ABSTRACT

Title of Dissertation: ADVANCING QUANTUM DIAMOND SENSING FOR DYNAMIC MAGNETIC IMAGING AND DIRECTIONAL RARE-EVENT DETECTION

Jiashen Tang
Doctor of Philosophy, 2026

Dissertation Directed by: Professor Ronald L. Walsworth
Department of Electrical and Computer Engineering

NV centers in diamond provide an optically accessible interface to microscopic physical information, including magnetic fields, crystal strain, temperature, and other local material or environmental properties. These capabilities have enabled a wide range of applications across the physical and life sciences. Motivated by the goal of extending what can be measured with NV diamond systems, this dissertation investigates how advanced quantum control, defect formation physics, and diamond material engineering can be used to access microscopic information in new regimes.

In the first part, I develop a pulsed quantum diamond microscope for wide-field imaging of static and dynamic magnetic fields. Motivated by the need to image weak, transient magnetic signals with both high sensitivity and high temporal resolution, this work combines double-quantum Ramsey magnetometry, spin-bath driving, and high-speed lock-in camera detection. The resulting platform demonstrates sub-millisecond magnetic imaging and stable averaging to resolve picotesla-scale fields while maintaining micrometer-scale spatial resolution. This provides a route toward measurements of

weak magnetic sources from electronic circuits, novel materials, and biological systems.

In the second part, I use NV centers to study particle-induced damage tracks in diamond. This work is driven by the development of quantum-diamond-based directional detectors for rare event interactions, while also being motivated by the need for improved damage-track modeling for paleodetectors and quantum material engineering. Using nitrogen-rich diamond and single ion implantation, I show that individual carbon ion damage events can produce spatially localized NV ensembles after annealing. These measurements provide an experimental benchmark for an implantation-and-annealing simulation framework, which is then used to assess how nanoscale directional information can remain encoded in, and be extracted from, the post-annealing NV distribution.

Finally, in the last part, I describe a collaborative effort with Element Six to engineer diamond materials for pulsed NV magnetometry. By advancing material synthesis and characterizing the tradeoffs among material properties, NV spin and optical properties, and experimental conditions, this work provides both new diamond sensors and guidance for material selection in NV quantum sensing. These results are especially relevant when sensor performance must be balanced against practical considerations such as available optical power, sensing volume, and experimental overhead.

ADVANCING QUANTUM DIAMOND SENSING FOR DYNAMIC MAGNETIC IMAGING AND DIRECTIONAL RARE-EVENT DETECTION

by

Jiashen Tang

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2026

Advisory Committee:
Professor Ronald Walsworth, Chair/Advisor
Professor John Cumings, Dean's Representative
Professor Edo Waks
Professor Anson Hook
Professor José Feijó

© Copyright by
Jiashen Tang
2026

Acknowledgments

Looking back on these years, I feel proud of the published papers and commercialization opportunities that came out of my PhD. I value these achievements both for the results themselves and for what the journey has taught me. It showed me that I can shape a scientific life that is distinct, diverse, and fulfilling. I learned that I can make academic contributions, carry ideas beyond the laboratory, and grow into a person who listens with empathy, connects openly, and creates opportunities for myself and others. I am grateful to the people and experiences that made this growth possible. Some shaped me through memorable turning points, changing how I think, work, and imagine what may be possible. Others carried me through the ordinary days of these years, giving them warmth, joy, and meaning.

My foremost gratitude goes to my PhD advisor, Ron Walsworth. Ron provided the foundation that touched many aspects of my PhD life. His support gave me the security to focus on the real work of graduate school. He created an environment that values freedom and trust: I was not simply assigned a narrow path to follow, but was given room to try things, find questions that interested me, and gradually build a direction of my own. Ron also built an influential platform and brought many good people together. Many of the opportunities I was able to pursue, from research collaborations to work beyond traditional academic research, were made possible because I could stand on that platform. At the same time, despite his seniority, he brought a youthful energy to the lab. His stories, banter, and genuine enjoyment of being around younger people made him not only an advisor I deeply respect, but also someone who could make an ordinary day feel livelier.

I would also like to express my appreciation for Jner Tzern Oon, whom I shared my first paper with. To this day, I remember less about the grind we went through together than about the empowering feeling it gave me: I realized that I could create something of value in academia. It was an early proof to others and to myself. Thank you, JJ, for opening the door to my academic independence.

Connor Hart became an important mentor to me – perhaps more than he realizes – as I led my first independent project. He brought a valuable combination of seriousness about science and lightness in the way he approached the work. When I was unsure whether I was capable of carrying the project forward, he offered practical ideas and advice that aroused in my heart a sense of steadiness that made the uncertainty feel manageable. Connor also shaped how I think about mentoring. His example showed me that one can hold high standards without making others feel small, and that encouragement can coexist with rigor and allows good work to emerge. These lessons will remain a touchstone for me: to be more patient, make room for uncertainty, and bring out energy and confidence in others.

A major turning point in my PhD came about a year and a half ago, when research funding suddenly felt less certain. The possibility of accelerated graduation made me think: What would distinguish me beyond the degree? What do I hope to build afterward? I knew I wanted work that felt meaningful to me – projects with tangible relevance to everyday life, and ones I could be passionate about. That search led to commercialization. I-Corps pushed me beyond academic research and into a world whose language and logic were initially unfamiliar. The process was demanding and often uncomfortable, but it also made me feel alive: conversations with people from different backgrounds taught me how to listen, ask questions, share an idea, and build goodwill. Finally, winning support from National I-Corps and the Maryland Innovation Initiative gave me a new sense of confidence that I could earn resources for an idea I believed in and create new opportunities.

I could not have made this journey alone. I am especially grateful to Daniel Ang and Niko Reed, who supported me in unique and important ways. Daniel's questions, advice, and steady reassurance often felt like the support of an older brother. When I felt defeated and frustrated, he reminded me that there was still a path forward and that I did not have to carry the whole weight of it by myself. Niko's genuine interest in the project and willingness to think alongside me made the work feel shared. Their

stories and enthusiasm gave me an encouraging sounding board.

I was fortunate to have Amanda Stein as a business mentor. She created important connections as I pitched the commercialization idea to funders. She brought humanity and kindness to a process that otherwise feels harsh. I also thank Prof. John Cumings, Dr. Thomas Templier, and Dan Nadash for the connections, guidance, and advocacy that helped carry the project forward.

These were some of the most important turning points of my PhD. But they were only milestones along the journey. What carried me from one milestone to the next were the ordinary days and weeks in between, made meaningful by people to whom I owe an equally deep gratitude.

To my mother and father, thank you for the many hours of long-distance calls that gave me a place to say what I could not express elsewhere.

To my sister, Yubi, thank you for your carefree energy, the example of your own excellence as a researcher, and the trust you gave me by asking me to officiate your wedding. Your presence in the U.S. over the past few years, and hearing you share both successes and hardships of your own PhD, have given me a sense of family. I am also grateful that you could visit our family when I could not.

To my high school friends, especially Xiaoying and Changlin, whose friendships have traveled with me not only through more than 12 years, but also across continents. Xiaoying, I treasure our many visits despite living 5-hours apart. I am grateful for the sense of care and closeness I feel when you are around. Changlin, you have somehow found a way to visit me at so many graduation moments, from high school through the end of my PhD. Thank you for always showing up – almost magically.

To my Walsworth labmates, past and present, thank you for the shared scientific growth, camaraderie, and the many forms of support that filled these years. I especially treasure the many outings with Katrijn, Andrew, Mason, Xiechen, Zechuan, and JJ. I enjoyed learning about your research as much as getting to know you personally.

To my housemates – Chenghao, Haowei, Zhiyin, Zechuan, Haoyi, and Huading – thank you for the shared food, movie nights, trips, and celebrations that filled our home, as well as your tolerance for the goofy side of me. To Yihui, Yaxin, and my other physics friends, thank you for the many trips, conversations, and moments of life shared alongside research.

To Sorah, Changmin, Joohyung, Chaohan, and friends in other labs, thank you for widening my scientific curiosity and for giving me opportunities to contribute to your experiments.

Special thanks to P.V. for your mentorship in life and beyond. You gave me the confidence to enter unfamiliar spaces with an open mind and curiosity. It was a wonderful experience visiting embassies and research institutes across Washington, D.C. Thank you for connecting me to a wider world through your generosity.

Table of Contents

Acknowledgements ii

Table of Contents vi

List of Tables ix

List of Figures x

Abbreviations xx

Chapter 1: Introduction 1
- 1.1 Nitrogen-vacancy centers in diamond . . . 1
- 1.2 Broadband DC magnetometry using NV-diamond . . . 7
- 1.3 NV-based directional rare-event detection . . . 10
- 1.4 Organization of this Dissertation . . . 15

Chapter 2: Dynamic Magnetic-Field Imaging with a Pulsed Quantum Diamond Microscope 17
- 2.1 Motivation and chapter overview . . . 17
- 2.2 Experimental platform and camera-based acquisition . . . 21
- 2.3 Pulsed magnetic-imaging protocol . . . 25
  - 2.3.1 DQ 4-Ramsey and camera implementation . . . 25
  - 2.3.2 Spin-bath driving . . . 28
  - 2.3.3 Differential measurement protocol . . . 30
- 2.4 Spin dephasing and magnetic sensitivity . . . 31
- 2.5 Spatial magnetic noise floor . . . 34
- 2.6 Static magnetic-field imaging . . . 38
- 2.7 Dynamic magnetic-field imaging . . . 41
- 2.8 Discussion and outlook . . . 45
  - 2.8.1 Technical limitations and paths forward . . . 46

Chapter 3: NV-Based Characterization of Single-Ion Damage Tracks in Diamond 49
- 3.1 Motivation and chapter overview . . . 49
- 3.2 Experimental platform and methods . . . 52
  - 3.2.1 Ion implantation . . . 52
  - 3.2.2 Confocal microscopy and NV measurement . . . 55
  - 3.2.3 Implantation and annealing simulation for NV formation . . . 55
  - 3.2.4 Head-tail directionality and machine learning . . . 57
- 3.3 Optical detection of individual ion-implantation events . . . 59
- 3.4 Experimentally benchmarked model of implantation, annealing, and NV formation . . . 60
  - 3.4.1 Implantation and annealing simulation . . . 60
  - 3.4.2 Testing sensitivity of NV-yield uncertainty to displacement energy and nitrogen concentration . . . 62

3.4.3 Short-range defect reactions beyond the nearest-neighbor approximation 66
3.5 Directional information retained after annealing 67
3.5.1 Counting-based head-tail asymmetry 67
3.5.2 Simulation-based inference of track direction 69
3.6 Spin coherence of implantation-created NV centers 70
3.7 Summary and outlook 72

Chapter 4: Low-Nitrogen Quantum Diamond: Engineering and Characterization 76
4.1 Motivation and chapter overview 76
4.2 Methods for sample synthesis and characterization 79
4.2.1 Key equations and parameters relevant to sensitivity 79
4.2.2 Diamond synthesis and treatment 83
4.2.3 NV-based wide-field strain mapping 85
4.2.4 Confocal-level characterization of NV spin and charge properties 85
4.3 Birefringence and nitrogen concentration of synthesized diamonds 86
4.4 NV spin dephasing time and strain contribution 87
4.5 Photoluminescence contrast and overhead time 92
4.6 NV charge state fraction 97
4.7 Discussion 99

Chapter 5: Conclusions and Outlook 102
5.1 Summary of dissertation contributions 102
5.2 Outlook for biomagnetic-field imaging 104
5.2.1 Prospects for biomineralization and electrically active cell imaging 104
5.2.2 Prospects for pollen tube studies 107
5.3 Outlook for NV-based damage-track reconstruction 110

Appendix A: Appendix for Dynamic QDM 113
A.1 Estimate of NV-NV interaction limited Ramsey dephasing time 113
A.2 Estimate of shot noise and quantization noise 114
A.3 Denoising of Picotesla-Scale Magnetic Image 118

Appendix B: Appendix for damage track study 120
B.1 SIMS measurement of nitrogen concentration 120
B.1.1 TOF-SIMS depth profiling 120
B.1.2 NV background PL as an independent estimate 121
B.1.3 Combined estimate and propagation into the simulations 122
B.2 Cluster decomposition and KMC annealing simulation details 122
B.3 NV photoluminescence (PL) analysis details 128
B.3.1 Data analysis procedure for confocal scans 128
B.3.2 Poisson analysis of multi-ion impacts 129
B.4 Improving diamond surface preparation 131
B.5 DFT calculations of defect interactions 134
B.5.1 Methods 134
B.5.2 Results 135

B.6 Additional spin echo measurement data 139

Appendix C: Appendix for low nitrogen quantum diamond 142
C.1 Ramsey measurement at high-birefringence location 142
C.2 NV spin dephasing time at nitrogen spin-bath limit 142
C.3 Full strain and stress maps 145

Appendix D: Interfacing Sample with Diamond in Ramsey QDM 147

Appendix E: Germinating Pollen Grains in the Ramsey QDM 152

List of Tables

3.1 The 13 summary features computed for each track and passed to the MNPE model. Here N is the number of defects and z_i is the position of the i-th defect along the track axis, measured relative to the track's centroid (each track is re-centered to $z = 0$). $N_\pm$ and $z_\pm$ denote the counts and positions of defects with $z_i > 0$ and $z_i < 0$, respectively. . 59

4.1 Principal material and measurement parameters characterized in this chapter, together with their relevance to practical pulsed Ramsey magnetometry. 80

B.1 Ar/Cl_2 ICP-RIE etching recipes and O_2 surface finishing step. 132

List of Figures

1.1 a) Crystal structure and energy levels of the NV center. (b) NV energy-level diagram. The enlarged view shows the electronic spin triplet ground state with zero-field splitting D and bias magnetic field B aligned to the sensing NV axis. 4

1.2 Schematics comparing NV magnetometry protocols based on continuous-wave optically detected magnetic resonance (CW-ODMR) and pulsed Ramsey. 8

1.3 Schematic of Ramsey magnetometry. A Ramsey free-induction decay measurement determines the dephasing time T_2^*, while a fixed-τ magnetometry curve converts small frequency or magnetic-field shifts into changes in normalized PL near the linear operating point. 9

1.4 Spin-independent dark matter–nucleon cross section as a function of dark matter mass, illustrating the neutrino floor for several detector target materials. Directional information provides a route to distinguish dark matter induced nuclear recoils from neutrino induced backgrounds in this regime. Adapted from Ref. [27]. 12

