the No-Cloning Theorem. This theorem states that it is impossible to take a mystery qubit in some unknown su-

perposition and make an exact copy of it [5]. Were it possible to make such copies, most decoherence issues would be easy to overcome by making many copies of a qubit and using measurements of those copies to correct any errors that creep into the system. However, with the No-Cloning Theorem rendering that approach moot, there appears to be no better way to avoid decoherence than isolating qubits from the rest of the world. Fortunately for quantum computing, while reducing intrinsic decoherence is still always worthwhile, theoretical physicists developed a scheme for correcting any errors that occur in a tightly-controlled quantum system without needing to measure the target qubits directly, known as quantum error correction.

Quantum error correction (QEC) is one of the key enabling ideas for the hope of creating a practical quantum computer. In QEC, a target register of computational qubits to protect from errors is augmented with additional qubits, known as ancilla qubits, which are entangled with the target register. The advantage of using these ancilla qubits is that they can be measured without disturbing the target qubits. To identify errors, the ancilla qubits are entangled in such a way that their measurements change based on whether or not an error has occurred on the important computational qubits. Finally, quantum logic gates can be applied to the computational qubits to fix the errors, without ever having measured these qubits (and thus pre-emptively terminated the calculation) [6]. QEC thus provides a clever way to skirt the No-Cloning Theorem handcuffing quantum measurements. However, it is not a silver bullet for all of quantum computing's ills – the error correction circuits impose a computational overhead, both in terms of additional computing steps and additional required qubits. In fact, in most proposed QEC approaches with realistic error rates, several orders of magnitude more qubits will be required to correct the errors in a quantum computer than to carry out the calculations themselves. This means that a few thousand qubits – a few quantum kilobytes – might require quantum megabytes of error correction capacity. Finally, the error correction circuits might fail altogether if the error rates of the underlying qubits aren't low enough to start with. Again, landmark results in quantum information

theory have helped establish exactly how low is “low enough,” in what has become known as “threshold theorems.” These theoretical results indicate that if the error rate on the individual qubits of a quantum computer are below a given threshold, then that computer will be able to carry out computations of any length without being lost to decoherence [7]. This threshold varies significantly depending on the specific connections that exist between computational qubits and their corrective ancilla qubits, and so an ongoing area of research involves designing improved qubit architectures that make the best use of the fewest qubits to protect coherence. Some recent designs have pushed the threshold to 1% error rates or higher [8], which is in range of most of the popular qubit platforms, which in principle makes increasing the number of qubits a viable path to realizing QEC.

While developing QEC and proving threshold theorems have provided major boosts to the viability of quantum computing, there is plenty of theoretical work remaining. Indeed, one of the central questions to address is precisely what advantages a fully-realized quantum computer will have over a classical computer. This seems like the first question that needs to be answered before embarking on the heroic quest to build a quantum computer; how could it not already be thoroughly investigated and soundly answered? There are a few factors at work. In the first place, nascent quantum computers have to compete against fully mature classical supercomputing opponents. With billions of operations per second and storage space for trillions of bits of information, even cheap classical computers today have capabilities that it’s hard to imagine quantum computers ever will. Thus, in order to be competitive, quantum computers must find their niche with problems where they have a significant advantage in algorithmic complexity – in other words, problems where the quantum solution simply requires a lot fewer steps than the classical solution, and where the advantage gets bigger as the problems get bigger. Identifying and proving which problems fall into this category of quantum advantage has been a slow process; one with a few headlining successes, but with very little low-hanging fruit. In some cases, progress even seems to move backward: it may appear clear that quantum computers have a clear advantage on a problem, until a

new classical algorithm starts to close the gap. Sometimes, the classical algorithm is even inspired by the once vastly-superior quantum algorithm [9]. While such cross-pollination of algorithm design is great news for enhancing the general power of computing, it does occasionally erode the apparent advantages of quantum computers. Finally, it can be much easier to convince oneself that one device or algorithm is superior to another than it is to prove it. This seems to be the case for quantum computers. While in an abstract sense they have clear advantages over their classical counterparts, rigorously certifying an advantage in a specific case is no simple task. This sticking point occurs in one of the most well-known open questions of classical computing, referred to as $\mathbf{P}=\mathbf{NP}$ for short. The $\mathbf{P}=\mathbf{NP}$ problem is a good place to begin discussing the exciting opportunities of quantum computing, in section 1.1.3.

In summary, quantum computers have faced towering challenges on both hardware and software fronts. A long list of engineering advances have been required to observe and stabilize qubits for use in computing, and more still will be needed to drive down decoherence rates, develop control protocols, and scale to large enough numbers of qubits for useful quantum devices. On the theoretical front, constant effort has been required to understand how qubits operate differently than bits, to invent quantum algorithms, and to prove that quantum computing advantages exist for those algorithms. The challenges are large, but in retrospect the progress has been rapid and continuous, offering hope for future breakthroughs and the eventual advent of quantum computers with truly incomparable performance.

1.1.3 Why are quantum computers exciting?

Quantum computers are exciting in part because of their novelty and strangeness, but they have only secured significant private funding because of sound reasons to expect useful technological developments. However, along with venture capital backing and media attention has come an abundance of “hype” and many misunderstandings regarding the potential impacts of quantum computers. It is easy to argue that quantum computers should be more

powerful than classical computers; the comparisons between bits and qubits in section 1.1.1 show that qubits can do everything that bits can, plus some things they cannot. Thus, there may be a huge number of quantum calculations which aren't possible on classical computers. On the other hand, there is no guarantee that any of these calculations are useful. Without a scientific understanding of computer algorithms, it is hard to say what the potential impact of quantum computers might be. Fortunately, QISE has received some milestone contributions from the field of theoretical computer science, especially complexity theory, which have helped clear up some of the questions around how exactly quantum computers are superior to classical.

Complexity theory is a subfield of computer science which tries to evaluate how hard different types of computational problems are. The “hardness” of a computational problem is made precise by estimating the number of steps that a computer would have to take in order to solve it. For example, consider the challenge of sorting a list of 20 items (of names, grocery items, phone numbers, etc.) in alphabetical or numerical order; this is a job which computers and humans both frequently do. A reasonable procedure, or *algorithm*, might be to scan the list for what should be the first entry (“Aaron,” a 202 area code, or similar) and write that at the top of a new list. This is done by checking every item in the list to see if it's earlier than all the other items, and so the first pass takes 20 steps. Then search the list again for the next item (another 19 steps), the third item, and so on. It takes 20 passes and 210 total steps to fully sort the list. Now, why care about the total number of steps? After all, a computer could take millions of steps in the amount of time it would take a person to do ten steps. Several reasons: first, it is important to make huge calculations take as few steps as possible, even for supercomputers. Second, consider how the number of steps increases as the size of the problem increases. What if the list was twice as long – 40 items? In this case, the sample algorithm would require 840 steps. Even though the problem size doubled, the number of steps quadrupled. This is common for most problems in computing – the number of steps increases much faster than the size of

the problem. The relationship between problem size and problem difficulty is predictable, and allows scientists to group problems and their algorithms into different *complexity classes* which contain problems of similar hardness. While it seems concerning that the example sorting algorithm got four times harder when the list got twice as long, sorting is actually one of the relatively easy problems in computing. Sorting a list of one thousand elements might take a modern computer only around one millisecond, so the size of the list can be doubled many times before it slows down to a noticeable level. Problems of this type are considered “efficiently solvable,” and grouped into the class **P**, for “polynomial time.”

Not all problems fall into **P**, however, and many important problems are far more difficult. For example, a canonical hard problem is the Traveling Salesman Problem, in which a person or vehicle wants to calculate the quickest route to visit a group of cities or points of interest on a map. What makes this problem so hard is that, while it’s fairly easy to do for a handful of cities, each additional destination explodes the number of possible routes to consider. The scaling is so extreme that adding just *one* additional city more than *doubles* the number of steps required to calculate the shortest route [10]. This rapid growth is characteristic of the exponential function, and so the Traveling Salesman Problem algorithms are termed “exponential-time” algorithms. Exponential-time algorithms quickly become too hard to run even on supercomputers, which is an obstacle for the many applications which would benefit from exact solutions. In some cases, a rough approximation to the best answer can be found more quickly, but software engineers still know exponential-time algorithms as a pitfall for any problem which needs to be solved at scale. The Traveling Salesman Problem, and many others of similar difficulty, are grouped together into the class **NP**, for “nondeterministic polynomial time,” a class which also contains all the problems in **P**¹. The relationship is depicted graphically in Fig. 1.1, along with **NP-complete**, a subset of **NP** which contains

1. The exact meaning of this name refers to technical characteristics of the problems it contains. Interestingly, despite the apparently clear differences between the hardest problems of **NP** and those in **P**, they have not yet been rigorously proven to be separate classes. This is the famous **P=NP** problem of complexity theory.

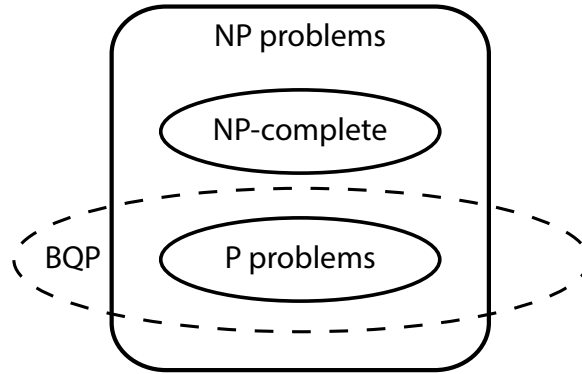


Figure 1.1: The hierarchy of quantum-relevant complexity classes. *Bounded-error quantum polynomial time* (**BQP**) represents the range of problems which can be efficiently solved by a quantum computer. While its span isn't fully understood yet, it completely contains the problems which a classical computer can efficiently solve (**P**), and some which it cannot (**NP**). It is still unknown whether or not **BQP** is completely contained in **NP**. **NP-complete**, which contains the provably hardest problems in **NP**, is not expected to overlap with **BQP**.

problems that are proven to be the hardest in **NP**.

The key question for QISE is how a quantum computer fares on these classes of problems when compared against classical computers. Encouragingly, a number of problems have been identified where a quantum computer can be proven (theoretically) to scale better with problem size than the best classical algorithms. Thus, quantum computers have been given their own complexity class, termed **BQP** for “bounded-error quantum polynomial time,” which contains all the problems which are efficiently solvable on a quantum computer. Fig. 1.1 also depicts the proposed relationship of **BQP** to other complexity classes. As expected from the qubit/bit relationship, **BQP** completely contains **P**, since a qubit does everything a bit can do. **BQP** is not believed to contain **NP-complete** at all, which means that computers will not be a magic bullet for contemporary hard computer problems. However, it still may extend beyond **NP**, which would be a remarkable discovery and is an active area of research. So far, several algorithms have been identified for problems which, by all appearances, are in **BQP** (quantum-feasible) but not in **P** (classical-feasible). These algorithms are the heart of the excitement around quantum computing, and there is optimism that these algorithms are just the tip of the iceberg, a hint that as computer scientists gain

familiarity with the unusual principles of quantum algorithm design, they will discover a host of new algorithms which continue to expand the region of quantum superiority over classical.

Quantum computing's advantage lies entirely in the superior scaling of these known quantum algorithms with no classical analogues. Qubits currently operate much more slowly than bits, so there are likely few computing gains to be had on small scales. However, as with the sorting and Traveling Salesman problems, quantum algorithms can win out in the end for large enough problem sizes, where classical computers would require an unreasonable number of steps. The benefit of the widening gap in the quantum versus classical steps to complete an algorithm will eventually outweigh the slowdown from classical bit to qubit. A principal goal in QISE research has been to seek out more algorithms with advantages over classical algorithms, especially where the classical algorithm is exponential-time, i.e. it scales very poorly.

The quantum list-searching algorithm, known as "Grover's algorithm," is one which illustrates how unusual quantum algorithms can be. The problem is to find a certain element in an unsorted list. For example, the task might involve finding someone with the first name "Francis" in a phone book, or checking if there are any books in a library catalog that have "dinosaur" in the title. Like the sorting problem, there is a very straightforward classical approach: start moving through the list from the beginning and check each item. On average, this algorithm has to check half the items in the list before it finds the target item. Thus, if the size of the list is doubled, the search also takes twice as long. Grover's algorithm, on the other hand, can take advantage of qubit superpositions and reach the correct answer far faster. In fact, it only takes a number of steps roughly equal to the square root of the number of list elements [11]. This means that the size of the list would have to be *quadrupled* before Grover's algorithm needs to take twice as many steps. On its face, this seems impossible. Doesn't an algorithm, quantum or not, still need to check every element in the list? That is classical thinking, and it is a challenge to all the researchers in the

field to develop “quantum thinking” in the pursuit of new algorithms. Superposition and entanglement allow quantum algorithms to approach problems in unusual ways, and for the list-search problem it results in an advantage over classical algorithms. However, it remains to be seen how widely applied Grover’s algorithm ultimately becomes. List-searching is a prevalent problem, both on its own and as a stepping stone for more involved algorithms, so there may be a lot of opportunities to apply a quantum search. On the other hand, it doesn’t achieve an exponential-time advantage over its classical counterpart, so the lists (and the quantum computers searching them) may have to be quite large before a quantum search becomes faster than a classical search.

Fortunately, another quantum algorithm does provide an exponential-time advantage – Shor’s factoring algorithm. The factoring problem is simple to describe: given a number, what are its prime factors? This is simple to do for small numbers, thanks to experience and a limited space of possibilities. 21 is 3×7 , and 527 is 17×31 , but what about 34,687,528,507,845,055,319? It would take significant effort for either a human or a computer to figure out that the two prime factors are 8,085,263,101 and 4,290,216,419. The factoring problem is hard enough that the only known solutions are exponential-time algorithms. The difficulty of the factoring problem has made it a key cog in a popular data encryption algorithm called RSA, which uses a large number to encrypt a message and one of its prime factors (known to the receiver, but not to the public) to decrypt it. Shor’s algorithm, however, does not take exponential time, only polynomial time. The drastically improved scaling for Shor’s algorithm suggests that it might be able to overtake a classical computer on relatively small problem sizes. The discovery of this algorithm and its low threshold for superiority has led to one of the most-discussed potential applications of quantum computing: quantum code-breaking. A sufficiently large quantum computer could efficiently find the prime factors of the large number used to encrypt large volumes of data, undoing the encryption for the quantum computer operator. The current RSA standards employ encryption keys a few

thousand bits long, so at least as many qubits would be required to break such an encryption². Given the prevalence of RSA encryption, including for storing and transmitting important data, this would place a large fraction of the world’s currently encrypted data at risk. It’s perhaps no wonder that governments and intelligence agencies have focused on the national security implications of a working quantum computer.

As an aside, while Shor’s algorithm appears to be a matter of serious public concern and one which fates quantum computers to be used for malicious purposes, there are two mitigating factors. The first is that a range of encryption options exist which do not rely on the difficulty of the factoring problem. Intelligence and standards agencies are proactively establishing an updated set of “post-quantum” cryptographic standards, which will greatly reduce the potential threat from a future quantum computer [12]. If and when it is ultimately constructed, a quantum computer will undoubtedly be tested against encryption standards, but by then the quantum-vulnerable options are likely to be phased out. The second is that the core element of Shor’s algorithm, the Quantum Fourier Transform, is useful for a wide range of problems besides attacking RSA encryption – cryptography applications are simply the ones which have garnered the most media attention thus far.

The quantum algorithms presented so far constitute some concrete reasons for optimism about the prospects of quantum computing, but there are several abstract reasons, as well. Nature, at its fundamental level, obeys the laws of quantum mechanics and is not simply classical. Building a quantum computer, then, requires scientists to understand and control matter and energy at its very base level. It is hard to imagine that this dramatic increase in our ability to engineer the natural world will *not* at some point lead to revolutionary technological developments. In pursuit of rigidly controlled quantum interactions, researchers are working on building atomically perfect materials, seeking the limits of precision measurement, and miniaturizing electronic components to unprecedented scales. From this position

2. These qubits would need to be error-corrected, which balloons the number of required qubits to the hundreds of thousands, millions, or more, depending on the quantum gate error rates and error correction paradigm.

at the emergence of a fledgling field, it may be hard to predict exactly what industries and wells of innovation may result from these new capabilities. The uncertainty is both enticing and cautionary, and suggests that a prudent course of action for the future of QISE involves striking an appropriate balance of risk-mitigating and exploratory research. Risk mitigation would entail continuing to improve estimates of the hardware and software requirements for functional quantum computers, to best assess the feasibility and timeline of their development. On the other hand, exploratory basic research will help open the door to as many breakthrough technologies as possible, casting a wide net in search of the unique and unusual applications of quantum phenomena.

One final, esoteric angle on quantum computing applications is worth mentioning – the potential integration of quantum algorithms into artificial intelligence. Across all of the natural world, scientists have discovered only two types of information (and corresponding information processors): classical and quantum. Despite large disparities in size and speed, all classical computers ranging from handheld calculators to supercomputers process their classical information according to the same fundamental rules. This “substrate independence” is what enables the results of computational complexity theory to be so widely applicable. Likewise, all quantum information processing systems follow the rules of quantum information. In every information processing system found in nature so far, the information is either classical or quantum – and that includes human beings! That implies that human-level intelligences could be perfectly emulated with a classical or quantum computer of sufficient complexity. However, it isn’t yet known for sure whether humans are entirely classical or quantum; while the most likely scenario is that human intelligences are classical computers, formed from literal neural networks and with bits implemented via chemical and electrical signaling, there have also been a few proposals for how coherent quantum information could be processed in the brain [13]. Regardless of the eventual resolution of this curious question of human nature, quantum information would contribute an intriguing twist to the quest to build human-level artificial intelligence. Either the human brain is a quantum information

processor, and quantum computing serves as the key ingredient for creating a comparable AI, or brains are merely classical computers, and a quantum computer serves as the only known route to surpass the fundamental limitations of the realm of brains and supercomputers.

1.1.4 Are there other important technologies in QISE?

The majority of the discussion so far, not to mention the bulk of mass media publicity surrounding QISE, has focused on quantum computing specifically. However, just as there are classical devices of interest beyond computers, there are quantum technology applications beyond quantum computers. The quantum technologies other than computing have been grouped into three general categories: quantum simulation, quantum communication, and quantum sensing.

Quantum simulation sounds, and is, related to quantum computing. In the field of quantum simulation, the objective is to outperform classical computers at the task of simulating large quantum mechanical systems, such as molecules or high-temperature superconductors [14]. This premise harkens back to the origins of the field of QISE, where Feynman supposed that a computer obeying quantum principles would have a much easier time simulating quantum systems than a classical computer [4]. The main difference between a general quantum computer and a quantum simulator is that a simulator is designed to model a specific class of physical systems, such as an atomic lattice, a set of similar molecules, or the nucleus of an atom. A quantum computer is intended to be completely reconfigurable, and able to run arbitrary quantum calculations. In this sense, quantum simulators are limited, highly specialized realizations of quantum computers. From a classical computing perspective, accurately simulating quantum mechanical systems is difficult, due to the extensive correlations between qubits that come along with superposition and entanglement. Perfectly simulating the evolution of a quantum system is an “exponential-time” problem, and so is a very slow process for even just a few atoms at a time. With classical computers, scientists must make approximations or enforce size limitations on the quantum systems they’d like to simulate.

Qubits make a far more natural basis for simulating other quantum systems. By leveraging the natural analogies between many different qubit platforms and focusing on a limited set of problems, researchers hope to more rapidly attain advantages over classical computing than by developing a complete, error-corrected quantum computer. Especially with this approach, quantum simulator development is likely to split in many different directions and use many different qubit platforms, each chosen based on their ability to simulate specific analogous quantum systems. Quantum simulators may even jump start a virtuous feedback cycle, where early quantum simulators are used to gain greater understanding of important materials or physical laws that are used to build better quantum simulators and computers.

Quantum simulators are already a reality, to an extent. No sharp threshold exists for what “is” and “is not” a quantum simulator, since even a single qubit can technically be considered a simulator for other single qubits. However, the threshold for an *effective* quantum simulator is outperforming classical simulations at the same task, or even simulations which are too hard to do classically. This milestone is more challenging, but has been reached in at least a few cases [15]. The threshold for a *useful* quantum simulator is higher still, as the simulation should help answer a research problem, but there are several promising directions and it seems only a matter of time before quantum simulators are a common computational resource in scientific disciplines. Given their highly specialized nature, it seems unlikely that quantum simulators will enjoy widespread commercial appeal. Eventually, when they reach a sufficient level of sophistication, quantum simulators could be used by materials and pharmaceutical companies in performing semiconductor and chemistry calculations to serve the development of new drugs and materials. In the shorter term, once the “useful” threshold has been convincingly cleared, quantum simulators are likely to serve as research resources for making and validating quantum calculations for research problems, much like the supercomputing clusters already available at national laboratories and universities.

Quantum communication is a fairly self-descriptive term, and refers to the capability of transmitting and receiving coherent quantum information across long distances. Of partic-

ular interest is the case of sending one of a pair of entangled qubits across a network, to connect the quantum systems of the sender and receiver. On its face, it sounds completely outside the realm of possibility – somehow qubit states, which are already difficult to protect in low-temperature vacuum chambers, are supposed to be maintained in transit across countries and oceans? Indeed, it is only because qubits based on individual particles of light (photons) have different properties than most other, “matter” qubits, that quantum communication is even plausible. Whereas matter qubits are relatively easy to keep stationary, photon qubits always travel at the speed of light. And while matter qubits have strong interactions with their surroundings that lead to rapid decoherence, photon qubits can have long coherence times if they are not absorbed. Thus, photons are a logical choice as a vehicle for carrying quantum information from one place to another. The principal challenges lie in mediating the conversion from stationary qubit to photon qubit and back again, and in ensuring the photons are not absorbed. For the first challenge, success appears contingent on a combination of perfecting fabrication processes, improving qubit architectures, and developing new optics components. A variety of methods have been shown to successfully transfer stationary qubit information into a traveling photon, and leading research is now on increasing efficiency of the transfer. For the second challenge, the telecommunications industry already has decades of experience in transmitting light from point to point with low absorption (sometimes called low “loss”). However, not only are the loss requirements of quantum computing technologies more stringent than those of classical fiber optic communication, but the standard signal amplification techniques cannot be applied to quantum communication. Classical laser signals can be measured and perfectly retransmitted to avoid errors, overcome loss, and extend the network range. Photon qubit transmissions, on the other hand, are converted back to classical information if they are ever measured, as described in section 1.1.1. New techniques have been proposed to reliably extend qubit transmissions: relay stations called “quantum repeaters” would share entangled qubits with the sender and receiver separately, and can then perform quantum operations to entangle the qubits held

by the sender and receiver directly. In this way, each quantum repeater can halve the distance required for qubit transmission. While full operation of a quantum repeater has not yet been demonstrated, precursor experiments have shown promising results in transmitting entanglement over short distances [16].

There is an alternative to using fiber optics to transmit quantum information: satellite-to-ground communication. A satellite-based approach circumvents the inevitable losses of even the best fiber optics by sending the photon qubits into the extreme vacuum of space, where there is virtually no chance of absorption. While loss by absorption is solved in space, transmitting to a satellite introduces the problem of ensuring the photons are correctly aimed at their target, which is trivial in fiber optics. Recent demonstrations on satellites equipped with quantum photon emitters have overcome the targeting hurdles and successfully distributed entangled photons to two earthbound receivers separated by 750 miles [17]. This provides satellite quantum communication an early advantage in terms of distance, but as of yet, satellite quantum communication has yet to demonstrate conversion back to stationary qubits at the receivers. The respective advantages of fiber optics and satellites may lead to the adoption of both for specific applications.

Quantum communication has several distinct application areas, and some have already resulted in commercial products. The earliest and most mature quantum technology, quantum key distribution (QKD), uses the correlated randomness of entangled qubits to generate encryption keys for large-scale, secure classical communication. This technology aims to replace RSA and other public-key encryption schemes, which are only secured by the presumed difficulty of reverse-engineering a mathematical problem like factoring. In QKD, by delivering qubits of an entangled pair to both the sender and receiver, the two parties can generate shared, unpredictable strings of bits to encrypt their classical communications. If a hacker attempted to intercept the quantum information transmissions used to share entanglement, it would again convert the transmission to a classical signal and destroy the entanglement. Thus, any interference in the communication channel is (theoretically) detectable, and the

encryption secure against attacks with either a classical or quantum computer. However, whether such an extreme degree of security is needed in most cases is questionable; the National Institute for Standards and Technology is developing a set of classical encryption standards which are expected to be secure against attacks by either supercomputers or full-fledged quantum computers, providing sufficient data protection at least in practice, if not in principle, for most applications. Short-range QKD networks are already in use in some national laboratories and metropolitan areas, and the aforementioned satellite-to-ground quantum communication channels have successfully implemented QKD over much longer scales. QKD networks could reasonably be deployed for secure networks that require the highest assurances of long-term security, such as internal networks in the defense, finance, and healthcare sectors.

Blind quantum computation is a prospective application which combines unique features of quantum computing and quantum communication, and similarly attempts to address concerns around data privacy. Cloud computing has become an enormous industry, but clients of cloud computing firms must be comfortable handing control of their data to third parties, exposing confidential information to additional security risks. In the classical arena, encrypted computation (also referred to as “confidential computing”) can protect data even on untrusted servers. Blind quantum computation provides a quantum analogue to this service, which will be in demand given the likelihood of the most powerful quantum computers being centralized resources. Once again leveraging the No-Cloning Theorem, blind quantum computing allows a quantum cloud computing client to transmit quantum information (the details of which are known only to the client) to the quantum cloud server, which is then used in the specified quantum computations. Since the transmitted information is secret and cannot be determined by the cloud service, the service provider can know and correctly implement the desired calculations without knowing their purpose or results [18].

Entangling disparate quantum systems via quantum communication channels is a fundamental operation for both computing and sensing. Especially in the early stages of quantum

computing, when individual devices have limited numbers of qubits, it may be easier to entangle multiple devices than to grow the size of a single device. Mature quantum communication technology might therefore provide a shortcut to scaling up quantum computers. Entangled sensors can also provide advantages over non-entangled sensors [19], and if the sensors are spatially separated, quantum communication techniques will likely be the means of entangling them.

Quantum sensing is the final major area of quantum technology. It is also the most heterogeneous of the application areas, covering the widest range of qubit platforms and as many different use cases as there are types of sensors. The unifying principle behind all quantum sensors is to leverage the quantum properties of a sensor to outperform purely classical sensors. In some cases, quantum sensors even make measurements that no classical sensor has been able to. For example, quantum sensors can detect individual microwave photons [20], which is far below the detection threshold of classical microwave detectors.

Sensing may not seem like a particularly noteworthy application for quantum technologies, but the importance of sensing is often understated. Smartphones include a wide range of sensors which enable convenient and by-now-expected features, such as GPS location readings, touchscreen and fingerprint interfaces, and gyroscopes for orientation-dependent interactions. Cameras and distance sensors are critical for the computer vision systems that enable self-driving vehicles, and the acquisition rate of images and environmental data directly factors into how safe and reliable such a vehicle can be. Medical diagnoses depend on accurate readings of tests for concentrations of chemicals or the presence of specific viruses and bacteria. These are just a few of the roles that sensors play in making technological progress possible – new measurement capabilities often lead to engineering innovation.

Improved sensing takes many forms. In some cases, a quantum sensor can take the same measurements in less time, or make higher-resolution measurements of the same target. In other cases, a quantum sensor can be miniaturized beyond classical capabilities. The quantum sensing systems under study are useful for a range of physical quantities, including force,

motion, magnetic and electric fields, and temperature [21]. The quantum systems themselves range from clouds of atoms to superconducting circuits to nanometer-sized vibrating membranes.

Of course, not all sensors see widespread commercial success. Some sensors are so specialized or require such extensive hardware support that they are never deployed outside of laboratory settings. This doesn't mean that the device is useless; many discoveries can be made in a lab and applied elsewhere. Most quantum sensors are restricted to the lab because of the difficulty of maintaining quantum coherence under uncontrolled conditions. An emerging quantum sensor, the nitrogen-vacancy (NV) center in diamond, has already been employed in numerous laboratory experiments, but also shows promise for deployment outside the lab because of its exceptional stability and coherence properties. The NV center is the central quantum system of this thesis.

* * *

Four decades after the first quantum computers were proposed, the field of QISE looks different than at its inception, and in many ways that could not have been predicted at the time. The major areas of application only came into focus along with significant research advances, and have been shaped by the parallel progress in classical computing, communication, and sensing. While it's a safe bet that the next forty years will be equally unpredictable, the current trajectory of quantum R&D seems to be accelerating and enthusiasm is high that commercial success is imminent. Like most new developments, the course of QISE will fall somewhere between the most optimistic and pessimistic projections. Unpredictable breakthroughs will occur and launch the field forward, but classical information industries aren't remaining stagnant either. No matter how much venture capital is committed, quantum computers are still far from implementing their envisioned quantum algorithms, and the intervening cycles of test and development will all take time. New hurdles will be discovered and surmounted, new applications will appear and others will become obsolete, and developmental estimates will have to be revised many times. But after decades of trying, no

fundamental obstacles to the creation of a quantum computer have been found. The future looks bright for the development of remarkable quantum technologies.

Chapter 2

The Nitrogen-Vacancy Center

This chapter will introduce the nitrogen-vacancy (NV) center in diamond, including its structure, the rich physics of what at first glance is a simple system, and the relevant applications in and outside of physics. The rest of this dissertation will assume the reader has familiarity with undergraduate-level quantum mechanics.

2.1 Physics of the Nitrogen-Vacancy Center

2.1.1 Structure and Formation

The NV center is an atomic defect contained within the lattice of a diamond crystal. It forms when a substitutional nitrogen impurity atom is adjacent to a vacancy, each replacing a carbon atom. Fig. 2.1 displays a diamond unit cell containing an NV center. The dangling bonds provide five non-bonding electrons, two from the nitrogen and one each from the three nearest-neighbor carbon atoms. The NV center has multiple stable charge states, and only the negatively charged NV center will be discussed from now on. With the capture of this additional electron, the NV center is formed as these six electrons fill up the molecular orbitals composed of the local dangling bonds. These molecular orbitals create an electronic level structure that resembles that of an isolated atom, with sharp excited state and hyperfine




Figure 2.1: Diamond lattice with NV center. A diamond unit cell contains tetrahedrally-coordinated carbon atoms. Here, a nitrogen impurity (N) and a vacancy defect (V) combine to form the NV center. A typical coordinate system is also shown.

transitions in spite of the surrounding crystal lattice. Diamond’s band gap of 5.5 eV is among the largest yet identified for a semiconductor, and leaves a pure diamond crystal transparent in the visible spectrum. The NV’s excited state transition falls within this experiment-friendly visible range, well within the band gap, which makes it possible to observe isolated NV centers within a diamond chip under a microscope when simultaneously excited with a laser. While single NV centers are too dim to register on a photodiode, and require single-photon detectors, ensembles of NV centers clearly register on CCD cameras, and particularly dense ensembles – to the extent the host diamond appears nearly black – can even be seen with the naked eye under laser illumination. Both single and ensemble NV centers have unique applications, and their relative strengths will be discussed in later sections. Since single NVs can be optically resolved, there is a natural transition from “single” to “ensemble” NV samples at the point where the NV centers become too dense to be individually resolved. In normal confocal microscopes, single NV samples host concentrations below 10^{12} NVs/cm³, while ensemble samples may have densities of 10^{16} NVs/cm³ or even higher.

Each lattice site in the diamond lattice has four nearest neighbor lattice sites, configured in a tetrahedral arrangement. Thus, for any substitutional nitrogen atom, there are four equivalent sites for a vacancy to occupy and create an NV center. Crystallographically, these

are the $\langle 111 \rangle$ directions. It is most common for rectangular diamond chips to be created with a $\langle 100 \rangle$ -oriented surface, for which all NV axes have a 54.7° angle to the surface. Under most circumstances, the four orientations are randomly distributed. If a specific NV orientation is desired for a single NV, several centers may need to be searched. NV center ensembles are composed of an even mixture of the four orientations, although there is statistical variation from region to region of the sample. While all four orientations are equivalent in the absence of external fields, an applied magnetic or electric field can break that degeneracy and single out one of the four NV sub-ensembles. In this case, experiments may be designed to probe only the desired NV orientation, with the remaining NV centers treated as background noise.

NV centers are formed both naturally and intentionally in diamond crystals. Nitrogen is the most common impurity in naturally-formed diamonds, and as an impurity is responsible for the yellow tint of class Ib diamonds¹. Vacancies are also plentiful in natural diamonds, which form under high temperature and pressure conditions that introduce many microscopic imperfections. Randomly, these nitrogen impurities and vacancies can coincide, naturally forming NV centers. For many years, high-quality, natural, class IIa diamonds provided the best substrates for single NV experiments. In addition to NV centers, many other species of defect tend to form in natural diamonds, but NV centers can usually be distinguished due to the high concentrations of nitrogen compared to other impurities, the relatively bright photoluminescence (PL) of NV centers, and their characteristic spectral signatures. Today, however, the diamond samples which host the highest-quality NV centers are lab-grown, not mined, in a process called chemical vapor deposition (CVD). CVD growth begins with a single-crystal diamond substrate, and takes place in a high-temperature vacuum chamber. Methane (CH_4) and hydrogen gases are released into the growth chamber, and enough rf power is injected to break the chemical bonds and create a carbon and hydrogen plasma above the substrate. With the right balance of conditions, the carbon atoms grow in epitaxial layers

1. In gemological classification of diamonds, Ib are distinguished by high nitrogen concentration and have a yellow tint. IIa diamonds have low impurity densities and are transparent.

of high-quality, low-nitrogen diamond. Growth techniques have advanced to the point where CVD samples often have so little residual nitrogen and vacancy concentration that few NV centers will be visible immediately, and they must be deliberately engineered in the sample.

Two methods exist to create NV centers: implantation and growth. In implantation, nitrogen ions are accelerated to energies of 1 - 200 keV and directed at a diamond substrate, where they penetrate to depths of up to several hundred nanometers. The collisions with lattice carbon atoms create a trail of vacancies along the path of the ions, and the stochastic series of collisions creates a few nanometers of variance in the final nitrogen impurity positions in the crystal. After implantation, the vacancies are mobilized in the lattice by annealing at temperatures above 800 °C, and some of them combine with nitrogen impurities to form stable NV centers. Alternatively, nitrogen can be introduced during the CVD growth process. As the crystal is grown one atomic layer at a time, nitrogen gas can be released alongside the methane and hydrogen. Some of the nitrogen is incorporated alongside the carbon in the diamond, and with proper calibration, this produces nitrogen-enriched layers of diamond as thin as two nanometers [22]. Unlike ion implantation, the growth process does not produce a significant number of vacancies, and so additional vacancies are then created in the sample, typically with electron irradiation. Finally, the sample is annealed until the vacancies reach the nitrogen-rich layers and form NV centers. Both sample preparation methods have advantages, subtleties, and variations which determine the correct choice for each application. In general, CVD growth is used to create single NVs with the best optical properties and in isotopically purified material (see section 2.1.4), while implantation is preferred for ensembles.

2.1.2 Spin and Optical Properties

The basic spin and optical properties of the NV center emerge from the overall C_{3v} symmetry of the center and its six electrons. Four of these electrons pair into lower-energy orbitals, while the remaining two form a spin-1 system across two degenerate orbitals, with triplet and




Figure 2.2: NV center ground state level structure. The ground state spin triplet degeneracy is broken by the zero field splitting, $D_{gs} = 2\pi \times 2.87$ GHz. The states further split in the presence of a magnetic field parallel to the NV axis according to the gyromagnetic ratio $\gamma_e = 2\pi \times 2.8$ MHz/G.

singlet subspaces. The ground state is the triplet with 3A_2 symmetry, containing $m_s = 0, \pm 1$ angular momentum states, which will be labeled $|0, \pm 1\rangle$ respectively. The crystal field breaks the degeneracy of the triplet states, contributing a $D_{gs}S_z^2$ term to the Hamiltonian, where $D_{gs} = 2\pi \times 2.87$ GHz is the zero-field splitting². This splitting creates a natural coordinate axis for the NV center, with the z axis along the vector connecting the nitrogen and vacancy lattice sites. The NV center couples to magnetic fields with a g tensor that appears isotropic and almost identical in magnitude to that of a free electron. Due to the ground state orbital singlet structure, the coupling to electric fields is much weaker, and so the principal Hamiltonian of the NV center ground state is [23]

$$H_{gs} = D_{gs} \left[\hat{S}_z^2 - S(S+1)/3 \right] + \gamma_e \mathbf{B} \cdot \mathbf{S}, \quad (2.1)$$

with the electron gyromagnetic ratio $\gamma_e = 2\pi \times 2.8$ MHz/G. At room temperature, the excited state spin sublevels of the NV center largely resemble its ground state, with a smaller zero-field splitting $D_{es} = 2\pi \times 1.42$ GHz.

2. Throughout this dissertation, \hbar will be set to 1, and energies expressed in units of angular frequency.




Figure 2.3: NV center photoemission spectrum. The emission from an ensemble of NV centers at room temperature is shown under 532 nm excitation. The zero-phonon line (ZPL) is visible at 637 nm, comprising approximately 3% of the total emission. The remaining photoluminescence (PL) corresponds to phonon-assisted transitions. A longpass filter at 600 nm is placed in the collection path to remove the excitation laser.