1.5 Schematic of the proposed solid-state directional dark matter detector readout pipeline. A rare event interaction is first detected and localized at the millimeter scale using phonon, charge, or photon readout. The corresponding damage track is then localized at the micrometer scale and reconstructed at the nanoscale using quantum-defect or X-ray-based imaging methods. Adapted from Ref. [30]. 14

2.1 (a) Imaging phantom magnetic fields using a quantum diamond microscope (QDM). 532-nm excitation laser light excites NV centers at the diamond chip surface. A planar, gold omega-loop delivers microwaves (MWs) to the diamond chip for NV spin-state control. Optical components collect NV photoluminescence onto either a photodiode or a high-speed camera (not shown). A bias magnetic field (not shown) is applied along one NV axis. An RF signal delivered by a nearby coil drives paramagnetic bath spins in the diamond to increase the ensemble NV spin dephasing time (T_2^*). A driven current in the phantom produces a magnetic-field pattern to be imaged with the QDM. An example image of simulated phantom magnetic signals, projected along the direction of a sensing NV axis with $5\,\mu$m standoff from the source to the NV surface, is shown in the inset. Scale bar: $50\,\mu$m. (d) Experimental apparatus. A flip mirror can route the NV photoluminescence to a photodiode for rapid measurements of ensemble NV spin properties and experimental optimization. 20

2.2 Optical image of the fabricated wire phantom. NV photoluminescence intensity under continuous laser illumination is collected using the QDM's camera. Semi-transparent white dashed traces of the phantom perimeter are added as a guide to the eye. The measured photoluminescence is reported in the camera unit, which is linearly proportional to the photon number. Scale bar: $50\,\mu$m. 23

2.3 (a) Sequence for magnetic imaging using the DQ 4-Ramsey protocol for NV interrogation together with spin-bath driving. The camera internal exposures alternate between the I and Q channel; on-chip circuitry accumulates differential results between two successive measurements in both channels. For example, the output in the I channel using the sequence in (a) will be $S_{I1} - S_{I2} + S_{I3} - S_{I4}$. (b) Left: Dual-tone MW resonant with $|m_s = 0\rangle \leftrightarrow |m_s = +1\rangle$ and $|m_s = 0\rangle \leftrightarrow |m_s = -1\rangle$ NV spin transitions is used to implement DQ pulses. Middle: Phase alternations of individual MW tones in a DQ 4-Ramsey sequence. The $\pm$X directions are defined in the effective Bloch sphere formed by the two NV spin levels $|m_s = 0\rangle$ and $|m_s = +1\rangle$ (or $|m_s = -1\rangle$). The duration of the DQ pulse is the time to change the NV spin state from $|m_s = 0\rangle$ to a superposition state with equal population between $|m_s = +1\rangle$ and $|m_s = -1\rangle$. Right: The permutation of MW-phase alternation in the Q channel results in NV fluorescence dynamics of opposite phase to that of the I channel. Differential results between these two measurements are output to an external frame. 26

2.4 (a) NV double electron-electron resonance (DEER) technique for measuring the $^{15}\mathrm{N}_s^0$ and free-electron spin resonance spectrum in the QDM diamond sensor. The total free evolution intervals between the initial and final single-quantum (SQ) $\pi/2$ pulse on the NV spins is chosen to match the ensemble NV spin-echo T_2 for optimal sensitivity to bath-spin resonances. The frequency of the SQ π pulse on the bath spins is swept to measure the NV DEER spectrum by the effect on the NV ensemble fluorescence contrast (see Ref. [57] for details). (b) Energy level diagram of the coupled electron-nuclear spin system in $^{15}\mathrm{N}_s^0$. (c) Simulated $^{15}\mathrm{N}_s^0$ and free-electron ($g = 2$) spin resonance spectrum (red) overlaid with NV DEER measurement results (blue) and offset vertically for clarity. The simulated resonant frequencies are broadened by a Lorentzian function with linewidth and relative contrast chosen to approximate the measurement data. 30

2.5 Measured magnetic sensing performance of the Ramsey-based QDM. Median per-pixel T_2^* is 2.2(1) µs and 1.2(1) µs with and without spin-bath driving. Median per-pixel η is 4.1(1) nT/$\sqrt{\mathrm{Hz}}$ and 9.4(1) nT/$\sqrt{\mathrm{Hz}}$ with and without spin-bath driving. 32

2.6 Spatial magnetic noise floor. Measurement data are obtained using the differential measurement protocol. Top: Images of magnetic noise with different acquisition times and pixel binning (see the bottom plot). Here, the acquisition time accounts for the duration of time spent acquiring measurements, but excludes the implementation-dependent time required to transfer and store data from the camera to a host computer. The magnetic noise is distributed randomly without significant spatial correlation. Field of view: $270 \times 270\,\mu\mathrm{m}^2$. Scale bar: 50 µm. Bottom: Spatial magnetic noise floor $\sigma_{spatial}$ as a function of acquisition time T_{acq} for no pixel binning and 3×3 pixel binning. $\sigma_{spatial}$ is computed as the standard deviation of magnetic-field measurements using the entire image after any pixel binning. A dashed black line depicts power law scaling behavior $\propto T_{acq}^{-1/2}$ as a guide to the eye. 36

2.7 Additional pixel binning results of spatial magnetic noise floor in QDM images. Measurement data are obtained using the differential measurement protocol. Upper six panels: Images of the background magnetic noise with different acquisition times and pixel binning (values indicated in lower panel). Scale bar: 50 µm. Lower panel: Spatial magnetic noise floor $\sigma_{spatial}$ as a function of acquisition time T_{acq}. A dashed black line depicts power law scaling behavior $\propto T_{acq}^{-1/2}$ as a guide to the eye. 37
2.8 (a) Simulated (Left) and measured (Right) phantom static magnetic-field patterns using the Ramsey-based QDM and a differential measurement protocol. The measurement is conducted with a 25 mV voltage (436 nA current) applied to the phantom and time-averaged for about 200 s. To compute SNR, the peak-to-peak amplitude difference of the magnetic-field pattern (solid box) is divided by the spatial magnetic noise obtained from a source-free region (dashed box). The simulated phantom magnetic-field pattern is obtained using the applied current and geometry of the phantom. To mimic the parameters of the imaging system, the simulated magnetic fields are further averaged over a depth of $10\,\mu$m (the NV-layer thickness) and binned with a pixel area of $1.9 \times 1.9\,\mu\text{m}^2$. NV-phantom standoff distance is the only free parameter, and is tuned to a value of 5 µm to best match the measured magnetic-field pattern. (b) Demonstration of a picotesla-scale QDM magnetic image. Left: After about $T_{acq} = 3600$ s ($T_{wall} \approx 7$ h) time-averaging, a magnetic-field pattern with a peak-to-peak amplitude difference of about 300 pT is resolved with $\mathrm{SNR} \approx 3$. Right: Applying a non-local mean denoising technique enhances image SNR by about $10\times$. All images are not binned and scale bars are 50 µm for (a, b). 39
2.9 Time-resolved imaging of dynamic magnetic fields using the Ramsey-based QDM and a differential measurement protocol. A voltage trace mimicking a synthetic, broadband human cardiac signal is applied to the phantom. Sets of QDM magnetic images are collected with temporal resolution of about 2 ms and binned with 3×3 pixels. The experiments are repeated 400 times to allow signal averaging. (a) Temporal variation of the peak magnetic field from the signal-averaged images is overlaid with the applied stimulus. A voltage-to-magnetic-field scaling factor of 274.26 nT/V is applied to the entire voltage trace (see main text). A subset of peak magnetic-field amplitudes from the acquired images are displayed, indicating good agreement between the applied voltages and magnetic measurements. (b) Selected magnetic-field images at time points labeled in (a). Each image has been signal averaged. The spatial magnetic noise at time (vi) is about 3 nT. 42

2.10 QDM magnetic imaging with sub-millisecond time resolution. Camera external frame rate is set to $F_s \approx 1.7$ kHz (i.e., time step between frames of about 0.6 ms). Magnetic signals are produced by applying single-tone AC signal voltages to the phantom. AC frequencies of 200 Hz, 500 Hz, and 750 Hz are used for different runs of the demonstration experiment. (a, b) Example time-domain magnetic-measurement data (500 Hz signal) from a single pixel in the QDM image, selected to have significant amplitude variation over a signal cycle. (c) Time-domain data from the single pixel is Fourier transformed and displayed, here for all three AC frequencies used in the demonstration experiments. The resulting spectrum is in good agreement with the applied signal frequencies; and with a simulation of the QDM measurement protocol that leads to a reduction in spectral peak amplitude near the Nyquist frequency ($F_s/2$). 44
2.11 Spatial magnetic noise floor $\sigma_{spatial}$ as a function of acquisition time T_{acq}. The differential measurement protocol is not employed and thus the utility of time averaging is limited. No pixel binning is applied. 46

3.1 Schematic of damage track formation and multi-scale readout in nitrogen-rich HPHT diamond. An incoming particle induces a carbon nuclear recoil, generating a cascade of vacancies and interstitials. The recoil also produces phonons that enable identification of candidate interaction events at the millimeter scale (single diamond chip). During annealing, vacancies diffuse and either recombine with interstitials, form vacancy clusters, or combine with substitutional nitrogen to create optically-active NV centers near the original track. The resulting NV photoluminescence (PL) enables localization at the micrometer scale, while the spatial distribution of NV centers and residual defect-induced strain enables nanoscale reconstruction of track morphology and inference of the incoming particle direction. 53
3.2 Schematic of ion implantation geometry. The diamond surface is tilted by $\sim 11^\circ$ relative to the incident carbon ion beam to suppress channeling effects. 54
3.3 Optical characterization of NV formation. a) Wide-field confocal NV PL map acquired at coarse spatial resolution ($\sim$500 nm). Bright PL spots correspond to NV ensembles formed at individual implantation sites, arranged in a regular grid consistent with the implantation pitch. Candidate sites selected for further high-resolution 3D confocal scans are marked with green crosses. The NV number is estimated from the PL intensity relative to calibrated single NV. Background NV counts in the HPHT sample are $\lesssim 2$. b) Representative high-resolution confocal scans of four implantation sites (fixed Z slice). Spatially separated PL spots are observed and are attributed to multi-ion implantation events, consistent with statistical analysis. c) Distribution of the number of NV centers per implanted ion (N_{NV}), extracted from integrated PL intensity (see main text). Data are collected from 222 sites, yielding a mean NV count of (17.5 ± 0.38), where the uncertainty denotes one standard error of the mean; the standard deviation is 5.7. 61

3.4 Modeling NV formation. a) Example damage track before and after annealing. Prior to annealing, vacancies (V) generated by an 800 keV carbon ion implantation are shown. During annealing, vacancies are consumed through recombination with interstitials, aggregation into vacancy clusters, or capture by substitutional nitrogen (not shown) to form NV centers. The resulting NV distribution after annealing is shown. b) Histogram of the number of NV centers per simulated implantation event. A mean NV number of 23.9 ± 1.1 (standard error of mean; standard deviation: 11.5) is obtained from 112 KMC-annealed tracks. A correction accounting for second-nearest-neighbor (2NN) defect interactions (DFT Corr.) reduces the predicted yield to 16.6 ± 0.7 NVs per event, in close agreement with experiment (Fig. 3.3(c)). 63
3.5 Sensitivity of vacancy production to displacement energy E_d. Mean number of vacancies generated per 800 keV carbon recoil as a function of the threshold displacement energy used in SIIMPL simulations. Each data point represents the average over 5000 simulated damage tracks. 64
3.6 Impact of nitrogen concentration on annealing simulations. (a) Mean NV yield as a function of nitrogen concentration $[N_s^0]$, showing weak dependence over a broad range (150–5000 ppm). (b) Representative spatial distribution of vacancies (gray) and resulting NV centers (red) at $[N_s^0] = 200\,\mathrm{ppm}$ and 5000 ppm. Higher nitrogen concentration leads to NV formation closer to the original damage track, while the overall NV yield remains largely unchanged. 65
3.7 Head-tail directional information from simulated damage tracks. a) Cumulative distribution of the asymmetry metric for vacancy (before annealing) and NV (after annealing) tracks. The asymmetry is defined as the ratio of defect counts in the tail and head regions of the track. For a threshold $x \geq 1$, the shaded region between $1/x$ and x indicates the indeterminate range. Tracks to the left of the shaded region are incorrectly classified and counted as false positives, while tracks to the right are correctly classified and counted toward the classification efficiency, following Ref. [24]. For both vacancy and NV tracks, the thresholds are chosen to give a $\lesssim$5% false positive rate. The resulting classification efficiency decreases from 93% for vacancy tracks to 74% for NV tracks, due to annealing-induced diffusion and reduced NV yield. b) Comparison of head-tail classification performance using NV tracks with the counting-based asymmetry metric and a machine learning (ML) approach. The ML model outputs a posterior probability over track direction, allowing tracks to be similarly classified as correct, incorrect, or indeterminate (see main text). At the same $\lesssim$5% false positive rate, the ML method improves classification efficiency from 74% (asymmetry metric) to 95%, demonstrating enhanced recovery of directional information from NV spatial distributions. 68
3.8 Cumulative plots of NV track classification using the counting-based asymmetry metric and machine learning approach. Top row: cumulative distributions of the asymmetry metric for 800 keV NV tracks. Bottom row: cumulative distributions of the machine learning posterior probability for the same tracks. Shaded regions indicate indeterminate ranges chosen to give a $\lesssim$5% false positive rate. The machine learning approach increases the classification efficiency relative to the asymmetry metric. 71

3.9 Spin coherence of ion-implanted NV centers. a) Distribution of NV numbers in the implantation sites selected for coherence measurements. b) Representative Hahn echo measurement (site 4). The signal is fitted to $C_0 \exp[-(\tau/T_2)^p]$ to extract the coherence time T_2 and decay exponent p. c) Hahn echo pulse sequence used for the measurements. The phase of the final microwave $\pi/2$ pulse is alternated between signal and reference shots to obtain the NV PL contrast C. d) Extracted coherence time T_2 (top) and decay exponent p (bottom) for nine single ion implantation sites. Error bars represent the standard deviation from fitting. The measured average values ($T_2^{\mathrm{avg}} \approx 1.47\,\mu$s, $p_{\mathrm{avg}} \approx 1.55$) are consistent with nitrogen spin-bath-limited ensemble decoherence. . . . 73

4.1 (a) NV ensemble DC magnetic field sensitivity estimated using the simplified metric $1/\tilde{\eta} = \sqrt{[\mathrm{N_s^0}] \times T_2^*}$, assuming spin-bath-limited T_2^* and 99.995% isotopic enrichment of ^{12}C. Corrections accounting for experimental overhead time (t_O) are included by multiplying with $1/\sqrt{(T_2^* + t_O)/T_2^*}$ (cf. Eq. 4.2). With zero overhead (blue line), sensitivity plateaus for $[\mathrm{N_s^0}] \gtrsim 1$ ppm. However, incorporating an overhead of 10μ s, typical in pulsed Ramsey protocols, shifts the optimal $[\mathrm{N_s^0}]$ towards lower values (red line). A preliminary comparison between the current low-$[\mathrm{N_s^0}]$ samples (~ 0.8 ppm) and a previously characterized high-$[\mathrm{N_s^0}]$ sample (~ 14 ppm suggests a nearly threefold sensitivity improvement. (b) Estimated optimal nitrogen concentration, and sensitivity improvement between 0.8 ppm and 14 ppm samples as functions of overhead time (t_O). 81

4.2 Optical birefringence images of ten low-$[\mathrm{N_s^0}]$ diamond samples from a single CVD synthesis run, prior to irradiation and annealing. Sample 10, exhibiting the strongest birefringence features, is selected for subsequent NV-based measurements after irradiation and annealing. Magnified views are shown of two regions of sample 10 with higher levels of birefringence (labeled A and B), with example values of high birefringence indicated. 88

4.3 Single-quantum (SQ) Ramsey measurements on diamond sample 10, far from the high birefringence regions. (a) Ramsey measurements performed at three representative locations of low strain, overlaid on the birefringence image from Fig. 4.2. (b) Ensemble NV spin dephasing at location 2. Here the microwave frequency is detuned from the center of the hyperfine-split NV spin transitions, allowing the decay envelope to be separated from hyperfine-induced beating. The extracted ensemble NV dephasing time is $T_{2,\mathrm{SQ}}^* = 17.7(4)\,\mu$s. 89

4.4 Wide-field imaging of diamond strain for sample 10. (a, b) Spatial maps of strain-induced frequency shifts in NV spin transitions across the entire diamond sample and within a selected sub-area. Strain fields shown are projections onto one NV crystallographic orientation. Each spatial map is mean-subtracted so that the global average is zero, emphasizing relative strain variations relevant to NV spin dephasing. (c, d) Histogram distributions of the strain-induced frequency shifts extracted from the maps in (a) and (b), respectively. Lorentzian fits to the histograms are also shown, from which linewidths (FWHM) Δ are determined. 91

4.5 (a) Lorentzian-fitted FWHM (Δ) of strain-induced frequency shifts as a function of sensor size L_{sensor} for sample 10. Sensor size is systematically varied by partitioning the full-diamond strain map (Fig. 4.4(a)) into smaller square sub-regions. Statistical analysis of Δ across sub-regions is presented, including minimum and median values. When more than five sub-regions are available, interquartile ranges are included: 25-75% (vertical solid lines) and 10-90% (vertical dotted lines). (b) Scaling of effective Ramsey DC magnetometry sensitivity metric with increasing sensor size. In the ideal scenario, where increasing the sensing volume introduces no additional NV ensemble dephasing – and thus no reduction in Ramsey T_2^* – the sensitivity metric scales inversely with sensor size. Here, fitting a power-law function to the effective sensitivity metric vs. L_{sensor} yields nearly inverse linear scaling behavior, indicating minimal strain limitation on the Ramsey dephasing time for bulk NV ensemble sensing using the present low-[N_s^0] diamonds. The overall low-strain effect of these diamonds applies even to sample 10 with its isolated regions of relatively high strain. 93

4.6 NV spin polarization dynamics, initialization time, and PL contrast as functions of laser excitation intensity. (a) Measured NV spin polarization dynamics and schematic measurement protocol (inset). A 1-ms laser pulse (not shown) initializes the NV ensemble into the $m_s = 0$ state prior to a microwave (MW) π pulse that initializes the NV ensemble into $m_s = -1$. A laser pulse then begins to repolarize the NV spins, with PL readout (Sig) acquired after a delay time t_{delay}. The measurement protocol is then repeated without the MW π pulse, and a reference PL measurement (Ref) is acquired. Example PL contrast measurement (Sig/Ref) at a laser intensity of 0.24 mW/μm^2, shown as a function of t_{delay}. Spin initialization time t_I (green cross) is extracted by fitting an exponential decay to the PL contrast data following its peak; we set t_I as the delay at which contrast decays to $1/e^3$ of its peak value, corresponding to repolarization of $\sim$95% of the ensemble NV spins. With this measurement protocol, the readout time is included in the reported initialization time. (b) Measured NV initialization time (t_I) as a function of laser intensity. The shaded blue region represents the range of initialization times predicted based on a five-level model of NV photophysics, with saturation intensities of 1 mW/μm^2 and 3 mW/μm^2 used to produce the lower and upper bounds, respectively. (c) Measured peak PL contrast for low- and high-[N_s^0] diamonds as a function of laser intensity. 95

4.7 Schematic of the simplified five-level NV photophysics model used in spin polarization dynamics simulations. 96

4.8 Experimentally-determined NV charge-state fraction ψ, defined as [NV^-]/([NV^-]+[NV^0]), as a function of laser intensity for both low- and high-[N_s^0] diamond samples. NV PL spectra used for decomposition analysis (see the main text) are collected using a 550 nm long-pass filter. 98

4.9 Estimated volume-normalized photon shot-noise limited DC magnetic field sensitivity for low- and high-[N_s^0] diamond samples as functions of optical excitation intensity. We assume a photon emission rate of 30 kcps per single NV at 1 mW/μm^2 illumination intensity. Sensitivity-relevant parameters (see Eq. 4.2), including the Ramsey dephasing times (T_2^* for both single- and double-quantum measurements), spin-state-dependent PL contrast (C), and experimental overhead time (t_O), are directly measured as functions of laser intensity. Nitrogen concentration and NV charge-state fraction, experimentally determined using the UV-Vis absorption method (Section 4.2.2), are used to estimate the number of NV^- centers contributing to sensing at a given laser intensity. The inset shows the ratio of estimated sensitivity between high-[N_s^0] and low-[N_s^0] samples as a function of optical excitation intensity, for SQ (solid) and DQ (dashed) Ramsey protocols, highlighting the relative advantage of low-nitrogen material at moderate to low optical excitation intensities. 101

5.1 Cumulative plots of NV track classification using the counting-based asymmetry metric and machine learning approach. Top row: cumulative distributions of the asymmetry metric for 100 keV NV tracks. Bottom row: cumulative distributions of the machine learning posterior probability for the same tracks. Shaded regions indicate indeterminate ranges chosen to give a $\lesssim$5% false positive rate. The machine learning model trained on 20,000 100 keV vacancy tracks.. 112

A.1 Denoised QDM picotesla-scale magnetic-field images after about 3600 s signal acquisition time. Original measurement data is presented in the main text (Sec. 2.6). Here, two denoising methods, Non-local mean (NLM) and Gaussian smoothing, are applied. Both algorithms are successful at improving SNR. However, NLM denoising, as shown in (a), additionally excels at restoring fine structures (e.g., contours, patterns, and edges where the Laplacian is non-zero) from a noisy image that can be otherwise distorted by a Gaussian smoothing, as shown in (b). Scale bar: 50 µm 119

B.1 **Nitrogen-concentration measurement. (a)** Representative TOF-SIMS depth profile of the CN^-/C^- ion ratio, showing a uniform signal over the profiled depth. **(b)** EAG–Rice TOF-SIMS calibration: CN^-/C^- ratio measured at Rice University plotted against the EAG reference nitrogen concentration for five regions of Sandia 2 (solid line: linear fit). **(c)** NV background PL calibration: confocal PL intensity plotted against the EAG reference nitrogen concentration for the same five regions (solid line: linear fit). 123

B.2 Cluster decomposition and lattice domain size selection for annealing simulations. (a) Probability of vacancy–vacancy interaction as a function of initial separation r, obtained from two-vacancy simulations at different nitrogen concentrations. The dashed theoretical line is computed using the 3D hitting probability, assuming an effective interaction radius of $1.54\,\text{Å}$. Solid lines show power-law fits of the form $1.54\,\text{Å}/r^p$. For the experimental nitrogen concentration $[\mathrm{N_s^0}] = 253\,\mathrm{ppm}$, the interaction probability at $154\,\text{Å}$ is $\sim$0.3%, supporting negligible cross-cluster interactions. (b) Example decomposition of an 800 keV implantation damage track into spatially separated clusters. The minimum pairwise defect spacing between clusters exceeds $154\,\text{Å}$. Vacancy (left) and carbon interstitial (right) clusters are shown. The legend indicates the number of defects in each cluster. (c) Root-mean-square (RMS) displacement of a diffusing vacancy as a function of nitrogen concentration, evaluated at the point where the number of unique sites visited equals the number of available capture sites ($N_{uni} = N_{cap}$), corresponding to a capture probability of $\sim$64%. (d) Vacancy capture probability as a function of RMS displacement for a nitrogen concentration of $[\mathrm{N_s^0}] = 253\,\mathrm{ppm}$. For each cluster, the simulation domain is defined by padding 100 unit cells in each direction from all vacancy defects. This ensures that the probability of vacancy escape is $\lesssim 0.01\%$. 124
B.3 Representative damage tracks before and after annealing simulations. Gray dots indicate the initial vacancy distribution, and red dots indicate the resulting NV distribution after annealing. The vacancy-to-NV conversion yield and head-tail asymmetry values, computed using a defect counting method, are shown for both vacancy and NV tracks. . 127
B.4 **NV-count extraction pipeline and multi-ion statistics**, illustrated for a representative 800 keV implantation square (6 μm pitch). **(a)** Raw coarse confocal map (500 nm pixels), in units of calibrated single-NV PL. **(b)** After region-of-interest and scratch masking and intensity thresholding, isolating the implantation-site PL. **(c)** Elliptical grid mask ($\pm 3\sigma$ of the beam profile) overlaid on the thresholded map; green crosses mark the analyzed sites and cyan boxes the three example sites shown in (d). **(d)** High-resolution nanovolume scans of the three boxed sites, resolving one, two, and three PL spots. **(e)** Distribution of the number of resolved spots per site over all $N = 222$ analyzed 800 keV sites (bars), with a maximum-likelihood truncated-Poisson fit (red; $\lambda = 0.99$, $\chi^2 = 4.1$, 2 d.o.f., $p = 0.13$). 130

B.5 Mitigation of polishing-induced NV background. (a) Confocal PL map of a scaife-polished HPHT diamond measured at a depth of $\sim 1\,\mu\mathrm{m}$ after annealing. The diamond is predominantly $\{100\}$-oriented, with four high-PL "petal" regions arising from growth sectors. Strong and highly non-uniform background PL is observed due to polishing-induced damage. Red arrows indicate regions of severe damage that can overwhelm implantation signals, black arrows highlight streak-like features consisting of isolated high-NV spots, and green arrows indicate regions with relatively low background. Three corners have been previously etched, showing indented features. (b) Confocal PL maps at increasing surface depths following high-power (250 W) Ar/Cl_2 ICP-RIE etching. The near-surface region ($\lesssim 1\,\mu\mathrm{m}$) exhibits elevated PL corresponding to $\sim$5-7 NV centers per confocal spot, indicating etch-induced damage. (c) Confocal PL maps at increasing depths following low-power (20 W) Ar/Cl_2 ICP-RIE etching. The near-surface PL is significantly reduced, with NV densities $\sim$1 per confocal spot, comparable to as-received HPHT diamond without annealing. 133
B.6 NEB-DFT calculations of Reaction 1. Geometries of the (a) initial, (b) transition geometry, and (c) final states along the reaction path. V_C and C_i defects are formed by the ejection of a C_C atom from its lattice site. Carbon atoms are shown in brown, while the resulting vacancy is shown in yellow. (d) MEP of Reaction 1. The transition state occurs at a V_C-C_i spacing of 1.75 Å corresponding to E_B = 15.5 eV. 136
B.7 NEB-DFT calculations of Reaction 2. Geometries of the (a) initial, (b) transition geometry, and (c) final states along the reaction path. Two V_C defects are formed by the dissociation of constituent V_C from the V_{2C} defect. Carbon atoms are shown in brown, while the resulting vacancy is shown in yellow. (d) MEP of Reaction 2. The transition state occurs at a V_C-V_C spacing of 2.53 Å corresponding to E_B = 5.95 eV. . . 137
B.8 Additional NV spin echo measurements. Spin echo measurements from 14 additional implantation sites in the HPHT sample. Each trace is fitted to a stretched exponential decay (red curve) with fixed exponent $p = 1.55$. The final panel shows the distribution of NV PL counts for the characterized sites. 140
B.9 Spin properties versus NV number. (a) Scatter plots of spin-state-dependent PL contrast C_0, coherence time T_2, and NV number per site, showing no clear correlation between spin properties and NV count. (b) Distribution of C_0 and T_2 across measurement sites. T_2 remains relatively constant around $1.4\,\mu$s, while C_0 exhibits a larger spread ($\sim$1-3%). 141

C.1 Single-quantum (SQ) and double-quantum (DQ) Ramsey measurements and fits yield the ensemble NV spin dephasing times $T^*_{2,\mathrm{SQ}}$ and $T^*_{2,\mathrm{DQ}}$ at a high-birefringence region (area B of sample 10, see main text Fig. 2). The reduced $T^*_{2,\mathrm{SQ}}$ reflects the strong local strain variation, while the DQ protocol mitigates these effects. ΔC is the change in NV PL contrast. 143
C.2 Spatial maps of experimentally determined NV spin-strain coupling terms (M_z) for each of four NV orientations and derived diamond stress tensors components σ across the entire diamond sample 10. Scale bars: $500\,\mu$m 146

D.1 Schematic of the Ramsey QDM sensor head assembly, including the wafer, microwave circuit, wire phantom, and diamond. 148