While the NV center is now primarily studied for its spin properties, its discovery and development have been facilitated by its characteristic optical properties. The excited state transition occurs directly at 637 nm – the zero-phonon line (ZPL) – and in a broad spectrum of phonon-assisted optical transitions extending out to around 800 nm, as shown in Fig. 2.3. These phonon-assisted transitions exist for absorption processes as well, allowing the defect to be excited with much shorter wavelengths than the resonant 637 nm. In many cases, widely-available 532 nm lasers are preferred for excitation, allowing full collection of the NV center photoluminescence while spectrally filtering the excitation laser. The excited state lifetime is 10-15 ns for most NV centers, which makes it a moderately bright center capable of emitting nearly 10^8 photons/s. However, most experimental setups which use microscopes to examine single NV centers will observe below 10^5 counts per second (cps) on their single-photon counters. A number of factors contribute to the disparity, including ionization and electron recapture cycles, non-radiative pathways, and low collection efficiency. Diamond has a high index of refraction of $n = 2.4$, which leads to the majority of subsurface emission being trapped in the sample due to total internal reflection. NV

centers tend to have high quantum yield³, measured above 90%. On the other hand, the ratio of ZPL emission to the total emission, known as the Debye-Waller factor, is quite low for the NV center, approximately 3%. This leads to the broad phonon sideband and can be a hindrance or a help, depending on the experimental goals. Collectively, these properties make it straightforward to observe single NV centers in diamond at room temperature using cost-efficient visible-spectrum microscopy equipment.

Unlike the ground state, the NV center excited state is an orbital doublet, with 3E symmetry. The transition is often modeled as two perpendicular electric dipoles aligned with the x and y axes of the NV coordinate system. Due to this dipolar configuration, the NV center emission angular distribution has only a mild isotropy along its axis, and it absorbs and emits in most linear polarizations. At room temperature, the ZPL is broadened by thermal phonons to a linewidth of a few nanometers (several terahertz, at this wavelength).

In addition to the electron spin, the nonzero nuclear spin of the nitrogen atom is a critical part of the NV center's spin properties. Nitrogen has two stable isotopes, ${}^{14}\text{N}$ and ${}^{15}\text{N}$, which have nuclear spin $I = 1$ and $I = 1/2$, respectively. The ${}^{14}\text{N}$ isotope is naturally 99.4% abundant, and so is predominant in natural diamonds. Despite the low natural abundance, the ${}^{15}\text{N}$ isotope is often used in ion implantation and CVD growth to distinguish created NVs from naturally occurring, and so its unique properties cannot be neglected. With the addition of the nuclear spin, the ground state Hamiltonian becomes

$$H_{gs} = D_{gs} \left[\hat{S}_z^2 - S(S+1)/3 \right] + \gamma_e \mathbf{B} \cdot \mathbf{S} + \mathbf{S} \cdot \mathbf{A} \cdot \mathbf{I} + Q \left[\hat{I}_z^2 - I(I+1)/3 \right]. \quad (2.2)$$

Here, \mathbf{A} is the hyperfine tensor, which is well approximated by a diagonal tensor with components $(A_{\parallel}, A_{\perp}) = 2\pi \times (-2.16 \text{ MHz}, -2.7 \text{ MHz})$ for ${}^{14}\text{N}$ and $2\pi \times (3.0 \text{ MHz}, 3.7 \text{ MHz})$ for ${}^{15}\text{N}$. Q is the nuclear quadrupolar interaction, which is present for ${}^{14}\text{N}$ only with a strength of $2\pi \times -4.95 \text{ MHz}$.

3. Quantum yield is often defined as the fraction of excited state population which decays radiatively instead of non-radiatively.

2.1.3 State Preparation, Manipulation, and Measurement

One of the most distinct aspects of the NV center is the interplay between its spin and optical degrees of freedom. The triplet excited states can decay via electron-phonon interactions to the singlet subspace, in a process known as intersystem crossing (ISC). The singlet states are themselves metastable, and decay back to the triplet ground states. This process is important because it is strongly spin selective: the $|\pm 1\rangle$ excited states undergo the ISC far more often than the $|0\rangle$ excited state. Two major consequences follow from this spin selectivity: ground state spin polarization and optical spin state readout.

Preferential coupling to the ISC provides the mechanism for a high degree of room temperature spin polarization. Initially, an NV center might start in a thermal ground state mixture, with no detectable polarization. An off-resonant laser is applied, exciting the entire ground state population. The $|0\rangle$ state decays radiatively, but some of the $|\pm 1\rangle$ population instead enters the ISC and the singlet states. The singlet states decay relatively evenly into all three triplet ground states, which leads to some residual radiative and non-radiative repopulation of the $|\pm 1\rangle$ ground states, but the net result is a transfer of population from $|\pm 1\rangle$ to $|0\rangle$. If the off-resonant laser is kept on, the cycle will repeat, leading to further polarization. The timescale of this process is set by the approximately 300 ns lifetime of the singlet state, and so after a few microseconds of laser excitation, the $|0\rangle$ ground state can be statistically polarized to over 90%. This simple optical polarization mechanism is the basis for nearly all NV center experiments.

The non-radiative ISC decay pathway also enables all-optical readout of the ground state spin state. In the absence of spin-selective processes, all three ground states would exhibit indistinguishable PL. The $|\pm 1\rangle$ states' increased coupling to the non-radiative ISC leads to lower PL, and thus photon counting can be used to differentiate $|\pm 1\rangle$ from $|0\rangle$. Technically, the singlet pathway also emits photons, but the transition lies in the infrared and can easily be filtered out of the visible frequency collection band. State measurement is accomplished by applying laser excitation and measuring PL. The PL contrast between the $|0\rangle$ and $|\pm 1\rangle$ states

can vary depending on a multitude of factors, including laser power, collection duration, and background filtering, but is optimally around 25-30%. Since the readout procedure is the same as the spin initialization procedure, the measurement gradually erases the spin state information and reinitializes it into $|0\rangle$. Thus, an important step in performing optical spin measurement of an NV center is calibrating the optimal collection window where spin state PL differences are maximized. PL contrast provides a means of determining the probability that an arbitrary state is in $|0\rangle$ versus $|\pm 1\rangle$. In other words, if a qubit is defined in the $\{|0\rangle, |-1\rangle\}$ basis, PL contrast determines $\langle S_z \rangle$ of a pure state, or the diagonal density matrix elements of a mixed state. While this measurement alone is not enough to fully determine a qubit state, the ability to apply coherent phase rotations enables successive measurements to determine $\langle S_x \rangle$ and $\langle S_y \rangle$. This approach provides full information on the density matrix of any qubit which can be mapped to the $|0\rangle$ and $|\pm 1\rangle$ states.

With a large magnetic dipole, the NV center ground state spin is typically manipulated with microwave fields. The zero-field splitting sets a natural frequency scale in the few gigahertz regime, comfortably in the range of microwave signal generators and common analog electronics components. These microwave drives can be delivered to the sample through a variety of stripline and waveguide geometries, ranging from simple straight wires across the sample surface to lithographically-defined coplanar waveguides. Different geometries can offer advantages in terms of frequency bandwidth, the ratio of drive strength to applied microwave power, and ease of fabrication. In all cases, the B_x or B_y component of the applied field (in the NV coordinate frame) drives corresponding spin rotations when it is in resonance with allowed NV spin transitions. This makes it possible to observe Rabi oscillations, demonstrating coherent quantum control and the ability to perform arbitrary rotations around the x and y axes.

The most common measurement for NV centers is optically detected magnetic resonance (ODMR). The optical readout of the NV center makes it easy to track the magnetic resonances which correspond to the ground-state level splittings. In the simplest form of ODMR,

an off-resonant readout laser continuously illuminates the NV while a CW microwave drive is stepped in frequency. The NV emits steady PL until the microwaves reach resonance with a ground-state transition. The increased population in the $|\pm 1\rangle$ states leads to a dip in PL which matches the line shape of the transition. A pulsed measurement, alternating between laser readout and microwave π pulses, improves the contrast of the measurement and is used when narrow frequency bandwidths are desired.

With calibration and compensation for pulse errors, microwave pulses have been demonstrated with single qubit error rates as low as 5×10^{-5} [24], among the best achieved for any qubit system. Following angular momentum selection rules, magnetic dipole transitions are allowed between the $|0\rangle$ and $|\pm 1\rangle$ states. While operations in the $\{|+1\rangle, |-1\rangle\}$ basis can be mediated by the $|0\rangle$ state, it is also possible to directly couple the states via electric field or strain interactions which are present, though more weakly allowed, in the Hamiltonian [25].

2.1.4 Interactions and Coherence Properties

While the ground state Hamiltonian of the NV center in Eq. 2.2 describes the majority of NV spin dynamics, there are additional contributions from fields and temperature dependences which play a role in shaping experiments and determining coherence times. After all, despite the NV's relative isolation in the diamond band gap, it is still affected by its host crystal, including distortions and vibrations in the lattice.

Strain and electric fields have similar effects on the level structure. Their combined influences on the NV ground state are described by [23]

$$\begin{aligned}
 H_{E, strain} = & d_{\parallel}(E_z + \delta_z) \left[\hat{S}_z^2 - S(S+1)/3 \right] + d_{\perp}(E_x + \delta_x) \left[\hat{S}_y^2 - \hat{S}_x^2 \right] \\
 & + d_{\perp}(E_y + \delta_y) \left[\hat{S}_x \hat{S}_y + \hat{S}_y \hat{S}_x \right].
 \end{aligned}
 \tag{2.3}$$

In this expression, \mathbf{E} and $\boldsymbol{\delta}$ are the static electric and strain fields, respectively, at the NV location. d_{\parallel} and d_{\perp} are the coupling parameters with corresponding values of 17 Hz/V·cm

and $0.35 \text{ Hz/V}\cdot\text{cm}$, where strain is measured in effective electric field units. These terms indicate that electric and strain fields aligned with the NV axis directly modify the zero-field splitting, while the transverse fields create couplings between the $|\pm 1\rangle$ states (anti-diagonal components of the Hamiltonian). The magnitudes of the coupling terms are small, and require fields of thousands of volts per centimeter to compare to a few milligauss of magnetic field. These small couplings follow from the orbital symmetry, which prevent a first-order electric dipole and require higher-order couplings to the excited state via spin-orbit and spin-spin terms.

Temperature has relatively minor effects on the ground state level structure. Temperature primarily affects the zero-field splitting D_{gs} , which has been fit to the Varshni equation or similar polynomial relationships [26]. D_{gs} decreases from $2\pi \times 2.877 \text{ GHz}$ at 0 K to $2\pi \times 2.870 \text{ GHz}$ at room temperature, where it locally varies with temperature at a rate of -74 kHz/K . This dependence allows for temperature sensing, described in section 2.2.3.

Temperature plays a much larger role in populating the phonon modes which lead to spin depolarization and decoherence. There are three phonon-mediated processes which cause transitions between the spin sublevels. The direct process couples the levels with a single phonon of appropriate symmetry, the Orbach process mediates spin flips through two-phonon excitation to a higher-lying state, and the Raman process induces spin flips via two-phonon excitation to a virtual state [27]. The Raman process is strongly temperature dependent, proportional to T^5 and T^7 for the NV center, and dominates the depolarization rate at room temperature. The NV center's T_1 is approximately 7 ms at room temperature, which can be viewed as short compared to some nuclear spin or trapped atom qubits, or exceedingly long relative to most electron qubits in solid state environments, particularly above cryogenic temperatures. This spin lifetime is largely consistent across different sample environments and in ensembles up until high densities, where spin-spin interactions finally exceed the thermal-phonon-driven depolarization. NV centers near surfaces (including in nanodiamonds [28]) or in proximity to sources of high magnetic noise show reduced T_1

times, but this effect can be leveraged as a sensing method (see section 2.2.3).

The NV center’s coherence times are often far shorter than the spin lifetime. A typical single NV center in a natural IIa diamond might have coherence time $T_2 \sim 500 \mu\text{s}$ and dephasing time $T_2^* \sim 2 \mu\text{s}$. This indicates that for high-quality single NVs, while there are relatively few processes which exchange energy with the NV, there are greater influences from pure dephasing sources. In the Hamiltonian, dephasing is induced by uncontrolled effects that alter a qubit’s level splitting, while coupling terms are responsible for reducing qubit lifetimes. Several dephasing culprits have already been described: temperature, electric field, and strain fluctuations all directly affect the zero-field splitting. These fluctuations can be external, arising from natural variations in laboratory conditions, or internal to the sample, such as the random redistribution of electrical charges residing on pervasive impurities and crystal defects. However, these interactions are relatively small influences compared to magnetic field fluctuations, which are the dominant source of decoherence for most NV centers.

In most samples, the majority of decohering magnetic field noise comes from ^{13}C spins in the diamond lattice. Pure diamond is entirely carbon, which has two stable isotopes: the spinless ^{12}C with 98.9% natural abundance, and the spin- $\frac{1}{2}$ ^{13}C which makes up the remaining 1.1% of carbon atoms. Even at this relatively low abundance⁴, there are likely to be thousands of nuclear spins within just a few nanometers of each NV center. The ^{13}C nucleus has a gyromagnetic ratio of $2\pi \times 1.07 \text{ kHz/G}$, and so the aggregate effect of the spin bath becomes significant due to the NV center’s strong interactions with magnetic fields. In natural abundance samples, the fluctuating fields of the bath spins can reach several hundred milligauss, setting a lower bound on the spin linewidth of a few hundred kilohertz and limiting the corresponding dephasing time to under ten microseconds. Flip-flop interactions between the electron and nuclear spins are suppressed by the energy differences and large

4. Compare against silicon, which has roughly 5% ^{29}Si nuclear spin density, or III-V materials with 100% nonzero nuclear spins.

zero-field splitting, but dephasing alone is significant enough to pose the primary constraint on coherence time. Because of the relatively slow quantum jumps of each individual nuclear spin, dynamical decoupling techniques are effective, leading to the orders of magnitude improvement from T_2^* to T_2 . It is unsurprising, then, that a large fraction of NV center research has focused on strategies to mitigate the decohering influence of the nuclear spin bath. While there are a wide range of control techniques that reduce the effect of the nuclear spins, including dynamical decoupling, decoherence-free subspaces, and bath driving, the most robust approach is to remove the ^{13}C spins altogether via isotopically purified growth.

Isotopic purity is a unique and potent advantage of CVD growth over natural diamonds. While ^{13}C atoms are unavoidable for natural diamonds, CVD growth presents an opportunity to grow new crystal with reduced abundance of ^{13}C spins. CVD-grown diamond reflects the isotope distribution of the precursor methane gas, so in principle diamond can be grown with zero nuclear spins. Isotopic purities have reached 99.998%, driving down the nuclear spin density by two orders of magnitude [29]. The resulting coherence times are impressive, showing T_2^* exceeding $30\ \mu\text{s}$ for ensembles [30] and $100\ \mu\text{s}$ for single NVs [31]. With decoupling, the coherence time can be extended beyond a millisecond, approaching the room temperature T_1 time. These drastic improvements in coherence time show the current and future importance of high-quality isotopically purified diamond for NV center experiments, where long spin coherence is possible without any sophisticated decoupling.

In the limit of low nuclear spin density, weaker interactions become the dominant dephasing mechanism. Ongoing research continues to drive down the various sources of decoherence, with indications that defect charge stability is an important factor. By doping diamond with phosphorous, dephasing times even reach $1.5\ \text{ms}$ [29], the longest observed in room-temperature solid state spins. As sample engineering advances, one goal will be to push the coherence time to the ultimate limit imposed by the spin lifetime, $T_2 \leq 2T_1$ [27].

2.1.5 Low-Temperature Properties

Remarkably, the properties described thus far have been for the NV center at room temperature, and these generally improve at cryogenic temperatures. The excited state level structure qualitatively changes below ~ 20 K, and so the cryogenic regime for NV center experiments entails liquid helium cryostats; sub-kelvin temperatures are not required to investigate most low-temperature NV center phenomena.

The NV's ground state physics are largely unchanged at low temperatures. The zero-field splitting becomes insensitive to temperature, although the electric field and strain interactions remain. However, freezing out the thermal phonon bath increases the spin coherence times significantly. Phonon-mediated spin relaxation becomes a far slower process than flip-flops with bath spins [32], and so T_1 is no longer a limiting factor. Without phonon-induced depolarization, high-order dynamical decoupling can extend the coherence time to over a second even in diamonds with natural abundance of ^{13}C [33].

Thermal phonons play a major role in shaping the room-temperature excited state, and its level structure can be fully resolved at low temperatures once the phonon bath is depleted. As a spin triplet, orbital doublet, the excited state is comprised of six levels. At elevated temperatures, thermal phonons average out the two orbitals, leaving an effective triplet structure similar to the ground state [34]. At low temperatures, all six levels can be resolved. Due to spin-orbit coupling, these levels are no longer eigenstates of \hat{S}_z , and are instead labeled according to their respective symmetry: $\{A_2, A_1, E_x, E_y, E_1, E_2\}$ [35]. In the absence of electric fields or strain, $|E_x\rangle$ and $|E_y\rangle$ are degenerate, as are $|E_1\rangle$ and $|E_2\rangle$. With the application of transverse fields, the two orbitals split in energy, yielding upper (3E_x) and lower (3E_y) branches. The three lower states exhibit several anticrossings, causing optical spin flips when the transitions are excited. The three upper states are relatively unmixed and serve as more efficient cycling transitions. Axial electric field or strain does not change the relative level structure of the excited state, but does provide a Stark shift of the excited state relative to the ground state – i.e., the optical transition frequency.

Given the gigahertz scale of the spin-orbit and spin-spin interactions of the excited state, the optical linewidths must be far smaller than at room temperature to optically resolve the individual levels. This proves to be the case, as under cryogenic conditions the linewidths approach the lifetime limit of ~ 13 MHz. Saturating this lower bound is not simply a matter of cooling the sample, however. As a result of the electric-field-induced Stark shifts, the optical transition frequency will jump if electric charges near the NV are suddenly redistributed. Crystal defects, such as nitrogen impurities or vacancy clusters, host weakly-bound electrons that can be ionized when laser pulses are applied. Thus, when excitation lasers are applied to the NV center, they may inadvertently mobilize a bath of electric charges that quickly settle into local potential minima. This process causes spectral jumps and is termed “spectral diffusion.” The magnitude of spectral diffusion depends on each NV’s local environment, and can range from tens of megahertz for bulk NVs to over a gigahertz for NVs near surfaces or nanostructures, where loosely-bound electric charges are plentiful. Reducing the effect of spectral diffusion is an active area of research, in service of applications where stable, narrow optical lines are important, such as quantum communication (section 2.2.2).

Along with the narrowing optical lines, the selection rules for resonant optical coupling become nontrivial at low temperatures. The $|A_2\rangle$, $|A_1\rangle$, $|E_1\rangle$, and $|E_2\rangle$ states have $m_s = \pm 1$ character and can couple to the $|\pm 1\rangle$ ground states via absorbing or emitting circularly polarized photons. The $|E_x\rangle$ and $|E_y\rangle$ states are $m_s = 0$ states and couple to the $|0\rangle$ ground state through linearly polarized photons. However, as noted above, transverse fields mix the excited state spin character and allow optically driven spin flips to occur uncontrolled [35]. The joint $m_s = \pm 1$ character of states such as $|A_2\rangle$ does create spin-photon entanglement following spontaneous emission, which could be used in schemes to remotely distribute spin-based entanglement, and creates a nearly closed Λ system for optical control of the ground state spin (see chapter 3).

For low-strain, low-spectral-diffusion NV centers, the relatively pure optical transitions enable enhanced optical spin initialization and readout methods. A desired ground state can

be initialized with high fidelity by resonantly pumping optical transitions from the other two ground states. Within a few microseconds, population in the unwanted ground states decays via the ISC or optical spin flip to the desired state, exceeding 99% state preparation fidelity [36]. Since the resonant excitation is spin-selective, this technique avoids repopulating the unwanted states. Optical initialization works most efficiently for the $|0\rangle$ state, due to the stronger ISC coupling of the $m_s = \pm 1$ excited states, but can be used to directly initialize any of the three ground states.

Similarly, higher readout fidelity is possible when resonant, spin-selective optical transitions are pumped instead of off-resonant excitation. In this method, the $|E_x\rangle$ state is nearly always chosen as the optical transition, since it is the best cycling transition. By driving only this transition with a narrowline laser, near unity contrast between the $|0\rangle$ and $|\pm 1\rangle$ states is possible. As $|E_x\rangle$ couples only weakly to the ISC or $|\pm 1\rangle$ states, photons can be collected until the spin finally flips or the defect ionizes. With sufficiently high collection efficiency, resonant readout can even enable single-shot readout, where the spin state is determined in a single, non-destructive measurement [36]. Single-shot readout is the prerequisite for a plethora of fundamental quantum physics experiments, as well as executing spin-photon entanglement protocols to distribute entanglement remotely. These and other applications will be covered in more detail in the next section.

2.2 NV Center Applications

The NV center exhibits a unique set of properties among the myriad qubit systems that are subjects of research today. No other system combines a solid-state host, optical readout, and long room-temperature coherence. While it has its own limitations, such as challenges around deterministic placement and the difficulty of entangling multiple NV centers together, a thriving set of applications has emerged from leveraging its strengths. These include as small quantum simulators, for quantum sensors of a broad set of physical parameters, and

as the basis for distributing remote, spin-spin entanglement over quantum networks.

2.2.1 Quantum Simulation

In many ways, the NV center is a model quantum system. It features a nearly ideal set of triplet states: the level structure is highly tunable, from the MHz to the GHz regime, using applied fields. In addition to long coherence times, NV centers exhibit rapid initialization and spin manipulation, allowing tens of thousands of unitary operations to be applied before the spin decoheres. A wide range of control is possible, originating from the combination of microwave signal generator IQ modulation and rf striplines created via a full suite of lithography techniques. The flexibility of the qubit system makes it a popular choice for modeling single-qubit or few-spin dynamics.

The NV center is capable of implementing unitary evolution of any single-qubit Hamiltonian. The Hamiltonian of any two-level quantum system can be regarded as a pseudospin- $\frac{1}{2}$ system in an effective magnetic field [37]. This statement can be reversed: an actual spin in a real magnetic field can model the dynamics of any two-level system. As a result, the NV center is useful for exploring such diverse phenomena as topology in Hopf insulators [38] and the dynamics of qubits beyond the rotating wave approximation [39]. Demonstrations have even extended beyond Hermitian evolution, encompassing PT-symmetric evolution [40] and exceptional points [41] that result from dissipation in the system. Chapter 3 presents several experiments which use the NV center as a model three-level system, demonstrating superadiabatic spin dynamics and geometric unitary operations.

The NV center is also an excellent candidate to test central-spin models of quantum systems, due to its optically accessible electron spin surrounded by a bath of dark nuclear and electron spins. This model is one of the most-studied theoretical frameworks for decoherence, and theoretical study works in a virtuous cycle with the NV center: NV experiments validate theoretical models for decoherence, which in turn are refined and generate approaches to mitigate decoherence in further NV experiments. Sample engineering presents a fruitful

opportunity to test models across a wide range of parameters space. Diamond samples can be doped with varying concentrations of several species of bath spins, including nitrogen donors and ^{13}C spins, that call for different theoretical treatments. Closely related is the development of sophisticated dynamical decoupling sequences, which are relatively easy to implement given the flexibility of NV center spin control. Numerous studies have explored the relationship of the NV to its spin bath [42], the efficacy of dynamical decoupling, or both [43].

Indisputably, the NV center is more challenging to arrange into large qubit networks than some other popular qubit systems, but NV centers have still be used to demonstrate important multi-qubit protocols. While coupling separate NV centers is difficult, each NV may naturally have access to multiple long-lived nuclear spins via hyperfine interactions. The intrinsic nitrogen nuclear spin provides a dependable two-qubit register, and strong coupling to nearby ^{13}C spins can provide a number of additional, well-resolved qubits. These accessible and readily controllable qubit registers have enabled early demonstrations of the building blocks of quantum error correction [44], quantum algorithms [45], and discrete time crystals [46], among many others.

The use of NV centers as quantum simulators may wane as other platforms improve their coherence times and scale to larger qubit registers. Undoubtedly, it will continue to be useful demonstration system, but few-qubit simulations are easy for classical computers and research will push toward larger systems with classically intractable dynamics. The breadth of quantum simulation applications for NV centers may thus depend on the success of quantum networking efforts for NV centers. Remote entanglement efforts for quantum communication applications could dovetail with larger NV-based quantum simulators. There are still efforts to develop on-chip NV center networks at room temperature [47], and while conceptual and engineering hurdles stand in the way, the payoff of a successful room-temperature, solid-state quantum network would be tremendous for all prospective quantum applications.

2.2.2 Quantum Communication

Fortunately, the NV center is also leading candidate for quantum communication and quantum networking. Protocols for creating distributed quantum networks require traveling qubits (all but mandating the use of optical photons) to interface with long-lived memory qubits at repeater or end nodes. Most qubit systems lack one or both of these characteristics, but both aspects are natural for NV centers – spin-dependent photoemission ferries entanglement between spatially separated spins, and entangled electron spin states can be stored in proximal nuclear spin memories. In recent years, NV-based entanglement testbeds have achieved major milestones in developing quantum networks and repeaters. NV center spins have been entangled via intermediary photon interference [48] in two- and three-node networks [49] at laboratory scales⁵. These demonstrations represent rapid progress toward metropolitan-scale quantum networks, but even using the natural suitability of the NV center for remote entanglement, a number of physics and engineering challenges must be solved to increase the distance, rate, and fidelity of the transmitted entanglement.

Quantum repeaters will be essential for distributing entanglement beyond ~ 100 km through fiber optics. As a single photon’s state cannot be copied or amplified, the exponentially increasing losses with distance in an optical fiber eventually make direct photon transmission unfeasible. A quantum repeater (designated point C) would operate between a sender (A) and receiver (B), and create entanglement separately between each endpoint. The repeater would then swap the entanglement from A-C and B-C to A-B via unitary operations, entangling the endpoints while only having to create entanglement across half the distance. Nested sets of repeaters could bring an arbitrary distance down to a manageable one. In addition, repeaters will take multiple copies of imperfect entanglement and combine them to generate a higher-fidelity entangled state, in a process called “entanglement distil-

5. Along the way, researchers have also used remote entanglement of NV centers to demonstrate a loophole-free violation of Bell’s inequality, thus proving that local, entanglement-free classical theories of physics are inadequate to describe the physical world [50].

lation.” Entanglement distillation [16] and entanglement swapping [49] have now both been demonstrated using NV centers, as key steps in a fully functioning quantum repeater. The resulting entanglement fidelities are high enough to demonstrate violations of Bell’s inequality, but must still be improved before they are viable as quantum network nodes. To that end, current research seeks to improve on the photon-collection interface for single emitters, which limits the rate of entanglement generation and the fidelity of the heralding process.

While there are several protocols to create remote NV-NV entanglement, a typical example (drawn from Ref. [49]) will illustrate both how most protocols function and how photon collection efficiency assists with entanglement generation. Two NV centers in separate chips are arranged so that their collection paths are the two inputs to a beamsplitter, with a single-photon counter at each output. The samples are cooled to low temperature to address the individual optical transitions, and a readout laser is tuned to the $|0\rangle \leftrightarrow |E_x\rangle$ transition of each NV. Both ground states are prepared in the superposition

$$|\alpha\rangle = \sqrt{\alpha}|0\rangle + \sqrt{1-\alpha}|-1\rangle,$$

and a readout pulse is applied to each NV, which ideally triggers a photon if the spin is in the $|0\rangle$ state. This creates entanglement between each spin and the presence or absence of a photon. The photon modes interfere on the beamsplitter, erasing the which-path information, so if a single detector click is measured, the spins are likely to be in the joint state

$$|\psi^\pm\rangle = \frac{1}{\sqrt{2}} \left(|0, -1\rangle \pm e^{i\theta} |-1, 0\rangle \right).$$

Here, θ is the phase difference of the two optical paths, and the \pm sign corresponds to which of the two detectors clicked. This protocol creates entangled states with a maximum fidelity of $1 - \alpha$, due to the possibility of creating the $|0, 0\rangle$ state, which also triggers one detector click. The entangled states are generated at a rate $\propto \alpha p$, where p is the probability of collecting a photon emitted by an NV. Thus, if photon collection efficiency is low, α may

need to be increased to create an appreciable entanglement rate, at the cost of lowered fidelity. Improvements to photon collection efficiency benefit both fidelity and entanglement rate in this and other entanglement protocols.

One of the more popular approaches to improve photon collection and readout fidelity is through a photonic cavity. A photonic cavity consists of two reflective interfaces which create constructive and destructive interference for interior photonic modes. The cavity is constructed to strongly enhance a target mode or modes, and a point emitter can be positioned to couple to the resonant modes. In the regime of weak coupling between cavity and emitter, an overlap between the emitter's spectrum and the cavity's resonant modes produces a *Purcell enhancement* of the emitter. This increases the number of emitted photons in the resonant modes and correspondingly decreases the optical lifetime. For an isotropic, spectrally broad emitter like the NV center, Purcell enhancement of the ZPL has two beneficial consequences: it increases the fraction of NV center emission in the ZPL relative to the phonon sideband, resulting in more photons useful for entanglement; and it directs more of the emission into spatial modes in the collection path, as opposed to being trapped in the diamond sample, so fewer photons are lost. Photonic cavities thus purport to address two of the major challenges in NV-NV remote entanglement. However, fabricating such cavities in diamond is difficult due to the limited repertoire of nanoscale fabrication techniques which are effective in the material. With persistence, photonic crystals have successfully been created in diamond membranes [51], but this leads to two integration challenges: positioning an NV center at the point of maximal coupling to the cavity and ensuring its spin and optical properties remain stable. NV centers are not yet deterministically formed with nanometer precision, so the cavities must either be formed around pre-selected defects or a large number of devices must be created in the hopes of probabilistic success [52]. Due to its permanent excited-state electric dipole, NV centers are often optically unstable when in proximity to diamond surfaces. This restricts the space of suitable designs, but some encouraging progress has been made by fabricating external cavities around diamond membranes [53]. Ongoing

research continues to pursue an efficient photonic interface for single NV centers with reliable fabrication yields.

Frequency conversion is a final factor in the development of NV-based quantum networks. Modern fiber optics have the best transmission around 1550 nm, but coherent NV photons are emitted at 637 nm, where attenuation is much higher. In the interest of squeezing the most out of every photon, it may be more efficient to convert NV emission into the telecom band, through a process called spontaneous four-wave mixing. This technique involves superimposing the ZPL photons on an infrared pump laser in a nonlinear crystal, where the beams are converted into sum and difference outputs. The difference of the photon energies is chosen to be at the desired telecom wavelength. If the conversion efficiency between 637 nm and ~ 1550 nm wavelengths can be improved sufficiently, better entanglement generation rates can be achieved by frequency conversion than by transmitting visible photons directly.

While NV centers are well-suited for quantum communication applications in general, they may also develop into an excellent platform for distributed sensing. The combination of remote entanglement suitability and sensing prowess is a rare one, but the NV center boasts both. However, it will remain to be seen if cryogenic temperatures, which are sub-optimal for many sensing applications, will always be required for NV-NV remote entanglement, or if a reliable room-temperature protocol can be developed.

2.2.3 Quantum Sensing

At the present moment, the NV center's primary application is in quantum sensing. Once again, its properties are favorable to the task at hand: long coherence times that lead to high sensitivities, atomic scale which facilitates high resolution measurements, and room temperature operation that allows it to be deployed on the widest possible range of sensing scenarios. Since sensing applications include addressing research problems on the laboratory scale, it has been relatively easy to find use cases for NV-based sensors. Many sensing demonstrations have been published, covering a broad swath of experimental parameters and

sensor designs, and many patents have emerged from the mix. Small companies are already producing NV-based research instruments, and while it remains to be seen if any products will see commercial success beyond the research community, several good candidates exist.

The NV sensing paradigm is generally based around measuring shifts in two or more of the ground state levels. A typical example might set a baseline qubit frequency and then monitor its shifts up or down to track an external variable of interest. This approach is valid for a wide range of parameters, including electric field, strain, temperature, and magnetic fields, since they all enter into the NV Hamiltonian directly. First, a reference point is taken, with the NV levels under initial environmental conditions. Then, any changes in level splitting can be correlated to changes in the environment. For example, at room temperature, the zero-field splitting changes at a rate of -74 kHz/K. A shift of -350 kHz in the $|0\rangle \leftrightarrow |-1\rangle$ transition thus indicates a 4.7 K increase in temperature.

It can be a downside that so many parameters affect the NV levels, as it creates potentially confounding influences. Electric field and strain have very similar effects, but fortunately are rarely varying simultaneously. When examining the $\{|0\rangle, |-1\rangle\}$ qubit, it may not be obvious if the qubit energy is changing due to magnetic field, temperature, or a mix of parameters. In this case, knowing both the $|0\rangle \leftrightarrow |-1\rangle$ and $|0\rangle \leftrightarrow |+1\rangle$ transition energies allows the different influences to be separated.

While the NV can potentially measure anything that influences its Hamiltonian, it is most often used for magnetic fields. Under most conditions, the NV center is far more sensitive to magnetic phenomena than any other effects – in fact, it is among the most sensitive magnetometers yet designed [54]. Because the B_z and $B_{x,y}$ field components create different NV level shifts, some magnetometer designs take advantage of the four possible NV orientations to perform full vector magnetometry, measuring up to eight distinct transitions that allow the vector magnetic field to be reconstructed [55]. Of course, this requires NV ensembles, and opens the possibility of performing wide-field imaging. In this mode, due to the high PL of an NV ensemble, the typical single-photon counter can be replaced with

a CCD camera. In addition, instead of focusing the laser excitation to a single point, a high-power beam is expanded to illuminate the whole sample. This arrangement allows NV centers across the entire sample to be measured simultaneously, providing a snapshot of an entire magnetic field distribution. This concept is in development to produce what should be, at the very least, a powerful research instrument: the quantum diamond microscope [56].

To compare different approaches to sensing, a standard metric is the sensor’s sensitivity. Sensitivities are expressed in units of the parameter per $\sqrt{\text{Hz}}$ (e.g. $\text{T}/\sqrt{\text{Hz}}$), and are often interpreted as the smallest signal which can be measured in 1 s. When comparing two sensors, a smaller sensitivity indicates a better sensor, with the ability to resolve an equivalent signal in a shorter amount of time. The $\sqrt{\text{Hz}}$ factor also gives a correct intuition for how measurements improve with time – in order to improve the signal-to-noise ratio of a measurement by a factor of two, the measurement must be integrated four times as long⁶. One of the key advantages of quantum-enhanced sensors (ones which leverage entanglement or other uniquely quantum effects) is that they may be able to surpass this usual limit, reaching a scaling of $1/T$ instead of $1/\sqrt{T}$ [21].

Sensitivities of quantum devices have an important relationship with the coherence time of their constituent levels or qubit states. Of course, if a signal is proportional to the coherence of a qubit, no signal will be acquired if the qubit decoheres due to a long measurement protocol. On the other hand, increasing the measurement time benefits the sensitivity. There is an optimal balance point between coherence and sensitivity, which corresponds to measuring the qubit at approximately its coherence time, T_2^* . When measured at this point, sensitivity is $\propto 1/\sqrt{T_2^*}$ [21]. Thus, an important dimension of improving quantum sensor performance is extending the coherence time of the sensor qubits or energy levels.

ODMR provides a straightforward means to measure dc or quasi-static influences on the NV, but more sophisticated protocols enable the detection of ac fields all the way up to the

6. Equivalently, four successive measurements of the original duration can be averaged together. This reflects the assumption of a time-invariant measurement, where measurement precision scales as $1/\sqrt{NT}$ for N measurements that each take T time.

gigahertz regime [21]. In a Ramsey measurement, a resonant $\pi/2$ pulse on, for example, the $|0\rangle \leftrightarrow |-1\rangle$ transition initializes the spin state into a superposition. Next, a signal which creates a static detuning or a slow-varying modulation is applied to the qubit, causing its spin state to precess on the Bloch sphere equator. After some time T , the spin state is measured, and the phase it acquired is interpreted based on the sensing mode. For example, if a B_z field was applied, the phase is rescaled by the gyromagnetic ratio of $\gamma_e = 2\pi \times 2.8$ MHz/G; if it was an E_z field, by the coupling of $d_{\parallel} = 2\pi \times 0.35$ kHz/(V/cm). Ramsey measurements are limited by the dephasing time T_2^* , similar to the ODMR limit, which has a maximum resolution equal to the transition linewidth of $\Gamma \propto 1/T_2^*$. Ramsey measurements can detect multiple subensembles with differing static detunings, however, such as the equally populated nitrogen nuclear spin hyperfine sublevels that are each detuned by 2 MHz. While they are optimal for static detunings, Ramsey measurements can resolve ac effects up to $3\pi/T_2^*$, or $\lesssim 1$ MHz for most NV centers [57].

At high frequencies, an oscillating signal will tend to average to zero during a dc measurement. Fortunately, spin echo techniques effectively rectify a synchronized signal, enabling more sensitive measurement approaches. The simplest spin echo sequence places a microwave π pulse in the center of a Ramsey measurement, simultaneously “refocusing” any dephasing on the spin superposition and inverting its evolution relative to an ac signal. As an example, consider a sinusoidal signal of frequency $2\pi/T$, such that it completes one full oscillation during a Ramsey measurement. The integrated signal is zero, but if a spin echo occurs at $t = T/2$, the second half of the signal adds instead of subtracting from the accumulated phase. To detect higher frequencies, more spin echoes can be added over the same duration T such that the pulse spacing $\tau = \pi/\omega$. The upper bound of this technique is set by how short the microwave π pulses can be made, and is currently ~ 100 MHz for NV centers. Chapter 4 presents a demonstration of adapting this technique to detect ac photocurrents in 2D materials.