D.2 Schematic of a wafer opening that enables direct sample interfacing with the diamond in the Ramsey QDM. 148
D.3 Top-view image of a glass pillar carrying magnetic beads fixed with S1813 photoresist. 149
D.4 Ramsey QDM setup with a mechanical arm holding the glass pillar. 150
D.5 Magnetic image of the pillar sample measured using the Ramsey-based QDM. The pillar is positioned within the square imaging area. 151

E.1 Lily pollen grain inside the square imaging area above the diamond surface. 152
E.2 (Top) Images of the bio-mount designed to integrate with the Ramsey QDM sensor head and create a perfusion environment for pollen-tube growth. A peristaltic pump is shown here, but in the final experiment it was replaced by an HPLC pump for improved stability when injecting culture medium into the bio-mount. (Bottom) Bio-mount inside the Ramsey QDM setup. The zoomed-in panel shows a schematic cross-sectional view through the central imaging area. 153
E.3 TEC-based water bath system used to regulate the temperature of the culture medium. . 154
E.4 Example temperature profile of culture medium submerged in the water bath (top) and an example of laboratory temperature variation (bottom). 155
E.5 Tobacco pollen grains germinating under normal NV magnetic-imaging operation. Top: before continuous magnetic imaging, pollen grains are loaded above the diamond surface. Bottom: after approximately 1 hour of continuous magnetic imaging, multiple pollen tubes are visible across the image. These images were acquired with a $20\times$ 0.35 NA objective for quick viewing of pollen grains from above the diamond surface (cf. Fig. D.4), rather than for NV photoluminescence imaging or optimized optical imaging of pollen tubes. 156

List of Abbreviations

AC	Alternating Current
AOM	Acousto-optic Modulator
ASD	Amplitude Spectral Density
CLSM	Confocal Laser Scanning Microscopy
CPW	Coplanar Waveguide
CVD	Chemical Vapor Deposition
CW	Continuous Wave
DAQ	Data Acquisition
DEER	Double Electron-Electron Resonance
DQ	Double Quantum
DQ4R	Double Quantum Four-Ramsey
FEM	Finite-Element Method
FFT	Fast Fourier Transform
FOV	Field of View
HPHT	High Pressure Hight Temperature
ISC	Intersystem Crossing
MW	Microwave
NA	Numerical Aperture
NMR	Nuclear Magnetic Resonance
NV	Nitrogen-vacancy
ODMR	Optically Detected Magnetic Resonance
PL	Photoluminescence
POP	Physical Optics Propagation
PSD	Power Spectral Density
QDM	Quantum Diamond Microscope
RF	Radio-frequency
SEM	Scanning Electron Microscopy

SNR	Signal-to-noise Ratio
SQ	Single Quantum
SQUID	Superconducting Quantum Interference Device
SWaP	Size, Weight, and Power
XRD	X-Ray Diffraction
ZFS	Zero-field Splitting
ZPL	Zero-phone Line

Chapter 1: Introduction

This chapter provides the essential background needed to understand the work in Chapters 2-4. The goal is not to provide an exhaustive review of nitrogen-vacancy (NV) centers in diamond, NV magnetometry, or directional rare-event detection, as these topics have already been covered extensively in the literature. Instead, I focus on the basic spin and optical properties of the NV center, the microwave and optical protocols underlying broadband DC magnetic sensing, and the proposed use of NV quantum-diamond systems for directional readout of particle-induced damage tracks. I also point to additional references that I found particularly useful for more detailed treatments of these topics.

1.1 Nitrogen-vacancy centers in diamond

The nitrogen-vacancy (NV) center [1–4] is a point defect in diamond consisting of a substitutional nitrogen atom adjacent to a carbon vacancy. The line connecting the nitrogen and vacancy sites naturally defines the NV symmetry axis, as shown schematically in Fig. 1.1. In the diamond lattice, there are four crystallographic $\langle 111 \rangle$ directions, giving rise to four possible non-orthogonal NV orientations. This structural property underpins the vectorial sensing capability of NV diamond. For example, in magnetic-field sensing, each NV orientation measures the projected field along its own axis, and measurements from multiple orientations can be combined to reconstruct the full magnetic-field vector [5].

The NV center can exist in several charge states, including the negatively charged NV^-, the neutral NV^0, and the positively charged NV^+ charge states [6]. The charge state most relevant for quantum sensing is NV^-, which has the spin and optical properties used for magnetometry in this

dissertation. However, NV^0 is often present alongside NV^- and typically acts as an optical background that degrades the effective NV readout contrast. Mitigating charge-state conversion is therefore an important practical consideration. For example, in Chapter 4, I discuss how charge-state conversion becomes a limitation for pulsed magnetometry at high optical excitation intensity. At the same time, NV^0 can be useful in other experimental regimes, such as spin-to-charge conversion readout [7] and diamond voltage sensing [8]. Hereafter, I use "NV" to refer to the negatively charged NV^- center unless otherwise stated.

NV has an electronic spin triplet ground state ($S = 1$), in which the $m_s = 0$ sublevel is separated from the nearly degenerate $m_s = \pm 1$ sublevels by the zero-field splitting $D \approx 2.87$ GHz at room temperature (Fig. 1.1). The nitrogen nucleus also carries a nuclear spin, with $I = 1$ for $^{14}\mathrm{N}$ and $I = 1/2$ for $^{15}\mathrm{N}$. This nuclear spin produces hyperfine structure in the NV spin resonance spectrum, further splitting the electronic spin transitions by approximately 2-3 MHz, depending on the nitrogen isotope [1, 9]. Most of the sensing principles and applications discussed in this thesis can be understood from the electronic spin structure alone.

Green optical excitation, commonly provided by a free-running 532 nm laser, enables both initialization and readout of the NV electronic spin (Fig. 1.1). The green laser excites the NV center from the ground-state triplet to an excited-state triplet in a largely spin-conserving fashion. From the excited state, the NV center can decay radiatively back to the ground state, producing NV photoluminescence (PL) with the zero-phonon line at 637 nm. However, coupling to lattice vibrations produces a broad phonon sideband extending to approximately 800 nm at room temperature. This broad vibronic structure is also why NV optical excitation does not require a narrow-linewidth resonant laser; wavelengths near 532 nm, and in some cases nearby green excitation wavelengths, can efficiently drive the optical cycle.

Beyond the spin-conserving radiative decay pathway, the optical cycle also contains a spin-dependent nonradiative pathway through intermediate singlet states, commonly referred to as the intersystem crossing (ISC) [10]. In the excited state, population associated with $m_s = \pm 1$ has a higher probability of entering this nonradiative pathway than population associated with $m_s = 0$. As a result, NV centers prepared in $m_s = 0$ emit more PL than those prepared in $m_s = \pm 1$, providing the foundational spin-dependent optical readout contrast. Because decay through the singlet manifold returns population to the ground-state triplet without preserving the original spin projection, so repeated optical cycling polarizes the NV spin into $m_s = 0$. Therefore, green laser excitation both initializes the NV spin into $m_s = 0$ and allows the final spin population to be read out optically. In addition to optical excitation, microwave (MW) irradiation can coherently drive transitions between $m_s = 0$ and $m_s = \pm 1$ when resonant with the corresponding spin transition, producing Rabi oscillations. This combination of optical initialization, microwave spin control, and spin-dependent PL readout forms the foundation of optically detected magnetic resonance (ODMR).

For the experiments in this dissertation, the most relevant ground-state NV Hamiltonian can be approximated as follows, written in a coordinate system with the z axis aligned to the NV symmetry axis:

$$\hat{H}/h \approx (D + M_z)\hat{S}_z^2 + \gamma_e B_z \hat{S}_z. \tag{1.1}$$

Here, $\hat{S}_z$ is the dimensionless spin-1 operator, D is the zero-field splitting, which has a temperature dependence, and M_z represents the longitudinal crystal-strain-induced shift. More generally, M_z is a linear combination of components of the crystal strain tensor [11]. $\gamma_e \approx 28.024\ \mathrm{GHz/T}$ is the NV electron gyromagnetic ratio, and B_z is the projection of the magnetic field along the NV axis, including contributions from both the applied bias field and any additional signal fields. This Hamiltonian gives

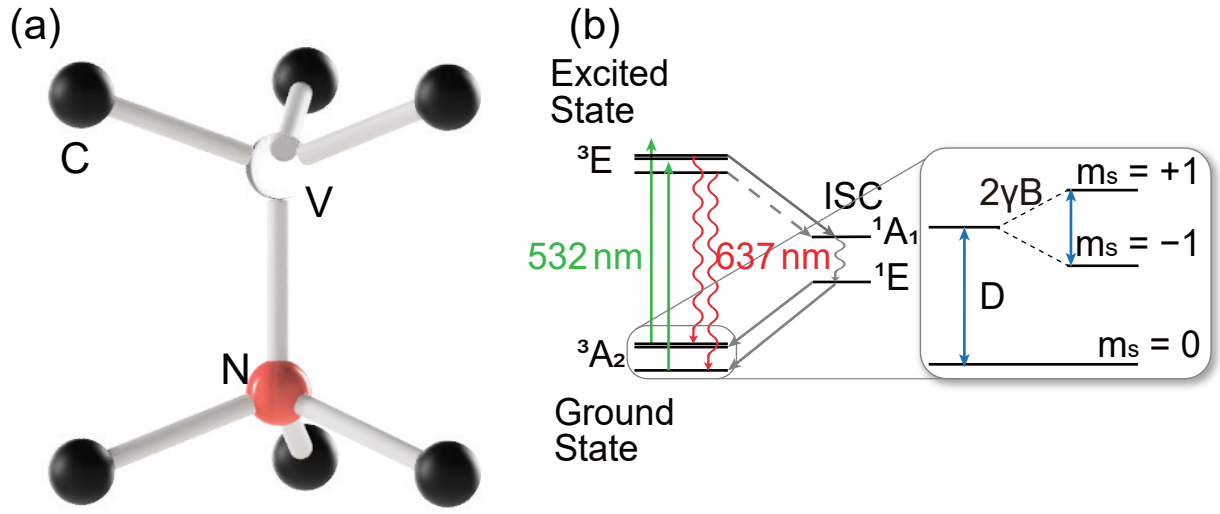


Figure 1.1: a) Crystal structure and energy levels of the NV center. (b) NV energy-level diagram. The enlarged view shows the electronic spin triplet ground state with zero-field splitting D and bias magnetic field B aligned to the sensing NV axis.

the transition frequencies

$$\nu_{\pm} = D + M_z \pm \gamma_e B_z, \tag{1.2}$$

corresponding to the $m_s = 0 \rightarrow m_s = +1$ and $m_s = 0 \rightarrow m_s = -1$ transitions. The magnetic field therefore shifts the two transitions in opposite directions, while temperature and longitudinal strain shift them in common mode. The full NV Hamiltonian contains additional terms, including hyperfine coupling to the nitrogen nuclear spin, transverse magnetic fields, electric fields, and both axial and transverse strain. These terms are important in many NV experiments and are treated in detail in several canonical reviews, for example Refs. [1, 3]. For the DC broadband magnetometry experiments considered here, the bias magnetic field is typically aligned close to one NV axis. Together with the large zero-field splitting D, this operating condition suppresses many transverse strain and electric field effects to higher order [12, 13].

NV centers can be created in diamond by combining substitutional nitrogen with mobile vacan-

cies. In many quantum diamond samples, nitrogen is first incorporated during growth, and vacancies are then introduced by irradiation, commonly using either electrons or ions. Subsequent high-temperature annealing allows vacancies to diffuse through the diamond lattice and become trapped by substitutional nitrogen atoms to form NV centers. The diamond samples used in Chapters 2 and 4 were grown and processed by Element Six. In Chapter 3, this same basic formation process is studied more directly, where implantation-induced vacancies form NV centers that act as optical markers of ion-induced damage tracks.

Two broad classes of synthetic quantum diamond are relevant here: chemical-vapor-deposition (CVD) diamond and high-pressure, high-temperature (HPHT) diamond. In CVD growth [14], a diamond substrate is placed inside a plasma reactor, where a carbon-containing gas, typically methane diluted in hydrogen, supplies carbon to the growing diamond surface. Nitrogen can be introduced during growth so that substitutional nitrogen is incorporated into the lattice. By controlling when nitrogen is introduced, one can grow either a bulk nitrogen-doped diamond sample or a thin near-surface nitrogen-doped layer. After irradiation and annealing, such layers can be converted into near-surface NV ensembles, which are commonly used in quantum diamond microscopes (QDMs) for wide-field imaging experiments, as in Chapter 2. CVD growth also allows control over the carbon isotope composition through the choice of methane source; high $^{12}\mathrm{C}$ enrichment reduces the residual $^{13}\mathrm{C}$ nuclear spin bath and can extend NV spin coherence.

HPHT diamond [15] is grown under high pressure and high temperature using a carbon source and a metal solvent catalyst. Compared with the CVD diamonds, HPHT diamonds often contain much higher concentrations of substitutional nitrogen, including the $\sim 200\ \mathrm{ppm}$-level nitrogen concentration used in Chapter 3. This high density of nitrogen capture sites is useful for forming NV centers near ion-induced damage tracks after annealing. One practical complication of HPHT diamond is that a single

as-grown crystal can contain multiple growth sectors with different impurity and strain properties, so the material often needs to be carefully cut to obtain a suitable single-crystal sector.

Crystal strain is an important material property for ensemble NV sensing. In CVD diamond, strain can arise from substrate polishing damage, dislocations threading from the substrate into the growth layer, and other growth imperfections [16]. Because strain shifts the NV spin transition frequencies, spatially varying strain can broaden ensemble resonances and reduce the effective spin dephasing time. For this reason, low-strain growth and careful substrate selection are important for producing high-quality quantum diamond for sensing applications, a topic explored further in Chapter 4.

NV-based sensing can be performed using either single NV centers [17] or ensembles of NV centers [1]. Single NV centers are useful for nanoscale sensing because the defect itself is atomic in size and can be positioned within nanometers of a sample. In contrast, ensemble NV sensors, particularly those used in QDM wide-field magnetic imaging, are typically limited to optical or near-optical spatial resolution, but gain signal-to-noise ratio by averaging over many nominally identical spins. This ensemble regime also introduces additional factors that influence sensing performance. Increasing the NV density increases the number of sensing spins, but it typically requires higher nitrogen concentration and irradiation dose, which can introduce additional paramagnetic spin baths and shorten NV spin coherence times. In addition, an ensemble diamond contains NV centers along all four crystallographic orientations. When only one orientation is used for sensing, the other orientations contribute PL that does not carry the desired magnetic signal, reducing the observed contrast. These tradeoffs in ensemble diamond engineering are considered further in Chapter 4.

1.2 Broadband DC magnetometry using NV-diamond

Equation 1.2 shows that magnetic-field information is encoded in the transition frequencies between NV spin sublevels. In practice, NV magnetometry uses microwave and optical control to convert these frequency shifts into a measurable change in NV PL. The simplest implementation is continuous-wave optically detected magnetic resonance (CW-ODMR), in which laser excitation and microwave driving are applied simultaneously while the microwave frequency is swept across the NV spin resonance (Fig. 1.2). A magnetic field shifts the resonance frequency and therefore changes the detected PL at a fixed microwave frequency. CW-ODMR is experimentally simple and widely used, but simultaneous optical and microwave excitation introduces power broadening and limits the achievable sensitivity for a fixed number of NV centers [1, 18]. For the high-sensitivity DC broadband magnetometry considered in this dissertation, Ramsey-type pulsed protocols are therefore central [1, 13].

In Ramsey magnetometry, the goal is usually to detect a weak signal field, ranging from picotesla to hundreds of nanotesla, on top of a much larger applied bias field. The bias field, typically a few millitesla, is aligned close to one NV orientation and splits the otherwise nearly degenerate $m_s = \pm 1$ states. This spectral separation allows microwave pulses to address the relevant spin transitions individually. The Ramsey sequence is shown schematically in Fig. 1.2. A first green laser pulse initializes the NV spin population into $m_s = 0$. A resonant microwave $\pi/2$ pulse then prepares a coherent superposition of spin states. In a single-quantum (SQ) Ramsey measurement, the microwave pulse addresses either the $|0\rangle \leftrightarrow |-1\rangle$ transition or the $|0\rangle \leftrightarrow |+1\rangle$ transition. For example, driving the $|0\rangle \leftrightarrow |-1\rangle$ transition prepares a state of the form

$$\frac{1}{\sqrt{2}} \left(|0\rangle + |-1\rangle \right). \tag{1.3}$$

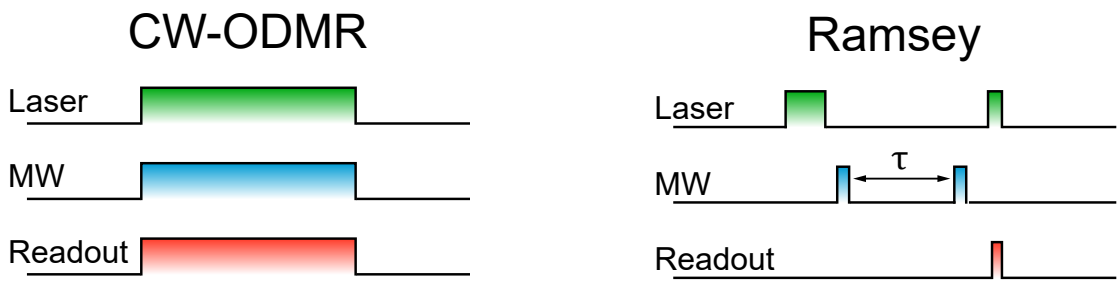


Figure 1.2: Schematics comparing NV magnetometry protocols based on continuous-wave optically detected magnetic resonance (CW-ODMR) and pulsed Ramsey.

During the subsequent free-precession interval τ, an additional signal field B_{sig} projected along the NV axis shifts the spin transition frequency and causes the superposition state to accumulate a relative phase

$$\phi_B^{SQ} = \gamma_e B_{\text{sig}} \tau. \tag{1.4}$$

A second microwave $\pi/2$ pulse maps this accumulated phase onto a population difference between the spin states. The final green laser pulse then reads out this population difference through spin-dependent NV PL.

Experimentally, Ramsey magnetometry is often understood through two related measurements (Fig. 1.3). First, one can sweep the free precession time τ in the absence of an applied signal field to measure the Ramsey free induction decay. The decay envelope gives the Ramsey dephasing time T_2^*, which helps to determine the optimal timescale for phase accumulation [1]. Oscillations in the Ramsey signal can arise from microwave detuning, hyperfine splitting, or other controlled MW frequency offsets. Second, after choosing a fixed value of τ, one can generate a magnetometry curve by varying the microwave detuning or equivalently by applying a known calibration magnetic field. This produces an optical response curve relating NV PL to effective frequency shift, and therefore to magnetic field

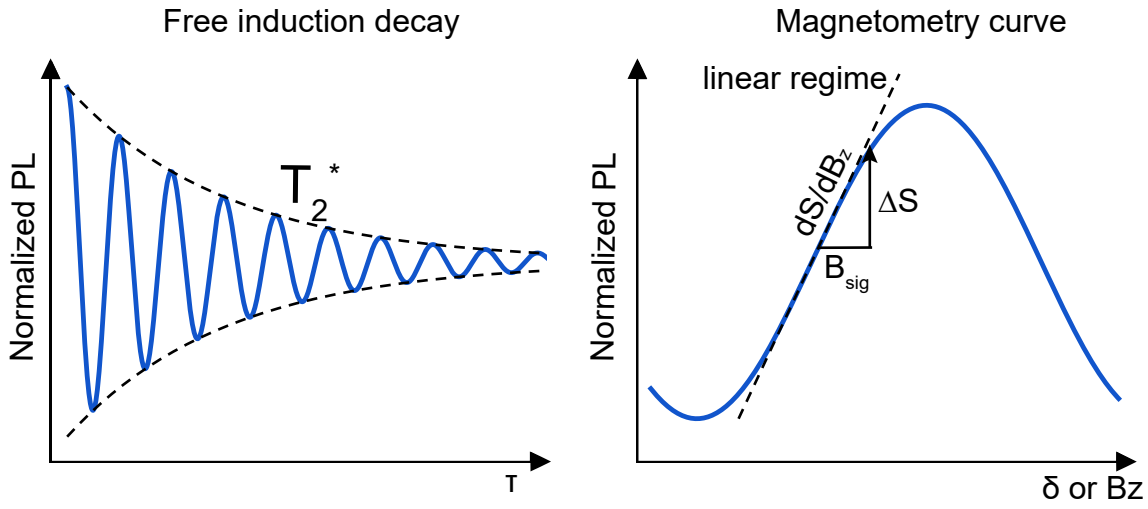


Figure 1.3: Schematic of Ramsey magnetometry. A Ramsey free-induction decay measurement determines the dephasing time T_2^*, while a fixed-τ magnetometry curve converts small frequency or magnetic-field shifts into changes in normalized PL near the linear operating point.

through γ_e. In the linear operating regime, the signal can be written approximately as

$$\Delta S \approx \frac{dS}{dB_z} B_{\text{sig}}, \tag{1.5}$$

where S denotes the measured Ramsey PL signal and dS/dB_z is determined by calibration. Magnetic imaging with a QDM applies this procedure pixel by pixel: the measured PL change in each pixel is converted into a magnetic field value using the local magnetometry response. This is the basic measurement principle used for the Ramsey QDM in Chapter 2.

The sensitivity of a Ramsey magnetometer [1] is determined by how efficiently a small magnetic-field-induced phase shift can be distinguished from measurement noise. Important parameters include the number of participating NV centers, spin-state readout contrast, collected PL photon number, chosen Ramsey free precession time τ (typically close to T_2^*), and the experimental overhead associated

with spin initialization and readout. The photon shot noise limit provides an important benchmark for the best sensitivity achievable with conventional optical readout. The full photon-shot-noise-limited sensitivity expression, including the influence of NV spin, optical, charge-state, and material parameters, is discussed in Chapter 4.

The Ramsey protocol can also be implemented in a double-quantum (DQ) configuration [13, 19]. In this case, the microwave pulses contain two frequency tones with a defined phase relationship, simultaneously resonant with the $|0\rangle \leftrightarrow |+1\rangle$ and $|0\rangle \leftrightarrow |-1\rangle$ transitions. These dual-tone pulses prepare a coherent superposition between the $|+1\rangle$ and $|-1\rangle$ states, rather than between $|0\rangle$ and one of the $m_s = \pm 1$ states. Because the two states shift in opposite directions under a magnetic field, the accumulated magnetic phase is

$$\phi_B^{\mathrm{DQ}} = 2\gamma_e B_{\mathrm{sig}} \tau. \tag{1.6}$$

In this case, frequency shifts that are common to both $|+1\rangle$ and $|-1\rangle$, such as changes in D from temperature or longitudinal strain shifts M_z, are largely rejected. This makes DQ Ramsey magnetometry especially useful for wide-field imaging experiments, where different regions of the diamond can experience different strain or temperature shifts. The Ramsey QDM developed in Chapter 2 uses this DQ approach, together with spin-bath driving and camera-based lock-in detection, to perform high-sensitivity DC broadband magnetic-field imaging.

1.3 NV-based directional rare-event detection

NV centers can be used not only for magnetic-field sensing; Equation 1.2 also shows that they are sensitive to crystal strain. Although the longitudinal strain shift M_z enters the transition frequencies in the same way as the zero-field splitting D, temperature variations are often spatially smoother

across a diamond chip, whereas crystal strain can have localized nanoscale or microscale features [20]. This strain sensitivity, together with the ability to optically resolve individual NV centers [3] and to use quantum-sensing-enabled super-resolution techniques for nanoscale NV position imaging [21–23], motivates the use of NV centers as readout tools for particle-induced damage tracks. One particularly important application is directional detection of rare events, including dark matter, and especially weakly interacting massive particle (WIMP) dark matter [24].

In conventional WIMP direct-detection experiments, an incoming dark matter particle or neutrino can scatter from a target nucleus and produce a low-energy nuclear recoil. As detector exposures improve, coherent scattering from solar and atmospheric neutrinos becomes an increasingly important background [25]. This background is difficult to remove because neutrinos can produce recoil-energy spectra that overlap with those expected from dark-matter interactions. This limitation is often described as the neutrino floor or neutrino fog (Fig. 1.4), a regime in which simply increasing exposure no longer provides efficient signal-background discrimination [26, 27].

Directional information provides a powerful way to distinguish these signals. The expected dark matter flux in the laboratory frame has a preferred direction set by the motion of the Solar System through the Galactic halo, while solar neutrinos arrive from the direction of the Sun. If the recoil direction, including both orientation and head-tail information, can be measured, then dark matter induced recoils and neutrino induced recoils could be separated [27, 28], allowing one to cut through the neutrino fog. Gas-based time projection chambers have demonstrated mature directional readout capabilities, but their low target density makes scaling to large detector masses challenging [29]. Solid-state detectors, in contrast, provide high target density, but the directional information is encoded on much shorter length scales.

Quantum diamond may provide a useful platform for addressing this directional readout chal-

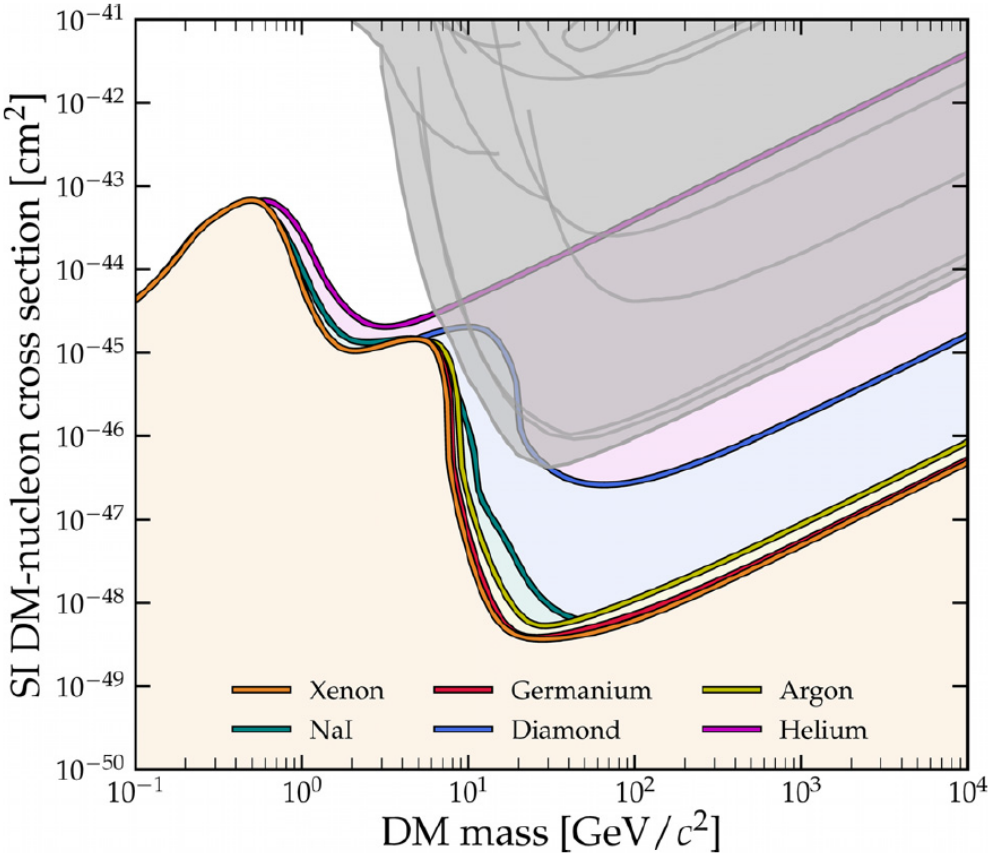


Figure 1.4: Spin-independent dark matter–nucleon cross section as a function of dark matter mass, illustrating the neutrino floor for several detector target materials. Directional information provides a route to distinguish dark matter induced nuclear recoils from neutrino induced backgrounds in this regime. Adapted from Ref. [27].

lenge in a solid-state system. Because recoil damage tracks in diamond are expected to have characteristic length scales of tens of nanometers, the proposed diamond-based directional detector (Fig. 1.5) naturally separates into a multiscale localization and readout problem [24, 27, 30–32]. A dark matter or neutrino interaction induces a nuclear recoil [24], generating a nanoscale cascade of lattice defects, including vacancies and interstitials, whose spatial distribution encodes information about the incoming particle direction. Candidate events are first identified at the millimeter scale, corresponding roughly to an individual diamond detector chip, using conventional phonon, charge, and/or photon readout [33–35]. At the micrometer scale, the damage track could be localized using high-throughput, high-sensitivity wide-field NV strain imaging with pre-existing NV ensembles [20], or by detecting newly formed NV centers created near the damage track after annealing. At the nanoscale, the morphology of the damage track may then be probed using techniques such as high resolution X-ray strain imaging [36, 37] or Fourier gradient imaging of individual NV centers [22]. Ultimately, the recoil direction could be inferred from the nanoscale spatial distribution of damage-related signals.