To measure signals in the gigahertz frequency range, the measurement technique actually

becomes more simple, not more complicated. Suppose the signal of interest is magnetic noise emanating from a ferromagnet. A qubit formed from two NV levels is tuned (via static external field) to the target frequency. The noise spectral density is known to be directly related to the transition rate between resonant, magnetically coupled energy levels like in the NV [58]. Since transitions out of the $|0\rangle$ state dim the NV's PL, a measurement of the relevant noise can be made just by applying constant laser excitation and monitoring the brightness relative to the absence of magnetic noise. This technique is known as T_1 relaxometry, and provides a simple method to investigate noise without microwave control.

Before concluding this overview of the NV center, it is worth highlighting a few intriguing applications of NV sensing presently under study. These examples illustrate the breadth of NV center sensing utility and how the properties of diamond-based single spins can be leveraged for creative measurement techniques.

- NV centers have been implanted into diamond anvil cells, which are used for applying pressures beyond 10^{11} Pa to samples. Remarkably, the NV centers still operate up to these very high pressures, and can be used to make localized magnetic measurements of samples under pressure [59]. This is otherwise impossible, given the extreme on any probes in the volume and the bulky surrounding apparatus. This research thread represents the convenient integration of nitrogen-vacancy centers with existing diamond-based devices.
- Fiber magnetometry seeks to develop NV centers into a miniaturized, portable probe by leveraging the NV's optical readout properties. Nanodiamonds containing NV centers are affixed to the end of an optical fiber, which is used both to deliver laser excitation and collect PL [60].
- NV centers can be incorporated into custom-fabricated, diamond scanning probe tips. These tips, with widths of a few tens of nanometers, contain single NV centers that can be rastered across a sample surface using the scanning probe piezoelectric con-

trollers [61]. The close proximity of the NV and the nanoscale precision of scan-probe positioners allow for magnetic mapping of samples with nanometer resolution. This application best leverages the atomic scale of the NV center to outperform comparable magnetometry techniques.

- Sensing applied to biological samples (“biosensing”) is an entire application field unto itself. Once again, the NV’s unique properties open avenues for novel research. Made entirely of carbon, diamond is non-toxic to organisms and chemically inert. The NV center’s optical interface and room temperature coherence are well-matched for biological sample conditions. A primary goal is developing nanodiamonds containing stable, long-coherence NVs for use measuring at the cellular level [62]. Already, demonstrations have shown that temperature can be measured inside a living cell [63]. Alternatively, nanodiamond surfaces could be functionalized and used as biomarkers with excellent photostability [64].
- Taking a slightly different tack, NVs can be used indirectly in NMR measurements as a source of polarization. Whereas most NMR measurements must rely on weak statistical polarization of target nuclear species, even at very high magnetic fields, the NV center can be rapidly polarized with laser excitation, and this polarization transferred to target spin populations to enhance contrast. While approaches to distribute this enhanced polarization throughout a macroscopic target sample are still in development [65], proof-of-principle demonstrations have shown orders of magnitude improvement in ^{13}C NMR signal relative to its normal signature in diamond [66].

The measurement protocols and NV center applications presented here are merely an overview of the rich world of NV quantum sensing. As a result of its exceptional stability and long coherence under a broad set of experimental conditions, its sensitivity to many physical quantities of interest, and the convenience of optical initialization and readout, the NV center has found plentiful opportunities to advance cutting-edge research. These applications will

only become more prevalent as material quality improves, as diamond fabrication techniques mature, and as quantum sensing protocols develop. In the future, the NV center may be a leading candidate to demonstrate some of the more exotic quantum sensing phenomena in a solid state probe, including squeezed-state sensing and distributed, entangled sensors. These breakthrough applications would only serve to accelerate the adoption of NV center sensors in many corners of the quantum research field and bolster their prospects for commercial deployment.

Chapter 3

All-Optical Single-Qubit Control

This chapter and figures are adapted from work published in references [1] and [2].

On the whole, microwave control of the NV center spin is more widespread than optical control. Microwaves are more simple to generate and apply evenly to a sample, the NV responds similarly at room and low temperatures, and the energy level splittings are high enough to make the rotating wave approximation (RWA) very good under most conditions. However, optical control can provide a number of advantages to low-temperature NV configurations. The optical spot localization is far superior, due to optical wavelengths under one micron and the simplicity of focusing optics, compared to microwave wavelengths greater than centimeters. This is advantageous when trying to independently address multiple defects in the same diamond sample – microwaves will produce far more crosstalk. In addition, the strong optical dipole and diffraction-limited focusing allow optical control to perform spin operations with far lower power input, which can be especially advantageous at low temperatures if thermal load to the sample is a concern. Furthermore, the greater energy difference of the optical regime has several beneficial effects. While the upper limits of microwave driving do cause significant deviations from the RWA [39], the $> 10^{11}$ Hz optical transitions make the RWA as exact as can reasonably be achieved in experimental settings. This enables driving with rates limited only by the control electronics, not by

optical bandwidths or RWA constraints.

This chapter presents two developments in the state of the art in optical control of the NV center spin. In section 3.2, pulse modulation of an adiabatic state transfer technique is used to coherently exchange population between the $|\pm 1\rangle$ ground states of a single NV center. In section 3.3, control over nanosecond-scale resonant optical pulses allows a complete set of geometric single-qubit unitary operations to be implemented. Both of these demonstrations utilize the same low-temperature lambda (Λ) system of the NV center.

3.1 Lambda Systems

A lambda (Λ) system refers to a generic three-level quantum system in which two ground states are both coupled to the same excited state, and without a direct coupling between them. This level arrangement is widely found in physical systems ranging from trapped atoms and single molecules to superconducting qubits and optomechanics. Λ systems enable a well-studied set of interesting phenomena that are not available to two-level systems, including electromagnetically-induced transparency, coherent population trapping, and adiabatic dark state evolution (the STIRAP protocol, section 3.1.1). Implementing Λ systems in the NV center has previously led to demonstrations of geometric phases [67] and coherent population trapping [68], among others.

As described in section 2.1.5, there are multiple combinations of NV ground and excited states that can be treated as a Λ system. Depending on the applied electric, strain, and magnetic fields, many of the excited states couple to two or even all three of the ground states, to varying degrees. In ref. [68], the excited state was magnetically tuned to an anticrossing of the $|E_y\rangle$ and $|E_1\rangle$ states. At this anticrossing, either excited state couples to the $|0\rangle$ and $|-1\rangle$ ground states, which form the Λ system used to demonstrate coherent population trapping.

For all the experiments to follow, the Λ system is constructed from the $|A_2\rangle$ excited

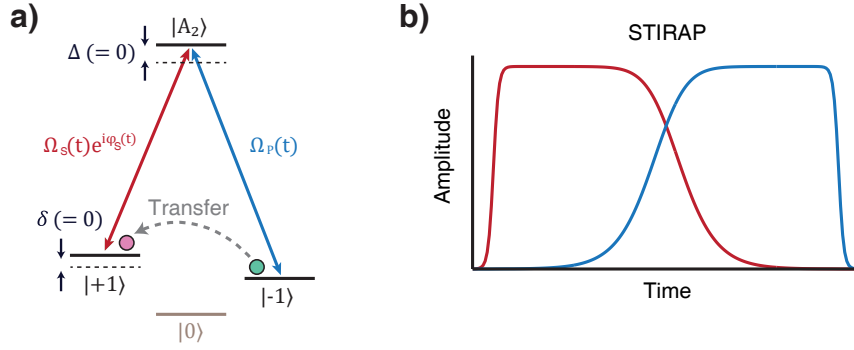


Figure 3.1: A low-temperature Λ system of the NV center. (a) The $|\pm 1\rangle$ ground states couple through strong optical transitions to the $|A_2\rangle$ state, forming the desired Λ system. The direct transition $|-1\rangle \leftrightarrow |+1\rangle$ is magnetically forbidden. At low temperature, the $|A_2\rangle$ level has a weak rate of nonradiative intersystem crossing transitions, and at low strain and applied fields, the decay rate to the $|0\rangle$ state is much less than the $|\pm 1\rangle$ states. Under these conditions, the Λ system is nearly ideal. (b) The STIRAP protocol, which coherently transfers a quantum state between two levels in a Λ system, involves adiabatically turning on the laser addressing one transition (blue) while turning off the laser addressing the other (red).

state and the $|\pm 1\rangle$ ground states. Unlike the $|E_{1,2}\rangle$ states, the $|A_2\rangle$ state has relatively weak coupling to the ISC despite its $|\pm 1\rangle$ character [69]. It also has a relatively weak decay pathway to the $|0\rangle$ state, which makes this subspace a nearly ideal Λ system. The level structure is shown in Fig. 3.1a.

3.1.1 STIRAP

One of the most widely-used Λ -system techniques is stimulated Raman adiabatic passage, or STIRAP [70]. This technique mediates population transfer between the two ground states without occupying the intermediate excited state. STIRAP achieves this through a counterintuitive ordering of two driving fields. Suppose the system is initialized in the $|-1\rangle$ state of the NV Λ system (Fig. 3.1a). The field which couples the initial and excited states is referred to as the pump pulse, with intensity Ω_P , and the field coupling the final and excited states is the Stokes pulse, with intensity Ω_S . In a dynamic state transfer, a π pulse with Ω_P

would excite the population to the intermediate level, and a second π pulse with Ω_S would complete the transfer. In STIRAP, the Ω_S drive is turned on first, despite there being no population in the destination state. Ω_S is slowly extinguished while Ω_P is ramped on, until the population transfer is fully complete (Fig. 3.1b). The key property of STIRAP is that the system will never occupy the intermediate state during this transfer. This property is advantageous if the intermediate state should be avoided, such as if it is dissipative in some manner. While this transfer protocol sounds unusual, it is successful due to the evolution of the dark state in the Hamiltonian.

A “dark state” of a Hamiltonian is a useful property which is also responsible for coherent population trapping. In the rotating frame of the $\{|-1\rangle, |+1\rangle, |A_2\rangle\}$ basis, the interaction Hamiltonian is [70]

$$H(t) = \frac{1}{2} \begin{pmatrix} 0 & 0 & \Omega_P(t) \\ 0 & 2\delta & \Omega_S(t)e^{i\phi_S} \\ \Omega_P(t) & \Omega_S(t)e^{-i\phi_S} & 2\Delta \end{pmatrix} \quad (3.1)$$

In this expression, Δ and δ are the one- and two-photon detunings (indicated graphically in Fig. 3.1), while ϕ_S is the phase offset between the pump and Stokes drive fields. This Hamiltonian has three eigenstates, two of which include components of $|A_2\rangle$ – in other words, they are “bright” states, because $|A_2\rangle$ emits PL. The third eigenstate is

$$\begin{aligned} |D(t)\rangle &= \cos \vartheta(t) |-1\rangle - \sin \vartheta(t) |+1\rangle, \\ \tan \vartheta(t) &= \frac{\Omega_P(t)}{\Omega_S(t)}. \end{aligned} \quad (3.2)$$

$|D\rangle$ is referred to as the “dark” state, since it has no overlap with the excited state. When expressed in terms of the two drive fields, the purpose of the pulse ordering becomes clear. Initially, with $\Omega_P = 0$, the dark state is the same as the initial state, $|-1\rangle$. At the end, with Ω_P reaching a maximum and $\Omega_S \rightarrow 0$, the dark state is the target state, $|+1\rangle$. In

between, if the drives are changed smoothly, there exists a continuous path in the Hilbert space connecting $| -1 \rangle$ and $| +1 \rangle$. According to the adiabatic theorem, if the parameters are changed slowly enough, the system will stay in the initial eigenstate, even as the character of that eigenstate evolves from $| -1 \rangle$ to $| +1 \rangle$. In quantitative terms, the condition for adiabatic evolution is

$$\Omega(t) \gg \left| \dot{\vartheta}(t) \right| = \frac{\left| \Omega_S(t)\dot{\Omega}_P(t) - \Omega_P(t)\dot{\Omega}_S(t) \right|}{\Omega(t)^2}, \quad (3.3)$$

where $\Omega(t) = \sqrt{\Omega_P(t)^2 + \Omega_S(t)^2}$.

If the adiabatic condition is approximately satisfied, a system which starts in the dark state will end in the dark state, and the parameters must only be chosen to make the initial state a dark state. This achieves the goal of never entering the intermediate state, and avoids the losses which would ensue. In the NV center Λ system, if $| A_2 \rangle$ is never populated, losses due to spontaneous emission will not occur. Radiative loss is one of the principal nonidealities in NV Λ system experiments. Section 3.2 describes experiments designed to avoid it, and it is the primary source of error for the experiments in section 3.3.

3.2 Superadiabatic Transitionless Driving

Adiabatic evolution is a common tool in quantum simulation experiments [71], geometric quantum computation [67], and quantum annealing [72]. As a result of the slow evolution required by the adiabatic condition, adiabatic protocols are robust against experimental imperfections like errors in drive strength. On the other hand, slow operations are vulnerable to decoherence. A paradigm termed “shortcuts to adiabaticity” (STA) aims to reproduce the results of adiabatic evolution in a shorter time frame, while still preserving the error-robust characteristics [73]. This section describes the results of one such STA applied to the STIRAP technique. The STA, termed “superadiabatic transitionless driving” (SATD), achieves similar transfer efficiencies to STIRAP in less time. The novel protocol is robust against dissipation and can be used in systems where losses are a significant factor.




Figure 3.2: Concept of superadiabatic evolution. The initial state $|\Psi_i\rangle$ and target state $|\Psi_f\rangle$ lie within the system’s full Hilbert space, \mathcal{H} . These two states are connected via an ideal adiabatic evolution (dashed line), but the constraints of finite evolution time introduce errors in the evolution that deviate from ideal (blue line) and result in the system reaching an excited state, $|\varphi\rangle$. The superadiabatic shortcut (red line) is calculated by incorporating a finite evolution time, and reproduces the desired final and initial states via a different path in \mathcal{H} . For any evolution, dissipation (black lines) creates loss of fidelity by incoherently taking the system to a different point in \mathcal{H} .

3.2.1 Adiabatic and Superadiabatic Control

Adiabatic evolution is a useful tool for controlling quantum systems. However, the adiabatic theorem, on which the evolution is based, is only ever approximately satisfied. Real experimental systems can never fully implement an adiabatic evolution, and it is rarely expedient to use long evolution times. To attempt an adiabatic evolution, the infinite-time evolution is truncated to an experimentally feasible time. This truncating step introduces some level of nonadiabatic evolution. Especially in systems where decoherence is nontrivial on adiabatic timescales, there is a need to design ways for systems to evolve such that the adiabatic results are reproduced with tighter time constraints.

STA is a class of protocols that does manage to fulfill both constraints. By introducing tailored control fields, the exact dynamics of a real, partly-adiabatic system can match a fully adiabatic evolution. One of the more straightforward STA techniques is known as “counter-adiabatic” or “transitionless” driving, in which modified or additional control fields precisely cancel the nonadiabatic components in an approximately-adiabatic evolution [74, 75]. As a

general strategy, counter-diabatic driving has proven effective at accelerating state transfer in two-level systems [76, 77] and in speeding up the manipulation of trapped atoms [78, 79]. However, the implementation of counter-diabatic driving may be complex, especially for higher-dimensional systems. There is no guarantee that the counter-diabatic variation of an adiabatic process will use the original set of control fields and level couplings. In particular, for STIRAP, the counter-diabatic evolution requires a coupling directly between the initial and final states [74]. This coupling is not typically present in STIRAP, and indeed is not assumed to exist in the most general form of a Λ system. Thus, other STA methods are required to accelerate the operation of STIRAP.

Superadiabatic transitionless driving, or SATD, has been developed to fill this need [80]. SATD introduces a dressed state basis that matches the initial and final conditions of an adiabatic evolution, but not necessarily at intermediate times. This concept is illustrated in Fig. 3.2, which shows a set of possible trajectories in the Hilbert space \mathcal{H} . The adiabatic path smoothly connects the initial and final states. As the speed of the traversal increases, the trajectory shifts, eventually resembling the non-adiabatic path, which reaches a different final state due to errors. In a counter-diabatic evolution, additional control fields allow the system to exactly reproduce the adiabatic path, but with shorter evolution times. And finally, the superadiabatic shortcut connects the correct initial and final states, but may trace a different path through \mathcal{H} to do so. This superadiabatic path follows a “dressed dark state” that maintains many of the same qualities of the dark state in the adiabatic evolution.

The SATD pulses are formulated as corrections to the original adiabatic pulses. Experimentally, STIRAP must be implemented in a finite time L , with uncorrected pulses

$$\begin{aligned}
 \Omega_S(t) &= \Omega \cos(\theta(t)), \\
 \Omega_P(t) &= \Omega \sin(\theta(t)), \\
 \theta(t) &= \frac{\pi}{2} \frac{1}{1 + e^{-\nu(t-L/2)}}, \\
 \nu(L, \epsilon) &= \frac{2}{L} \log \left(\frac{\pi}{2} \frac{1}{\sin^{-1}(\epsilon)} - 1 \right).
 \end{aligned} \tag{3.4}$$

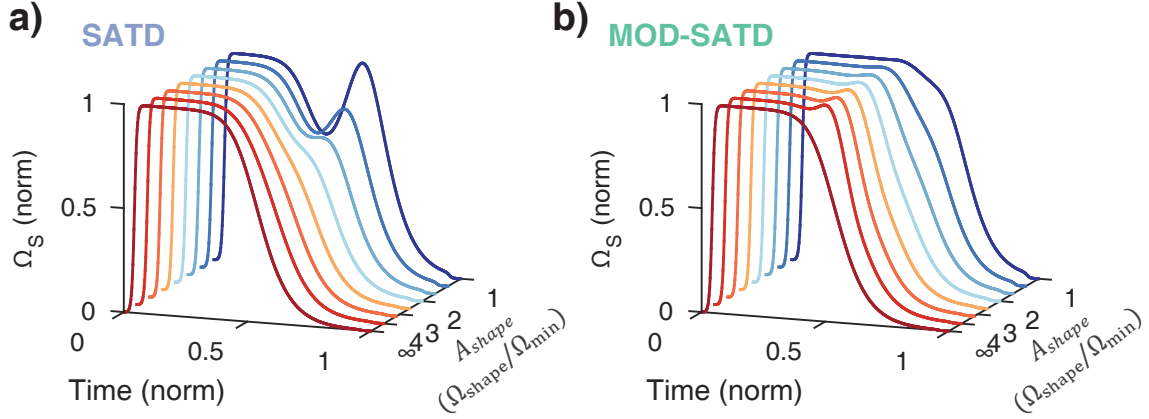


Figure 3.3: Superadiabatic STIRAP pulse shapes. The analytical variations on the STIRAP Stokes pulse ($\Omega_S(t)$) are shown for both superadiabatic transitionless driving (SATD, a) and a variant which reduces excited state occupation (MOD-SATD, b). For a given pulse duration, $A_{shape} \rightarrow \infty$ reduces to the original adiabatic protocol. At $A_{shape} = 1$, the correction is the maximum which can be implemented without exceeding the original drive strength. In both cases, the pump pulse (Ω_P) is equal to the time-reversed Stokes pulse.

This description introduces several parameters. ν determines the rate of the sweep, and ϵ accounts for the inevitable truncating of the pulses (since infinite time is required in the adiabatic ideal) by setting $\Omega_P(0) = \Omega_S(L) = \epsilon \cdot \Omega$. These definitions maintain a constant Ω over the duration of the pulse. In addition, note that the adiabaticity condition for these pulse shapes, given by Eq. 3.3, is simply $\Omega \gg |\dot{\theta}(t)|$, and increasing the optical power for a given L improves the adiabatic approximation.

With the adiabatic pulses established, the superadiabatic pulses can be calculated according to the procedure of ref. [80]. These are

$$\begin{aligned}\Omega_S^{SATD}(t) &= \Omega \cos(\theta(t)) - \frac{4\Omega \sin(\theta(t))\ddot{\theta}(t)}{\Omega^2 + 4\dot{\theta}(t)^2} \\ \Omega_P^{SATD}(t) &= \Omega \sin(\theta(t)) + \frac{4\Omega \cos(\theta(t))\ddot{\theta}(t)}{\Omega^2 + 4\dot{\theta}(t)^2}\end{aligned}\tag{3.5}$$

In order to compare the performance of adiabatic and corrected pulses, the amplitude of the corrected pulses should not exceed the adiabatic pulses; that is, $\Omega_S^{SATD}, \Omega_P^{SATD} \leq \Omega$ for all

t. Otherwise, SATD could have access to more adiabatic regimes simply by increasing the pulse power. When this condition is enforced, a minimum optical power emerges for a given pulse length:

$$\Omega_{min}^{SATD} = \frac{\nu}{1.315}. \quad (3.6)$$

This constraint is used to define a dimensionless adiabaticity parameter $\Omega/\Omega_{min}^{SATD}$. It will be referenced in two cases. First, when deriving theoretical pulse shapes for an intended drive strength, it will be called the shape parameter $A_{shape} = \Omega_{shape}/\Omega_{min}^{SATD}$. Second, for describing the experimental conditions, it will be called the experimental adiabaticity $A = \Omega_{exp}/\Omega_{min}^{SATD}$. The family of SATD curves is shown in Fig. 3.3a, ranging from maximum to minimum correction at the corresponding shape parameter. As $A \rightarrow \infty$, the evolution becomes fully adiabatic, the correction vanishes, and the original adiabatic pulse shapes are recovered. Alternatively, as $A \rightarrow 1$, the correction reaches its maximum for constant Ω . Different pulse shapes (A_{shape}) can be implemented at the same optical power (A), although *a priori* it would seem that optimal performance should be achieved for $A_{shape} = A$.

Unlike the adiabatic ideal, the SATD evolution deliberately introduces some excited state population in order to accelerate the evolution. This exposes the system to the losses of the excited state. A variation of SATD, termed MOD-SATD, can be derived by reducing the contribution of the excited state to the dressed dark state. The corresponding family of pulse shapes is shown in Fig. 3.3b. Both SATD and MOD-SATD are implemented in the NV Λ system, with results described in the following sections. The derivations of both SATD and MOD-SATD pulse shapes are presented in Appendix A.

3.2.2 Experimental Setup

These experiments are performed on a single NV center in a commercial, CVD-grown diamond sample obtained from Element Six. The diamond is (100)-oriented, with the selected NV center axis at a 54.7° angle to the surface normal. The sample is mounted in




Figure 3.4: Experimental hardware setup. A single NV center in a diamond chip is held at 5.5 K in a cryostat with optical access. A 532-nm (Initialize) and two 637-nm (Readout, Interact) lasers are integrated into the optical path using dichroic mirrors (DC), and phonon-assisted NV emission in the 650-800 nm band is collected into a fiber leading into an avalanche photodiode (APD). Phase- (PEOM) and amplitude- (AEOM) electro-optic modulators split the interaction laser into coherent sidebands to address both Λ -system transitions. The AEOM is driven by a 10-GHz arbitrary waveform generator (AWG), and a second AWG performs IQ modulation on a pair of signal generators (SG) which drive the PEOM and perform resonant spin manipulations on the NV center via on-chip waveguides.

a closed-cycle cryostat manufactured by Montana Instruments, which maintains a stable base temperature of 5.5 K. An external permanent magnet splits the $| -1 \rangle$ and $| +1 \rangle$ levels by $2\pi \times 1.414$ GHz. Multiple optical and microwave control fields are applied, as indicated in Fig. 3.4. Two microwave signal generators (SRS SG396, “SG”) output at the ground state transition frequencies of $2\pi \times 2.171$ GHz and $2\pi \times 3.585$ GHz, and are directed to the sample through coplanar waveguides. A 532-nm laser initializes the charge state and polarizes into the $|0\rangle$ state with $>90\%$ fidelity. Two tunable 637 nm lasers are calibrated to the $| -1 \rangle \leftrightarrow |A_2\rangle$ and $|0\rangle \leftrightarrow |E_Y\rangle$ transitions for Λ system interactions and resonant readout, respectively. Their wavelengths are stabilized by a High Finesse WS6-200 wavemeter. In order to address both transitions of the Λ system, the interaction laser is passed through a phase electro-optic modulator (Jenoptik PM 635, “PEOM”) that can modulate the phase of the optical waveform. The difference frequency of the two microwave drives, equal to the splitting between the $| \pm 1 \rangle$ levels, is applied to the PEOM. This periodic modulation induces harmonic sidebands on the laser beam, such that the 0th and 1st harmonics simultaneously

address the two branches of the Λ system. The overall amplitude of the lasers, Ω , is an important parameter, and is precisely controlled by a second, amplitude-EOM (Jenoptik AM 635, “AEOM”). Both phase and amplitude modulations are controlled by a fast 10-GHz arbitrary waveform generator (Tektronix 7102, “AWG”). A secondary, synchronized 1-GHz AWG (Tektronix 5014C) operates the acousto-optic modulators and rf switches that implement slower gating of the three lasers and ground-state microwave pulses. This AWG is also responsible for triggering the photon-counting logic switches which bin collected photons into useful data output categories.

Optical measurements of the NV center are performed through a custom-built confocal microscope setup. PL is collected through a 100X objective with a numerical aperture of 0.9. This objective, with a working distance of 1 mm, is held at room temperature inside the cryostat by a cryo-optic heater, which provides a long thermal lag between the objective and the cryostat platform. The optical path contains 550 nm and 650 nm dichroic elements that integrate the green and red lasers; the red lasers are combined on a non-polarizing beamsplitter before entering the main path. Half-waveplates in each of the laser source paths allow for independently adjustable linear polarization. Long-pass (647 nm) and short-pass (830 nm) filters attenuate any scattered laser light prior to the collected PL entering the single-photon counting module (PerkinElmer SPCM-AQR-16), which has a maximum detection efficiency of 70% at 700 nm and fewer than 50 dark counts per second. To steer the excitation beams, a fast steering mirror (Newport FSM 300-01) is placed at the $4f$ point of the two confocal lenses. This arrangement allows the input angle of the excitation lasers at the back aperture of the objective to be smoothly varied without changing their lateral positioning, which in turn allows rastering of the laser focal point in the sample plane.

Several calibration and characterization measurements are required before implementing SATD. By taking repeated scans of the $|E_X\rangle$ line, spectral diffusion of $\sigma \sim 2\pi \times (31 \pm 6)$ MHz is measured. A wide scan of the excited state transitions indicates that $|E_X\rangle$ and $|E_Y\rangle$ are separated by 3.4 GHz, corresponding to a low transverse strain of $2\pi \times 1.7$ GHz.




Figure 3.5: Optical Rabi oscillations on a single NV center. A 100-ns laser pulse is applied, resonant with the $| -1 \rangle \leftrightarrow | A_2 \rangle$ transition. Photoluminescence (PL) collected during the pulse measures the excited state population, and reveals decaying, coherent oscillations. A fit to the decaying oscillations yields the Rabi frequency $\Omega = 2\pi \times 171$ MHz and the dephasing time $T_{orb} = 18$ ns. The background exponentially decays due to spontaneous emission to off-resonant ground states.

Next, to characterize $| A_2 \rangle$, a π pulse is calibrated on the $| -1 \rangle \leftrightarrow | A_2 \rangle$ transition (Fig. 3.5). After applying a π pulse, time-resolved PL is collected using a time-correlating card (FAST ComTec P7889) with 0.1 ns resolution. The resulting decay curve fits to a single exponential with a time constant of $T_1 = 11.1 \pm 0.2$ ns. This optical lifetime provides the main source of dissipation in the experiments. Next, optically-driving Rabi oscillations on the isolated $| 0 \rangle \leftrightarrow | E_Y \rangle$ transition are used to calibrate the ratio of optical power in the 0th and 1st harmonic sidebands of the interaction laser – i.e., Ω_P/Ω_S . Finally, separate sets of optically-driven Rabi oscillations on the $| -1 \rangle \leftrightarrow | A_2 \rangle$ and $| +1 \rangle \leftrightarrow | A_2 \rangle$ transitions are used to correct for differences in the response of the two transitions to the same optical power, which vary with polarization. This data also ensures accurate estimation of Ω_P and Ω_S . Drive rates on the order of 100 MHz are achieved for input optical powers on the order of 5 μ W.

3.2.3 Superadiabatic Performance

First, the relationship between Ω and L is explored. An initial pulse length of $L = 16.8$ ns is chosen, for which the minimum drive strength that allows an SATD correction is $2\pi \times 72.6$ MHz. The system is initialized into the $| -1 \rangle$ state, and then the three sets of

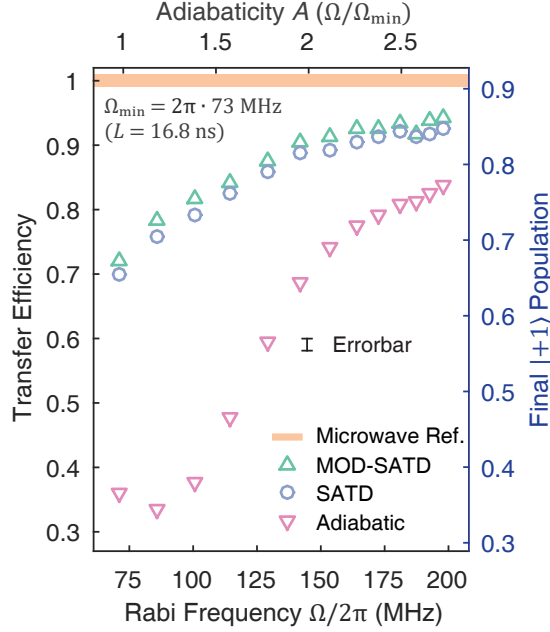


Figure 3.6: Improved STIRAP efficiency due to SATD and MOD-SATD. The system is initialized in $|-1\rangle$, and a pair of STIRAP pulses is applied to transfer it to $|+1\rangle$. For constant $L = 16.8$ ns pulses, the uncorrected adiabatic protocol rapidly falls off in efficiency as optical power is decreased, whereas the superadiabatic variants maintain high transfer efficiency and increasingly outperform the adiabatic case. Transfer efficiency is normalized against the maximum population transfer achievable with microwave drives (yellow bar), as a result of imperfect state initialization. The slight uptick in efficiency for the adiabatic case at $\Omega = 2\pi \times 73$ MHz results from the increasing unintended $|A_2\rangle$ excitation and incoherent decay to the $|+1\rangle$ state.

STIRAP techniques are employed to evaluate their transfer efficiencies: adiabatic, SATD, and MOD-SATD. The superadiabatic pulses are each selected according to the corresponding experimental adiabaticity: $A_{shape} = A$. After each pair of pulses, the resulting populations in the $|\pm 1\rangle$ ground states are sequentially measured by projecting them to the $|0\rangle$ state and normalizing the collected PL on the $|0\rangle \leftrightarrow |E_Y\rangle$ transition. The calculated transfer efficiencies are shown in Fig. 3.6. Superadiabatic protocols demonstrate a clear advantage in efficiency over the adiabatic equivalent, especially as the adiabatic approximation breaks down in the $A \rightarrow 1$ limit. The efficiencies still show imperfections, as the excited state lifetime is comparable to the length of the pulse and will degrade the performance of protocols which involve $|A_2\rangle$. However, these losses are minimal compared to excitations out of the

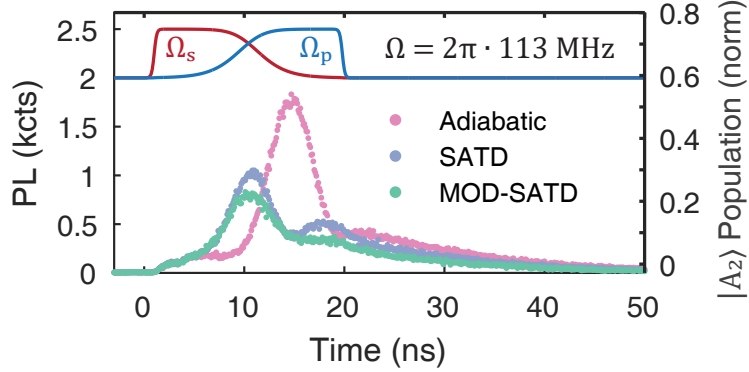


Figure 3.7: Photoluminescence during the STIRAP protocol. Photoluminescence (PL) results from $|A_2\rangle$ occupation (right axis), and reveals the differences between adiabatic and superadiabatic approaches. The adiabatic pulses (red and blue lines, offset above) unintentionally populate $|A_2\rangle$ once the resonant pump pulse is applied. The SATD variants coherently occupy the excited state in the middle of the protocol and drive the population to the final state. MOD-SATD shows reduced PL compared to SATD, consistent with its design to reduce $|A_2\rangle$ occupation. The long decay tail beyond 20 ns results from residual excited state population, which is nonideal in all cases, but worst for the adiabatic protocol.

protocol’s dark state, as in the adiabatic case. SATD outperforms the adiabatic transfer over the full experimental range, and as expected, MOD-SATD improves further by decreasing the detrimental excited state occupancy. The apparent uptick in adiabatic transfer efficiency at $\Omega = 2\pi \times 72.6$ MHz is due to incoherent decay from the unintentionally excited population, and hides the fact that coherent population transfer is trending to zero. If the dominant dissipation in $|A_2\rangle$ were to states outside of the qubit basis, this illusory recovery would not appear.

An examination of the PL collected during the STIRAP protocol shows that the evolution in each case matches expectations. With $|A_2\rangle$ the only state involved which emits PL, the collected emission is proportional to the excited-state occupancy at that point of the protocol. As illustrated in Fig. 3.7, all three pulse pairs show significant population in $|A_2\rangle$. However, SATD and MOD-SATD are designed to coherently occupy and recover from the excited state, which is reflected in the peaks at the center of the symmetric pulse sequence. For the adiabatic pulses, on the other hand, all excited state population is erroneous, and its more significant losses are reflected in the large spike in PL as the pump laser is activated.

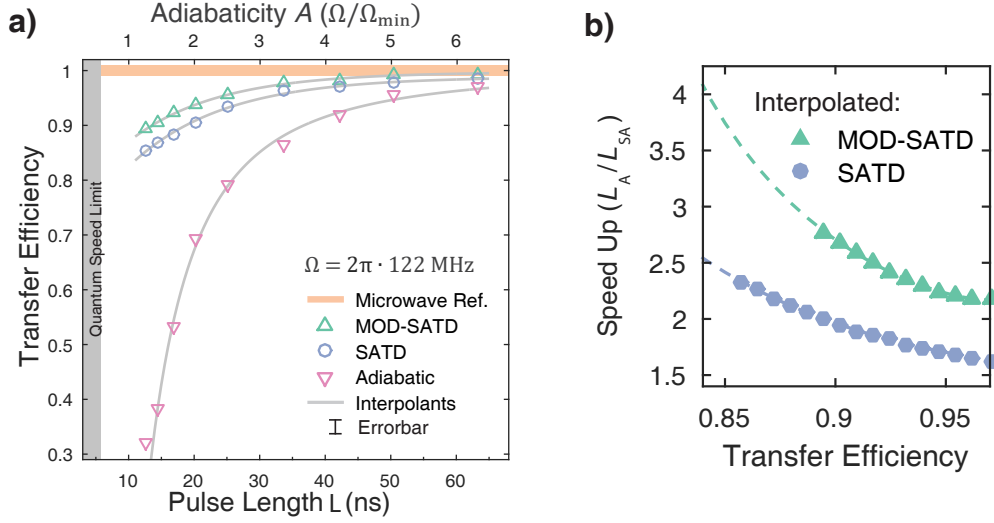


Figure 3.8: Speedup of SATD and MOD-SATD protocols. (a) Fig. 3.6 showed how transfer efficiency changed with Ω for constant L ; here, Ω is fixed at $2\pi \times 122$ MHz while L is varied. Similarly, as $A \rightarrow 1$, the adiabatic protocol becomes inefficient while SATD and MOD-SATD perform relatively well. The SATD techniques maintain high efficiency even as the pulse width approaches the quantum speed limit (gray vertical bar). (b) By interpolating along each data set (gray lines), it is possible to calculate the speedup which the superadiabatic protocols offer over the adiabatic approach (inset). To reach the same target transfer efficiency, the two superadiabatic protocols offer a speedup between 1.5X and 3X.

Consistent with its theoretical origins, MOD-SATD yields $\sim 20\%$ lower PL than SATD.

Next, $\Omega = 2\pi \times 122$ MHz is held constant while L is varied between 13 and 65 ns, to examine the relative speedup that superadiabatic protocols can provide. Since $\Omega_{\min} \propto L^{-1}$, this set of experiments also explores a wide range of A . As shown in Fig. 3.8a, the SATD and MOD-SATD pulses maintain a higher transfer efficiency than adiabatic pulses for all L . Based on the smooth trend of the data, it is possible to fit phenomenological curves to each set of data and invert the relationship between the transfer efficiency and pulse length. This allows an estimation of the speedup to achieve a target fidelity when applying SATD or MOD-SATD instead of adiabatic evolution. The corresponding speedup factors are shown in Fig. 3.8b. For example, SATD reaches 90% transfer fidelity using a pulse 2.0 times shorter

than similar adiabatic transfer, and for MOD-SATD the speedup is 2.7.