In this proposed detector concept, NV centers primarily contribute to the micrometer-to-nanometer readout stages. For NV-based strain localization, the diamond sample would likely be a CVD diamond, because lower nitrogen concentrations and longer spin-coherence times are advantageous for high-sensitivity strain sensing and high-throughput searches across a full diamond chip with micrometer-scale resolution. For detection of newly formed NV centers, the diamond sample would instead likely be nitrogen-rich HPHT diamond. The high concentration of substitutional nitrogen reduces the average spacing between nitrogen capture sites, increasing the probability that damage-induced vacancies form NV centers near the original recoil track during annealing.

However, the optimal sensing modality at each length scale remains an open question. Several issues must be addressed before such a detector can be realized. It is necessary to know whether low-

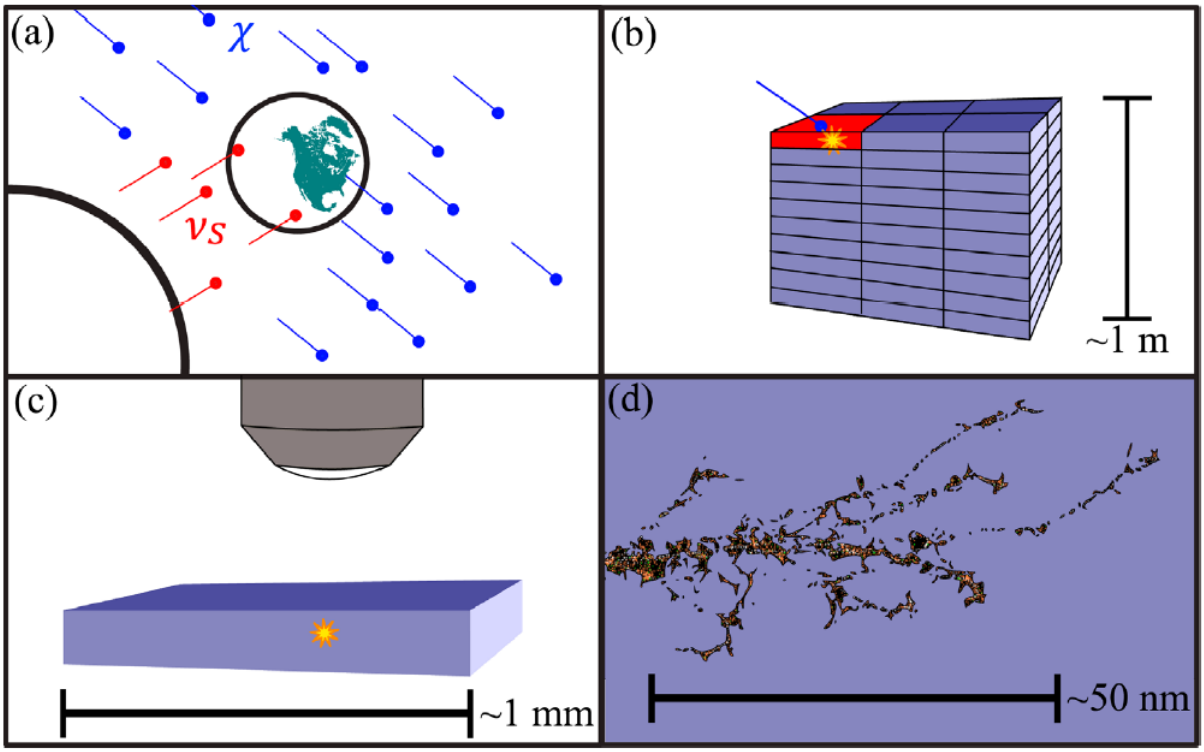


Figure 1.5: Schematic of the proposed solid-state directional dark matter detector readout pipeline. A rare event interaction is first detected and localized at the millimeter scale using phonon, charge, or photon readout. The corresponding damage track is then localized at the micrometer scale and reconstructed at the nanoscale using quantum-defect or X-ray-based imaging methods. Adapted from Ref. [30].

energy recoils produce enough detectable NV or strain signal for micrometer-scale localization, and how this signal scales with recoil energy. In the case of newly formed NV centers in HPHT diamond, it is also necessary to determine whether annealing can preserve directional information in the original damage track, and whether the final NV or strain distribution can be reconstructed with sufficient spatial resolution to infer the recoil direction. Answering these questions requires experimentally calibrated models of damage-track creation, defect evolution during annealing, NV formation, and the associated strain distribution. These questions motivate the work in Chapter 3, where single carbon-ion implantation events in nitrogen-rich diamond are used as a benchmarking platform for studying NV formation and retained directional information in particle-induced damage tracks.

Directional detector capability is not only useful for dark matter searches. Directional neutrino detection could also improve tests of the Standard Model at high-flux, energetic neutrino sources [38]. In addition, diamond detectors can probe sub-GeV light dark matter in a non-directional mode, capitalizing on diamond's favorable properties for conventional signal readout [33, 35]. More broadly, improved modeling of damage track creation and annealing evolution could be useful for studying track formation and survivability over geological timescales relevant to mineral-based dark-matter detection [39, 40], and for guiding implantation strategies in quantum device fabrication, including cases involving surface boundaries or non-uniform dopant distributions such as delta-doped nitrogen layers [41].

1.4 Organization of this Dissertation

This dissertation investigates approaches for extending the microscopic information accessible with NV-diamond systems.

Chapter 2 develops a pulsed quantum diamond microscope for wide-field imaging of static and dynamic magnetic fields. By combining Ramsey magnetometry, spin-bath driving, and high-speed camera detection, this work demonstrates improved sensitivity and sub-millisecond magnetic imaging capability.

Chapter 3 investigates NV-based readout for particle-induced lattice damage. Individual carbon-ion implantation events are localized through spatially confined NV photoluminescence after annealing, and an experimentally calibrated implantation-and-annealing model is used to assess the survival of directional information in the resulting NV distribution.

Chapter 4 examines the diamond material engineering for pulsed magnetometry. By synthesizing and characterizing low-nitrogen, low-strain, isotopically enriched diamond, this work identifies the practical tradeoffs among coherence time, NV density, charge stability, and optical overhead that determine pulsed sensing performance.

Chapter 5 synthesizes these results and discusses future opportunities in biological magnetic imaging using pulsed quantum diamond microscopy, as well as outlook for low-energy recoil detection and nanoscale damage-track reconstruction.

Chapter 2: Dynamic Magnetic-Field Imaging with a Pulsed Quantum Diamond Microscope

Provenance and contributions. This chapter is adapted from J. Tang, Z. Yin, C. A. Hart, J. W. Blanchard, J. T. Oon, S. Bhalerao, J. M. Schloss, M. J. Turner, and R. L. Walsworth, "Quantum Diamond Microscope for Dynamic Imaging of Magnetic Fields" [42]. I led the experimental implementation, data analysis, curation, and visualization, and writing of the original manuscript draft, and contributed to subsequent revision and editing with R. L. Walsworth and the other authors. Z. Yin contributed to simulations of the phantom magnetic-field patterns, assisted with the implementation of the spin-bath-driving measurements, and contributed to data analysis. Related early development of the pulsed QDM imaging platform is described in Refs. [12] and [43].

2.1 Motivation and chapter overview

As introduced in Chapter 1, a quantum diamond microscope (QDM) employs a diamond chip with a near-surface layer of nitrogen-vacancy (NV) centers to enable wide-field magnetic imaging. QDMs have garnered increasing interest because they combine micron-scale spatial resolution, millimeter-scale fields of view, and compatibility with diverse samples across the physical and life sciences. The measurements considered in this chapter target DC and broadband magnetic fields with relevant frequency components extending into the low-kHz range. QDM measurements have already enabled applications in geoscience [44, 45], 2D condensed matter systems [46], electronic circuits [47, 48], and biological systems [49–52].

A majority of QDMs realized to date utilize a continuous-wave optically detected magnetic

resonance (CW-ODMR) sensing protocol. In CW-ODMR, optical and microwave (MW) fields are applied simultaneously during NV interrogation: the optical field polarizes and reads out the NV spin ensemble, while the MW field drives spin transitions whose resonance frequencies encode magnetic information. This simultaneous optical and MW irradiation produces competing effects that constrain the achievable magnetic sensitivity for a fixed number of NV centers per imaging pixel [1]. As a result, measurements of nanotesla-scale magnetic fields can require averaging intervals of hours [49, 50, 53], limiting the imaging of weaker magnetic sources, as well as sample throughput. This sensitivity constraint is especially critical when both the spatial and temporal resolution are required for imaging dynamic magnetic fields, such as those generated from action-potential (AP) currents in electrically-active cells. To date, dynamic NV-based measurements have detected magnetic fields from neuron [51], cardiac [54], and muscle [55] AP currents, but these measurements relied on spatially integrating ensemble NV photoluminescence over lateral length scales $> 100\,\mu\mathrm{m}$ onto a single photodiode and signal averaging over repeated AP measurements. In particular, the temporal resolution achieved in previous dynamic QDM biomagnetic imaging experiments [50, 56] are inadequate for resolving individual AP currents on the sub-millisecond timescales.

QDM per-pixel magnetic sensitivity and temporal resolution can be improved by implementing pulsed Ramsey magnetometry with acquisition using a high-frame-rate camera [12, 13, 43]. Pulsed Ramsey sensing protocol separates the intervals of NV optical preparation and readout, MW control, and magnetic-field sensing, affording optimization of each NV-interrogation stage for improved sensitivity compared to CW-ODMR. Recent Ramsey-based magnetic-imaging advances have demonstrated robustness against errors arising from spatially heterogeneous MW control fields, diamond strain, and temperature variations [13], enabling spatially uniform sensitivity improvements relative to CW-ODMR-based measurements [47].

In this chapter, I describe the development and benchmarking of a pulsed Ramsey-based QDM for wide-field imaging of static and dynamic magnetic fields. The measurement protocol integrates double-quantum 4-Ramsey magnetometry (DQ4R), spin-bath driving [57], high-speed lock-in camera acquisition, and differential measurement protocols. Spin-bath driving is used to extend the NV spin dephasing time and thereby improve magnetic sensitivity, extending this approach from previous confocal-volume demonstrations to wide-field imaging. The QDM is benchmarked using a fabricated wire phantom (Fig. 2.1(a)), whose geometry generates nontrivial, micron-scale spatial patterns of magnetic fields.

The following sections first describe the experimental platform and integrated pulsed imaging protocol. The performance of the magnetic-imaging system is then evaluated through measurements of NV spin dephasing, magnetic sensitivity, and the spatial magnetic noise floor. Static magnetic-field imaging experiments establish micron-scale spatial resolution and demonstrate capability to resolve a picotesla-scale magnetic-field pattern by time averaging and signal modulation. The utility of image denoising for improving the measured image signal-to-noise (SNR) ratio is also examined. Finally, dynamic magnetic fields are imaged by applying to the phantom a broadband synthetic waveform that mimics a human cardiac signal. These measurements demonstrate time-resolved reconstruction of spatial magnetic-field patterns, while additional single-frequency measurements further establish sub-millisecond sampling capability. The chapter concludes by discussing technical limitations and prospects for further QDM improvements toward imaging weaker and more transient biomagnetic signals.

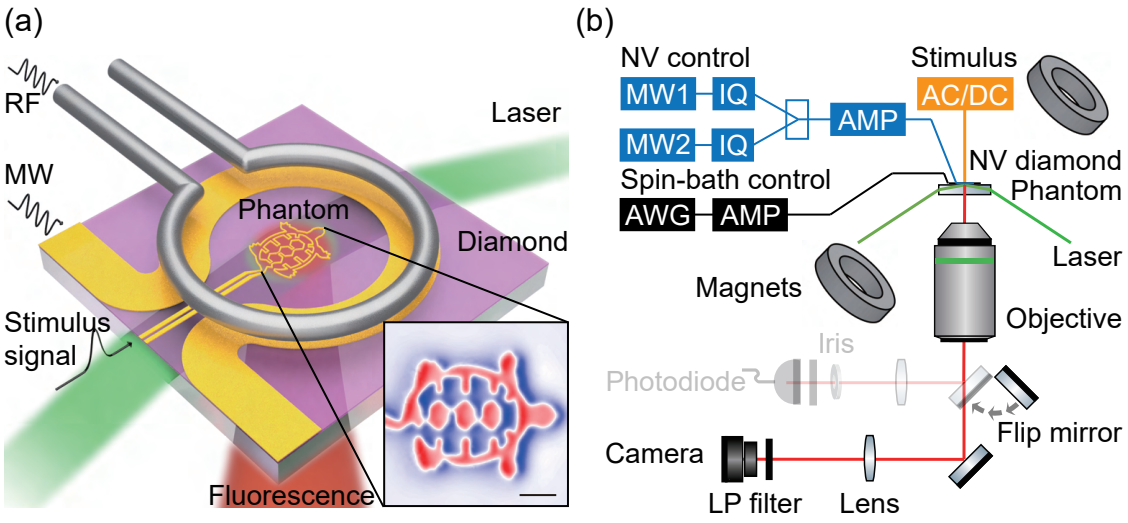


Figure 2.1: (a) Imaging phantom magnetic fields using a quantum diamond microscope (QDM). 532-nm excitation laser light excites NV centers at the diamond chip surface. A planar, gold omega-loop delivers microwaves (MWs) to the diamond chip for NV spin-state control. Optical components collect NV photoluminescence onto either a photodiode or a high-speed camera (not shown). A bias magnetic field (not shown) is applied along one NV axis. An RF signal delivered by a nearby coil drives paramagnetic bath spins in the diamond to increase the ensemble NV spin dephasing time (T_2^*). A driven current in the phantom produces a magnetic-field pattern to be imaged with the QDM. An example image of simulated phantom magnetic signals, projected along the direction of a sensing NV axis with $5\,\mu$m standoff from the source to the NV surface, is shown in the inset. Scale bar: $50\,\mu$m. (d) Experimental apparatus. A flip mirror can route the NV photoluminescence to a photodiode for rapid measurements of ensemble NV spin properties and experimental optimization.

2.2 Experimental platform and camera-based acquisition

The experimental setup is shown schematically in Fig. 2.1(b). The QDM sensor consists of a $10\ \mu$m-thick, ^{15}N-enriched CVD diamond layer with a nitrogen concentration of $[\mathrm{N}] = 17\ \mathrm{ppm}$ and $> 99.995\%$ ^{12}C enrichment, grown by Element Six Ltd. on a $2 \times 2 \times 0.5\ \mathrm{mm}^3$ high-purity diamond substrate. Post-growth treatment via electron irradiation and annealing increases the NV concentration to about $2.4\ \mathrm{ppm}$. Defect concentrations were reported by Element Six Ltd., using characterization techniques outlined in Ref. [58].

The excitation path is not shown in full in the schematic of Fig. 2.1(b). A diode-pumped solid-state laser (Lighthouse Photonics, Sprout-D5W) generates $532\ \mathrm{nm}$ excitation laser beam. The laser beam is focused onto an acousto-optic modulator (AOM; Gooch & Housego, Model 3250-220) using an $f = 500\ \mathrm{mm}$ spherical convex lens; and the first-order diffracted beam from the AOM is used for the measurements. Optical pulses for NV initialization and readout are created by gating the AOM radio frequency (RF) driver with switches (Mini-Circuits, ZASWA-2-50DR+) using transistor-transistor logic (TTL) pulses from a pulse generator (SpinCore Technologies, PBESR-PRO-500). After the AOM, the first-order beam is shaped by an $f = 130\ \mathrm{mm}$ cylindrical convex lens and an $f = 100\ \mathrm{mm}$ spherical convex lens to produce an elliptical beam profile incident on a polished side surface of the diamond sample. Approximately $2.5\ \mathrm{W}$ of optical laser power enters the diamond and undergoes total internal reflection at the surface of the NV layer, illuminating an area of approximately $300 \times 600\ \mu\mathrm{m}^2$.

NV photoluminescence is collected with a $20\times$ 0.75 NA Nikon objective and passed through a $647\ \mathrm{nm}$ long-pass filter (Semrock). A flip mirror directs the collected photoluminescence either to a photodiode (Thorlabs, PDA36A2), used for rapid ensemble-spin characterization and experimental optimization, or to a lock-in camera (Heliotis, heliCam C3) for magnetic imaging. Photodiode

voltage measurements are acquired using a data acquisition (DAQ) system (National Instruments, NI USB6363).

Dual-tone microwave (MW) pulses for the double-quantum 4-Ramsey protocol described in Section 2.3.1 are generated by two signal sources with built-in IQ mixers (Stanford Research Systems, SG384) and gated by switches controlled with TTL pulses. A fabricated gold Ω-shaped coplanar waveguide, with nominal thickness of $300\ \mathrm{nm}$ and width of $800\ \mu\mathrm{m}$, delivers MW fields to the NV layer. The fabricated palladium wire phantom, patterned in the shape of a terrapin, has a trace width of approximately $4\ \mu\mathrm{m}$, spans an area of approximately $220 \times 180\ \mu\mathrm{m}^2$, and has a thickness of approximately $10\ \mathrm{nm}$. To generate phantom magnetic-field signals, stimulus voltages produced by a waveform generator (Keysight, 33622A) are applied to the phantom to induce currents. Both the phantom and the MW waveguide are fabricated on a silicon-carbide wafer, which provides mechanical support and heat dissipation. The diamond is glued to the surface of the MW waveguide; thus giving a geometric minimum phantom-to-NV separation set by the approximately $300\ \mathrm{nm}$ MW structure thickness. The actual standoff distance is assessed later by comparison of the measured phantom fields with finite-element simulations.

Multi-tone, resonant radio-frequency (RF) pulses for spin-bath driving are generated by an arbitrary waveform generator (Keysight, M8190) and gated by switches controlled with TTL pulses. The RF fields are delivered to the bath spins by a shorted coaxial loop, about $1\ \mathrm{mm}$ wide, positioned above the MW waveguide. The role of spin-bath driving within the pulsed imaging protocol is described in Section 2.3.2.

A bias magnetic field of approximately $4.3\ \mathrm{mT}$ is produced by two pairs of temperature-compensated permanent samarium-cobalt ring magnets. The field is aligned along one of the four NV ensemble axes, such that the QDM measures the projection of the signal magnetic field along the selected NV axis.

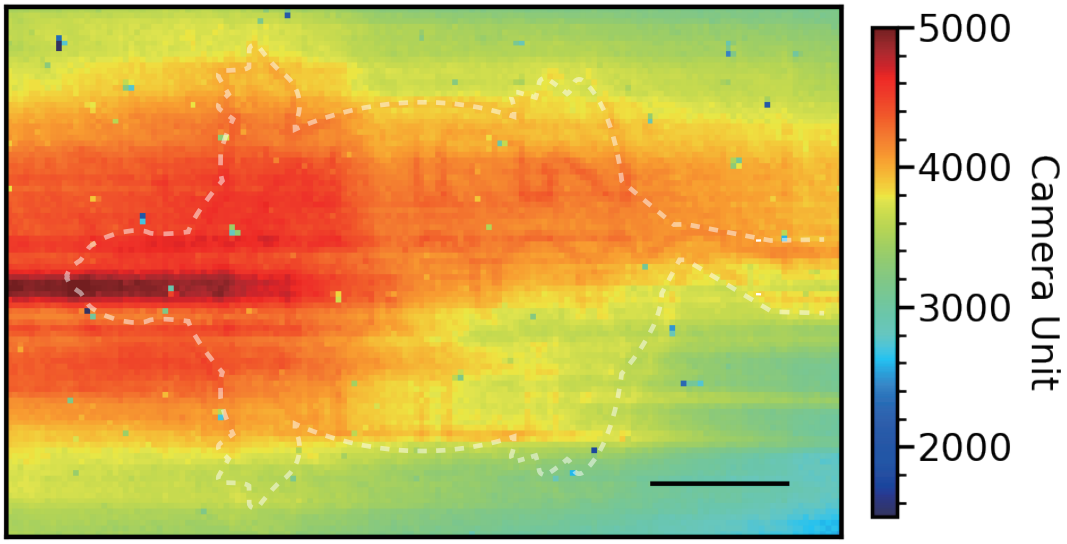


Figure 2.2: Optical image of the fabricated wire phantom. NV photoluminescence intensity under continuous laser illumination is collected using the QDM's camera. Semi-transparent white dashed traces of the phantom perimeter are added as a guide to the eye. The measured photoluminescence is reported in the camera unit, which is linearly proportional to the photon number. Scale bar: $50\,\mu\mathrm{m}$.

This bias-field magnitude was convenient for the present apparatus; but similar QDM performance should be possible for a wide range of magnetic fields, from millitesla to tesla scales.

The size of a single camera pixel in the objective focal plane is determined using the known geometry of the phantom and its optical image as visualized using NV fluorescence. The metallic phantom traces reflect a portion of the incident excitation light, locally increasing the optical illumination and producing enhanced NV photoluminescence that follows the phantom structure (Fig. 2.2). The corresponding area of the NV layer sampled by each camera pixel is approximately $1.9 \times 1.9\ \mu\mathrm{m}^2$. This calibration agrees with the estimate using the $20\times$ objective magnification and the physical camera pixel dimensions of $39.6 \times 39.6\ \mu\mathrm{m}^2$. A $270 \times 270\ \mu\mathrm{m}^2$ region of the NV layer is selected as the field of view for the measurements described below.

The lock-in camera has a controllable internal exposure frequency of up to $1\ \mathrm{MHz}$ and an external frame rate of up to $3.8\ \mathrm{kHz}$. During acquisition, successive internal exposures are stored alternately in the in-phase (I) and quadrature (Q) channels. A unit cycle consists of four internal exposures, photoluminescence is accumulated sequentially in exposures I_1, Q_1, I_2, and Q_2. The camera performs on-chip subtraction between successive exposures in each channel, yielding differential signals of the form $I_1 - I_2$ and $Q_1 - Q_2$. This so-call "demodulation" cycle can be repeated multiple times before the accumulated channel signals are digitized and output as an external frame. The number of accumulated cycles therefore sets the external frame rate, ensuring it is below the maximum value. Section 2.3.1 describes how the MW phase cycle of the DQ 4-Ramsey sequence is synchronized with this I/Q acquisition scheme to produce a magnetic signal in each external frame.

The camera buffer stores a continuous acquisition of up to 500 external frames, after which the data are transferred to the control computer. Transfer from the camera buffer requires approximately $5\ \mathrm{s}$. Additional software-related overhead associated with data conversion and storage contributes approximately $2\ \mathrm{s}$ for static magnetic-field imaging, where the average of 500 frames is stored, and approximately $17\ \mathrm{s}$ for dynamic magnetic-field imaging, where individual frames are retained. No sensing is performed during these transfer and processing intervals. Accordingly, the signal acquisition time T_{acq} denotes only the time spent actively acquiring magnetic measurements, whereas the wall-clock time T_{wall} includes the full elapsed experimental time, including buffer transfer, data conversion, and storage.

2.3 Pulsed magnetic-imaging protocol

2.3.1 DQ 4-Ramsey and camera implementation

The DQ 4-Ramsey sequence for magnetic imaging utilizes dual-tone MW pulses and a MW-phase alternation scheme to achieve uniform magnetic sensitivity over the field of view [13]. The dual-tone MW pulse generates a double-quantum (DQ) coherence between the $|m_s = +1\rangle$ and $|m_s = -1\rangle$ ground state spin triplet sublevels of the NV center. Compared to using a single-quantum (SQ) coherence between $|m_s = 0\rangle$ and $|m_s = +1\rangle$ or $|m_s = -1\rangle$, the DQ scheme is twice as sensitive to magnetic fields and can isolate magnetic signals from spurious strain and temperature artifacts. In practice, however, spatial inhomogeneity of MW control fields, variations in diamond strain, and temperature can cause errors in DQ pulses across the NV ensemble and generate a residual SQ coherence. Thus a single DQ Ramsey measurement is still not fully immune to these spurious signals contributions. To mitigate them, the phases of the two MW tones are alternated across four individual DQ Ramsey measurements [13]. If we denote the fluorescence signals from the i-th measurement as S_i, the phase-alternation scheme generates DQ magnetic signals in S_2 and S_4 that are opposite to S_1 and S_3, while residual SQ signals in the sum $S_1 + S_3$ are also observed in $S_2 + S_4$. A final normalized signal

$$S_{\text{norm}} = S_1 - S_2 + S_3 - S_4 \tag{2.1}$$

therefore preserves the DQ component while suppressing residual SQ contributions arising from pulse errors and common-mode magnetic artifacts. This phase-alternated measurement protocol is also well suited to lock-in camera detection, which further suppresses laser intensity fluctuations during signal

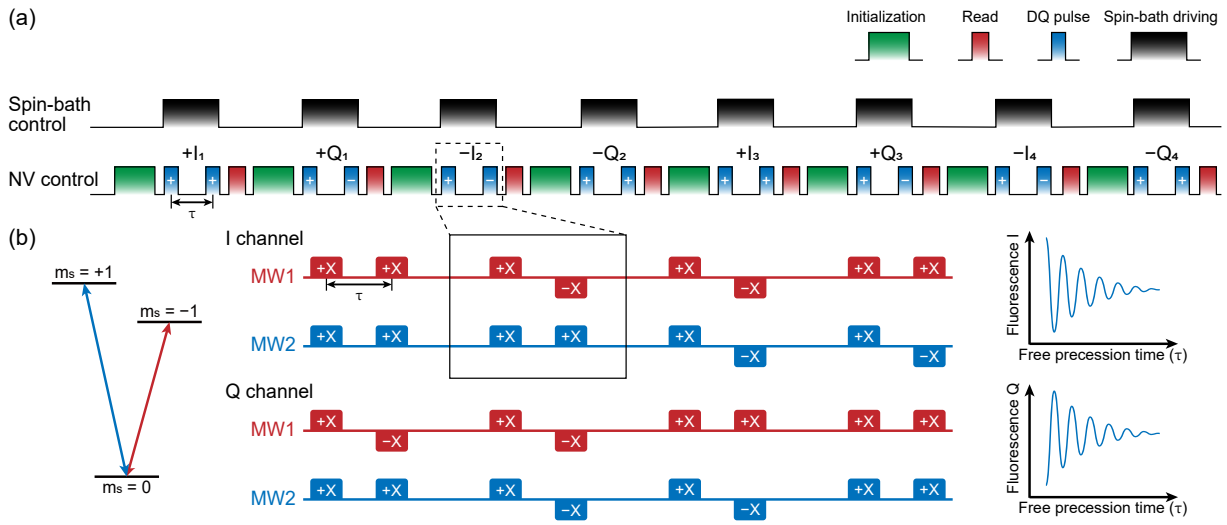


Figure 2.3: (a) Sequence for magnetic imaging using the DQ 4-Ramsey protocol for NV interrogation together with spin-bath driving. The camera internal exposures alternate between the I and Q channel; on-chip circuitry accumulates differential results between two successive measurements in both channels. For example, the output in the I channel using the sequence in (a) will be $S_{I1} - S_{I2} + S_{I3} - S_{I4}$. (b) Left: Dual-tone MW resonant with $|m_s = 0\rangle \leftrightarrow |m_s = +1\rangle$ and $|m_s = 0\rangle \leftrightarrow |m_s = -1\rangle$ NV spin transitions is used to implement DQ pulses. Middle: Phase alternations of individual MW tones in a DQ 4-Ramsey sequence. The $\pm$X directions are defined in the effective Bloch sphere formed by the two NV spin levels $|m_s = 0\rangle$ and $|m_s = +1\rangle$ (or $|m_s = -1\rangle$). The duration of the DQ pulse is the time to change the NV spin state from $|m_s = 0\rangle$ to a superposition state with equal population between $|m_s = +1\rangle$ and $|m_s = -1\rangle$. Right: The permutation of MW-phase alternation in the Q channel results in NV fluorescence dynamics of opposite phase to that of the I channel. Differential results between these two measurements are output to an external frame.

accumulation.