In addition to slow adiabatic protocols, the superadiabatic transfers can also be compared against the fastest possible quantum evolution. For the same Λ coupling structure, assuming no direct coupling between the ground states, the “quantum speed limit” corresponds to a transfer protocol which simply applies the maximum coupling on both transitions until the state is transferred. In other words, $\Omega_P^{QSL} = \Omega_S^{QSL} = \Omega$. For $\Omega = 2\pi \times 122$ MHz, the corresponding L_{QSL} is 5.8 ns. However, the quantum speed limit protocol includes significant occupancy of $|A_2\rangle$, reaching 50% probability at the midpoint of the transfer. Thus, the fastest possible transfer would show low transfer efficiencies. The shortest SATD pulses are 12.6 ns, which is roughly a factor of two slower than the quantum speed limit, but still achieve $> 85\%$ transfer efficiency. The superadiabatic protocols strike a balance between speed and efficiency among the possible quantum trajectories through the Hilbert space.

Finally, two extensions of STIRAP show the flexibility of superadiabatic techniques for quantum state control: coherent phase transfer and fractional STIRAP. To demonstrate phase transfer, the superposition state $|\psi_I\rangle = (|0\rangle + e^{i\phi}|-1\rangle)/\sqrt{2}$ is initialized with a microwave pulse. Next, STIRAP or the superadiabatic variant is applied to evolve the system to a target final state of $|\psi_F\rangle = (|0\rangle + e^{i\phi}|+1\rangle)/\sqrt{2}$. Last, the X and Y projections of the superposition are measured to confirm that the phase ϕ is correctly transferred. This procedure is illustrated schematically in Fig. 3.9a, and the results for several values of ϕ are plotted in Fig. 3.9b. Both adiabatic and superadiabatic processes coherently transfer the initial phase, but as with the full state transfer, SATD and MOD-SATD are significantly more efficient. The solid arcs on the polar plot mark the square root of the corresponding transfer efficiencies (suitable for comparing to state amplitudes), and make clear that a greater fraction of the superadiabatic transfer efficiency is due to coherent evolution. For adiabatic pulses, on the other hand, a large portion of the transferred population is not phase coherent, and results from incoherent decay out of the excited state.

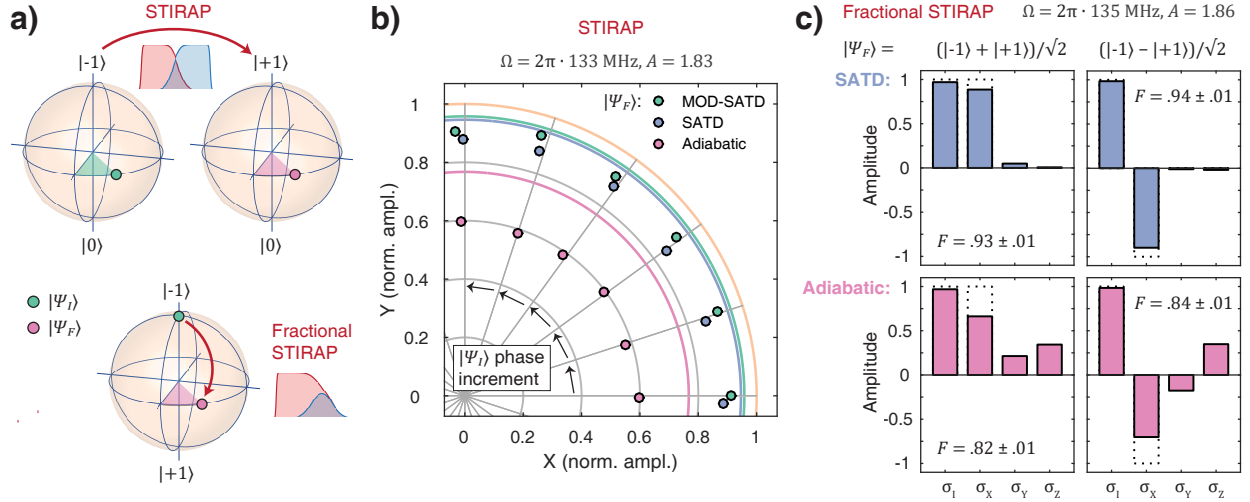


Figure 3.9: Phase coherence of STIRAP and fractional STIRAP. (a) Top, the procedure to confirm the coherence of STIRAP is shown. A superposition is initialized in the $\{|0\rangle, |-1\rangle\}$ subspace and STIRAP is used to transfer it to a superposition in the $\{|0\rangle, |+1\rangle\}$ subspace. The same phase relation should exist between basis states in both initial and final superpositions. Bottom, the pulse shapes and result of fractional STIRAP are shown. By truncating the STIRAP pulses, the initial state of $|-1\rangle$ can be transferred to a superposition of the $|\pm 1\rangle$ states. (b) The X and Y projections of the final superposition are shown for various initial phase relations (gray lines). All three protocols maintain phase coherence, but the superadiabatic protocols are closer to the overall transfer efficiency (solid color arcs), indicating primarily coherent transfer. The adiabatic state projections indicate a large incoherent contribution to the final state. (c) Quantum state tomography is performed on the the final superposition states resulting from fractional STIRAP. The adiabatic transfer achieves an average fidelity of 0.83 ± 0.01 , whereas SATD improves to 0.94 ± 0.01 .

In fractional STIRAP (“f-STIRAP”), the same principles of dark state evolution are applied to create a superposition of initial and final states instead of completing full state transfer. The pulses of f-STIRAP resembles the full STIRAP pulses, except that they adiabatically ramp to zero partway through the evolution (shown in Fig. 3.9a). The superadiabatic derivation can be applied to create SATD variations of f-STIRAP as well. Fig. 3.9c shows quantum state tomography of two different superpositions of the $|\pm 1\rangle$ basis, created with $\Omega = 2\pi \times 135$ MHz. Once again, SATD demonstrates an advantage over uncorrected adiabatic pulses, averaging a fidelity of 0.94 ± 0.01 , against the adiabatic fidelity of 0.83 ± 0.01 .

3.2.4 Influence of Dissipation

The fundamental assumption of STIRAP is that the intermediate state is lossy and needs to be avoided. Otherwise, if the state is lossless, there is no reason to deviate from the quantum speed limit protocol. As an adiabatic protocol, STIRAP is inherently robust against noise and control errors. Inefficiencies and dissipation can be reduced by improving the adiabatic approximation, either by increasing optical power or extending the pulse sequence. However, STA protocols are usually derived assuming unitary (lossless) evolution and ideal experimental implementation. Despite the improved speed and efficiency of SATD demonstrated previously, it is an open question to what degree it is robust to the remaining dissipation and any control errors.

This NV center's dissipation is broadly categorized by $T_1 = 11.1 \pm 0.2$ ns, but its specific decay paths to the ground state can be more finely characterized. Nominally, $|A_2\rangle$ couples to $|\pm 1\rangle$ only, which should manifest as dissipation that preserves population but contributes to loss of coherence. However, the excited state mixing induced by strain and the magnetic field weakly allows transitions to the $|0\rangle$ state. This loss does not contribute to incoherent transfer, but rather leads to leakage from the qubit basis. Furthermore, the optical transitions experience dephasing on fast timescales, which also reduces the coherence of transfers. The spontaneous emission branching ratios between pairs of ground states can be calculated by initializing either $| -1\rangle$ or $| +1\rangle$ and pumping its transition to $|A_2\rangle$. After the initial state is fully depleted and the excited state has fully decayed, the resulting populations in the other two ground states reflect the ratio of decay rates from $|A_2\rangle$. Performing these measurements for $|\pm 1\rangle$, the ratios Γ_{-1}/Γ_0 and Γ_{+1}/Γ_0 are found, where Γ_i describes the decay rate for the $|A_2\rangle \rightarrow |i\rangle$ transition. These two quantities can be combined with the relation $1/T_1 = \sum_i \Gamma_i$

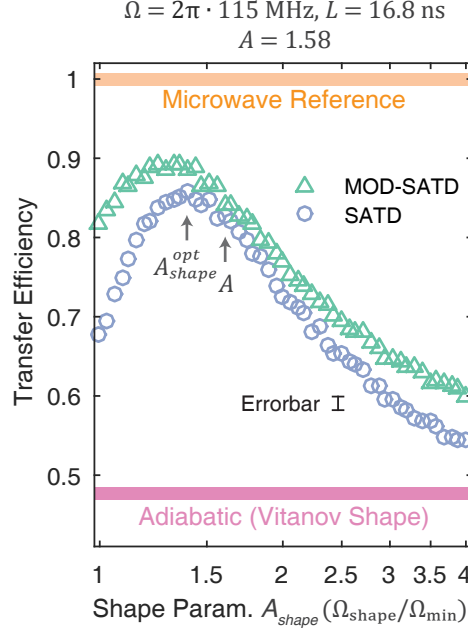


Figure 3.10: Robustness of SATD and MOD-SATD to deviations in pulse shape. The results of applying different pulse shapes, with L and Ω held constant, are shown. With $\Omega = 2\pi \times 115 \text{ MHz}$ and $L = 16.8 \text{ ns}$, the adiabaticity A of the transfer is 1.58. For this A , the adiabatic transfer efficiency is < 0.5 (magenta line), and the superadiabatic variants are an improvement no matter which pulse shape A_{shape} is used. Surprisingly, the optimal A_{shape} for both SATD and MOD-SATD is less than A , indicating that the system behaves as if more nonadiabatic than the pulse parameters predict.

to determine the decay rates:

$$\begin{aligned}\Gamma_{+1} &= 2\pi \times (8.5 \pm 0.3)\text{MHz}, \\ \Gamma_{-1} &= 2\pi \times (4.3 \pm 0.2)\text{MHz}, \\ \Gamma_0 &= 2\pi \times (1.6 \pm 0.1)\text{MHz}.\end{aligned}\tag{3.7}$$

Evidently, while $\Gamma_0 < \Gamma_{\pm 1}$, it is still appreciable on the timescales of STIRAP. In addition, by fitting the decay of multiple sets of optically driven Rabi oscillations and incorporating the measured lifetimes, the orbital dephasing rate of $\Gamma_{orb} = 2\pi \times (8.8 \pm 0.1) \text{ MHz}$ is extracted. These parameters more fully describe the intrinsic dissipation of the system than simply the optical lifetime.

To probe the robustness of the superadiabatic protocols, pulses with $A_{shape} \neq A$ are

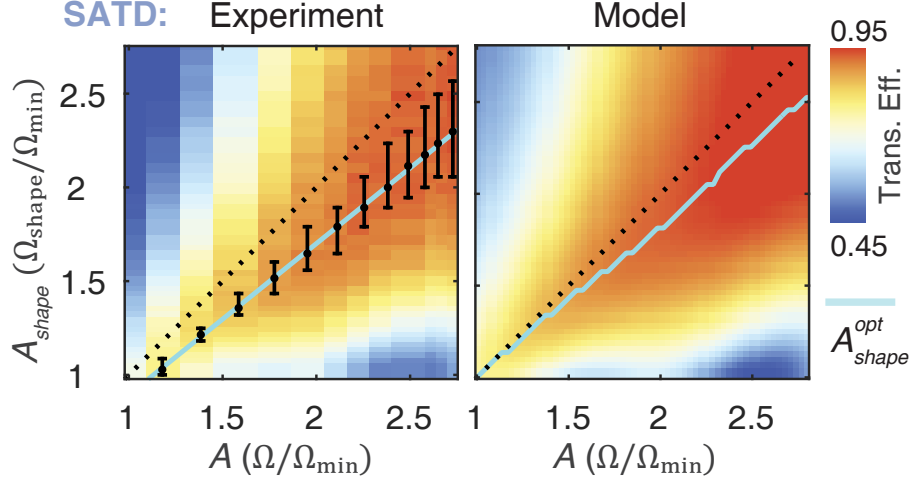


Figure 3.11: Experimental and theoretical data on the influence of dissipation on STIRAP. At left, experimental data shows the SATD transfer efficiency when independently varying A_{shape} and A . In the absence of dissipation, $A_{shape} = A$ is expected, given by the dotted line. However, the maximal transfer efficiency is observed for $A_{shape} < A$ (black data points, where interval indicates ± 0.01), and fits well to a linear relationship across a broad range of A (cyan line). At right, this deviation is reproduced by a master equation simulation that incorporates the dissipative features of optical lifetime T_1 , orbital dephasing Γ_{orb} , and spectral diffusion σ . A_{shape}^{opt} is again $< A$ for all A .

tested. Both $\Omega = 2\pi \times 115$ MHz and $L = 16.8$ ns are fixed, corresponding to $A = 1.58$, while a range of shape parameters is implemented. In Fig. 3.10, the transfer efficiency results show that both SATD and MOD-SATD continue to outperform adiabatic pulses even for large deviations from the nominal A_{shape} . The adiabatic comparison, which has no shape parameter, is indicated by the magenta bar. In addition, for both superadiabatic variants, $A_{shape} = A$ appears not to be the optimal working point. Instead, $A_{shape}^{opt} < A$ in both cases. While this deviation may be partly attributed to experimental errors in implementing the target A_{shape} , adding theoretical investigations can provide a more thorough explanation for the origins of this phenomenon.

Fig. 3.11 compares the results of both experimental investigation and theoretical modeling of the performance of SATD in the NV's dissipative Λ system. In experiment (left panel), the transfer efficiency is plotted for a range of adiabaticity A . The optimal transfer shape parameter at each laser power fits well to a linear trend (cyan line) described by

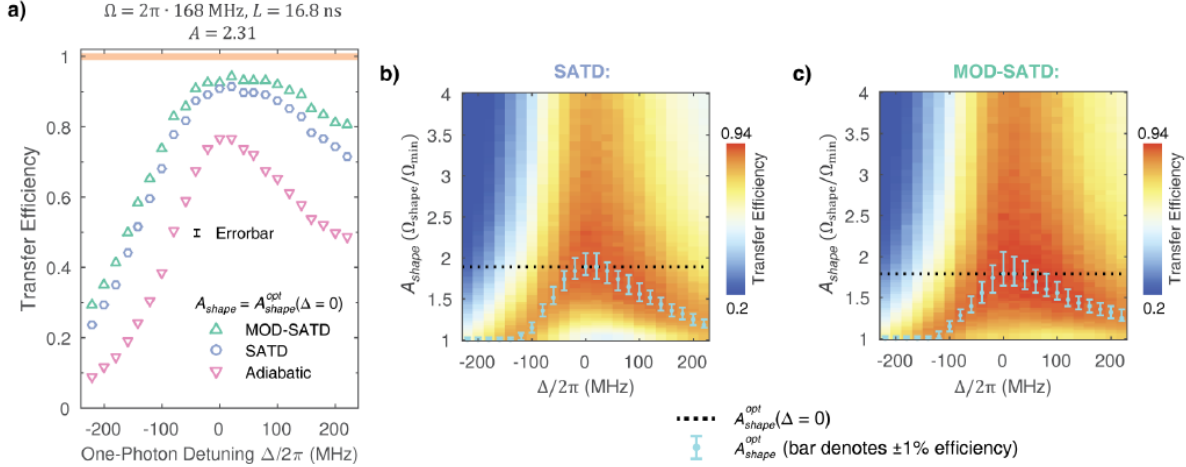


Figure 3.12: Superadiabatic transfer efficiency with deliberate one-photon detuning. (a) Transfer efficiency of all three STIRAP protocols is shown as a function of one-photon detuning Δ , for fixed pulse parameters Ω , L , and A_{shape} . All methods show decreases in efficiency with increased detuning, and superadiabatic protocols maintain high transfer efficiency over a wider range of detuning. (b) Sweeping A_{shape} while performing SATD with deliberate detuning at constant A shows that A_{shape}^{opt} decreases as Δ increases. This implicates spectral diffusion σ as part of the cause of $A_{shape}^{opt} < A$ in resonant experiments. The dashed line corresponds to the data points in (a). (c) Same as in (b) for MOD-SATD.

$A_{shape}^{opt} = (0.81 \pm 0.02) \cdot A + (0.09 \pm 0.03)$. With a slope < 1 , the data and fit clearly indicate deviations from the $A_{shape} = A$ trend (dotted line) which would be expected in the absence of dissipation. A master equation model (right panel) reproduces the data well, including the transfer efficiency and the qualitative trend of $A_{shape}^{opt} < A$. The maximum transfer efficiency predicted by the model is also highlighted by a cyan trendline. To match the experimental results, this model includes the measured T_1 , Γ_{orb} , and spectral diffusion values of the $|A_2\rangle$ level. The appearance of $A_{shape}^{opt} < A$ in the simulation confirms that it is not solely due to experimental imperfections. Physically, the presence of spontaneous emission and dephasing, as well as the off-resonance condition induced by spectral diffusion, decrease the effective coupling rate to the excited state. Thus, the required superadiabatic correction (which is designed to populate the excited state) increases to compensate. Data taken with deliberate one-photon detuning (in which the frequencies of both laser fields are detuned by the same amount) help support this interpretation (Fig. 3.12). In Fig. 3.12a, the deliberate detuning

decreases the transfer efficiency of all three methods, but the superadiabatic protocols are at least as robust as the adiabatic pulses. In Fig. 3.12b-c, both SATD and MOD-SATD show that A_{shape}^{opt} decreases as detuning increases while A is held constant. Thus, the measured spectral diffusion of $2\pi \times 31$ MHz can directly account for part of the discrepancy between A_{shape}^{opt} and A under nominally resonant conditions.

3.2.5 Conclusion

These experiments demonstrate that superadiabatic shortcuts can be effectively applied to STIRAP, and that SATD and MOD-SATD in particular are fast and robust quantum control techniques. The influence of dissipation and noise on the results emphasize the importance of accelerated control techniques for achieving high state fidelities, which are a crucial component of QISE technologies. These techniques are expected to be relevant for a variety of qubit systems, including optically active group-IV defect spins in semiconductors which have weak ground-state magnetic dipoles and in superconducting qubit designs with Λ system level structures. STIRAP can also be extended to the transfer of quantum states in systems where the singular intermediate level is replaced by a continuum of states, such as in a waveguide [81], and thus may be useful in implementing quantum networks.

3.3 All-Optical Holonomic Gates

Like STIRAP, unitary quantum operations are the basis for countless experiments. Conventionally, quantum operations are “dynamic,” and based on phases accumulated over time between two states with different energies. This includes quantum gates implemented by coupling two eigenstates with a resonant field, which create dressed states in the stationary basis that acquire phases with time. Dynamic gates are simple to implement with tunable couplings or resonant control fields, and calibration curves based on Rabi oscillations are standard in most experimental qubit systems. However, despite their simplicity and

ubiquity, dynamic gates may not be the optimal choice for quantum operations in the presence of significant noise, systematic control errors, or when process fidelities close to 1 are desired, such as in fault-tolerant quantum computing. The alternative to dynamic phases are geometric phases, which are prevalent throughout physics but have been overlooked in quantum control. This section describes the implementation of a complete set of geometric (holonomic) single-qubit gates in the NV center. The holonomic gates are implemented via a single non-adiabatic optical cycle in the NV $|A_2\rangle$ Λ system.

3.3.1 Holonomic Gates

Unlike dynamic phases, which accumulate over time, geometric phases instead arise as a system completes a cyclic evolution in its parameter space [82]. They are not a uniquely quantum phenomenon, as classical geometric phases abound. A canonical example is the Foucault pendulum, which precesses (acquires a phase) as the Earth rotates on its axis. More simply, the appearance of a geometric phase can be seen for a vector pointing parallel to the surface of a sphere. The vector starts at the sphere’s north pole, and is then smoothly transported to the equator. Next, it traces a quarter-rotation along the equator, before being brought back to the north pole. When it returns, it has rotated by 90° . This rotation is independent of how fast the path is traversed, and instead depends on the details of the path geometry – if instead the vector takes a half-rotation around the equator, the resulting *geometric phase* is 180° . This path dependence is an essential characteristic of both classical and quantum geometric phases, and the motivation for its “geometric” moniker. Geometric phases result from the topology of the evolution space; if the same vector were transported along any path on a flat surface, no geometric phase would appear.

This classical example of geometric phase has a quantum equivalent: evolution on the Bloch sphere. If the two-state system represented on a Bloch sphere traces out a closed path on its surface, the entire subspace will acquire a quantum phase which is proportional to the area the path encloses. For example, the path described on the surface of the classical sphere

– from pole to equator, along the equator, and back to the pole – can be described by a series of three unitary gates: $R_y(-\pi/2) \cdot R_z(\pi/2) \cdot R_x(\pi/2)$, where $R_i(\theta)$ indicates a rotation around axis i by angle θ . This path encloses $1/8 \cdot (4\pi) = \frac{\pi}{2}$ area of the Bloch sphere surface. On the classical sphere, $\frac{\pi}{2}$ (90°) matches the phase acquired by the translated vector. In the quantum regime, however, this relationship need not hold. For the Bloch sphere, due to the spin- $\frac{1}{2}$ equivalence of the two-level system, the subspace acquires phase $\gamma = -\frac{A}{2}$ for enclosed area A [82].

Because of the geometric phase’s independence of dynamical considerations, it has recently emerged as an alternative to dynamic quantum control in implementing unitary quantum operations [83]. The principle is to leverage the global qualities of geometric phases to perform high-fidelity gates that are relatively immune to the local fluctuations that are usually characteristic of noise, decoherence, and experimental error. The Bloch sphere evolution again provides a concrete example. As the system progresses along a path on the Bloch sphere, fluctuations in the evolution (perhaps resulting from qubit energy splitting fluctuations) that are parallel to the path have little impact on the enclosed area. Even transverse fluctuations may cancel out over the course of a complete cycle, and so the geometric evolution may be insensitive to multiple sources of error [67, 84]. This intrinsic robustness is reminiscent of adiabatic protocols such as STIRAP, and indeed geometric phase protocols are often implemented using adiabatic means.

Adiabatic geometric phases have been implemented in the NV center by leveraging the same Λ system described in section 3.2, the $|A_2\rangle \leftrightarrow |\pm 1\rangle$ subspace [67]. The system is initialized in the $|-1\rangle$ ground state, and then (fully adiabatic) STIRAP transfers the system to $|+1\rangle$ along a longitudinal line of the Bloch sphere. Once STIRAP is completed, a second application of STIRAP, with a different phase relationship, transfers the population back along a different longitudinal line. This completes a cyclic evolution on the $|\pm 1\rangle$ Bloch sphere, enclosing a deterministic area and acquiring the corresponding phase on the $|\pm 1\rangle$ states. Because this phase cannot be measured within the basis alone, the $|0\rangle$ state serves

as a reference and confirms the presence of the purely geometric phase.

Holonomic¹ quantum computation (HQC) aims to develop fully geometric quantum operations to perform fault-tolerant quantum logic gates [85]. Geometric phases on single qubits are a step in that direction, but strategies for arbitrary unitary and two-qubit operations, at minimum, will be required for universal quantum computing. Fortunately, non-Abelian (i.e. non-commuting, matrix-valued) geometric phases can result from cyclic evolution in larger state spaces than a two-level system. In addition, though adiabatic evolution enjoys the benefits of noise and control error robustness, it is vulnerable to decoherence and dissipation, as explored in section 3.2. Thus, a major goal of HQC is to develop protocols for non-adiabatic, non-Abelian holonomic gates which maintain some of the geometric phase’s noise robustness characteristics [86]. It is currently an open question whether non-adiabatic gates will exhibit protections against noise, but simulations suggest that the non-adiabatic geometric operations do feature more protection against control errors than dynamic gates [87].

One approach for implementing non-Abelian single-qubit gates involves three-level dynamics, and has been demonstrated in various forms on several qubit platforms. π -rotations around arbitrary Bloch sphere axes are possible by resonantly addressing a three-level system, and have been demonstrated in superconducting qubits [88], liquid-state NMR systems [89], and the full triplet ground state of the NV center [90, 91]. However, pairs of these π -rotations must be combined to create an arbitrary unitary operation. Alternatively, Abelian geometric phases or hyperbolic secant pulses can create arbitrary rotation angles in a single non-adiabatic loop, but they each introduce dynamic phases as well. A recent proposal [92] combines the desirable features of both approaches, enabling arbitrary single-qubit unitary operations with a single, non-adiabatic holonomic gate. The key addition of the proposal is detuning of the previously resonant-only control pulses. The follow sections present the results of implementing these novel holonomic gates on an optical Λ system of the NV center.

1. Borrowing the term “holonomy” from the field of differential geometry.

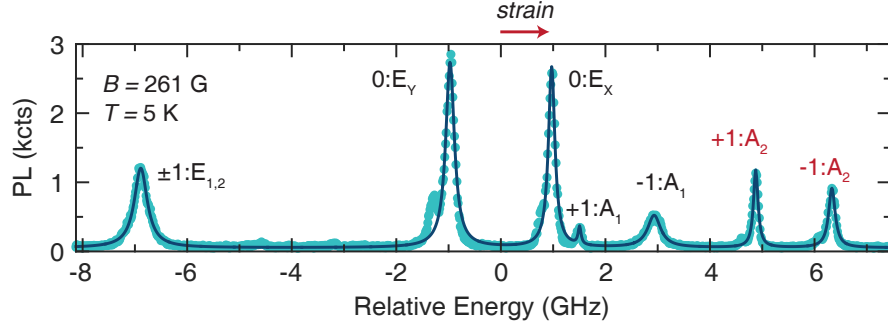


Figure 3.13: Photoluminescence excitation sweep of a single NV center. At 5 K, a 637-nm laser is tuned across the NV center excited-state resonances while off-resonant photoluminescence (PL) is collected. During the scan, two microwave drives resonant with the ground-state $|0\rangle \leftrightarrow |\pm 1\rangle$ transitions are applied to ensure the system does not decay into off-resonant ground states. With a $B = 261$ G magnetic field applied, the $|\pm 1\rangle$ ground states split in energy, and their respective excited-state transitions become resolved. The two highest-lying transitions, $|\pm 1\rangle \leftrightarrow |A_2\rangle$ (in red), form the Λ system used for optical control experiments. The splitting of the $|0\rangle \leftrightarrow |E_{Y,X}\rangle$ transitions provides a measurement of the transverse strain at the NV center, in this case a relatively small $\delta = 2\pi \times 0.94$ GHz.

3.3.2 Experimental Implementation

The experiments detailed in the following sections use the same sample and NV center described in section 3.2.2. Slight variations in the precise parameters include a temperature of 5.0 K and a magnetic field of 261 G which splits the $|\pm 1\rangle$ ground states by $2\pi \times 1.461$ GHz. The optical properties show some variation from cryogenic cycling, as this NV shows a transverse excited-state strain of $2\pi \times 0.97$ GHz (Fig. 3.13) and spectral diffusion of $2\pi \times 15$ MHz. These parameters are more ideal for Λ system experiments, as they introduce less noise and less mixing of the excited states.

Holonomic gates are implemented by a very similar Hamiltonian to the STIRAP Hamiltonian (Eq. 3.1), which is given in the rotating frame of the $\{|-1\rangle, |+1\rangle, |A_2\rangle\}$ basis:

$$H(t) = \frac{1}{2} \begin{pmatrix} 0 & 0 & \Omega_P(t) \\ 0 & 0 & \Omega_S(t)e^{-i\phi} \\ \Omega_P(t) & \Omega_S(t)e^{i\phi} & 2\Delta \end{pmatrix} \quad (3.8)$$

The primary distinction here is that the two-photon detuning δ is set to 0. Once again,

the qubit basis of interest is the $|\pm 1\rangle$ subspace. When calculating unitary operations on this subspace, there is no definite initial or final state, as in the case of STIRAP. Thus, the “pump” and “Stokes” labels become ill-defined. However, the convention of Ω_P addressing the $|-1\rangle \leftrightarrow |A_2\rangle$ transition and Ω_S addressing the $|+1\rangle \leftrightarrow |A_2\rangle$ transition is maintained for consistency. A similar parameterization in terms of polar angles is made here:

$$\begin{aligned}\Omega_P(t) &= \sin\left(\frac{\theta}{2}\right)\Omega(t), \\ \Omega_S(t) &= -\cos\left(\frac{\theta}{2}\right)\Omega(t).\end{aligned}\tag{3.9}$$

As with STIRAP and SATD, the θ parameter describes the relative amplitude of the two resonant laser fields. Unlike the preceding experiments, however, θ has no time dependence for non-adiabatic gates. The time dependence is solely in the total power, $\Omega(t)$.

Geometric operations appear in this Hamiltonian under certain conditions. First, considering the case where $\Delta = 0$, a qubit starting in the $|\pm 1\rangle$ subspace will return to it at time $t = \tau$ if $\int_0^\tau \Omega(t)dt = 2\pi$, regardless of the details of the optical pulse $\Omega(t)$. Due to the resonance condition, the eigenstates are degenerate in the rotating frame, and so the evolution is purely geometric. At $t = \tau$, the unitary evolution operator on the qubit states is [92]

$$\hat{U} = |D\rangle\langle D| - |B\rangle\langle B|.\tag{3.10}$$

$|D\rangle$ and $|B\rangle$ are the dark and bright qubit states, respectively, given by

$$\begin{aligned}|D\rangle &= \cos\left(\frac{\theta}{2}\right)|-1\rangle + \sin\left(\frac{\theta}{2}\right)e^{i\phi}|+1\rangle, \\ |B\rangle &= \sin\left(\frac{\theta}{2}\right)|-1\rangle - \cos\left(\frac{\theta}{2}\right)e^{i\phi}|+1\rangle.\end{aligned}\tag{3.11}$$

Examining these equations, it becomes clear that \hat{U} represents a π rotation around the Bloch sphere axis defined by the dark and bright states – note the resemblance of \hat{U} to the Pauli matrix σ_z . Furthermore, this rotation axis is fully determined by the experimental

parameters θ and ϕ . Analytically, \hat{U} can be rewritten as

$$\hat{U}(\theta, \phi) = \mathbf{n} \cdot \boldsymbol{\sigma}, \quad (3.12)$$

where $\mathbf{n} = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$ describes an axis on the Bloch sphere and $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ are the Pauli operators². Thus, π rotations around arbitrary axes are possible in a resonant-only Λ system in general, and can be implemented in the NV center $|A_2\rangle$ Λ system.

For more general dynamics, nonzero Δ can be considered. Here, due to the $\Delta |A_2\rangle \langle A_2|$ term, the Hamiltonian no longer is guaranteed to commute with itself at different times. Thus, the details of the pulse shape $\Omega(t)$ become significant. By choosing a square pulse with $\Omega(t) = \Omega$ for $0 \leq t \leq \tau$, the Hamiltonian can be made time-independent and purely geometric evolution is recovered. While the energy of $|A_2\rangle$ is nonzero in the rotating frame, any state which begins in the zero-energy $|\pm 1\rangle$ subspace will remain at zero energy throughout the time-independent holonomic gate pulse, thus ensuring that the dynamic phase vanishes. With $|\Delta| > 0$, τ must be adjusted to ensure cyclic evolution as $\tau = 2\pi/\sqrt{\Omega^2 + \Delta^2} \equiv T_{2\pi}$. At $t = \tau$, the resulting unitary operation on the $|\pm 1\rangle$ subspace is

$$\begin{aligned} \hat{U}(\theta, \phi, \Delta, \Omega) &= |D\rangle \langle D| + e^{i\gamma} |B\rangle \langle B| \\ &= e^{i\gamma/2} e^{-i(\gamma/2)\mathbf{n} \cdot \boldsymbol{\sigma}}, \end{aligned} \quad (3.13)$$

where now γ generalizes to

$$\gamma = \pi \left(1 - \frac{\Delta}{\sqrt{\Omega^2 + \Delta^2}} \right). \quad (3.14)$$

This operation describes a rotation by angle γ around axis \mathbf{n} on the $\{|-1\rangle, |+1\rangle\}$ Bloch sphere (along with a global phase). Unlike the resonant case, $\gamma \in (0, 2\pi)$, so both the axis

². If two of these arbitrary π rotations are concatenated, the result is $\hat{U}(\theta_m, \phi_m)\hat{U}(\theta_n, \phi_n) = \mathbf{m} \cdot \mathbf{n} - i\boldsymbol{\sigma} \cdot (\mathbf{n} \times \mathbf{m})$. This describes a rotation of $2 \cos^{-1}(\mathbf{n} \cdot \mathbf{m})$ around the axis $\mathbf{n} \times \mathbf{m}$, which is a general SU(2) operation.




Figure 3.14: Evolution of the NV center Λ system during holonomic gates. The Bloch sphere representation of the $\{|B\rangle, |A_2\rangle\}$ subspace (at right) shows how the initial state $|B\rangle$ completes a 2π rotation around an axis offset from the pole by angle $\alpha = \tan^{-1}(-\Omega/\Delta)$. This completed cycle encloses an area -2γ , resulting in geometric phase γ on the subspace. In the $\{|D\rangle, |A_2\rangle\}$ basis (at left), the evolution is trivial. This corresponds to a rotation in the $\{|B\rangle, |D\rangle\}$ ground-state basis, effecting the desired unitary operation through proper choice of $|D\rangle$.

and the angle of the rotation are fully controlled. Thus, by choice of Ω_P , Ω_S , and Δ , the gate parameters can be set to any desired single-qubit unitary operation. The specific θ , ϕ , and γ parameters used for the gates in these experiments are recorded in Appendix B.

Geometrically, this evolution can be represented on two Bloch spheres, shown in Fig. 3.14. On the $\{|B\rangle, |A_2\rangle\}$ Bloch sphere (right), the state completes a cycle of evolution which begins in the ground state and revolves around an axis set at polar angle $\alpha = \tan^{-1}\left(-\frac{\Omega}{\Delta}\right)$. This cycle encloses an area $A = -2\gamma$, adding a phase γ to both $|B\rangle$ and $|A_2\rangle$ (see Eq. 3.13). On the $\{|D\rangle, |A_2\rangle\}$ Bloch sphere, the evolution is trivial since there is no coupling between the dark and A_2 states.

3.3.3 Resonant Gates

To implement these gates in the NV center, several calibration steps are necessary. The phase electro-optic modulator (PEOM) reprises its role as creating the phase-controlled laser sidebands which simultaneously address the two Λ system transitions. The amplitude




Figure 3.15: Measurement of square optical pulse rise and fall times. A 637-nm laser pulse is created using square inputs and directed to the sample. A small number of reflected photons penetrate the dichroic mirror and enter the collection path. The photon counts are binned by a time-correlating card with 0.1 ns resolution. The collected counts fit well to trapezoidal shapes (solid lines), with a rise time of 1.2 ± 0.1 ns and a fall time of 1.3 ± 0.1 ns.

EOM (AEOM) is used to tailor the pulse shape, which in this case is ideally a square pulse. Fig. 3.15 shows measurements of the pulse shape based on residual laser photons which are reflected off the diamond surface, and while not focusing on the NV center. The detected pulse shape is well-approximated by a trapezoid, and the fit values of the 0-100 rise and fall times are 1.2 ± 0.1 ns and 1.3 ± 0.1 ns, respectively. This creates a good approximation to a square pulse for the holonomic gates to be performed. While some residual optical power is apparent outside of the square pulse, this leakage only persists for ~ 10 ns before and beyond the ideal pulse, since the AEOM transmission is windowed with an additional acousto-optic modulator with a 45 dB extinction ratio.

Next, the optical coupling Ω must be calibrated on each of the Λ system transitions. As PL is proportional to $|A_2\rangle$ occupation, time-resolved PL measurements during constant laser excitation reveal Rabi oscillations between the initial and excited state, as in Fig. 3.16. These optical Rabi oscillations can be taken at fixed detuning, and demonstrate that coherent oscillations can be resolved out to $|\Delta| > 2\Omega$. Multiple sets of measurements, on both the $|\pm 1\rangle \leftrightarrow |A_2\rangle$ transitions, enable $T_{2\pi}$ to be calibrated as a function of both Ω and Δ .

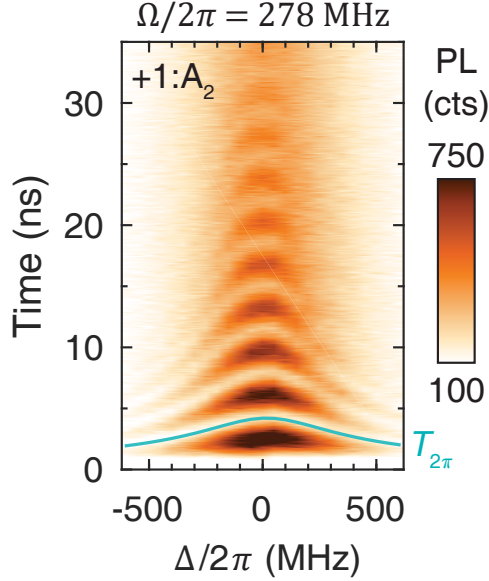


Figure 3.16: Rabi oscillations between the NV center ground and excited states. A near-resonant optical pulse is applied to the $|+1\rangle \leftrightarrow |A_2\rangle$ transition, and as in Fig. 3.5, PL is collected during the pulse proportional to $|A_2\rangle$ occupation. Data is taken as a function of static detuning Δ , revealing generalized Rabi oscillations where frequency increases and amplitude decreases with $|\Delta|$. This data set, and equivalent ones for the $| -1\rangle \leftrightarrow |A_2\rangle$ transition, are used to calibrate the 2π pulse time $T_{2\pi}$ for a single holonomic gate (cyan line).

While resonant holonomic gates alone do not constitute a complete set of single-qubit gates, as they only perform π rotations, they are nevertheless frequently-used operations and create several elements of the important Clifford gate set. Fig. 3.17 shows the results of initializing the $|z\rangle$ ($\equiv |-1\rangle$) state, performing gates with the full range of θ values, and measuring the resulting $|z\rangle$ and $|-z\rangle$ states (Fig. 3.17a). These operations are described by $\hat{U}(\theta, 0, 0, 2\pi \times 150 \text{ MHz})$, which include the ubiquitous σ_x (X , $\theta = \pi/2$) and Hadamard (H , $\theta = \pi/4$) gates, and are completed in $T_{2\pi} = 6.7 \text{ ns}$. As θ increases, the final state undergoes oscillations between $|z\rangle$ and $|-z\rangle$. These oscillations deviate from the ideal, dissipation-free evolution, but are recreated by a master equation model which includes optical lifetime (T_1) and dephasing (T_ϕ) effects. Despite the short duration of the gates, T_1 is still of the same order as $T_{2\pi}$, and so dissipation has a noticeable effect.

The deviations introduced by dissipation lead to an experimental asymmetry in place

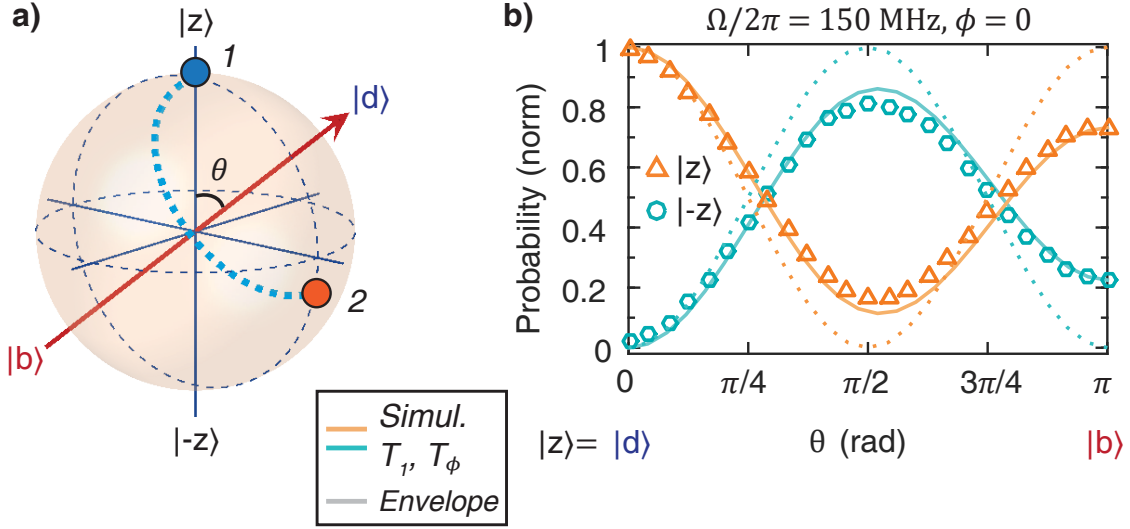


Figure 3.17: Resonant holonomic gates as θ is varied. (a) The system is initialized in $|\pm z\rangle$, and resonant (π rotation) holonomic gates are applied around axes with a range of θ values. (b) The $|\pm z\rangle$ populations after the holonomic gate are measured. At $\theta = \pi/2$, the gate ideally swaps $|\pm z\rangle$, and at $\theta = \pi$, the gate is an identity operation for the initial states, although they are fully aligned to the bright state. The increasing bright state overlap is responsible for the loss of contrast with increasing θ , departing from the ideal gate (dotted lines). A master equation simulation incorporating T_1 and T_ϕ reproduces observations well (solid lines).

of a theoretical symmetry. In principle, $\hat{U}(\theta) = \hat{U}(\theta + \pi)$, as π and $-\pi$ operations are equivalent. However, as the bright and dark states are exchanged when $\theta \rightarrow \theta + \pi$, the gates no longer show the same performance if dissipation differently affects the two states. Indeed, Fig. 3.17b shows that $\theta = 0$, where the initial state is the dark state, experiences much less loss than at $\theta = \pi$, where the initial state is the bright state. This observation suggests that in systems where intermediate state dissipation is significant, and in scenarios where the input state is known, selecting the holonomic gate with a smaller $|\langle \psi_i | B \rangle|^2$ will demonstrate higher fidelity.

Next, θ is held fixed and ϕ is varied, to demonstrate full control over the axis of rotation. This set of gates, described by $\hat{U}(\pi/2, \phi, 0, 2\pi \times 220 \text{ MHz})$, includes the σ_y gate (Y) where $\phi = \pi/2$ (Fig. 3.18a). The higher optical power lowers the gate time to $T_{2\pi} = 4.5 \text{ ns}$. Fig. 3.18b shows the results with an initial state $|x\rangle = (|-1\rangle + |+1\rangle)/\sqrt{2}$. The X and Y projections of the Bloch sphere are measured after the gate, since the final state ideally lies

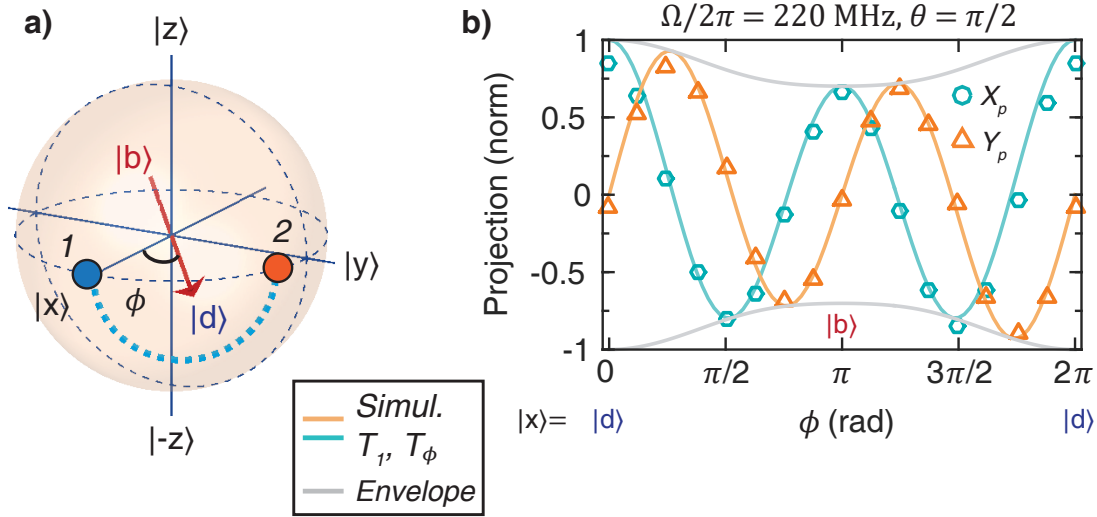


Figure 3.18: Resonant holonomic gates as ϕ is varied. The system is initialized to $|x\rangle$, and resonant holonomic gates with $\theta = \pi/2$ and a range of ϕ values are applied (left). The resulting $|x\rangle$ and $|y\rangle$ projections are then measured. At $\phi = \pi$, the gate is an identity operation on $|x\rangle$, while fully aligned to the bright state. As in Fig. 3.17, this bright state overlap reduces contrast, leading to the observed envelope (gray lines). A master equation simulation incorporating T_1 and T_ϕ reproduces observations well (cyan and yellow lines).

on the Bloch sphere equator. Two full periods of oscillation are observed, since a π rotation about an axis at $\phi = \phi_{ax}$ transfers the initial state at $\phi = 0$ to $\phi = 2\phi_{ax}$. As in Fig. 3.17, the visibility of the oscillations decreases as the bright state comes into alignment with the initial state, in agreement with simulation. However, due to the shorter gate, the losses are reduced and the minimum visibility is improved.

After demonstrating control over the axis of rotation and an understanding of the impact of dissipation, determining gate fidelities will help complete the assessment of the holonomic gate set's performance. Quantum process tomography (QPT) is an appropriate method to thoroughly characterize the effect of the holonomic gates. QPT follows from the description of a quantum gate as a transformation of the initial density matrix (ρ_i) to a final density matrix (ρ_f):

$$\rho_f = \sum_{j,k=1}^4 \chi_{jk} E_j \rho_i E_k. \quad (3.15)$$

In this decomposition, the basis operators are $E_j \propto \{\hat{I}, \boldsymbol{\sigma}\}$. χ is the process matrix, which

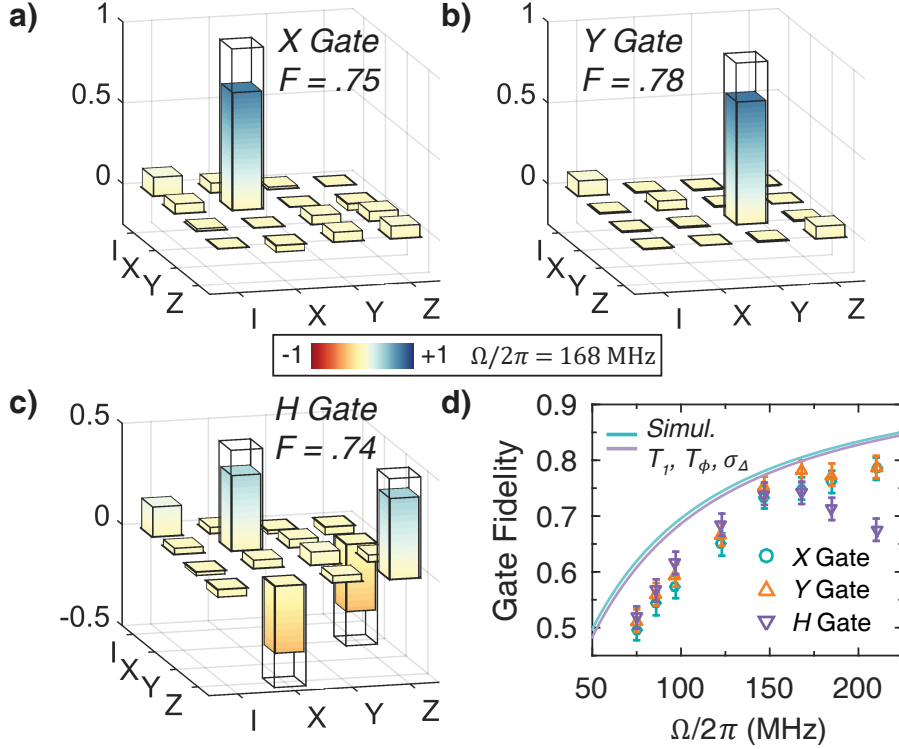


Figure 3.19: Quantum process tomography of resonant holonomic gates. (a-c) The real values of the experimentally determined process matrices for X , Y , and H gates are shown, along with the calculated process fidelities. Hollow bars indicate the ideal values for unity fidelity. The experimental and optimal imaginary values are negligible. (d) Experimental and simulated gate fidelities as Ω is increased. Dissipation is responsible for the majority of losses, as captured in simulation, so faster gates demonstrate higher fidelity. The remaining discrepancy is attributed to crosstalk between optical transitions, which may also explain the decrease in H fidelity at large Ω .

describes the independent basis operator contributions to the density matrix evolution. In the single-qubit case, χ has sixteen elements labeled by E_j and E_k , each complex-valued and with magnitude ≤ 1 . By measuring the initial and final density matrix elements, the relationship can be inverted to determine the elements of χ . The gate's process fidelity is computed via the differences between the experimental and theoretical process matrix elements. Additional information about QPT is available in Appendix B.

Fig. 3.19 shows the results of QPT for the X , Y , and H gates. The bars display the real components of all elements of χ for each gate at $\Omega = 2\pi \times 168$ MHz, as well as the theoretically optimal values. The imaginary components of each are negligible and

are theoretically zero. The resonant gates achieve maximum fidelity of between 0.74 and 0.78 ± 0.02 . Master equation simulations are again important to understanding the errors present, and incorporate the measured spectral diffusion as well as dissipative effects. As Ω is increased, the average gate fidelity rises from approximately 0.5 to its maximum, an improvement which is matched by the simulations. However, an additional decrease of ~ 0.05 is evident in the experimental data, and indicates the presence of other appreciable experimental errors, which may include unintentional excitation due to higher harmonic sidebands of the laser and crosstalk between the two transition. It is also noteworthy that while the X and Y gate fidelities continue to increase for $\Omega > 2\pi \times 168$ MHz, the fidelity of H begins to drop. This is not well-understood, and may result if the H gate is more susceptible to systematic errors such as the nonzero rise time as the optical coupling strength is increased.

3.3.4 Detuned Gates

In this section, following the theoretical guidelines proscribed in section 3.3.2, the unique features of detuned holonomic gates are demonstrated. A straightforward starting point is the phase gate $Z(\gamma)$, described by $\hat{U}(0, 0, \gamma, \Omega)$. This set includes the standard σ_z , S , and T operations, which corresponding to $\gamma = \pi$, $\pi/2$, and $\pi/4$, respectively. The gate time in this case is not constant, but decreases with increasing detuning, as marked in Fig. 3.16. Fig. 3.20 shows the experimental phase accumulated between the initial and final states as a function of the fixed detuning. This phase accumulation is measured from the X and Y projections of the Bloch sphere, and is averaged over $|x\rangle$ and $|y\rangle$ input states. Data for both $\Omega = 2\pi \times 134$ MHz and $2\pi \times 252$ MHz are plotted, along with curves that are not fits, but rather predictions of the accumulated geometric phase as a function of detuning from Eq. 3.14. The close agreement between predictions and observations, for two different values of Ω and extending to $|\Delta| \sim 4\Omega$, provides strong evidence for the geometric origin of the observed phases. It also suggests that the ~ 1 ns rise time does not have a large effect on the dynamics at these optical powers, as the data match predictions even for large detuning

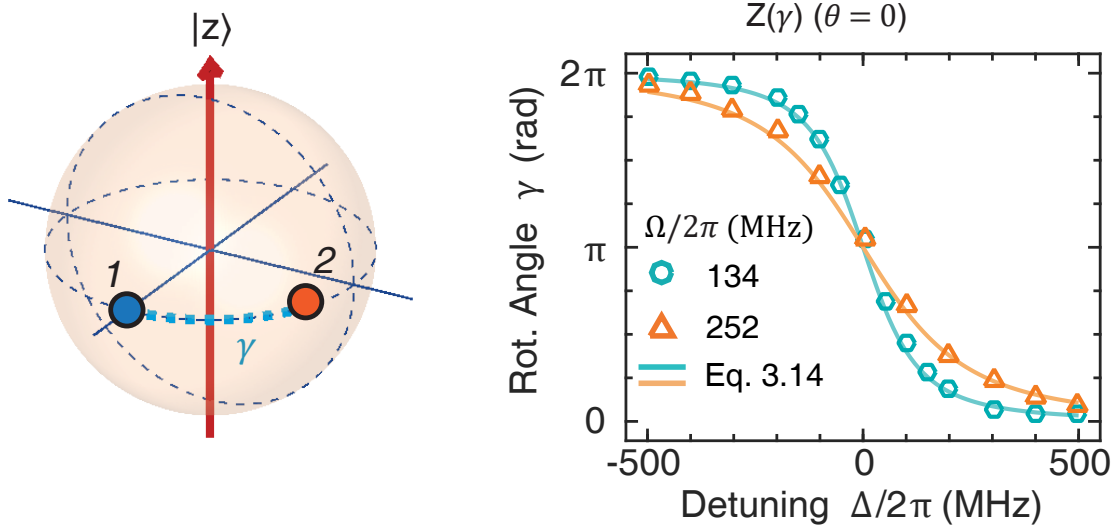


Figure 3.20: Detuned holonomic phase gates. The system is initialized in $|x\rangle$ and $|y\rangle$ states, after which a phase gate is applied, with rotation angle γ depending on Δ . The resulting X and Y projections are measured to determine the phase shift, for two different Ω . The observed phase shifts match predictions, as the solid lines show Eq. 3.14 given the experimentally measured Ω and Δ .

where the gate times become very short.

Quantum process tomography of the Z gate family reveals that the off-resonant Z gates achieve higher fidelities than on resonance. Fig. 3.21a shows the rapid improvement in fidelity with detuning for both optical coupling strengths, reaching ~ 0.95 for $\gamma \sim \pi/4$. As with errors in the previous protocols, dissipation is a likely culprit, a hypothesis which is supported by the similarities between the fidelity dip and the optical linewidth, as well as the improvement of the resonance fidelity with increasing Ω . Once again, the master equation simulation captures the quantitative and qualitative trends of the experiment by incorporating the sources of dissipation. In Fig. 3.21b, the simulations separate out the relative contributions of each dissipative or noise factor, and by including spectral diffusion closely match the experimental results for resonant Z gates up to $\Omega = 2\pi \times 252$ MHz, where the fidelity peaks at 0.86 ± 0.02 . At the highest optical power, nearly $2\pi \times 400$ MHz, the fidelity begins to decrease again, presumably due to power-broadened driving of other nearby transitions and the increasingly non-ideal shape of the pulse.

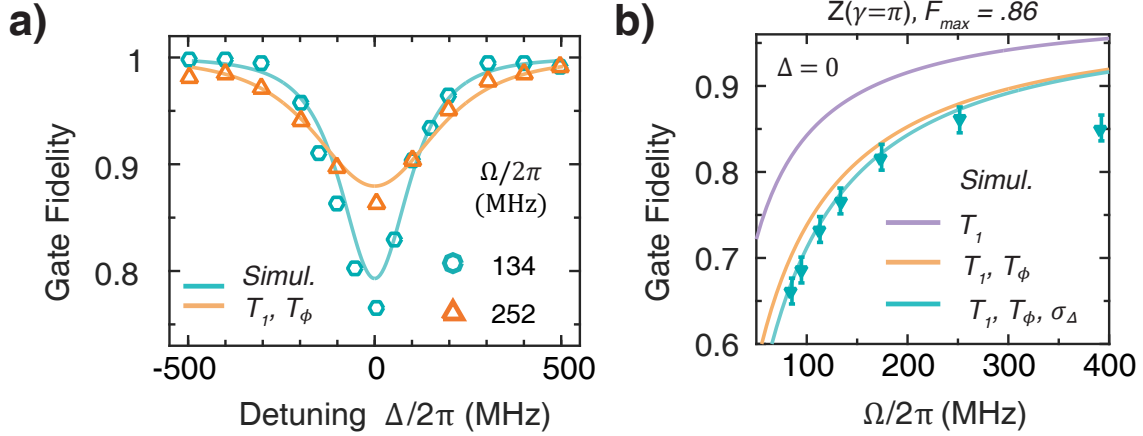


Figure 3.21: Fidelity of the detuned holonomic phase gates. (a) Quantum process tomography is used to determine the fidelity of Z gates as a function of Δ . As $|\Delta|$ is increased, excited state population decreases, resulting in higher fidelities. Faster gates also experience less dissipation, corresponding to higher fidelities for $\Omega = 2\pi \times 252$ MHz on resonance than $\Omega = 2\pi \times 134$ MHz. (b) This trend is observed in resonant gates of increasing Ω , as well as the master equation simulations (solid lines). Successively incorporating T_1 , T_ϕ , and σ_Δ effects matches observations, until large Ω is reached where the square pulse approximation begins to break down.

In the final set of operations, off-resonant gates around the x and y axes are demonstrated and characterized. These gates are described by $\hat{U}(\pi/2, 0, \gamma, 2\pi \times 152 \text{ MHz})$, and include the $X(\pm\pi/2)$, $Y(\pm\pi/2)$ operations that complete the set of Clifford gates for a single qubit. By choosing $|z\rangle$ as the initial state, the final state traces out increasing transfer to the opposite pole, peaking at a complete transfer on resonance (Fig. 3.22a). The final populations are shown in Fig. 3.22b, along with the expected results from dissipationless evolution (dotted lines) and matched well by the simulations incorporating the major sources of loss.

Performing QPT on the detuned X and Y gates provides more insight into the advantages of off-resonant gates in the presence of significant dissipation. Fig. 3.23 plots the process matrix values for $X(\pi/2)$ (including the imaginary components) and the real components of $Y(\pm\pi/2)$. These gates are realized at $\Delta/\Omega = \pm 1/\sqrt{3}$ for $\Omega = 2\pi \times 152$ MHz, and with $T_{2\pi} = 5.7$ ns. In contrast to the resonant gates, these off-resonant gates achieve average fidelities of 0.81 ± 0.02 . More than any other demonstration described, this data illustrates the superiority of off-resonant gates. The same $Y(\pi/2)$ could be implemented via two resonant

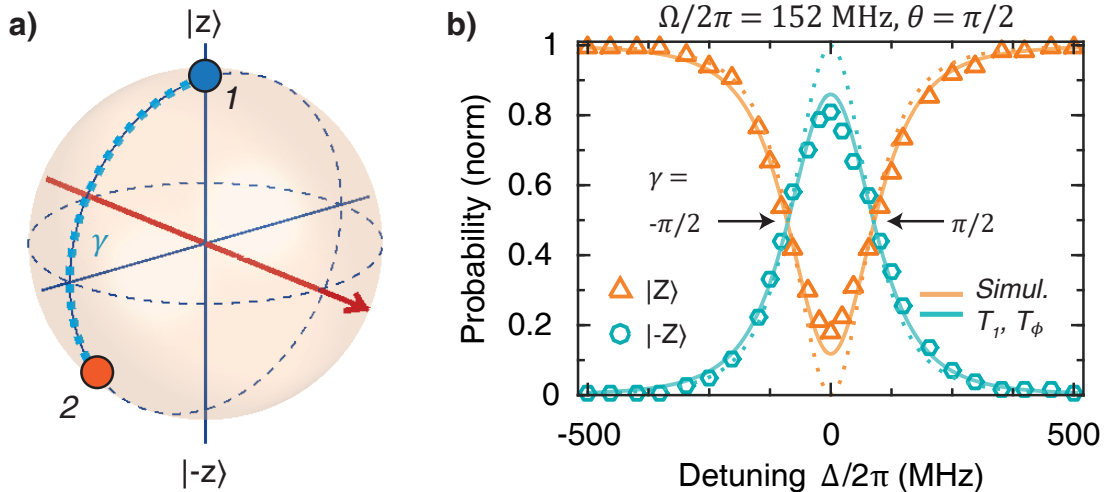


Figure 3.22: Detuned holonomic X gates. After initializing $|\pm z\rangle$, X gates, with $\theta = \pi/2$ and $\phi = 0$, are performed for various Δ , and the final populations are measured. On resonance, the holonomic gate performs a π rotation, with contrast reduced from the ideal (dotted lines) by dissipation. A master equation simulation (solid lines) matches observations. This set of gates includes the important $\pm\pi/2$ rotations, denoted in the plot.

gates, since $Y(\pi/2) = X(\pi) \cdot H(\pi)$. However, as observed in section 3.3.3, the combined fidelity of the two concatenated resonant gates would fall to $0.74 \cdot 0.75 = 0.56$, far below the single detuned gate.

For the all-optical gates presented here, the fidelities fall below 0.9, and below what is currently required for error-correction gate thresholds. Several experimental improvements could increase these fidelities. Faster pulses (generated by higher-bandwidth optical modulators) would reduce excited state occupation, but would also require faster modulator rise times and high extinction ratios to ensure the pulse shape is still approximately square. Simulations indicate that the dynamic phases resulting from nonzero rise and fall times will limit the fidelity when $\Omega > 2\pi \times 600$ MHz. Noting that fidelities at high optical powers are already exhibiting nonideal behavior in the data presented here, reducing these sources of error is a high priority. If two independent lasers were used instead of sidebands generated by the PEOM, the higher-harmonic sidebands could be eliminated, and the polarization selectivity of the $|\pm 1\rangle \leftrightarrow |A_2\rangle$ transitions could be leveraged to reduce crosstalk.

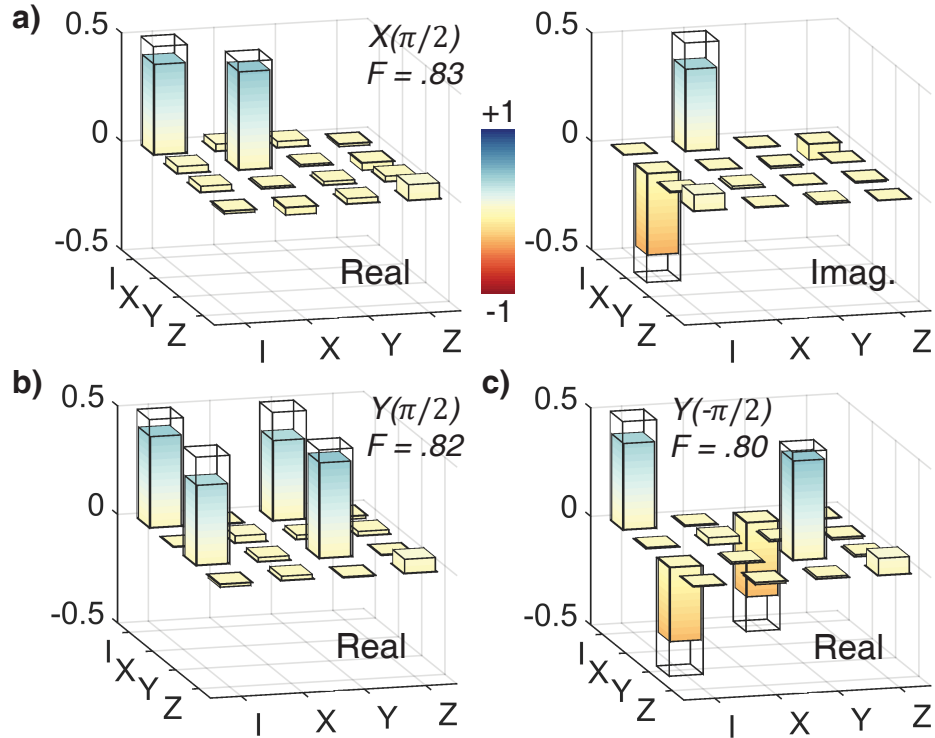


Figure 3.23: Quantum process tomography of detuned X and Y gates. (a) The real and imaginary experimental process matrix elements of the $X(\pi/2)$ gate are shown. Hollow bars indicate ideal values. (b-c) Experimental process matrices for the $Y(\pm\pi/2)$ gates; imaginary values are ideally zero.

3.3.5 Conclusion

This section presents an implementation and thorough exploration of the behavior and fidelities of non-adiabatic, non-Abelian geometric phases using the all-optical Λ system of a single NV center in diamond. The all-optical control enables short gate times of less than 10 ns and the capability of addressing a single defect among many in a sample. Unlike resonant-only gates previously demonstrated, adding control of detuning enables arbitrary single-qubit unitary gates with a single holonomic loop. Due to the rapid dissipation of the excited state, comparable even to the nanosecond-scale gate times, the resonant gate fidelity is reduced and detuned gates show improved fidelities. This provides a compelling motivation to explore detuned holonomic gates for systems with lossy intermediate states, such as hybrid quantum systems or states coupled via waveguides. These all-optical gates on spin qubits

can be extended to cavity-mediated two-qubit gates [93], or may be useful for optically-active defects with longer excited state lifetimes or that have difficulty implementing ground-state microwave control due to magnetically forbidden transitions.

Chapter 4

Spatiotemporal Photocurrent Mapping in Molybdenum Disulfide

This chapter and figures are adapted from work published in reference [3].

Even after decades of research, exciting new developments in NV center sensing are still emerging. While much of the research activity has been focused on the premier magnetometry applications, other sensing modes offer unique opportunities advance research capabilities and provide insights into important physical systems. This chapter covers a novel demonstration of NV centers in the detection of photocurrents in the monolayer semiconductor MoS₂. This technique combines well-established ac magnetometry sensing sequences with synchronized, pulsed photoexcitation to reach new sensitivity regimes in NV-based current detection. The advance in sensitivity enables the discovery of circulating photocurrent vortices in the monolayer sample, which is challenging for conventional photocurrent detection methods as a result of the zero net current of the feature. These improvements in NV photocurrent sensing protocols may lead to greater understanding of the light-matter interactions and material properties of the monolayer and thin-film materials under study in QISE and condensed matter physics, and the sensing protocol can be generalized to yet further novel approaches to NV-based sensing in future demonstrations.

4.1 Electric Current Sensing with NV Centers

Electric current is one of the most fundamental quantities across all of science and engineering. Innumerable inventions have been created to measure, induce, and direct currents, and it is in turns a tool and an object of study in scientific research. Currents can arise from many sources, driven by voltages that result from chemical storage, semiconductor band gap engineering, thermal gradients, or other phenomena. Each of these sources of current has been – and continues to be – studied to understand exactly how currents arise from the underlying materials and interactions. With sufficient understanding, current measurements can then be reversed, and themselves used to probe the host material properties and new device physics. Thus, an extensive toolkit of methods for measuring current is an important prerequisite for both fundamental and applied sciences.

While NV centers will never compete with conventional ammeters and oscilloscopes in terms of flexibility and sensitivity, their unique properties are useful for measuring current in certain geometries and samples. In the most common current measurements, a sample or device closes the connection between two electrical leads, and any generated current may be collected in the completed circuit for measurement by a lab instrument. This transport measurement approach is straightforward and flexible. However, it suffers a few drawbacks. The aggregate current does not include details of the local distribution of current within a sample, and so cannot measure net-zero-current features like eddies and vortices without creative electrical contact placement or sample engineering. Furthermore, it depends on being able to extract the total current from a sample. In some cases this is almost trivial, but not all samples readily support ohmic contacts for efficient carrier collection. Both these challenges take on greater importance in studies of nanoscale electronic devices and quantum materials. With miniaturized or nanofabricated components, it becomes difficult to electrically isolate segments of the sample to test for defects or nonideal behavior. In addition, introducing metallic probes into high purity quantum materials or nanostructures risks disrupting or overwhelming the very phenomena of interest with stray capacitances,

voltage and magnetic fluctuations, and quantum tunneling effects. As an alternative, NV center sensors can provide contactless, local current sensing based on the stray magnetic fields produced by flowing electric currents.

The NV current sensing technique is a specialized use of NV magnetometry. A current flowing in a sheet, wire, or nanostructure produces a corresponding magnetic field. As described in section 2.2.3, NV centers are most sensitive to magnetic fields and can even reconstruct vector magnetic fields. By performing these measurements in the vicinity of electric currents, NV centers can be used to infer the distribution of current which produced the measured fields [94, 95]. This concept alone is not new, and describes the working principle of Hall effect and other magnetic current sensors. However, NV centers provide higher spatial resolution than any other detection method. The most comparable technique¹, scanning SQUID magnetometry, does achieve high sensitivities and spatial resolution [96], but utilizes a metallic probe which must be cooled to cryogenic temperatures. NV center sensors can reach nanometer resolution under ambient conditions and provide an all-dielectric environment. In addition, NV centers are naturally compatible with optical excitation, while scanning SQUID magnetometers require exciting thicker samples from the back surface [97–99]. Recent experiments have used wide-field measurements on NV ensembles to map current flow through graphene nanoribbons [55] and used a single scanning NV center to observe hydrodynamic current flow in graphene [100]. These investigations have reported current densities down to $\sim 1 \mu\text{A}/\mu\text{m}$ with $\sim 50 \text{ nm}$ spatial resolution, demonstrating the impressive capabilities of NV centers for current sensing.

For any magnetometry-based current sensing technique, an essential step is inferring the originating current distribution from the measured magnetic field. Depending on the specifics of the NV sensor, this can be accomplished in one of several ways. For 2D (implanted) ensembles with a mix of NV orientations, ODMR measurements provide full vector magnetic

1. Ref. [54] provides a comparison of the advantages and disadvantages of several techniques which are relevant for nanoscale magnetic current sensing.




Figure 4.1: NV sensing concept. A superposition state, such as $|\psi_i\rangle=|x\rangle$, is initialized for some baseline qubit energy splitting. A static or time-varying signal that affects the energy difference of the qubit causes the state to accumulate a phase ϕ in the rotating frame (here, negative precession is portrayed). At the end of the sequence, the final state $|\psi_f\rangle$ is measured to determine the acquired phase. This phase is then correlated to the physical parameter of interest, such as via the gyromagnetic ratio, which relates external magnetic fields to qubit energy.

field information at the depth of the ensemble, which can be inverted to give the current distribution [55]. However, even if the magnetic field can only be measured along a single axis, a 2D map of the magnetic field will still be sufficient to recreate the current distribution if it is also two-dimensional [101]. This is the method employed by scanning single-NV sensors.

The remainder of this chapter describes the use of near-surface NV center ensembles to measure photocurrents in the monolayer semiconductor MoS₂. Photocurrents are electrical currents induced by the absorption of light, and form the basis of photodetector and photovoltaic technologies [102]. Photodetectors, which measure the intensity of light, are core pieces of fast fiber optic communication networks; photovoltaics are the basis for renewable solar energy. Beyond these critical modern technologies, photocurrents provide insight into material properties as a research tool. The response of a sample to light of various energy and polarization characterizes the substrate’s band structure, transport properties, and more novel phenomena such as valley degrees of freedom [103, 104]. Thus, photocurrent measurement techniques are a building block for designing novel materials and improving photodetector and photovoltaic devices.

The NV-based photocurrent measurements presented here are also magnetometry tech-

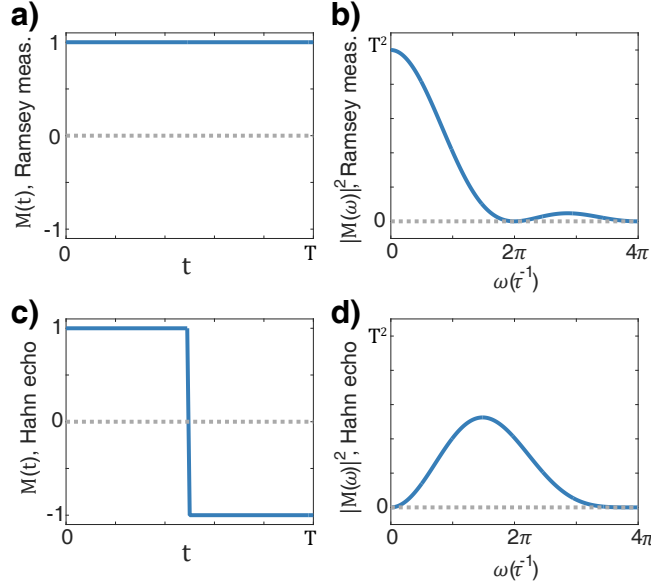


Figure 4.2: Time- and frequency-domain filter functions for sensing sequences. T is the total time of the sequence, and $\tau = T/2$. (a) Time-domain filter function for a Ramsey measurement; it is constant over the measurement duration. (b) Frequency-domain filter function for a Ramsey measurement. It is most sensitive to dc signals. (c) Time-domain filter function for a spin echo measurement; it inverts when a π pulse is applied. (d) Frequency-domain filter function for a spin echo measurement. Low-frequency signals are suppressed.

niques, but are not based on ODMR measurements of the NV level splittings. Instead, small shifts of the levels are resolved by accumulating a coherent phase as a result of Ramsey or spin echo pulse sequences. These pulse begin by initializing a superposition of two spin states, for example $|0\rangle$ and $| -1\rangle$, and measuring the phase difference between the initial and final states, as shown in Fig. 4.1. If the sample produces no magnetic field, the superposition remains in the initial state (in the rotating frame) and acquires no phase. However, a magnetic field will shift the qubit energy, causing the state to precess around the Bloch sphere equator. The accumulated phase ϕ is proportional, via the gyromagnetic ratio γ_e , to the average magnetic field throughout the measurement time. Quantitatively,

$$\phi = \int_T 2\pi\gamma_e B_{\parallel}(t)M(t)dt, \quad (4.1)$$

for a measurement of time T , and B_{\parallel} the component of the field aligned with the NV

center's axis. M is the time-domain modulation or filter function, which takes on values $\{-1, 1\}$ during the measurement depending on the spin operations applied to the state, and is 0 otherwise. For the Ramsey measurement, $M(t) = 1$ for all $0 \leq t \leq T$ (Fig. 4.2a). The filter function has a corresponding frequency-domain expression (Fig. 4.2b), obtained through the Fourier transform of the time-domain filter function [57]:

$$|\tilde{M}(\omega)|^2 = \left| \int e^{-i\omega t} M(t) dt \right|^2. \quad (4.2)$$

The corresponding Ramsey filter function shows that the sequence is most sensitive to dc signals, but does also respond to signals up to $\sim \frac{2\pi}{T}$. To resolve higher-frequency components, a spin echo can be added to flip the modulation function (Fig. 4.2c), which alters the frequency-domain filter function accordingly (Fig. 4.2d). By extrapolating to sets of periodic spin flips, the peak sensitivity frequency can be increased further, matching the periodicity of the spin flips.

The spin echo sequences offer multiple benefits for more sensitively detecting photocurrents. Repeated spin flips effectively implement dynamical decoupling, which suppresses low-frequency noise (see Fig. 4.2d, again) that simultaneously decoheres the superposition and obscures small photocurrent signals. At higher pulse numbers, the longer spin coherence time leads to a better sensitivity ($\propto 1/\sqrt{T_2}$) and a narrower detection bandwidth.

To take advantage of these characteristics, ac photocurrents must be generated instead of dc. In the novel photocurrent detection protocol, the excitation laser is pulsed at the periodicity corresponding to the sensing sequence maximum sensitivity. Unlike noise spectroscopy, which measures signals from noise sources that match in frequency but are incoherent, the resulting photocurrents can be phase-matched to the detection sequence, producing a larger coherent response and signal-to-noise. This technique is shown to be effective for resolving photocurrents in MoS₂.

MoS₂ is one of the class of transition metal dichalcogenides (TMD), which are 2D materi-

als with strong light-matter interactions. Bulk TMD samples consist of layers that are bound by van der Waals forces, but these individual layers, like graphene, are found to be stable as isolated monolayers. Unlike graphene, a direct band gap emerges for few-layer TMDs as a result of quantum confinement. At the band gap (~ 2 eV for most examples), strong absorption occurs at exciton resonances. These excitons have been measured to have large, ~ 0.3 eV binding energies due to the limited dielectric screening in the 2D host material. These intense interactions hosted by an atomically thin material have made TMDs a candidate for high-bandwidth photodetectors, high-efficiency photovoltaics, and flexible electronics. Resolving the details of photocurrent generation and behavior in TMDs is therefore pertinent to developing these applications.

4.2 Experimental Setup

In order to be an effective local probe of photocurrents, the NV center sensors must be brought into proximity with an MoS₂ sample. While this can be accomplished with an AFM tip hosting a single NV center, here a bulk diamond chip containing an ensemble of near-surface NV centers is used, similar to previous graphene measurements [55]. The MoS₂ is placed on the surface of the diamond, and coordinated photocurrent excitation and spin measurements can be made due to the suitable spectral behaviors of the sensor and sample (Fig. 4.5). This approach can be employed for other two-dimensional (2D) materials and thin films with relatively low PL in the NV emission band.

First, two diamond samples were prepared with NV ensembles. The starting substrates were 2x2x0.5 mm IIa diamond chips obtained from Element Six. These samples are (100)-oriented and initially contained < 5 parts per billion residual nitrogen. To create the NV ensembles, ¹⁵N ions were implanted with an area dose of 10^{12} cm⁻² (CuttingEdge Ions, LLC), which results in approximately 85 NV centers per optical spot size. The implantation energy was chosen to be 30 keV (sample 1) and 15 keV (sample 2), which produces a layer

of nitrogen roughly 40 nm and 20 nm below the surface, as calculated via the Stopping and Range of Ions in Matter program. These depths provide two options in the trade-off between proximity and coherence; shallower NVs are closer to the sensing target, but experience more noise from the surface. Ultimately, the deeper sample was selected to gather the majority of data presented here, since the current features vary over much longer lengthscales than the NV depth, which reduces the importance of proximity. Sample 2 showed similar results to sample 1, however. The implantation was also performed at a 7° angle to the surface normal, in order to avoid ions channeling along crystal planes and traveling deeper into the sample than anticipated.

After ion implantation, the NV centers were created by annealing. First, an anneal at 850°C for 4 hr under argon forming gas formed the NV centers by mobilizing vacancies to combine with implanted nitrogen. The NV ensemble spin coherence was measured to be $\sim 8\ \mu\text{s}$. Next, a second, higher-temperature anneal was performed with the goal of reducing the density of non-NV paramagnetic defects which contribute to decoherence [105]. This anneal entailed 6 hr at 800°C followed by 2 hr at 1100°C , again in argon forming gas. Finally, a 2 hr bath in a 1:1 mixture of sulfuric and nitric acids, heated to 225°C , oxygen-terminated the surface and helped remove graphitized surface material. The resulting average T_2 times increased fourfold, to $34\ \mu\text{s}$.

Though often obtained via exfoliation, MoS_2 can now be grown in high-quality monolayers at wafer scales. Here, just such a monolayer was grown via metal-oxide CVD on a SiO_2/Si substrate via previously-established techniques [106]. The monolayer, polycrystalline samples display uniform electrical properties and an average grain size of several microns. Next, poly(methyl methacrylate) (PMMA) was spun on the MoS_2 for 60 s at 3000 RPM, followed by a 10-min. bake at 180°C . Thermal release tape was applied and used to peel the monolayer from its substrate.

The MoS_2 was transferred to the diamond substrate via a vacuum transfer technique [107] to ensure a smooth, clean interface. The diamond was heated to 100°C in a vacuum chamber

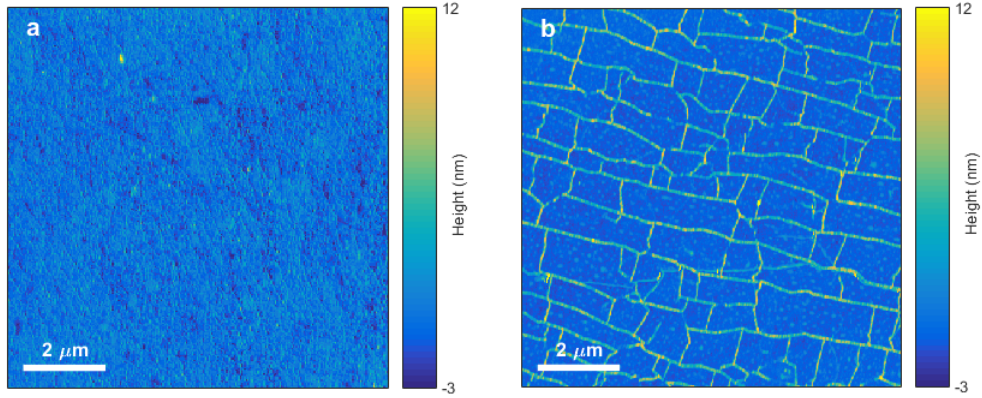


Figure 4.3: Atomic force microscope images of MoS₂ on a diamond substrate. (a) Atomic force microscope (AFM) image of an MoS₂ sample on diamond that had never been cooled down. Surface roughness is 1.2 nm, compared to 0.3 nm roughness for bare diamond. This sample had been exposed to air, with oxidation potentially increasing the roughness. (b) AFM image of a different MoS₂ sample on diamond, at room temperature, shortly after a cryogenic cycle. Nanoscale wrinkles appear at room temperature due to differences in thermal expansion coefficients during the warming process, but are not expected to be present at low temperatures. Surface roughness of the areas excluding wrinkles is 0.6 nm.

held below 70 mTorr, and then an intra-chamber manipulator pressed the PMMA/MoS₂ stack onto the diamond for five minutes to ensure good adhesion. The thermal release tape handle was removed, followed by a 1 hr bath in heated acetone to remove the PMMA. This resulted in a uniform monolayer spanning the diamond surface. Finally, to allow for control measurements, a second piece of thermal release tape was applied to half of sample 1, and used to remove the MoS₂ covering that half of the sample. This produced a sample which could easily be used to compare measurements on an NV ensemble with and without a surface covering of MoS₂.

An atomic force microscope (AFM) image confirms the quality of the MoS₂-diamond surface interface. As shown in Fig. 4.3, the regions covered by MoS₂ have surface roughness between 0.6-1.2 nm r.m.s., increasing as the sample ages and MoS₂ oxidizes. This compares to 0.3 nm r.m.s. for bare diamond, so the MoS₂ conforms nearly perfectly to the sample surface and no evidence of trapped contaminants or bubbles is observed. Fig. 4.3a shows a sample that has never been cooled down, while Fig. 4.3b shows a sample after a cryogenic cycle. Nanoscale wrinkles are evident, approximately 10 nm in height, which have previously

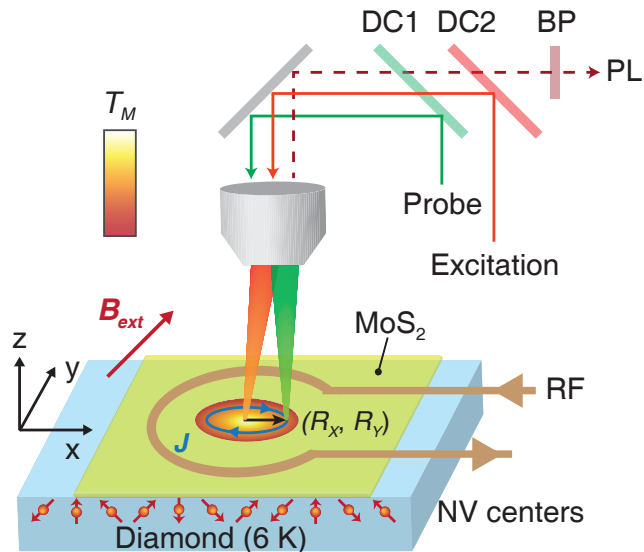


Figure 4.4: Experimental schematic for measuring photocurrents in MoS₂ with NV centers. MoS₂ is transferred to the surface of a diamond with a near-surface ensemble of NV centers. Two independently steerable laser beams are focused through the objective to the sample: a probe beam at 532 nm to read out NV center PL contrast, and an excitation beam at 661 nm to induce photocurrents in the MoS₂. These beams are integrated into the optical path with dichroic mirrors (DC), while bandpass filters (BP) restrict the PL collection to 685-830 nm. A suspended wire loop (RF) delivers microwave fields across the sample. An external magnetic field (\mathbf{B}_{ext}) is aligned with the [111] direction, addressing a subensemble of NV centers. The sample is cooled to 6 K in a helium cryostat.

been observed for 2D materials on substrates with significantly different thermal expansion coefficients [108]. The wrinkles appear at room temperature when warming from low temperature, as the MoS₂ expansion is much greater than diamond, but are not expected to be present during the initial cooldown process when the MoS₂ contracts more than its substrate.

After successfully transferring the MoS₂ monolayer to the diamond, the sample is placed in the liquid helium cryostat and cooled to 6 K. The same custom confocal microscope provides optical access, but instead of coplanar waveguides, microwaves are delivered via a wire coil approximately 700 microns in diameter, and suspended ~ 100 microns above the surface (Fig. 4.4). The permanent magnet, on a 1-in.-travel stage, provides magnetic fields between 100 and 500 G, with the misalignment to one of the four NV axes ranging from 1-4° and typically $\sim 2^\circ$. The 650 nm dichroic, previously used to couple in NV-resonant lasers at 637 nm, is replaced with a 685 nm dichroic to couple in optical excitation for the

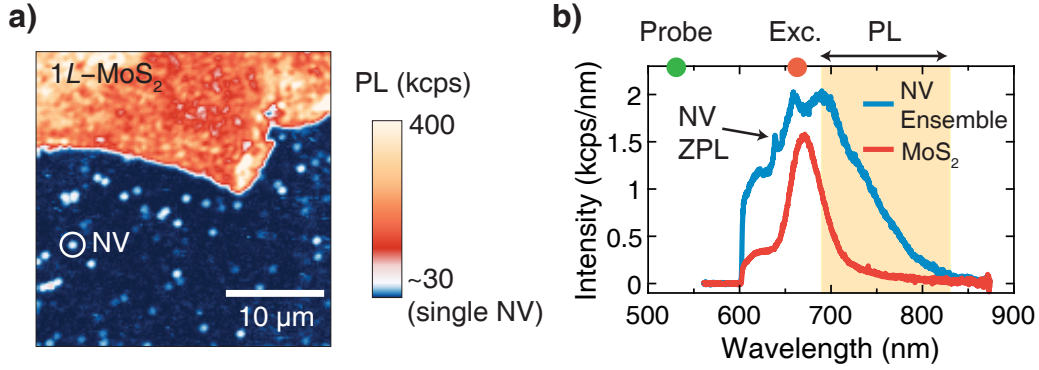


Figure 4.5: Optical characterization measurements of MoS₂ and NV centers. At left, a PL map of the edge of an MoS₂ monolayer on a diamond with sparse NV centers is shown. The MoS₂ PL overwhelms that of single NV centers, necessitating the use of ensembles. At right, spectrometer measurements of two samples are shown under 532 nm excitation: an NV ensemble with the same density as used to measure photocurrents, but with no MoS₂ on the sample, and an MoS₂ monolayer on a silicon substrate. The MoS₂ excitation wavelength (Exc.) is near the peak of MoS₂ PL, but longer than the 637 nm required to excite NV center optical transitions (NV ZPL). The PL measurement band from 685-830 nm is highlighted, showing that the NV phonon sideband dominates the collection. A longpass filter was used at 600 nm to remove 532 nm excitation. All measurements shown here are taken at room temperature.

MoS₂ at >660 nm. For this excitation pathway, an AOM (Gooch & Housego) provides initial gating, while a free-space EOM (ConOptics) with an 8-ns rise time performs fast switching in concert with a Glan-Thompson polarizer.

The optical properties of NV centers and MoS₂ overlap somewhat, but are different enough to provide independent means of addressing each. Fig. 4.5a shows a sample (not used for any later data) which juxtaposes single NV centers and the edge of monolayer MoS₂ on the diamond surface. PL collected under CW 532 nm excitation demonstrates that MoS₂ fluorescence, nearing 400 kcps, overwhelms the signal of single NV centers at ~30 kcps. This explains the need for ensemble samples, where the enhanced PL will not be dominated by broad MoS₂ emission. Fig. 4.5b shows separate spectra taken of the NV ensemble alone and an MoS₂ sample still on a silicon substrate. These spectra are taken at room temperature with 532 nm excitation, and indicate how each sample is addressed in the same setup. 532 nm light is partially absorbed by the MoS₂, but only up to a few percent of the total optical

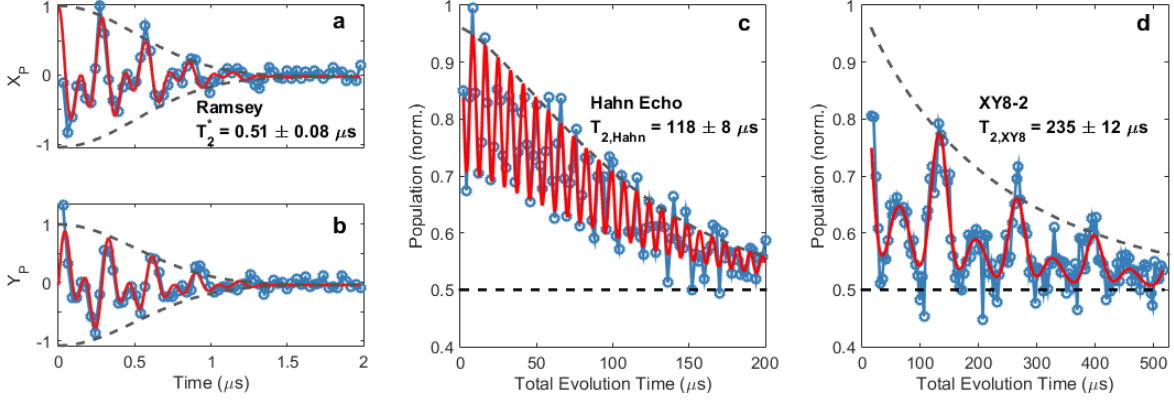


Figure 4.6: Coherence properties of near-surface NV center ensembles. At 6 K, spin measurements are performed on the NV ensemble in proximity to the MoS₂ monolayer. Fits are solid red lines, while the stretched exponential envelopes that extract the coherence time are dashed lines. (a-b) Ramsey measurement X and Y projections with a 5-MHz detuning. (c) Hahn echo measurement, with the fit accounting for ^{13}C interactions. (d) Coherence oscillations during XY8-2 sequence. The fit includes terms for ^{13}C interactions and field misalignment. Sensing measurements are generally performed at $T \sim 120 \mu\text{s}$ to maximize coherence.

power, which can be increased to compensate. The NV ensemble emission is much broader than the PL from the MoS₂ monolayer, so the PL collected in the 685 to 830 nm band is primarily NV fluorescence. This band is used to read out the spin-dependent NV PL. Finally, the MoS₂ sample emits strong PL in its A exciton peak centered around 660 nm, of which a small but manageable fraction extends into the collection window. More importantly, when this transition is pumped resonantly (to excite photocurrents), it should not disturb the NV center spins.

Fig. 4.6 shows the results of spin coherence measurements on NV centers at 6 K in proximity to the surface MoS₂. The Ramsey, Hahn echo, and XY8-2 measurements are fit with stretched exponential envelopes to determine coherence times of $T_2^* = 0.51 \pm 0.08 \mu\text{s}$, $T_{2,\text{Hahn}} = 118 \pm 8 \mu\text{s}$, and $T_{2,\text{XY8}} = 235 \pm 12 \mu\text{s}$, respectively. The XY8-2 sequence involves sixteen evenly-spaced spin echoes around alternating x and y axes. These measurements confirm that the NV spin properties are excellent in the presence of MoS₂, and T_1 measurements (not shown) while illuminating with constant 662 nm (MoS₂-resonant) light confirm that NV absorption of that wavelength is negligible at cryogenic temperatures. After plac-

ing an MoS₂ sample near an NV ensemble and verifying that both MoS₂ excitation and NV center spin readout can function independently, photocurrent measurements can proceed.

4.3 ac Photocurrent Measurements

As mentioned in section 4.1, the photocurrent sensing protocol adapts existing ac magnetometry methods, specifically the XY8 sequence (used in many experiments, for example [109, 110]). Except where otherwise specified, the data presented here are collected on the $\{|0\rangle, |-1\rangle\}$ qubit. Shown schematically in Fig. 4.7, the XY8 sequence is defined by the pattern

$$\left(\frac{\pi}{2}\right)_y - [\pi_y - \pi_x - \pi_y - \pi_x - \pi_x - \pi_y - \pi_x - \pi_y]^N - \left(\frac{\pi}{2}\right)_{\text{var}}, \quad (4.3)$$

where the subscript indicates the axis of rotation, each π pulse is separated by a time τ , and the time separation between the $\pi/2$ pulses and the chain of π pulses is $\tau/2$. The first pulse initializes the $|x\rangle$ superposition state, and the final pulse projects the $\pm X$, $\pm Y$ components of the final state of the system to determine the phase accumulation. The pseudo-alternating axes of rotation decouple the spin from noise in both X and Y components, and also make the sequence robust against pulse errors. The XY8 sequence is standard in performing ac magnetometry measurements, where often the pulse spacing τ is varied until a drop in the X projection is observed, indicating a magnetic signal at frequency π/τ . The number of pulses can be increased, usually in sets of the eight-pulse block of Eq. 4.3, and so the nomenclature XY8- N refers to a sequence with N repetitions of the basic XY8 pattern ($8N$ total π pulses).

To add in a detectable photocurrent signal, 662-nm MoS₂ excitation is pulsed at the same rate as the XY8 spin flips, with laser pulses of duration τ and a wait period of τ in between. For all measurements presented with varying τ , this indicates that both the spin echo separation and the laser pulse duration were matched and swept together. This coordinates the laser-induced signal with the XY8 modulation function, ensuring that it is possible to observe a signal; for example, if the pulse duration was equal to 2τ , the

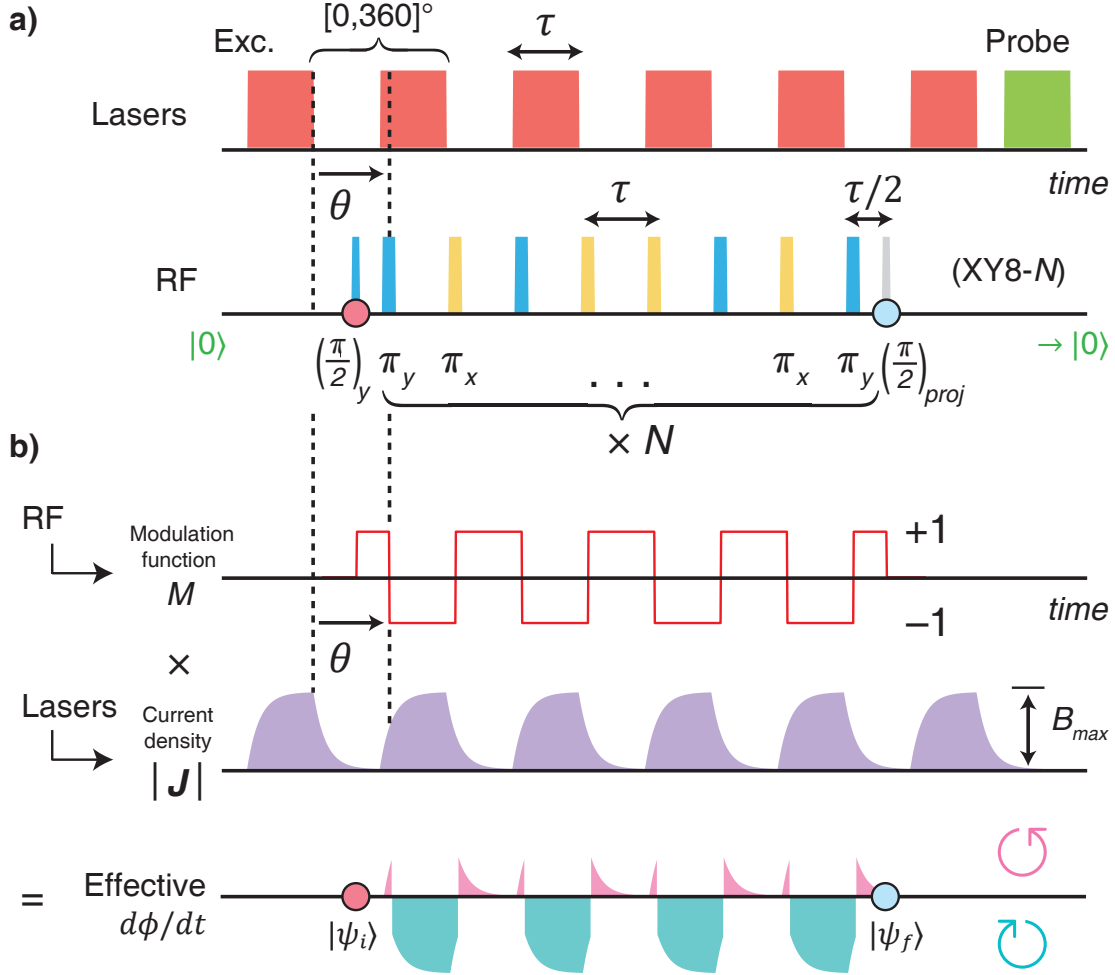


Figure 4.7: ac photocurrent measurement protocol. (a) The two lasers are gated by modulators, and the excitation (red) is synchronized to the spin flips of the NV sensing sequence (RF, blue and yellow bars). The relative phase between the lasers and rf is defined as $\theta = 0$ when the first laser pulse ends in sync with the first spin flip, as shown. At the end of the XY8 sequence, the X and Y projections are measured by projecting to the $|0\rangle$ state. (b) The sensing sequence corresponds to a time-domain filter function M which inverts at each spin flip. The laser excitation produces a periodic photocurrent signal \mathbf{J} , which the sensing sequence detects via its stray magnetic field. The product of the filter function and signal gives the instantaneous phase accumulation on the NV superposition, $d\phi/dt$. If the signal has a nonzero rise time, as shown here, it introduces an effective phase shift between the signal and sensing sequence, such that $\theta_{opt} > 0$. It also broadens the original square excitation pulses, such that the overlap with the filter function is both positive and negative (bottom line). This reduces the total ϕ acquired for a given B_{max} .

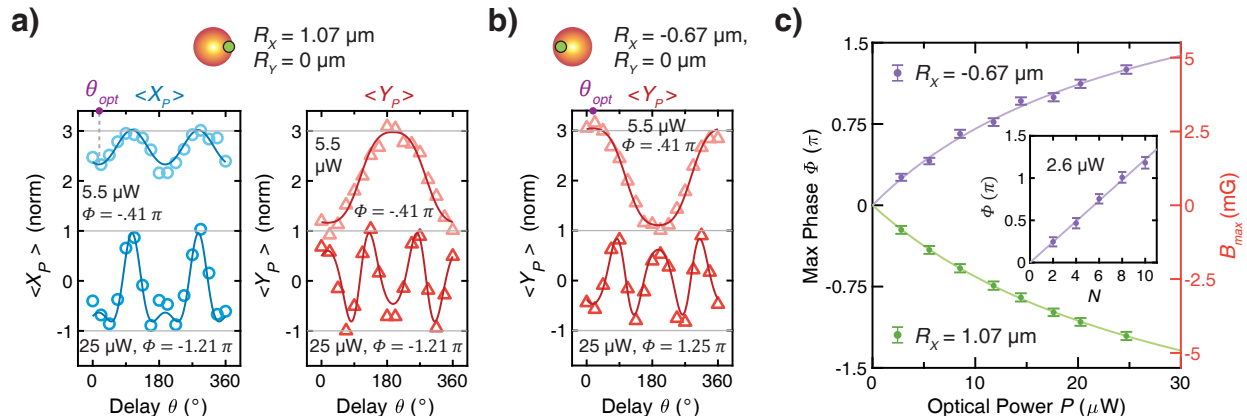


Figure 4.8: Detection of ac photocurrents and dependence on optical power. By varying the delay phase θ , the origin of the signal is unambiguously attributed to the laser and the behavior of X and Y is characterized. (a) X and Y with the probe laser displaced to $R_X = 1.07 \mu\text{m}$, for two different optical powers. Simultaneously fitting X and Y extracts Φ , the maximum phase accumulation. Φ increases with optical power. (b) If the probe is moved to the opposite side of the excitation, $R_X = -0.67 \mu\text{m}$, the X component is similar, but the Y component changes sign, indicating a reversal of photocurrents. (c) Power dependence of the signal at both spots, revealing a saturating signal. Φ can be directly related to a maximum magnetic field (B_{max}) of a few milligauss (right axis). Inset: the signal increases linearly with time, and so does not result from transients.

modulation function times pulse amplitude would always vanish and no signal could be observed. However, there is still a degree of freedom in the relative timing of the laser pulses and the spin flips. The *delay phase* θ parameterizes this variable (Fig. 4.7a, top line). It is termed a phase because it is periodic – if the laser pulses are shifted in time by $\pm 2\tau$, the train of pulses still appears the same. θ is defined such that $\theta = 360^\circ$ is equivalent to a shift of $\Delta t = 2\tau$, and at $\theta = 0$ the first laser pulses switches off exactly at the time of the first microwave π pulse. This means that the laser overlaps perfectly with the odd (1st, 3rd, 5th, ...) evolution periods between pulses, and if the laser causes a phase, it will be at its positive maximum for $\theta = 0$. For $\theta = 180^\circ$, the first laser pulse switches *on* in time with the first π pulse, the laser pulses overlap with the *even* evolution intervals, and any laser-induced phase should be at its most negative. This behavior is analogous to tuning the phase shift of a lock-in amplifier.

Initially, the sequence parameters are set to $N = 2$ and $\tau = 7.6 \mu\text{s}$. This sixteen-pulse

sequence thus has total time $T = 121.6 \mu\text{s}$, which coincides with a strong revival of coherence at the initial magnetic field of 226 G (see Fig. 4.6c). When performed on an uncovered area of diamond, no phase accumulation is observed for any θ , as expected, since the absorption of 662 nm light is negligible. However, when performed on a region covered by MoS₂, large phase oscillations appear, as observed in the X_P and Y_P projections of the final state (Fig. 4.8a-b). These oscillations appear sinusoidal in θ , and match the expected behavior of an observed sinusoidal phase as θ is varied:

$$\begin{aligned} X_P(\theta) &= A \cos(\Phi \cos(\theta - \theta_{opt})) + C, \\ Y_P(\theta) &= A \sin(\Phi \cos(\theta - \theta_{opt})) + C. \end{aligned} \tag{4.4}$$

A and C are the normalized amplitude and offset, expected to be exactly 1 and 0 respectively, but included to properly fit noisy experimental data. Φ is the maximum phase accumulation, and θ_{opt} is a phase offset corresponding to that maximum. These equations are derived from Eq. 4.1 assuming B_{\parallel} is sinusoidal and synchronized with the spin flips. As the excitation power is increased from 5.5 μW to 25 μW , the oscillations increase in amplitude, but jointly fit well to Eq. 4.4. This allows Φ to be extracted, and it is found to increase in magnitude from $\Phi = -0.41\pi$ at the lower optical power to $\Phi = -1.25\pi$ at the higher power. As the power is increased, a sublinear dependence of Φ on power P is observed (Fig. 4.8c). The data fit well to a saturation curve,

$$\Phi(P) = \frac{\Phi_{sat}P}{P + P_{sat}}, \tag{4.5}$$

where Φ_{sat} is the maximum achievable (saturated) phase accumulation at high optical power, and P_{sat} is roughly the power scale of saturation, defining the point where $\Phi_{sat}/2$ is reached. For the upper (positive-valued) curve in Fig. 4.8c, the fit parameters are $\Phi_{sat} \sim 2.5\pi$, $P_{sat} \sim 25 \mu\text{W}$. These values of Φ correspond to average B_{\parallel} on the scale of milligauss, which will be made more precise in later analyses.

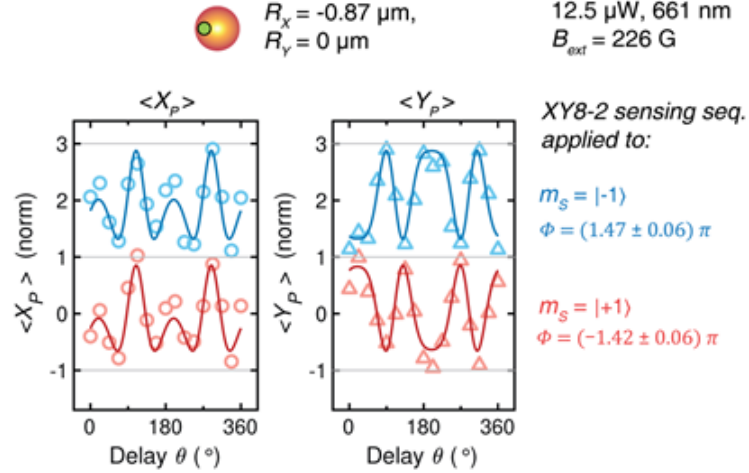


Figure 4.9: Comparison of photocurrent measurements between $| -1 \rangle$ and $| +1 \rangle$ branches. By measuring phase accumulation on both the $\{|0\rangle, | -1 \rangle\}$ and $\{|0\rangle, | +1 \rangle\}$ bases, it is possible to determine the symmetric and asymmetric shifts on the $|\pm 1\rangle$ levels. This distinguishes thermal and electric field effects from magnetic fields. The observed phase accumulation is equal and opposite to within error bars, indicating that thermal and electric field effects are negligible in these measurements.

Before proceeding with extensive characterization of the photocurrent features, several important control measurements must be discussed. As the order of the XY8 sequence is increased, Φ increases linearly (inset to Fig. 4.8c) as expected for a stable photocurrent signal. This suggests that the signal does not originate either from a transient or from a slowly-growing source. Next, Fig. 4.9 compares the phase accumulation in the $\{|0\rangle, | -1 \rangle\}$ and $\{|0\rangle, | +1 \rangle\}$ bases under identical conditions. The X_P data are similar, but the Y_P data appear to be reflected, indicating that $\Phi_{0,-1} = -\Phi_{0,+1}$. This reference measurement is important to establish the magnetic character of the signal, since temperature shifts and electric fields (both of which can arise from laser irradiation) both shift the qubit levels. However, these effects would symmetrically shift the $|\pm 1\rangle$ levels, so the asymmetric shift in Fig. 4.9 can only result from magnetic fields. The magnitude of the signal is approximately equal and opposite, providing no evidence to suggest contributions from the confounding variables. Additionally, 662 nm excitation is used throughout to generate photocurrents, but the effects persist when using 675 nm excitation. The resulting photocurrents are smaller,

consistent with lower MoS₂ absorption at that wavelength. Finally, the ideal expressions in Eq. 4.4 are modified to account for a distribution of signals within the NV ensemble. The photocurrent signal changes with a lengthscale comparable to the optical spot size of ~ 500 nm, and so different parts of the ensemble detect slightly different phase accumulation. To quantitatively account for this, Gaussian averaging is added to the signal dependence on θ :

$$\begin{aligned}\bar{X}_P(\Phi, \theta, \sigma) &= \frac{1}{\sqrt{2\pi}\sigma} \int X_P(\Phi, \theta) e^{-(\Phi' - \Phi)^2 / 2\sigma^2} d\Phi' \\ &= e^{-\cos(\theta - \theta_{opt})^2 \sigma^2 / 2} X_P(\Phi, \theta),\end{aligned}\tag{4.6}$$

and similarly, $Y_P(\Phi, \theta, \sigma) = e^{-\cos(\theta - \theta_{opt})^2 \sigma^2 / 2} Y_P(\Phi, \theta)$.

Introducing this correction tends not to change the fit value of Φ significantly, but does improve the quality of the fits and matches the qualitative expectations.

4.4 Photocurrent Rise Times

In the preceding measurements, the green NV measurement laser and red MoS₂ excitation laser were not directed at the same point, but instead intentionally misaligned by approximately one micron. The relative positioning of the two lasers determines the photocurrents that are observed at the measurement spot, and so independent control of the two beams enables spatial mapping of the photocurrent distribution. To systematically study these effects, the following positioning coordinate system is adopted. The external magnetic field is aligned with one of the four NV axes, all of which point out of the surface at a 54.7° angle to the surface normal (Fig. 4.4). The surface normal becomes the z axis, and the preferred NV axis is defined to be in the x - z plane. Equivalently, the x axis can be defined in the direction of the projection of the NV axis onto the surface. The y axis is then perpendicular to the NV axis along the surface.

The excitation lasers can be independently steered by adjusting their separate alignment paths into the microscope objective. Reflections off the sample, including focal spots of both




Figure 4.10: Phase shift of photocurrent signal as a function of location. Eq. 4.4 is used to fit X and Y data taken along the R_X axis. The optimal delay θ_{opt} increases monotonically, and agrees with a symmetric linear fit (red line). θ_{opt} indicates an increasing rise time τ_r as the probe and excitation are increasingly separated, which tracks the propagation of the photocurrent distribution across the sample. The large error bars near $R_X = 0.4 \mu\text{m}$ occur because the signal crosses through zero near these points and the phase offset cannot be fit well.

lasers, are imaged with a flip-mounted pellicle beamsplitter and CMOS camera so as to not interfere with spin measurements. The pixel intensity values for each laser are fit to 2D gaussian distributions to find their center locations, and a 1951 USAF resolution test target (Thorlabs) is used to calibrate the conversion from image size to physical distance. The image pixel size was found to be 95 nm, and beam centers could be reproducibly positioned with an accuracy of 50 nm. The coordinate convention (R_x, R_y) designates a separation between the red and green beam centers, taking the red focal spot as the origin.

Before mapping the magnitude of the photocurrent distribution, the spatially-resolved measurements provide insight into the photocurrent dynamics. Fits from multiple datasets taken along the x axis indicate that the phase offset, θ_{opt} , steadily increases with separation (Fig. 4.10). This suggests that the observed phase offsets are not due to calibration errors between the excitation pulse train and the sensing sequence, but rather arise from physical phenomena. In this system, a phase offset can appear naturally due to photocurrent rise times which are comparable to τ . If laser excitation produced an instantaneous photocurrent

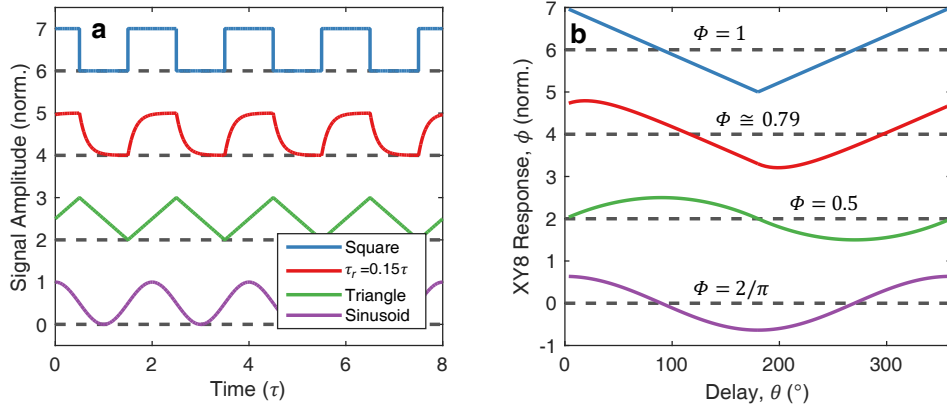


Figure 4.11: Example signal shapes with different relationships between maximum and average amplitude. (a) Square, exponential, and triangle waveforms represent (normalized) photocurrent signals in different rise time regimes: $\tau_r \ll \tau$, $\tau_r \sim \tau$, and $\tau_r \gg \tau$, respectively. The sinusoid waveform is used to derive Eq. 4.4. (b) Corresponding ϕ as a function of θ for each waveform in (a). While the square pulses produce a piecewise linear response that is not observed in experiment, the responses of exponential and triangle waveforms resemble a phase-shifted sinusoid response, justifying the use of the model. The maximum phase $\Phi = \max(\phi(\theta))$ depends on the rise time, despite all signals having normalized amplitude. This dependence is captured in the β parameter during data analysis.

response, θ_{opt} would remain zero. However, if an appreciable fraction of τ passes before the photocurrents reach their maximum, and it takes some time for them to dissipate after the excitation is removed, the net effect is to temporally shift the photocurrents relative to the laser (see Fig. 4.7, purple trace). Longer rise times should correspond to increased phase offsets, and so Fig. 4.10 indicates that the photocurrent distribution rises fastest in the center of the excitation spot, with increasing rise times at larger radii. This observation agrees with intuition, and is explored further in section 4.5.

The knowledge that the photocurrents have appreciable rise times on the scale of τ creates a need for more sophisticated data analysis to interpret observations, but along with improved physical models comes additional physical insight. The current density \mathbf{J} should not be modeled as a series of square pulses, as this model cannot have $\theta_{opt} > 0$ and exhibits sharp corners in the resulting phase accumulation that do not appear in observations (Fig. 4.11, blue trace). While a sinusoid (purple trace) is a good first approximation as the fundamental frequency of the laser modulation, and useful to derive Eq. 4.4, it does not

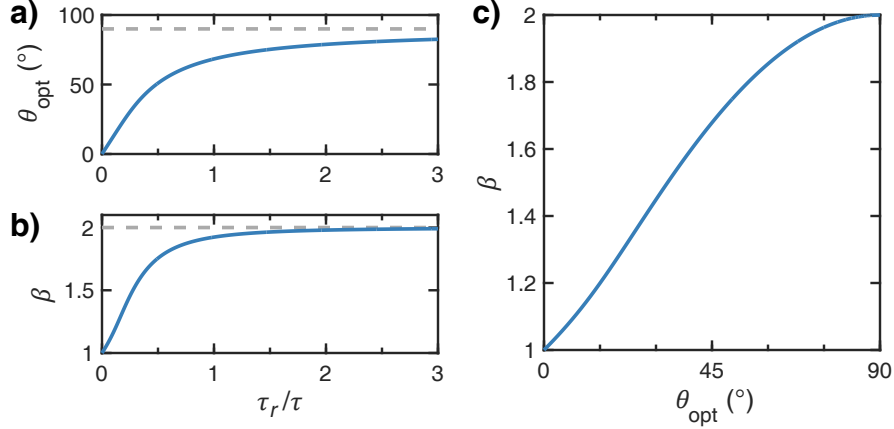


Figure 4.12: Amplitude correction β calibrated against phase shift θ_{opt} . (a-b) From the model Eq. 4.7, analytical forms for θ_{opt} and β as functions of the rise time τ_r are computed. θ_{opt} cannot surpass a 90° phase shift and β ranges from 1 to 2. (c) Combining (a) and (b), β can be determined by the experimentally-observed θ_{opt} . This enables a proper estimate of B_{max} from Φ .

correspond well to the sharp modulation of the laser. A more natural choice is to use exponential rise and fall times τ_r , as in the red trace of Fig. 4.11. In the limit of short rise times, $\tau_r \ll \tau$, the square wave is recovered. In the limit of long rise times, $\tau_r \gg \tau$, the signal approximates a triangle wave (green trace), which resembles a phase-shifted sinusoid and achieves the maximum θ_{opt} of 90° . The steady-state form of the exponential curves can be parameterized as follows:

$$B_{\parallel}(t) = \begin{cases} B_{max}(1 - A_0)^{-1}(1 - e^{-t/\tau_r}), & \text{for } 0 \leq t < \tau \\ B_{max}(1 - A_0)^{-1}(e^{-(t-\tau)/\tau_r} - A_0), & \text{for } \tau \leq t < 2\tau \end{cases} \quad (4.7)$$

The constant $A_0 = e^{-\tau/\tau_r}$ normalizes the functional form so that the amplitude is always B_{max} .

With time-dependent, non-square photocurrent responses, it becomes critical to determine the relationship between the accumulated phase and the maximum magnetic field, since the magnetic field is ultimately proportional to the current. The sensor is effectively performing an integrated readout, with limited sensitivity to fine temporal details of magnetic

field. For square pulses, the gyromagnetic ratio is a simple, context-independent conversion factor. For this signal with a finite rise time, dividing Φ by the gyromagnetic ratio will underestimate the true maximum current, since the signal is partly smoothed across the modulation function (see bottom line of Fig. 4.7). To account for this systematic error, a proportionality constant β is introduced, which relates the amplitude of the signal to the acquired phase:

$$\beta \equiv \frac{\gamma_e T B_{max}}{|\Phi|}. \quad (4.8)$$

The normalization by $\gamma_e T$ ensures that β is independent of the sequence length. For a square pulse, $\beta = 1$, while for a sinusoid $\beta = \frac{\pi}{2}$. An analytical expression for the exponential rise of photocurrent signal (Eq. 4.7) allows the calculation of β as a function of the rise time, shown in Fig. 4.12. Since the rise time is not known *a priori*, a further calculation of θ_{opt} as a function of τ_r is necessary. This relation is then inverted to derive the proper β once the phase offset is measured, as shown in the right panel of Fig. 4.12. The data analysis now produces an accurate analysis of the maximum magnetic field by accounting for the phase offset. For most data sets, $\beta \sim 1.25$.

4.5 Photocurrent Vortex Mapping

With proper data analysis, the spatially-resolved photocurrent measurements can be analyzed to determine the originating current distribution. Scans are taken along both the x and y axes relative to the excitation beam, with results shown in Fig. 4.13. In both cases, the detected signal changes sign for some regions of the sample. This reversal indicates that the direction of the current flow has also reversed. In equilibrium, the divergence-free condition $\nabla \cdot \mathbf{J} = 0$ and the rotational symmetry of the laser and sample system require that any current must flow as vortices around the point of rotational symmetry – in other words, the center of the laser focal point. A vortex pattern, like the one in Fig. 4.13, can explain the unusual, asymmetric shape of the photocurrent distribution measured in the surface

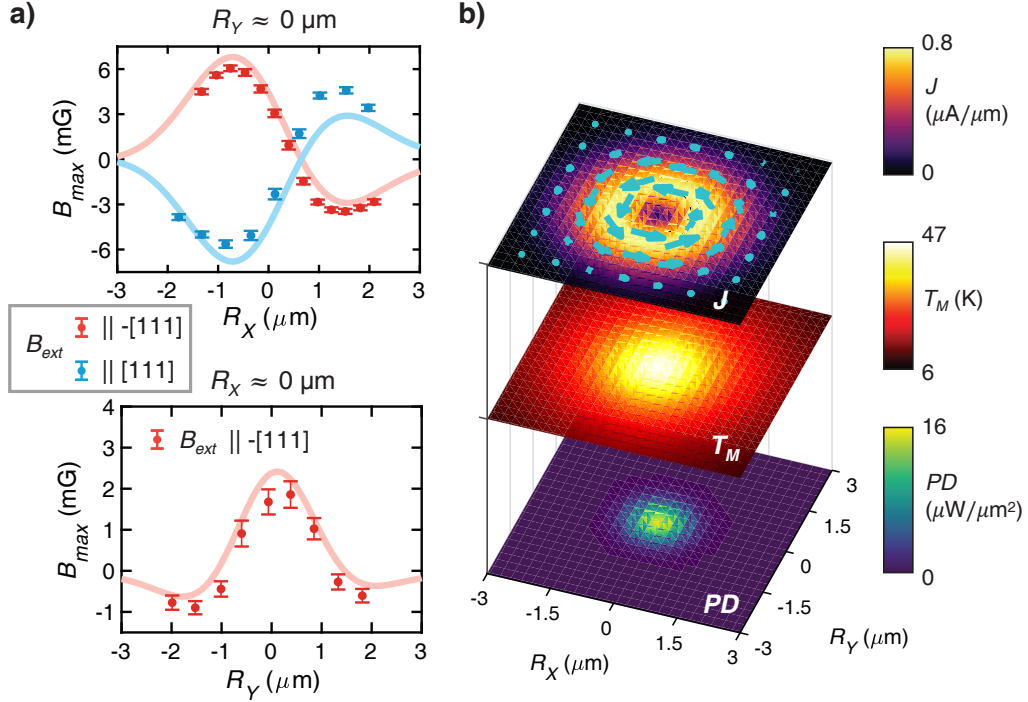


Figure 4.13: Spatial mapping of a photocurrent vortex in MoS₂. (a) The R_X and R_Y scans of the vortex. If the external field is flipped, the circulating currents reverse direction (blue data). Solid lines are fits to the data using the current distribution model in Eq. 4.10 and surrounding text. (b) The bottom layer shows the optical power density, assuming a symmetric spot, for total incident power $25 \mu\text{W}$ (PD); middle layer shows the simulated temperature distribution in the MoS₂ monolayer (T_M); top layer shows the current distribution fit to the data (J).

scans. As the currents circulate, they produce solenoid-like fringing fields. While the current distribution is rotationally symmetric around the z axis, the NV axes are not parallel to z . Since the NV centers respond to parallel magnetic fields, this creates an asymmetry in the resulting fields at the NV depth of 40 nm. The NV axis is fully perpendicular to the y axis, though, and so the y -axis distribution is symmetric still. This vortex extends out to several microns in size, larger than the excitation beam ($\sigma = 0.45 \mu\text{m}$).

The external magnetic field plays a significant role in the resulting photocurrent distribution. While symmetry and steady-state considerations restrict current to flowing in a vortex, the symmetry of left- and right-handedness would still result in zero net current. However, the out-of-plane magnetic field breaks the symmetry and enables a net current to

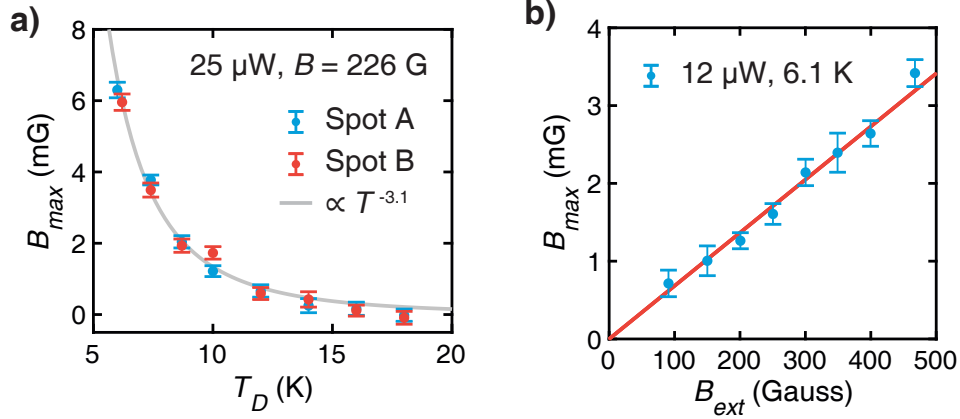


Figure 4.14: Temperature and magnetic field dependence of the observed MoS₂ photocurrent. (a) The photocurrent signal rapidly diminishes with increasing base temperature. This trend is consistent across the sample, and so does not appear to depend on the local MoS₂-diamond interface. A phenomenological trendline (gray) is shown as a guide to the eye. (b) The observed signal increases linearly with applied field B_{ext} . In both plots, data are collected at $R_X = -0.95 \mu\text{m}$.

be established. As an important corollary of this condition, the induced photocurrent should change handedness if the external magnetic field changes sign. To test this, the permanent magnet which applies the static magnetic field is removed, flipped, and replaced, resulting in approximately the same magnetic field magnitude but opposite direction. The R_X scan with the flipped magnetic field is shown in Fig. 4.13 in blue, and corroborates this hypothesis, as the B_{max} values appear negated². Further supporting the importance of the external field, the photocurrent is observed to increase linearly as the magnetic field strength is increased, despite the same applied optical power (Fig. 4.14b). Carrier mobility is expected to be excellent in the MoS₂ monolayer, and so with maximum magnetic fields of 0.05 T, the photocurrent dynamics are well within the low-field regime.

The system's temperature is also a key condition. The cryostat's base temperature of 6 K exhibits the highest signal, with the photocurrents dropping off precipitously (approximately

2. An important technical note here: when the direction of the external field is reversed, the NV axis also reverses, and so the raw data acquires the same phase. The direction of the lab-frame magnetic field can still be determined unambiguously, and for clarity it is this quantity (B_{max} relative to the surface) that is plotted in Fig. 4.13.

as T^{-3}) as the chamber and sample temperature is increased (Fig. 4.14a). To ensure that this dependence is not sensitive to local strain or surface conditions, data is taken from two different parts of the sample at each temperature. The data from both spots match closely. In general, this homogeneity is observed for all measurements; the outcomes are stable and repeatable across different parts of the MoS₂ samples, even when lasers have been directed at a spot for extended periods.

The posited origin of the photocurrents at a microscopic level is the photothermoelectric effect (PTE), in which a heating laser induces a temperature gradient, leading to optically-induced thermoelectric currents. PTE accounts for a multitude of observations to this point. The signal is observed to depend strongly on applied field and temperature, displays a significant rise time and diffusion behavior – the phase offset increasing with radius – and saturation behavior with increasing optical power. Furthermore, MoS₂ has previously been reported to have large thermoelectric and PTE responses [111, 112]. More specifically, given the presence of a static magnetic field, the origin of the currents may be attributed to the Nernst effect, in which electric current appears transverse to both a magnetic field and a temperature gradient in a material [113]. With the radially symmetric thermal gradient, Nernst currents can circulate around the center of the temperature distribution, as observed in the photocurrent measurements. To test the validity of a thermal photocurrent origin, thermal simulations are performed.

The simulation of thermal dynamics in MoS₂ start from a 2D, rotationally symmetric picture. The 2D heat equation is given by

$$\rho c_p \frac{\partial T_M(R, t)}{\partial t} - \kappa \frac{1}{R} \frac{d}{dR} \left[R \frac{dT_M(R, t)}{dR} \right] + \frac{G}{d} [T_M(R, t) - T_D] = f(R). \quad (4.9)$$

ρ is the MoS₂ density; $\kappa(T_M)$ is the MoS₂ thermal conductivity; G is the thermal interface conductivity between MoS₂ and diamond; d is the MoS₂ thickness; f is the laser heating source, assumed gaussian and with a measured beam radius; and T_M and T_D are the tem-

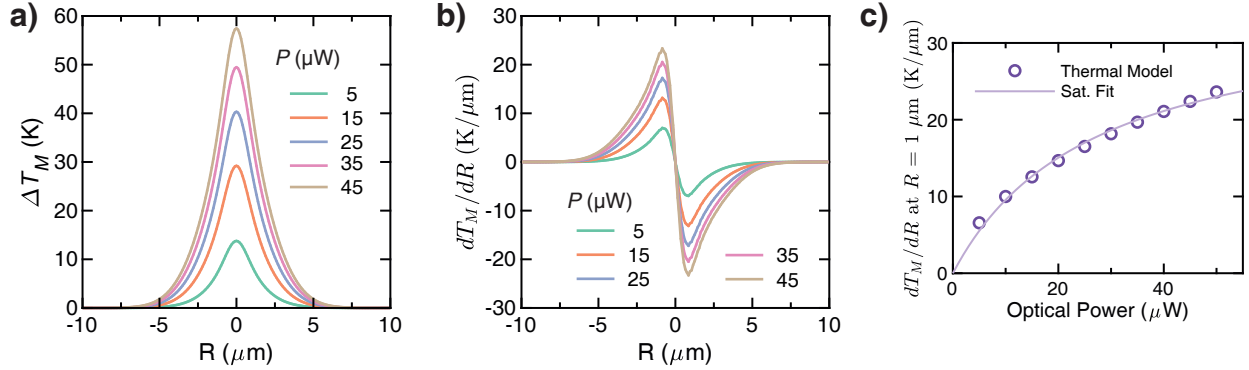


Figure 4.15: Temperature simulations in monolayer MoS₂. (a) Increase ΔT_M from a base temperature of 9 K after an optical pulse of duration 7.6 μs . (b) Spatial dependence of the thermal gradient for various optical powers. The incident power does not dramatically change the normalized distribution. (c) Dependence of the thermal gradient at $R = 1 \mu\text{m}$ on optical power. A sublinear, saturating behavior is observed, in agreement with experiment, with a fit line displayed.

temperatures of MoS₂ and diamond, respectively. T_D is assumed constant, due to diamond's high thermal conductivity and low absorption. The remaining parameters, except for c_p , are either taken from published values or measured in experiment (see Appendix C). T_D is set to be 9 K, slightly above the cryostat base temperature of 6 K, to account for steady-state heating. The MATLAB PDE toolbox is used to perform the simulation, and some results are displayed in Fig. 4.15. The model reproduces many of the experimental observations, including a roughly Gaussian thermal distribution (since PTE currents are proportional to temperature gradients) and a saturating rise time with additional power. c_p is chosen such that the system demonstrates a rise time of $\sim 1 \mu\text{s}$, which requires $c_p \approx 200 \text{ J}/(\text{kg K}^2) \times T_M$, whereas the room temperature heat capacity is measured in other experiments to be 400 J/(kg K) [111]. This discrepancy is still unexplained, but could arise from surface contamination such as residual PMMA or adsorbed gases during the cooling process. The overall good agreement in simulation supports the PTE origin of the photocurrents. The simulated equilibrium thermal distribution is represented in the middle layer of Fig. 4.13b.

A second model is used to calculate the magnetic field produced by a circulating current distribution, since the measurements provide no information about the magnitude or

temperature dependence of the MoS₂ Nernst coefficient which relates the temperature gradient to a transverse current density. As the temperature distribution is roughly gaussian, a current density proportional to the derivative of a gaussian is assumed,

$$J(R) = J_0 \frac{\sqrt{e}}{\sigma_J} e^{-\frac{R^2}{2\sigma^2}}, \quad (4.10)$$

where J_0 is the amplitude and σ_J the radius of maximum current density. The current is assumed to circulate around the origin due to the previously discussed symmetry considerations, and σ_J is taken to be $\sim 1.0 \mu\text{m}$, close to the maximum of the derivative of the simulated thermal distribution (Fig. 4.15, middle). This current distribution is displayed in the top layer of Fig. 4.13b, along with the thermal simulation results and the measured excitation beam power density. $\mathbf{B}(\mathbf{R})$ is then calculated by integrating the Biot-Savart law over the concentric infinitesimal current loops. The field is calculated for a depth of 40 nm below the MoS₂ plane. Finally, to match the experimental conditions, \mathbf{B} is projected onto the NV axis. The ideal NV axis in the specified coordinate system is $\hat{n}_N V = (\sqrt{2}, 0, 1)/\sqrt{3}$, but slight rotations around the y (-5°) and z (8°) axes are included to allow for experimental nonidealities. The R_X and R_Y scan locations are also allowed to have small offsets (< 250 nm, or about two pixels) to account for systematic beam positioning errors. This projected field is shown in Fig. 4.13 alongside the experimental data, and shows good agreement to observations.

The current density model approximates the experimental observations with $J_0 = 0.8 \mu\text{A}/\mu\text{m}$ and $\sigma_J = 1.0 \mu\text{m}$. Integrating $J(R)$ over $R > 0$, an estimated maximum of $1.3 \mu\text{A}$ flows in the vortex for an excitation power of $25 \mu\text{W}$. Accounting for the 85% transmission of the microscope objective, this indicates a photo-response of $\sim 60 \text{ mA/W}$, and in an applied field of 226 G along the NV center. Since only the out-of-plane component is expected to affect the photocurrent, the normal component of 130 G leads to a linear slope of $\sim 4.6 \text{ A}/(\text{W T})$ at low fields and optical powers. This strong photo-response corresponds to the




Figure 4.16: Spatially-resolved photocurrent rise time measurements. By sweeping τ (for both optical pulse duration and spin flip intervals), different values of τ_r/τ can be tested. For small τ , the signal diminishes and the rise time can be fit using Eq. 4.7. Data sets from two different R_X are displayed, documenting the increase in rise time for points farther from the center of the excitation beam.

high Seebeck coefficient and low thermal conductivity of monolayer MoS_2 .

A final set of measurements on the photocurrent vortex delves deeper into understanding photocurrent rise times by taking spatially and temporally resolved measurements. The thermal simulations predict that thermal rise times increase at farther distances from the excitation beam center, and some evidence for increasing rise times was already seen in the phase offset data in Fig. 4.10. To more directly investigate, measurements with different τ are taken with other conditions held constant. As θ_{opt} is only observed to vary with position, not optical power or XY8 order, these measurements explore different values of τ_r/τ . Fig. 4.16 shows the measured B_{max} for sweeps of τ at two different positions, one closer and one farther from the excitation spot center. As $\tau \rightarrow 0$ at constant τ_r , the maximum achieved magnetic field is reduced from B_{max} by the factor $(1 - e^{-\tau/\tau_r})$ (Eq. 4.7). This relationship is directly observed by measuring the field instead of θ_{opt} , and τ_r is fit to $0.8 \pm 0.2 \mu\text{s}$ for the closer spot and $1.4 \pm 0.2 \mu\text{s}$ for the farther spot, consistent with expectations. Thus, the ability to control θ , τ , and independently steer the green and red lasers enables spatiotemporal photocurrent measurements which investigate rise time and




Figure 4.17: Smallest resolved photocurrent. The phase-resolved measurements, enabling signal to accumulate in Y , result in sensitive signal detection. Fields as small as 0.14 ± 0.05 mG can be measured, corresponding to 20 nA/ μm current density. Very little information is carried in X measurements, so significant gains in measurement efficiency are possible by optimizing for the most sensitive data points.

diffusion phenomena, which are challenging to observe locally in 2D materials.

4.6 Measurement Sensitivity

In Fig. 4.17, a minimum signal of $B_{min} = 0.14 \pm 0.05$ mG is detected after ~ 2 hr averaging time. The Y_P measurements provide nearly all the signal information, as $Y_P \propto \sin \phi \approx \phi$ is more sensitive than $X_P \propto \cos \phi \approx 1$ for small ϕ . This highlights the advantage of performing phase-resolved ac measurements instead of either dc or random-phase ac measurements (typical for magnetic noise sensing). To convert the detection limit to a current density, an infinite sheet model is used, which produces constant field regardless of the sensor distance. This current sheet is assumed to flow along the y axis, so that its field is maximally projected along the NV axis. The minimum detectable current is then calculated via

$$\begin{aligned} \mathbf{B}_{sheet} &= \frac{\mu_0 J_{min}}{2} (1, 0, 0) \\ B_{min} &= \mathbf{B}_{sheet} \cdot \hat{n}_{NV} \\ J_{min} &= 20 \text{ nA}/\mu\text{m} \end{aligned} \tag{4.11}$$

This provides a factor of 50 improvement over the resolution limit of dc current sensing with ensembles or scan-probe tips demonstrated so far. In addition, this figure is calculated based on the full set of $X_P(\theta)$, $Y_P(\theta)$ measurements. As evident in Fig. 4.17, the X_P measurements provide no information on small signals, and need not be taken. Similarly, the information from Y_P data can be maximized by taking just the one or two points with the most signal per time (at $\theta = \theta_{opt}$). The magnetic sensitivity of the measurement is given by

$$\eta_B = \frac{\sigma}{dS/dB_0}, \quad (4.12)$$

where σ is the standard deviation of signal S , and B_0 is the field amplitude. For an instantaneous, square wave photocurrent signal with 50% duty cycle, the signal produced by the XY8- N sequence is

$$\Phi = (2\pi\gamma_e) \times TB_0/2. \quad (4.13)$$

Here, $T = 8N\tau$ is the sensing time. At small Φ , Y_P is approximately linear in Φ , and

$$\frac{dY_P}{dB_0} = CN_{ph}\pi\gamma_e T. \quad (4.14)$$

N_{ph} is the number of photons collected in the measurement, and $2C$ is the percent contrast between the bright and dark spin states. When photon counts are limited by Poissonian shot noise, as in these experiments, the signal standard deviation is $\sigma = \sqrt{N_{ph}}$. Finally, the magnetic field sensitivity is

$$\eta_B = \frac{1}{C\pi\gamma_e T \sqrt{\dot{N}_{ph}}}. \quad (4.15)$$

Under typical measurement conditions, $T \approx 150 \mu s$, the photon collection rate is $\dot{N}_{ph} \sim 0.75$ per measurement, and the contrast C , reduced by MoS₂ background, is ~ 0.01 . With these

figures, the sensitivities are estimated to be

$$\begin{aligned}\eta_B &= 110 \frac{\text{nT}}{\sqrt{\text{Hz}}}, \\ \eta_J &= 200 \frac{\text{nA}}{\mu\text{m}\sqrt{\text{Hz}}}.\end{aligned}\tag{4.16}$$

This measurement is still far from optimized. NV ensemble measurements have achieved ac magnetic sensitivities well below $1 \text{ nT}/\sqrt{\text{Hz}}$, which improve the current sensitivity accordingly. Several simple steps to increase sensitivity include optimizing the NV ensemble density, using isotopically purified material to increase coherence times, and improving collection efficiency.

4.7 Conclusion

This chapter has presented a novel technique for detecting quasistatic and ac photocurrents with NV centers in diamond. This technique is well-suited for 2D materials and thin films, as it locally detects photocurrents through their magnetic field signatures, avoiding the need for invasive electrical contacts. Synchronizing pulsed photoexcitation with ac magnetometry sequences improves the sensor coherence time and allows phase-resolved measurements of current densities as small as $20 \text{ nA}/\mu\text{m}$, 50 times smaller than the dc limit. The high sensitivity and optical diffraction-limited spatial resolution are used to map a photocurrent vortex in MoS_2 attributed to photothermoelectric origins. Temporal resolution provided by ac measurements at different frequencies chart a rise time of approximately $1 \mu\text{s}$ and demonstrate the propagation of the photocurrent distribution across the sample. The protocols and analysis presented here are expected to become useful methods for investigating electrical, thermal, and photocurrent phenomena in 2D materials.

Appendix A

Details of Superadiabatic Pulse

Shapes

Stimulated Raman adiabatic passage (STIRAP) is a means of transferring population between two states without a direct coupling, via an intermediate excited state. As an adiabatic technique, it is characterized by robustness to noise and experimental imperfections, but also by long evolution time to achieve a high degree of adiabaticity. Superadiabatic transitionless driving (SATD) is a modification of STIRAP which aims to reduce nonadiabatic errors introduced by finite evolution times. The key elements are that SATD intentionally deviates from the ideal adiabatic evolution, while still matching the desired final and initial states. It also only uses existing couplings instead of introducing new control fields or interactions, and seeks to maintain the adiabatic robustness even over shorter durations. Here, the original STIRAP protocol, SATD, and a variant SATD (MOD-SATD) are presented for the $\{|-1\rangle, |+1\rangle, |A_2\rangle\}$ Λ system of the NV center. In addition to full STIRAP, the modifications necessary to perform fractional STIRAP are included.

First, the details of STIRAP are recapped from section 3.1.1. The Λ system Hamiltonian

is

$$H(t) = \frac{1}{2} \begin{pmatrix} 0 & 0 & \Omega_P(t) \\ 0 & 2\delta & \Omega_S(t)e^{i\phi_S} \\ \Omega_P(t) & \Omega_S(t)e^{-i\phi_S} & 2\Delta \end{pmatrix} \quad (\text{A.1})$$

where Ω_P is the pump pulse, Ω_S is the Stokes pulse, ϕ_S is the phase difference between the two, and Δ and δ are the one- and two-photon detunings, respectively. In STIRAP, $\Delta = \delta = \phi_S = 0$ typically. When the drives are parameterized with an overall and relative amplitude,

$$\begin{aligned} \Omega_P(t) &= \Omega(t) \sin \theta(t), \\ \Omega_S(t) &= \Omega(t) \cos \theta(t), \end{aligned} \quad (\text{A.2})$$

the eigenstates of the system can be written as follows:

$$\begin{aligned} |B_{\pm}(t)\rangle &= \frac{1}{\sqrt{2}} (\sin \theta(t) |-1\rangle + \cos \theta(t) |+1\rangle \pm |A_2\rangle), \\ |D(t)\rangle &= \cos \theta(t) |-1\rangle - \sin \theta(t) |+1\rangle. \end{aligned} \quad (\text{A.3})$$

$|B_{\pm}\rangle$ are the two bright states, due to their $|A_2\rangle$ components, while $|D\rangle$ is the dark state. The STIRAP protocol utilizes the time dependence of the instantaneous dark state to evolve an initial state (here, $|D(t_i)\rangle = |-1\rangle$) to a different final state (here, $|D(t_f)\rangle = |+1\rangle$). To more suitably describe the adiabatic characteristics of the protocol, a change of basis can be applied to make the eigenstates time independent. Using the unitary

$$U(t) = |D\rangle \langle D(t)| + |B_+\rangle \langle B_+(t)| + |B_-\rangle \langle B_-(t)|, \quad (\text{A.4})$$

the Hamiltonian can be rewritten in the adiabatic basis as

$$\begin{aligned}
H_{ad}(t) &= U(t)H(t)U^\dagger(t) + i\frac{dU(t)}{dt}U^\dagger(t) \\
&= \frac{\Omega(t)}{2} (|B_+\rangle\langle B_+| - |B_-\rangle\langle B_-|) \\
&\quad + \frac{i\dot{\theta}(t)}{2} \left(\frac{|B_+\rangle + |B_-\rangle}{\sqrt{2}} \langle D| - |D\rangle \frac{\langle B_+| + \langle B_-|}{\sqrt{2}} \right) \\
&= \frac{\Omega(t)}{2} M_z + \frac{\dot{\theta}(t)}{2} M_y.
\end{aligned} \tag{A.5}$$

In the last line, a new set of spin-1 operators is introduced, M_i , which are defined on the $\{|B_+\rangle, |D\rangle, |B_-\rangle\}$ basis:

$$M_x = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad M_y = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & i & 0 \\ -i & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \quad M_z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \tag{A.6}$$

These operators obey the commutation relation $[M_i, M_j] = i\varepsilon_{ijk}M_k$, with ε the Levi-Civita symbol. In this formulation, it becomes clear that the system can only remain exactly in the dark state if $\dot{\theta} = 0$, and that the adiabatic approximation is satisfied if $\Omega \gg \dot{\theta}$ for all t .

In adiabatic shortcut methods like SATD, the Hamiltonian is augmented with some additional control, $H_c(t)$, which ideally reduces or eliminates the effect of the off-diagonal term in the adiabatic Hamiltonian. Nonadiabatic STIRAP errors can be completely suppressed by applying a direct coupling between $| -1 \rangle$ and $| +1 \rangle$, but this subverts the original assumption of STIRAP that the initial and final states are not necessarily coupled. In order to avoid introducing additional coupling terms, the adiabatic states can be dressed by a time-dependent unitary $V(t)$. In particular, the dressed dark state can be modified from the adiabatic dark state, expressed generally as

$$\left| \tilde{D}(t) \right\rangle = V^\dagger(t) |D(t)\rangle = c_+(t) |B_+(t)\rangle + c_-(t) |B_-(t)\rangle + c_D(t) |D(t)\rangle. \tag{A.7}$$

$V(t)$ is chosen to still satisfy the boundary conditions,

$$\begin{aligned} \left| \tilde{D}(t_i) \right\rangle &= V^\dagger(t) |D(t_i)\rangle = |-1\rangle, \\ \left| \tilde{D}(t_f) \right\rangle &= V^\dagger(t) |D(t_f)\rangle = |+1\rangle, \end{aligned} \quad (\text{A.8})$$

implying $V(t_i) = V(t_f) = I$. In this dressed state basis, the corrected Hamiltonian becomes

$$H_{dr}(t) = V(t)U(t) \left(H(t) + H_c(t) \right) U^\dagger(t) V^\dagger(t) + iV(t) \frac{dU(t)}{dt} U^\dagger(t) V^\dagger(t) + i \frac{dV(t)}{dt} V^\dagger(t). \quad (\text{A.9})$$

H_c can then be designed to eliminate couplings between the dressed bright and dark states. Due to the boundary conditions, the initial and final states match the adiabatic case, so in the absence of dissipation, perfect state transfer is expected.

The SATD protocol is derived by choosing $V(t)$ to diagonalize the original adiabatic Hamiltonian (Eq. A.5). Ref. [80] details the full derivation, with the resulting dressed state transformation given by

$$V_{SATD}(t) = \exp \left[-i \arctan \left(2 \frac{\dot{\theta}(t)}{\Omega(t)} \right) M_x \right]. \quad (\text{A.10})$$

The corresponding correction term is given by

$$\begin{aligned} H_c(t) &= -iU^\dagger(t) V_{SATD}^\dagger(t) \frac{dV_{SATD}(t)}{dt} U(t) \\ &= 2 \frac{\ddot{\theta}(t)\Omega(t) - \dot{\theta}(t)\dot{\Omega}(t)}{\Omega^2(t) + 4\dot{\theta}^2(t)} \left(\cos \theta(t) |-1\rangle \langle A_2| - \sin \theta(t) |+1\rangle \langle A_2| + h.c. \right). \end{aligned} \quad (\text{A.11})$$

When the adiabatic pulse parameters are applied, that is (see also Eq. 3.4),

$$\begin{aligned} \theta(t) &= \frac{\pi}{2} \frac{1}{1 + e^{-\nu(t-L/2)}}, \\ \Omega(t) &= \Omega_0, \end{aligned} \quad (\text{A.12})$$

the SATD pulse shapes in Fig. 3.3 are produced.

MOD-SATD is a further modification of the SATD protocol which reduces the excited state component of the dressed dark state. This is accomplished by reducing the degree of transformation in the dressed state, since the adiabatic dark state has no excited state component. The dressed state unitary $V(t)$ is modified from Eq. A.10 as follows:

$$V_{MOD}(t) = \exp \left[-i \arctan \left(2 \frac{\dot{\theta}(t)}{f(t)\Omega(t)} \right) M_x \right]. \quad (\text{A.13})$$

If $f(t) > 1$, the angle of the unitary transformation will be reduced toward an identity operation. A correction of the form

$$f(t) = 1 + Ae^{-t^2/T^2} \quad (\text{A.14})$$

was optimized in Ref. [80] to minimize $\langle A_2 | \tilde{D}(t) \rangle$ by tuning parameters A and T . Using the modified dressed state, the correction term becomes

$$\begin{aligned} H_c(t) &= -iU^\dagger(t)V_{MOD}^\dagger(t) \frac{dV_{MOD}(t)}{dt} U(t) \\ &= 2 \frac{\ddot{\theta}(t)\Omega(t)f(t) - \dot{\theta}(t)(\dot{\Omega}(t)f(t) + \Omega(t)\dot{f}(t))}{\Omega^2(t)f^2(t) + 4\dot{\theta}^2(t)} \times \\ &\quad \left(\cos \theta(t) |-1\rangle \langle A_2| - \sin \theta(t) |+1\rangle \langle A_2| + h.c. \right). \end{aligned} \quad (\text{A.15})$$

The corresponding MOD-SATD pulse shapes, also starting from the adiabatic pulses, are shown in Fig. 3.3.

Finally, fractional STIRAP [114] is also derived from the original adiabatic pulse shapes. The fractional STIRAP boundary conditions differ from STIRAP:

$$\begin{aligned} |\tilde{D}(t_i)\rangle &= |-1\rangle, \\ |\tilde{D}(t_f)\rangle &= \cos \theta_f |-1\rangle + e^{-i\phi_S} \sin \theta_f |+1\rangle. \end{aligned} \quad (\text{A.16})$$

To achieve the desired final state, θ is ramped to an intermediate angle θ_f instead of $\pi/2$.

The corresponding protocol parameters are

$$\begin{aligned}
\theta(t) &= \frac{\theta_f}{1 + e^{-\nu t}}, \\
\Omega(t) &= \frac{\Omega_0}{1 + e^{\xi(t-t_0)}}, \quad \xi > 0, \\
\Omega_P(t) &= -\Omega(t) \sin \theta(t), \\
\Omega_S(t) &= \Omega(t) \cos \theta(t) e^{i\phi_S}.
\end{aligned} \tag{A.17}$$

The SATD variant of fractional STIRAP is likewise found through Eq. A.11. The pulses employed in Fig. 3.9c use $\xi = \nu$, $t_0 = 3.73/\nu$, which correspond to the same maximal speed as full-STIRAP SATD: $\nu = 1.315 \Omega_{min}^{SATD}$ (Eq. 3.6).

Appendix B

Details of Holonomic Quantum Gates

B.1 Holonomic Gate Parameters

Gate	Loop Parameter θ	Loop Parameter ϕ	Detuning (Δ/Ω)	$U(\theta, \phi, \Delta, \Omega)$
X	$\frac{\pi}{2}$	0	0	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$
Y	$\frac{\pi}{2}$	$\frac{\pi}{2}$	0	$\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$
$Z(\gamma)$	0	-	$\frac{\Delta}{\Omega}$	$\begin{pmatrix} 1 & 0 \\ 0 & e^{i\gamma} \end{pmatrix}$
H	$\frac{3\pi}{4}$	0	0	$\frac{1}{\sqrt{2}} \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix}$
$X(\pi/2)$	$\frac{\pi}{2}$	0	$\frac{1}{\sqrt{3}}$	$\frac{1+i}{2} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix}$
$X(-\pi/2)$	$\frac{\pi}{2}$	0	$\frac{-1}{\sqrt{3}}$	$\frac{1-i}{2} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix}$
$Y(\pi/2)$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{1}{\sqrt{3}}$	$\frac{1+i}{2} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$
$Y(-\pi/2)$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{-1}{\sqrt{3}}$	$\frac{1-i}{2} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$

Table B.1: Experimental parameters for implementing holonomic gates. These choices reflect the control fields as defined in Eqs. 3.8 and 3.9, which perform the geometric operation U described in Eqs. 3.13 and 3.14. Relative to the conventional definition, H is performed on swapped basis states due to the experimental implementation, but is still a Hadamard gate.

B.2 Quantum Process Tomography

As described in section 3.3.3 and performed throughout section 3.3, quantum process tomography (QPT) is a method to thoroughly characterize the effect of a quantum process, even a dissipative one, on arbitrary quantum states [115]. QPT is performed by preparing a set of initial states in the $|\pm 1\rangle$ basis, applying the process under study, and measuring the resulting output states.

The basis states for the NV center ground state qubit are as follows:

$$\begin{aligned} |\pm z\rangle &= |\mp 1\rangle, \\ |\pm x\rangle &= \frac{1}{\sqrt{2}} (|+z\rangle \pm |-z\rangle), \\ |\pm y\rangle &= \frac{1}{\sqrt{2}} (|+z\rangle \pm i|-z\rangle). \end{aligned} \tag{B.1}$$

To perform QPT on a test process \mathcal{E} , the initial states $|\psi_i\rangle = \{|z\rangle, |-z\rangle, |x\rangle, |y\rangle\}$ are prepared, corresponding to initial density matrices $\rho_i = |\psi_i\rangle\langle\psi_i|$. Next, the operation \mathcal{E} is applied, and the final density matrix is measured through a standard quantum state tomography approach [116], which entails projecting the final state along the X , Y , and Z axes. After \mathcal{E} (or any other quantum process) is applied, the final density matrix ρ_f is determined by the initial density matrix and a process matrix χ :

$$\rho_f = \sum_{j,k=1}^4 \chi_{jk} E_j \rho_i E_k. \tag{B.2}$$

The E_j basis operators are chosen to be

$$\begin{aligned} E_1 &= I & E_3 &= Y = -i\sigma_y \\ E_2 &= X = \sigma_x & E_4 &= Z = \sigma_z \end{aligned} \tag{B.3}$$

Performing state tomography for the set of initial states allows χ to be determined by inverting Eq. B.2, following the method of ref. [115]. The estimated χ is constrained to be

a positive semidefinite matrix, to ensure it is physical. However, while residual population in $|A_2\rangle$ is negligible due to the short lifetime, loss can occur to the $|0\rangle$ state due to errors in microwave projection pulses and weaker radiative decay pathways. Thus, $\text{Tr}(\chi)$ is not guaranteed to be 1, and in experiment is observed to be slightly less than 1. The reported fidelity metrics are computed via $\mathcal{F} = \text{Tr}(\chi_{exp}\chi_{id})$, given experimentally determined χ_{exp} and theoretically ideal χ_{id} . To isolate the errors resulting from the holonomic gate alone, the calculated fidelities are normalized by the identity operation - that is, QPT is performed for an identity operation (a simple wait time instead of an optical pulse) to calculate its fidelity \mathcal{F}_I , and the fidelity of other gates G is $\mathcal{F}_G/\mathcal{F}_I$. The identity fidelity is $\mathcal{F}_I = 0.97 \pm 0.01$.

Appendix C

Thermal Modeling of Molybdenum Disulfide Laser Heating

Section 4.5 describes the 2D heat equation model which is used to estimate the heating of the MoS₂ monolayer on diamond as a result of photoexcitation. This appendix provides more details of the simulations and their relevance to the experimental results. The 2D heat equation (reproduced from Eq. 4.9) is

$$\rho c_p \frac{\partial T_M(R, t)}{\partial t} - \kappa \frac{1}{R} \frac{d}{dR} \left[R \frac{dT_M(R, t)}{dR} \right] + \frac{G}{d} [T_M(R, t) - T_D] = f(R). \quad (\text{C.1})$$

Due to the radial symmetry of the laser and the uniformity of both the MoS₂ sheet and the MoS₂-diamond interface, the 2D model can be reduced to the radially symmetric form shown above, and solved for the MoS₂ temperature T_M as a function of radius R and time t .

There are a large number of unknown or estimated physical quantities in this system, which prevent the use of the discussed measurements for precisely fitting thermal parameters. However, this does not preclude their future use for such ends, but rather reflects the difficulty of measuring the material properties of 2D materials, especially at low temperatures. For these simulations, the use of related (but not exactly matched) literature values

Symbol	Quantity	Value	Units	Source
ρ	MoS ₂ density	5.06×10^{-15}	$\frac{\text{kg}}{\mu\text{m}^3}$	[111]
c_p	Specific heat capacity of MoS ₂	$200 \cdot T_M$	$\frac{\text{J}}{\text{kg K}}$	See text
κ	MoS ₂ thermal conductivity	$(2 \times 10^{-7}) \cdot T_M$	$\frac{\text{W}}{\mu\text{m K}}$	[117] (expt.)
G	Thermal interface conductance	5×10^{-10}	$\frac{\text{W}}{\mu\text{m}^2 \text{K}}$	[118] (theory)
d	MoS ₂ thickness	6.5×10^{-4}	μm	[111]
T_D	Diamond substrate temperature	6	K	Measured

Table C.1: Simulation parameters for 2D heat equation model.

and assumptions proves necessary to obtain quantitative results. Monolayer MoS₂ thermal conductivity, κ , has been experimentally measured down to low temperatures, with $\kappa \sim T^{1.3}$ below 100 K [117]. For the simulation, this is approximated with a linear temperature dependence. In some cases, only theoretical studies of the parameter exist. For example, a key parameter is the thermal interface conductance, G , which in ref. [118] was examined for monolayer MoS₂ on a SiO₂ substrate. G is found to strongly decrease with temperature, $G \propto T^4$, and its approximate value at 10 K is adopted for the simulations. The temperature of the diamond substrate, T_D , is assumed to match the cryostat temperature measured at the base of the sample holder due to good thermal contact with the sample mount and diamond's excellent thermal conductivity. The thermal parameters incorporated into simulation, and their sources, are collected in Table C.1. Several laser parameters, which define

Symbol	Quantity	Value	Units	Source
f	Laser heat source	$\frac{P\eta}{2\pi R_0^2} \frac{\alpha}{t} e^{-R^2/2R_0^2}$	$\frac{\text{W}}{\mu\text{m}^3}$	TEM ₀₀ mode
P	Laser power at objective	$1 - 45 \times 10^{-6}$	W	Measured
η	Objective transmission	0.85	-	Specification
α	MoS ₂ absorption	3×10^{-2}	-	[119, 120] (expt.)
R_0	Laser beam width	0.45	μm	Measured

Table C.2: Simulation parameters for laser heating source.

the magnitude of the heating, are also measured. Total laser power, P , is measured at the back aperture of the confocal microscope objective, and the radius of the gaussian beam at the sample surface, R_0 , is measured via the reflected camera image. These parameters are listed in Table C.2. Due to the approximations that enter into the calculations, the simulation is not expected to be predictive, but rather to support the conclusions drawn from the experimental data and inform order-of-magnitude estimates.

The low-temperature heat capacity of MoS₂, c_p , is an important parameter for setting the timescale of the dynamics, but is not well-studied. Ref. [121] predicts $c_p \sim T^{1.1}$ below 50 K, and so a linear model is adopted. With the other parameters reported here, a value of $c_p = 200 \text{ J}/(\text{kg K}^2) \cdot T_M$ reproduces the microsecond-scale dynamics of the experiment. This estimate is significantly higher than the existing theoretical predictions for single-crystal, monolayer MoS₂ [121, 122]. While the MoS₂ sample is polycrystalline, and grain boundaries can increase c_p , the grain size of about $3 \mu\text{m}$ is \gtrsim the photocurrent effects. Other potential sources of increased heat capacity include defects in the sample, accumulated during storage and handling, PMMA residue from the initial transfer process, or adsorbates aggregating during the cooling process. The presence of some moderate MoS₂ photobleaching may suggest the presence of surface adsorbates. This photobleaching, which occurs only at low temperature, does not appear to affect the photocurrent measurements. Further experimental and theoretical work will be necessary to investigate the possibilities. However, the qualitative agreement provided by the simulations with available parameters support the photothermal origins of the effects and motivate these future studies.

References

- [1] Brian B Zhou, Alexandre Baksic, Hugo Ribeiro, Christopher G Yale, F Joseph Heremans, Paul C Jerger, Adrian Auer, Guido Burkard, Aashish A Clerk, and David D Awschalom. Accelerated quantum control using superadiabatic dynamics in a solid-state lambda system. *Nature Physics*, 13(4):330–334, 2017.
- [2] Brian B Zhou, Paul C Jerger, Vladislav O Shkolnikov, F Joseph Heremans, Guido Burkard, and David D Awschalom. Holonomic quantum control by coherent optical excitation in diamond. *Physical Review Letters*, 119(14):140503, 2017.
- [3] Brian B Zhou, Paul C Jerger, Kan-Heng Lee, Masaya Fukami, Fauzia Mujid, Jiwoong Park, and David D Awschalom. Spatiotemporal mapping of a photocurrent vortex in monolayer mos 2 using diamond quantum sensors. *Physical Review X*, 10(1):011003, 2020.
- [4] Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21:467–468, 1982.
- [5] William K Wootters and Wojciech H Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
- [6] Peter W Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4):R2493, 1995.
- [7] Dorit Aharonov and Michael Ben-Or. Fault-tolerant quantum computation with constant error rate. *SIAM Journal on Computing*, 2008.
- [8] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, 2012.
- [9] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 217–228, 2019.
- [10] Michael Held and Richard M Karp. A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics*, 10(1):196–210, 1962.

- [11] Lov K Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79(2):325, 1997.
- [12] William Barker, Murugiah Souppaya, and William Newhouse. *Migration to Post-Quantum Cryptography*. NIST, Aug 2021.
- [13] Matthew PA Fisher. Are we quantum computers, or merely clever robots. *Int. J. Mod. Phys. B*, 31(1743001):10–1142, 2017.
- [14] Iulia M Georgescu, Sahel Ashhab, and Franco Nori. Quantum simulation. *Reviews of Modern Physics*, 86(1):153, 2014.
- [15] Pascal Scholl, Michael Schuler, Hannah J Williams, Alexander A Eberharter, Daniel Barredo, Kai-Niklas Schymik, Vincent Lienhard, Louis-Paul Henry, Thomas C Lang, Thierry Lahaye, et al. Quantum simulation of 2d antiferromagnets with hundreds of rydberg atoms. *Nature*, 595(7866):233–238, 2021.
- [16] Norbert Kalb, Andreas A Reiserer, Peter C Humphreys, Jacob JW Bakermans, Sten J Kamerling, Naomi H Nickerson, Simon C Benjamin, Daniel J Twitchen, Matthew Markham, and Ronald Hanson. Entanglement distillation between solid-state quantum network nodes. *Science*, 356(6341):928–932, 2017.
- [17] Sheng-Kai Liao, Wen-Qi Cai, Wei-Yue Liu, Liang Zhang, Yang Li, Ji-Gang Ren, Juan Yin, Qi Shen, Yuan Cao, Zheng-Ping Li, et al. Satellite-to-ground quantum key distribution. *Nature*, 549(7670):43–47, 2017.
- [18] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 517–526. IEEE, 2009.
- [19] Zheshen Zhang and Quntao Zhuang. Distributed quantum sensing. *Quantum Science and Technology*, 6(4):043001, 2021.
- [20] Kunihiro Inomata, Zhirong Lin, Kazuki Koshino, William D Oliver, Jaw-Shen Tsai, Tsuyoshi Yamamoto, and Yasunobu Nakamura. Single microwave-photon detector using an artificial λ -type three-level system. *Nature Communications*, 7(1):1–7, 2016.
- [21] Christian L Degen, F Reinhard, and Paola Cappellaro. Quantum sensing. *Reviews of Modern Physics*, 89(3):035002, 2017.
- [22] Kenichi Ohno, F Joseph Heremans, Lee C Bassett, Bryan A Myers, David M Toyli, Ania C Bleszynski Jayich, Christopher J Palmstrøm, and David D Awschalom. Engineering shallow spins in diamond with nitrogen delta-doping. *Applied Physics Letters*, 101(8):082413, 2012.
- [23] Marcus W Doherty, Neil B Manson, Paul Delaney, Fedor Jelezko, Jörg Wrachtrup, and Lloyd CL Hollenberg. The nitrogen-vacancy colour centre in diamond. *Physics Reports*, 528(1):1–45, 2013.

- [24] Xing Rong, Jianpei Geng, Fazhan Shi, Ying Liu, Kebiao Xu, Wenchao Ma, Fei Kong, Zhen Jiang, Yang Wu, and Jiangfeng Du. Experimental fault-tolerant universal quantum gates with solid-state spins under ambient conditions. *Nature Communications*, 6(1):1–7, 2015.
- [25] Arne Barfuss, Johannes Kölbl, Lucas Thiel, Jean Teissier, Mark Kasperczyk, and Patrick Maletinsky. Phase-controlled coherent dynamics of a single spin under closed-contour interaction. *Nature Physics*, 14(11):1087–1091, 2018.
- [26] Marcus W Doherty, Victor M Acosta, Andrey Jarmola, Michael SJ Barson, Neil B Manson, Dmitry Budker, and Lloyd CL Hollenberg. Temperature shifts of the resonances of the nv- center in diamond. *Physical Review B*, 90(4):041201, 2014.
- [27] Gary Wolfowicz, F Joseph Heremans, Christopher P Anderson, Shun Kanai, Hosung Seo, Adam Gali, Giulia Galli, and David D Awschalom. Quantum guidelines for solid-state spin defects. *Nature Reviews Materials*, pages 1–20, 2021.
- [28] Aedan Gardill, Matthew C Cambria, and Shimon Kolkowitz. Fast relaxation on qutrit transitions of nitrogen-vacancy centers in nanodiamonds. *Physical Review Applied*, 13(3):034010, 2020.
- [29] ED Herbschleb, H Kato, Y Maruyama, T Danjo, T Makino, S Yamasaki, I Ohki, K Hayashi, H Morishita, M Fujiwara, et al. Ultra-long coherence times amongst room-temperature solid-state spins. *Nature Communications*, 10(1):1–6, 2019.
- [30] Erik Bauch, Connor A Hart, Jennifer M Schloss, Matthew J Turner, John F Barry, Pauli Kehayias, Swati Singh, and Ronald L Walsworth. Ultralong dephasing times in solid-state spin ensembles via quantum control. *Physical Review X*, 8(3):031025, 2018.
- [31] Gopalakrishnan Balasubramanian, Philipp Neumann, Daniel Twitchen, Matthew Markham, Roman Kolesov, Norikazu Mizuochi, Junichi Isoya, Jocelyn Achard, Johannes Beck, Julia Tissler, et al. Ultralong spin coherence time in isotopically engineered diamond. *Nature Materials*, 8(5):383–387, 2009.
- [32] A Jarmola, VM Acosta, K Jensen, S Chemerisov, and D Budker. Temperature-and magnetic-field-dependent longitudinal spin relaxation in nitrogen-vacancy ensembles in diamond. *Physical Review Letters*, 108(19):197601, 2012.
- [33] Mohamed H Abobeih, Julia Cramer, Michiel A Bakker, Norbert Kalb, Matthew Markham, Daniel J Twitchen, and Tim H Taminiau. One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment. *Nature Communications*, 9(1):1–8, 2018.
- [34] LJ Rogers, RL McMurtrie, MJ Sellars, and NB Manson. Time-averaging within the excited state of the nitrogen-vacancy centre in diamond. *New Journal of Physics*, 11(6):063007, 2009.

- [35] Jeronimo R Maze, Adam Gali, Emre Togan, Yiwen Chu, Alexei Trifonov, Efthimios Kaxiras, and Mikhail D Lukin. Properties of nitrogen-vacancy centers in diamond: the group theoretic approach. *New Journal of Physics*, 13(2):025025, 2011.
- [36] Lucio Robledo, Lilian Childress, Hannes Bernien, Bas Hensen, Paul FA Alkemade, and Ronald Hanson. High-fidelity projective read-out of a solid-state spin quantum register. *Nature*, 477(7366):574–578, 2011.
- [37] Dirk Dubbers and Hans-Jürgen Stöckmann. *Quantum physics: the bottom-up approach: from the simple two-level system to irreducible representations*. Springer Science & Business Media, 2013.
- [38] X-X Yuan, L He, S-T Wang, D-L Deng, F Wang, W-Q Lian, X Wang, C-H Zhang, H-L Zhang, X-Y Chang, et al. Observation of topological links associated with hopf insulators in a solid-state quantum simulator. *Chinese Physics Letters*, 34(6):060302, 2017.
- [39] GD Fuchs, VV Dobrovitski, DM Toyli, FJ Heremans, and DD Awschalom. Gigahertz dynamics of a strongly driven single quantum spin. *Science*, 326(5959):1520–1522, 2009.
- [40] Yang Wu, Wenquan Liu, Jianpei Geng, Xingrui Song, Xiangyu Ye, Chang-Kui Duan, Xing Rong, and Jiangfeng Du. Observation of parity-time symmetry breaking in a single-spin system. *Science*, 364(6443):878–880, 2019.
- [41] Wenquan Liu, Yang Wu, Chang-Kui Duan, Xing Rong, and Jiangfeng Du. Dynamically encircling an exceptional point in a real quantum system. *Physical Review Letters*, 126(17):170506, 2021.
- [42] R Hanson, VV Dobrovitski, AE Feiguin, O Gywat, and DD Awschalom. Coherent dynamics of a single spin interacting with an adjustable spin bath. *Science*, 320(5874):352–355, 2008.
- [43] G de Lange, ZH Wang, D Riste, VV Dobrovitski, and R Hanson. Universal dynamical decoupling of a single solid-state spin from a spin bath. *Science*, 330(6000):60–63, 2010.
- [44] Tim H Taminiau, Julia Cramer, Toeno van der Sar, Viatcheslav V Dobrovitski, and Ronald Hanson. Universal control and error correction in multi-qubit spin registers in diamond. *Nature Nanotechnology*, 9(3):171–176, 2014.
- [45] Jingfu Zhang, Swathi S Hegde, and Dieter Suter. Efficient implementation of a quantum algorithm in a single nitrogen-vacancy center of diamond. *Physical Review Letters*, 125(3):030501, 2020.
- [46] J Randall, CE Bradley, FV van der Gronden, A Galicia, MH Abobeih, M Markham, DJ Twitchen, F Machado, NY Yao, and TH Taminiau. Observation of a many-body-localized discrete time crystal with a programmable spin-based quantum simulator. *arXiv preprint arXiv:2107.00736*, 2021.

- [47] Marcus W Doherty, Carlos A Meriles, Audrius Alkauskas, Helmut Fedder, Matthew J Sellars, and Neil B Manson. Towards a room-temperature spin quantum bus in diamond via electron photoionization, transport, and capture. *Physical Review X*, 6(4):041035, 2016.
- [48] Hannes Bernien, Bas Hensen, Wolfgang Pfaff, Gerwin Koolstra, Machiel S Blok, Lucio Robledo, Tim H Taminiiau, Matthew Markham, Daniel J Twitchen, Lilian Childress, et al. Heralded entanglement between solid-state qubits separated by three metres. *Nature*, 497(7447):86–90, 2013.
- [49] Matteo Pompili, Sophie LN Hermans, Simon Baier, Hans KC Beukers, Peter C Humphreys, Raymond N Schouten, Raymond FL Vermeulen, Marijn J Tiggelman, Laura dos Santos Martins, Bas Dirkse, et al. Realization of a multinode quantum network of remote solid-state qubits. *Science*, 372(6539):259–264, 2021.
- [50] Bas Hensen, Hannes Bernien, Anaïs E Dréau, Andreas Reiserer, Norbert Kalb, Machiel S Blok, Just Ruitenberg, Raymond FL Vermeulen, Raymond N Schouten, Carlos Abellán, et al. Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682–686, 2015.
- [51] Noel H Wan, Sara Mouradian, and Dirk Englund. Two-dimensional photonic crystal slab nanocavities on bulk single-crystal diamond. *Applied Physics Letters*, 112(14):141102, 2018.
- [52] Noel H Wan, Brendan J Shields, Donggyu Kim, Sara Mouradian, Benjamin Lienhard, Michael Walsh, Hassaram Bakhru, Tim Schröder, and Dirk Englund. Efficient extraction of light from a nitrogen-vacancy center in a diamond parabolic reflector. *Nano Letters*, 18(5):2787–2793, 2018.
- [53] Daniel Riedel, Immo Söllner, Brendan J Shields, Sebastian Starosielec, Patrick Appel, Elke Neu, Patrick Maletinsky, and Richard J Warburton. Deterministic enhancement of coherent photon generation from a nitrogen-vacancy center in ultrapure diamond. *Physical Review X*, 7(3):031040, 2017.
- [54] Estefani Marchiori, Lorenzo Ceccarelli, Nicola Rossi, Luca Lorenzelli, Christian L Degen, and Martino Poggio. Nanoscale magnetic field imaging for 2d materials. *Nature Reviews Physics*, pages 1–12, 2021.
- [55] Jean-Philippe Tetienne, Nikolai Dontschuk, David A Broadway, Alastair Stacey, David A Simpson, and Lloyd CL Hollenberg. Quantum imaging of current flow in graphene. *Science Advances*, 3(4):e1602429, 2017.
- [56] Edlyn V Levine, Matthew J Turner, Pauli Kehayias, Connor A Hart, Nicholas Langelier, Raisa Trubko, David R Glenn, Roger R Fu, and Ronald L Walsworth. Principles and techniques of the quantum diamond microscope. *Nanophotonics*, 8(11):1945–1973, 2019.

- [57] Piotr Szańkowski, Guy Ramon, Jan Krzywda, Damian Kwiatkowski, et al. Environmental noise spectroscopy with qubits subjected to dynamical decoupling. *Journal of Physics: Condensed Matter*, 29(33):333001, 2017.
- [58] Aashish A Clerk, Michel H Devoret, Steven M Girvin, Florian Marquardt, and Robert J Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Reviews of Modern Physics*, 82(2):1155, 2010.
- [59] S Hsieh, P Bhattacharyya, C Zu, T Mittiga, TJ Smart, F Machado, B Kobrin, TO Höhn, NZ Rui, M Kamrani, et al. Imaging stress and magnetism at high pressures using a nanoscale quantum sensor. *Science*, 366(6471):1349–1354, 2019.
- [60] Akihiro Kuwahata, Takahiro Kitaizumi, Kota Saichi, Takumi Sato, Ryuji Igarashi, Takeshi Ohshima, Yuta Masuyama, Takayuki Iwasaki, Mutsuko Hatano, Fedor Jelezko, et al. Magnetometer with nitrogen-vacancy center in a bulk diamond for detecting magnetic nanoparticles in biomedical applications. *Scientific Reports*, 10(1):1–9, 2020.
- [61] Matthew Pelliccione, Alec Jenkins, Preeti Ovarthaiyapong, Christopher Reetz, Eve Emmanouilidou, Ni Ni, and Ania C Bleszynski Jayich. Scanned probe imaging of nanoscale magnetism at cryogenic temperatures with a single-spin quantum sensor. *Nature Nanotechnology*, 11(8):700–705, 2016.
- [62] Gopalakrishnan Balasubramanian, Andrii Lazariiev, Sri Ranjini Arumugam, and De-wen Duan. Nitrogen-vacancy color center in diamond—emerging nanoscale applications in bioimaging and biosensing. *Current opinion in chemical biology*, 20:69–77, 2014.
- [63] Georg Kucsko, Peter C Maurer, Norman Ying Yao, Michael Kubo, Hyun Jong Noh, Po Kam Lo, Hongkun Park, and Mikhail D Lukin. Nanometre-scale thermometry in a living cell. *Nature*, 500(7460):54–58, 2013.
- [64] Shu-Jung Yu, Ming-Wei Kang, Huan-Cheng Chang, Kuan-Ming Chen, and Yueh-Chung Yu. Bright fluorescent nanodiamonds: no photobleaching and low cytotoxicity. *Journal of the American Chemical Society*, 127(50):17604–17605, 2005.
- [65] J-P Tetienne, LT Hall, AJ Healey, GAL White, M-A Sani, F Separovic, and LCL Hollenberg. Prospects for nuclear spin hyperpolarization of molecular samples using nitrogen-vacancy centers in diamond. *Physical Review B*, 103(1):014434, 2021.
- [66] A Ajoy, A Sarkar, E Druga, P Zangara, D Pagliero, CA Meriles, and JA Reimer. Low-field microwave-mediated optical hyperpolarization in optically pumped diamond. *Journal of Magnetic Resonance*, 331:107021, 2021.
- [67] Christopher G Yale, F Joseph Heremans, Brian B Zhou, Adrian Auer, Guido Burkard, and David D Awschalom. Optical manipulation of the berry phase in a solid-state spin qubit. *Nature Photonics*, 10(3):184–189, 2016.

- [68] Christopher G Yale, Bob B Buckley, David J Christle, Guido Burkard, F Joseph Heremans, Lee C Bassett, and David D Awschalom. All-optical control of a solid-state spin using coherent dark states. *Proceedings of the National Academy of Sciences*, 110(19):7595–7600, 2013.
- [69] ML Goldman, A Sipahigil, MW Doherty, NY Yao, SD Bennett, M Markham, DJ Twitchen, NB Manson, A Kubanek, and MD Lukin. Phonon-induced population dynamics and intersystem crossing in nitrogen-vacancy centers. *Physical Review Letters*, 114(14):145502, 2015.
- [70] Nikolay V Vitanov, Andon A Rangelov, Bruce W Shore, and Klaas Bergmann. Stimulated raman adiabatic passage in physics, chemistry, and beyond. *Reviews of Modern Physics*, 89(1):015006, 2017.
- [71] K Kim, M-S Chang, S Korenblit, R Islam, EE Edwards, JK Freericks, G-D Lin, L-M Duan, and C Monroe. Quantum simulation of frustrated ising spins with trapped ions. *Nature*, 465(7298):590–593, 2010.
- [72] Arnab Das and Bikas K Chakrabarti. Colloquium: Quantum annealing and analog quantum computation. *Reviews of Modern Physics*, 80(3):1061, 2008.
- [73] Erik Torrontegui, Sara Ibáñez, Sofia Martínez-Garaot, Michele Modugno, Adolfo del Campo, David Guéry-Odelin, Andreas Ruschhaupt, Xi Chen, and Juan Gonzalo Muga. Shortcuts to adiabaticity. *Advances in Atomic, Molecular, and Optical Physics*, 62:117–169, 2013.
- [74] Mustafa Demirplak and Stuart A Rice. Adiabatic population transfer with control fields. *The Journal of Physical Chemistry A*, 107(46):9937–9945, 2003.
- [75] Michael V Berry. Transitionless quantum driving. *Journal of Physics A: Mathematical and Theoretical*, 42(36):365303, 2009.
- [76] Mark G Bason, Matthieu Viteau, Nicola Malossi, Paul Huillery, Ennio Arimondo, Donatella Ciampini, Rosario Fazio, Vittorio Giovannetti, Riccardo Mannella, and Oliver Morsch. High-fidelity quantum driving. *Nature Physics*, 8(2):147–152, 2012.
- [77] J Zhang, JH Shim, I Niemeyer, T Taniguchi, T Teraji, H Abe, S Onoda, T Yamamoto, T Ohshima, J Isoya, et al. Experimental implementation of assisted quantum adiabatic passage in a single spin. *Physical Review Letters*, 110(24):240501, 2013.
- [78] J-F Schaff, X-L Song, P Capuzzi, P Vignolo, and G Labeyrie. Shortcut to adiabaticity for an interacting bose-einstein condensate. *EPL (Europhysics Letters)*, 93(2):23001, 2011.
- [79] Shuoming An, Dingshun Lv, Adolfo Del Campo, and Kihwan Kim. Shortcuts to adiabaticity by counterdiabatic driving for trapped-ion displacement in phase space. *Nature Communications*, 7(1):1–5, 2016.

- [80] Alexandre Baksic, Hugo Ribeiro, and Aashish A Clerk. Speeding up adiabatic quantum state transfer by using dressed states. *Physical Review Letters*, 116(23):230503, 2016.
- [81] Hung-Shen Chang, Youpeng Zhong, Audrey Bienfait, Ming-Han Chou, Christopher R Conner, Étienne Dumur, Joel Grebel, Gregory A Peairs, Rhys G Povey, Kevin J Satzinger, et al. Remote entanglement via adiabatic passage using a tunably dissipative quantum communication system. *Physical Review Letters*, 124(24):240502, 2020.
- [82] Erik Sjöqvist. Geometric phases in quantum information. *International Journal of Quantum Chemistry*, 115(19):1311–1326, 2015.
- [83] Artur Ekert, Marie Ericsson, Patrick Hayden, Hitoshi Inamori, Jonathan A Jones, Daniel KL Oi, and Vlatko Vedral. Geometric quantum computation. *Journal of Modern Optics*, 47(14-15):2501–2513, 2000.
- [84] Simon Berger, Marek Pechal, Abdufarrukh A Abdumalikov Jr, Christopher Eichler, Lars Steffen, Alexey Fedorov, Andreas Wallraff, and Stefan Filipp. Exploring the effect of noise on the berry phase. *Physical Review A*, 87(6):060303, 2013.
- [85] Lara Faoro, Jens Siewert, and Rosario Fazio. Non-abelian holonomies, charge pumping, and quantum computation with josephson junctions. *Physical Review Letters*, 90(2):028301, 2003.
- [86] Erik Sjöqvist, Dian-Min Tong, L Mauritz Andersson, Björn Hessmo, Markus Johansson, and Kuldip Singh. Non-adiabatic holonomic quantum computation. *New Journal of Physics*, 14(10):103035, 2012.
- [87] Yuhei Sekiguchi, Naeko Niikura, Ryota Kuroiwa, Hiroki Kano, and Hideo Kosaka. Optical holonomic single quantum gates with a geometric spin under a zero field. *Nature Photonics*, 11(5):309–314, 2017.
- [88] Abdufarrukh A Abdumalikov Jr, Johannes M Fink, Kristinn Juliusson, Marek Pechal, Simon Berger, Andreas Wallraff, and Stefan Filipp. Experimental realization of non-abelian non-adiabatic geometric gates. *Nature*, 496(7446):482–485, 2013.
- [89] Guanru Feng, Guofu Xu, and Guilu Long. Experimental realization of nonadiabatic holonomic quantum computation. *Physical Review Letters*, 110(19):190501, 2013.
- [90] C Zu, W-B Wang, L He, W-G Zhang, C-Y Dai, F Wang, and L-M Duan. Experimental realization of universal geometric quantum gates with solid-state spins. *Nature*, 514(7520):72–75, 2014.
- [91] Silvia Arroyo-Camejo, Andrii Lazariev, Stefan W Hell, and Gopalakrishnan Balasubramanian. Room temperature high-fidelity holonomic single-qubit gate on a solid-state spin. *Nature Communications*, 5(1):1–5, 2014.
- [92] Erik Sjöqvist. Nonadiabatic holonomic single-qubit gates in off-resonant λ systems. *Physics Letters A*, 380(1-2):65–67, 2016.

- [93] Guido Burkard, Vladislav O Shkolnikov, and David D Awschalom. Designing a cavity-mediated quantum cphase gate between nv spin qubits in diamond. *Physical Review B*, 95(20):205420, 2017.
- [94] Steffen Steinert, Florian Dolde, Philipp Neumann, Andrew Aird, Boris Naydenov, Gopalakrishnan Balasubramanian, Fedor Jelezko, and Joerg Wrachtrup. High sensitivity magnetic imaging using an array of spins in diamond. *Review of Scientific Instruments*, 81(4):043705, 2010.
- [95] Linh My Pham, David Le Sage, Paul L Stanwix, Tsun Kwan Yeung, D Glenn, Alexei Trifonov, Paola Cappellaro, Philip R Hemmer, Mikhail D Lukin, Hongkun Park, et al. Magnetic field imaging with nitrogen-vacancy ensembles. *New Journal of Physics*, 13(4):045021, 2011.
- [96] Lior Ella, Asaf Rozen, John Birkbeck, Moshe Ben-Shalom, David Perello, Johanna Zultak, Takashi Taniguchi, Kenji Watanabe, Andre K Geim, Shahal Ilani, et al. Simultaneous voltage and current density imaging of flowing electrons in two dimensions. *Nature Nanotechnology*, 14(5):480–487, 2019.
- [97] J Beyer, T Schurig, A Lüdge, and H Riemann. Squid-nde of semiconductor samples with high spatial resolution. *Superconductor Science and Technology*, 13(5):532, 2000.
- [98] J Beyer, D Drung, and T Schurig. Squid photoscanning: an imaging technique for nde of semiconductor wafers and devices based on photomagnetic detection. *IEEE Transactions on Applied Superconductivity*, 11(1):1162–1167, 2001.
- [99] Y Nakatani, T Hayashi, and H Itozaki. Laser-squid microscope for noncontact evaluation of solar cell. *Physica C: Superconductivity and its Applications*, 471(21-22):1249–1252, 2011.
- [100] Mark JH Ku, Tony X Zhou, Qing Li, Young J Shin, Jing K Shi, Claire Burch, Laurel E Anderson, Andrew T Pierce, Yonglong Xie, Assaf Hamo, et al. Imaging viscous flow of the dirac fluid in graphene. *Nature*, 583(7817):537–541, 2020.
- [101] Bradley J Roth, Nestor G Sepulveda, and John P Wikswo Jr. Using a magnetometer to image a two-dimensional current distribution. *Journal of Applied Physics*, 65(1):361–372, 1989.
- [102] FHL Koppens, T Mueller, Ph Avouris, AC Ferrari, MS Vitiello, and M Polini. Photodetectors based on graphene, other two-dimensional materials and hybrid systems. *Nature Nanotechnology*, 9(10):780–793, 2014.
- [103] Kin Fai Mak, Kathryn L McGill, Jiwoong Park, and Paul L McEuen. The valley hall effect in mos2 transistors. *Science*, 344(6191):1489–1492, 2014.
- [104] Ziliang Ye, Dezheng Sun, and Tony F Heinz. Optical manipulation of valley pseudospin. *Nature Physics*, 13(1):26–29, 2017.

- [105] J-P Tetienne, RW De Gille, DA Broadway, T Teraji, SE Lillie, JM McCoe, N Dontschuk, LT Hall, A Stacey, DA Simpson, et al. Spin properties of dense near-surface ensembles of nitrogen-vacancy centers in diamond. *Physical Review B*, 97(8):085402, 2018.
- [106] Kibum Kang, Saien Xie, Lujie Huang, Yimo Han, Pinshane Y Huang, Kin Fai Mak, Cheol-Joo Kim, David Muller, and Jiwoong Park. High-mobility three-atom-thick semiconducting films with wafer-scale homogeneity. *Nature*, 520(7549):656–660, 2015.
- [107] Kibum Kang, Kan-Heng Lee, Yimo Han, Hui Gao, Saien Xie, David A Muller, and Jiwoong Park. Layer-by-layer assembly of two-dimensional materials into wafer-scale heterostructures. *Nature*, 550(7675):229–233, 2017.
- [108] Hichem Hattab, Alpha T N’Diaye, Dirk Wall, Claudius Klein, Giriraj Jnawali, Johann Coraux, Carsten Busse, Raoul van Gastel, Bene Poelsema, Thomas Michely, et al. Interplay of wrinkles, strain, and lattice parameter in graphene on iridium. *Nano Letters*, 12(2):678–682, 2012.
- [109] Stephen J DeVience, Linh M Pham, Igor Lovchinsky, Alexander O Sushkov, Nir Bargill, Chinmay Belthangady, Francesco Casola, Madeleine Corbett, Huiliang Zhang, Mikhail Lukin, et al. Nanoscale nmr spectroscopy and imaging of multiple nuclear species. *Nature Nanotechnology*, 10(2):129–134, 2015.
- [110] I Lovchinsky, JD Sanchez-Yamagishi, EK Urbach, S Choi, S Fang, TI Andersen, K Watanabe, T Taniguchi, A Bylinskii, E Kaxiras, et al. Magnetic resonance spectroscopy of an atomically thin material using a single-spin qubit. *Science*, 355(6324):503–507, 2017.
- [111] Michele Buscema, Maria Barkelid, Val Zwiller, Herre SJ van der Zant, Gary A Steele, and Andres Castellanos-Gomez. Large and tunable photothermoelectric effect in single-layer mos₂. *Nano Letters*, 13(2):358–363, 2013.
- [112] Kedar Hippalgaonkar, Ying Wang, Yu Ye, Diana Y Qiu, Hanyu Zhu, Yuan Wang, Joel Moore, Steven G Louie, and Xiang Zhang. High thermoelectric power factor in two-dimensional crystals of mo s₂. *Physical Review B*, 95(11):115407, 2017.
- [113] Kamran Behnia. *Fundamentals of thermoelectricity*. OUP Oxford, 2015.
- [114] NV Vitanov, KA Suominen, and BW Shore. Creation of coherent atomic superpositions by fractional stimulated raman adiabatic passage. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 32(18):4535, 1999.
- [115] M Howard, J Twamley, C Wittmann, T Gaebel, F Jelezko, and J Wrachtrup. Quantum process tomography and linblad estimation of a solid-state qubit. *New Journal of Physics*, 8(3):33, 2006.
- [116] Joseph B Altepeter, Daniel FV James, and Paul G Kwiat. 4 qubit quantum state tomography. In *Quantum state estimation*, pages 113–145. Springer, 2004.

- [117] Milad Yarali, Xufei Wu, Tushar Gupta, Debjit Ghoshal, Lixin Xie, Zhuan Zhu, Hatem Brahmi, Jiming Bao, Shuo Chen, Tengfei Luo, et al. Effects of defects on the temperature-dependent thermal conductivity of suspended monolayer molybdenum disulfide grown by chemical vapor deposition. *Advanced Functional Materials*, 27(46):1704357, 2017.
- [118] Zhun-Yong Ong, Yongqing Cai, and Gang Zhang. Theory of substrate-directed heat dissipation for single-layer graphene and other two-dimensional crystals. *Physical Review B*, 94(16):165427, 2016.
- [119] Dumitru Dumcenco, Dmitry Ovchinnikov, Kolyo Marinov, Predrag Lazic, Marco Gibertini, Nicola Marzari, Oriol Lopez Sanchez, Yen-Cheng Kung, Daria Krasnozhan, Ming-Wei Chen, et al. Large-area epitaxial monolayer mos2. *ACS Nano*, 9(4):4611–4620, 2015.
- [120] Jung Jun Bae, Hye Yun Jeong, Gang Hee Han, Jaesu Kim, Hyun Kim, Min Su Kim, Byoung Hee Moon, Seong Chu Lim, and Young Hee Lee. Thickness-dependent in-plane thermal conductivity of suspended mos2 grown by chemical vapor deposition. *Nanoscale*, 9(7):2541–2547, 2017.
- [121] Dipankar Saha and Santanu Mahapatra. Analytical insight into the lattice thermal conductivity and heat capacity of monolayer mos2. *Physica E: Low-dimensional Systems and Nanostructures*, 83:455–460, 2016.
- [122] Jie Su, Zheng-tang Liu, Li-ping Feng, and Ning Li. Effect of temperature on thermal properties of monolayer mos2 sheet. *Journal of Alloys and Compounds*, 622:777–782, 2015.