Figure 2.3 illustrates the implementation of the DQ 4-Ramsey protocol together with spin-bath-driving pulses for wide-field magnetic imaging using the Heliotis heliCam lock-in camera. The spin-bath drive is generated from continuous-wave multi-tone RF signals and gated such that it is applied during the Ramsey free-evolution interval; its physical role and calibration are described in the following subsection. The NV photoluminescence signals from successive DQ Ramsey measurements are synchronized with the internal camera exposures described in Section 2.2. The camera alternates

between the in-phase (I) and quadrature (Q) channels and performs on-chip subtraction between successive exposures in each channel. A DQ 4-Ramsey phase cycle is synchronized with the internal exposures in both channels. The MW phase permutation in the Q channel is chosen such that the resulting signal S^Q_{norm} has the opposite polarity to the corresponding signal S^I_{norm} accumulated in the I channel. Taking the difference between these two channel signals yields the final magnetic signal,

$$S_{\mathrm{norm}} = S^I_{\mathrm{norm}} - S^Q_{\mathrm{norm}}. \tag{2.2}$$

For all magnetic imaging experiments (including measurements with and without spin-bath driving, measurements of spatial magnetic noise, and phantom magnetic-field images), the NV photoluminescence signal in the camera external frame S_{norm} is converted to magnetic signal using a magnetometry calibration curve (as similarly described in Chapter 1.2). Briefly, with the Ramsey free evolution interval τ fixed, the MW frequency is swept around the NV electronic spin resonances to emulate the change of magnetic field B; the response of S_{norm} is then measured, producing a magnetometry calibration curve. The optimal MW frequency for magnetic measurements is then determined by maximizing the slope dS_{norm}/dB. For DQ magnetometry, the two MWs that are separately resonant with NV electronic spin transitions $|m_s = 0\rangle \leftrightarrow |m_s = +1\rangle$ and $|m_s = 0\rangle \leftrightarrow |m_s = -1\rangle$ are swept differentially (i.e., a detuning of δ is applied to one MW tone while a detuning of $-\delta$ is applied to the second MW tone). The Ramsey free evolution interval is chosen near the ensemble dephasing time T_2^*; in practice, τ is adjusted to maximize the median per-pixel magnetic sensitivity because hyperfine-induced Ramsey-fringe beating constrains the available operating points.

2.3.2 Spin-bath driving

Spin-bath driving [57] is used to decouple dipolar interactions between the NV centers and paramagnetic bath spins ($S = 1/2$) in the diamond. The relevant bath-spin species include neutral substitutional nitrogen, $^{15}\mathrm{N}_\mathrm{s}^0$, in the ^{15}N-enriched diamond used here, as well as free electron spins. A strong multi-tone radio-frequency (RF) drive is applied during the Ramsey free evolution interval to induce rapid Rabi oscillations of the bath spins. The effective dipole moments of these spins are thereby time averaged, reducing their contribution to inhomogeneous broadening of the NV spin resonance and extending the NV spin dephasing time, T_2^*.

The frequency components of the RF drive must be resonant with the relevant bath-spin transitions. These transitions are determined experimentally using NV double electron-electron resonance (DEER) spectroscopy [57, 59], as shown in Fig. 2.4(a). In a DEER measurement, the NV spin signal is monitored while an RF pulse is swept through the bath-spin resonance frequencies. To guide the frequency range for these measurements, the expected bath-spin transition frequencies are calculated. For free-electron spins, the transition angular frequency is $\gamma_e|\vec{B}|$, where γ_e is the electron gyromagnetic ratio, and $|\vec{B}|$ is the magnitude of the applied bias magnetic field. The transition frequency spectrum for $^{15}\mathrm{N}_s^0$ is calculated using its ground-state Hamiltonian [60]

$$\frac{\hat{H}}{h} = \frac{\mu_B}{h}\vec{B}\cdot\overline{g}\cdot\hat{\vec{S}} + \frac{\mu_N}{h}\vec{B}\cdot\hat{\vec{I}} + \hat{\vec{S}}\cdot\overline{A}\cdot\hat{\vec{I}}, \tag{2.3}$$

where μ_B is the Bohr magneton, μ_N is the nuclear magnetic moment, h is the Planck's constant, $\vec{B} = (B_x, B_y, B_z)$ is the applied bias magnetic field vector. Here, $\overline{g}$ is the electronic g-factor tensor, $\hat{\vec{S}} = (\hat{S}_x, \hat{S}_y, \hat{S}_z)$ is the dimensionless electronic spin vector, $\overline{A}$ is the hyperfine tensor, and $\hat{\vec{I}} =$

$(\hat{I}_x, \hat{I}_y, \hat{I}_z)$ is the dimensionless nuclear spin vector.

When the z-axis is chosen as one of the four NV crystal axis directions, tensors $\overline{g}$ and $\overline{A}$ are diagonal and can be written as

$$\overline{g} = \begin{bmatrix} g_\perp & 0 & 0 \\ 0 & g_\perp & 0 \\ 0 & 0 & g_{||} \end{bmatrix}, \overline{A} = \begin{bmatrix} A_\perp & 0 & 0 \\ 0 & A_\perp & 0 \\ 0 & 0 & A_{||} \end{bmatrix}, \tag{2.4}$$

where $g_\perp$, $g_{||}$, $A_\perp$, and $A_{||}$ are the gyromagnetic and hyperfine on and off-axis components, respectively. $^{15}\mathrm{N}_s^0$ has $S = 1/2$ and $I = 1/2$, leading to the four eigenstates $|m_S = \pm 1/2, m_I = \pm 1/2\rangle$. The two corresponding dipole-allowed transitions ($\Delta m_S = \pm 1$, $\Delta m_I = 0$, solid arrows), along with the two first-order forbidden transitions ($\Delta m_S = \pm 1$, $\Delta m_I = \pm 1$, dashed arrows) are shown in Figure 2.4(b). The simulated $^{15}\mathrm{N}_s^0$ resonance spectrum, using $g_\perp = g_{||} = 2$, $A_\perp = -559.7\,\mathrm{MHz}$, $A_{||} = -113.83\,\mathrm{MHz}$, and $|\vec{B}| = 4.278\,\mathrm{mT}$ aligned with an NV axis, is shown in Figure 2.4(c), along with the simulated free-electron spin resonance frequency ($\mathrm{g} = 2$) and the measured NV DEER spectrum.

After the RF frequency components are determined, the RF amplitudes are increased to achieve Rabi frequencies of approximately $2\ \mathrm{MHz}$ for the dipole-allowed $^{15}\mathrm{N}_\mathrm{s}^0$ and free electron transitions. The attainable Rabi frequencies are limited by the available RF amplifier power. Increasing the bath-spin Rabi frequency suppresses the corresponding bath-spin contribution to NV dephasing, such that T_2^* eventually approaches the limit imposed by the next dominant dephasing mechanism, here NV-NV dipolar interactions [57]. Because the shorted RF loop produces a spatially varying RF amplitude across the field of view, the bath-spin Rabi frequency and the resulting extension of T_2^* can vary

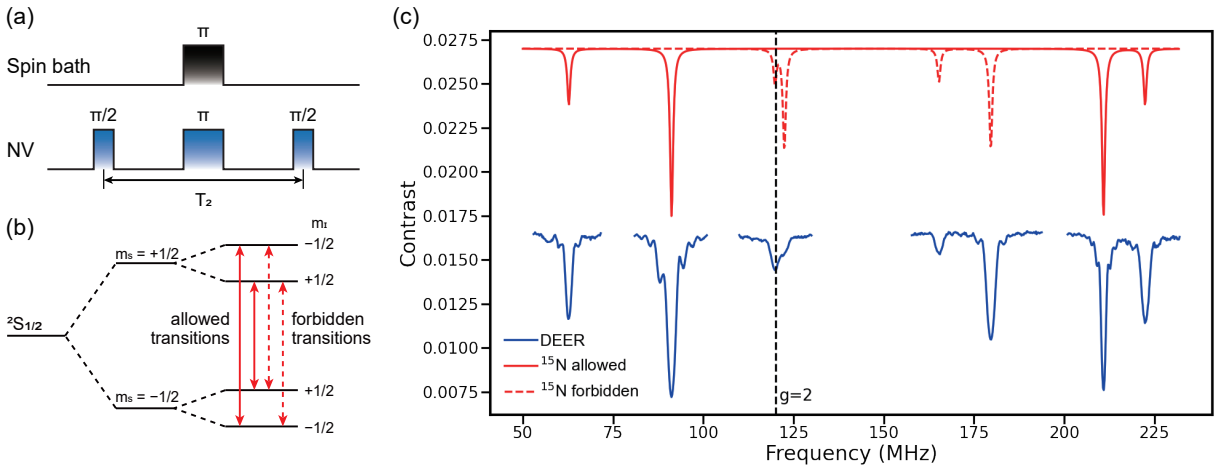


Figure 2.4: (a) NV double electron-electron resonance (DEER) technique for measuring the $^{15}N_s^0$ and free-electron spin resonance spectrum in the QDM diamond sensor. The total free evolution intervals between the initial and final single-quantum (SQ) $\pi/2$ pulse on the NV spins is chosen to match the ensemble NV spin-echo T_2 for optimal sensitivity to bath-spin resonances. The frequency of the SQ π pulse on the bath spins is swept to measure the NV DEER spectrum by the effect on the NV ensemble fluorescence contrast (see Ref. [57] for details). (b) Energy level diagram of the coupled electron-nuclear spin system in $^{15}N_s^0$. (c) Simulated $^{15}N_s^0$ and free-electron ($g = 2$) spin resonance spectrum (red) overlaid with NV DEER measurement results (blue) and offset vertically for clarity. The simulated resonant frequencies are broadened by a Lorentzian function with linewidth and relative contrast chosen to approximate the measurement data.

spatially.

2.3.3 Differential measurement protocol

Despite the DQ 4-Ramsey protocol and lock-in camera acquisition, several sources of spatially and temporally varying magnetic noise remain difficult to suppress. In particular, variations among helicam camera pixels can produce a pixel-to-pixel baseline pattern (this variation is not static). Vibration and temperature change can induce drift of the bias magnetic field, generating slowly varying spatial artifacts in magnetic images. To mitigate these effects, a differential measurement protocol is used. Each differential measurement pair consists of two sets of 500 camera frames acquired using

identical magnetic sensing sequences, except that the current applied to the phantom is modulated between the two frame sets. This approach is well suited to controllable magnetic signals, including those from electrical circuits, externally stimulated electrically active cells, or other sources whose magnetic signal can be reversed or gated.

In the following experiments, for static magnetic-field measurements, one set of frames is collected with positive current applied to the phantom, while the other frame set is collected with negative current. For dynamic magnetic-field measurements, the applied current is modulated between on and off. The first modulation scheme has identical magnetic sensitivity compared to measuring with steady (unmodulated) current flowing through the phantom, while the second modulation scheme costs approximately a factor of 2 in sensitivity.

After calculating the difference between the two frame sets of each measurement pair, the resulting magnetic image is further processed by subtracting a 2D Gaussian-filtered version of the image, with filter width $\sigma_{x,y}^{\mathrm{filt}} = 18\ \mu\mathrm{m}$. This procedure reduces residual magnetic gradients with spatial length scales larger than those of the phantom field pattern. The filter width was selected experimentally to exceed the characteristic spatial variation length scale of the phantom magnetic field. The processed magnetic images are then averaged over the acquisition time.

2.4 Spin dephasing and magnetic sensitivity

We first characterize the performance of the Ramsey-based QDM with and without spin-bath driving. The NV spin dephasing time T_2^* and magnetic-field sensitivity η are evaluated on a pixel-by-pixel basis over the field of view. For these measurements, a dual-tone MW implementation of the Ramsey sequence generates a DQ coherence that is robust against common-mode frequency shifts

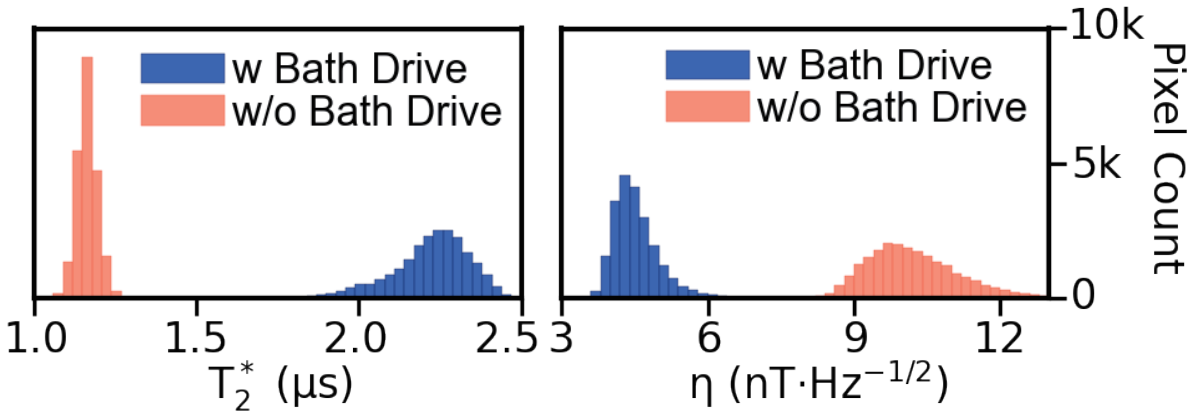


Figure 2.5: Measured magnetic sensing performance of the Ramsey-based QDM. Median per-pixel T_2^* is 2.2(1) µs and 1.2(1) µs with and without spin-bath driving. Median per-pixel η is 4.1(1) nT/$\sqrt{\text{Hz}}$ and 9.4(1) nT/$\sqrt{\text{Hz}}$ with and without spin-bath driving.

from diamond strain and temperature variation such that the NV ensemble T_2^* is limited by dipolar interactions with bath spins. Under this condition, spin-bath driving is expected to improve the NV spin dephasing time and magnetic sensitivity as shown in Figure 2.5.

With spin-bath driving, the median per-pixel T_2^* is extended by about 1.8× to 2.2(1) μs, approaching the estimated NV-NV interaction limit of 2.5 µs for this sample (See Appendix A.1). The median per-pixel η is enhanced by about 2.4× to 4.1(1) nT/$\sqrt{\text{Hz}}$ in the absence of pixel binning. The measured sensitivity with spin-bath driving is comparable to the combined camera-quantization- and photon-shot-noise-limited estimate of 3.2 nT/$\sqrt{\text{Hz}}$, to which the two noise sources make approximately equal contributions (See Appendix A.2).

The experimental per-pixel magnetic sensitivity results are computed from

$$\eta = \delta B \sqrt{T_{\text{acq}}}, \tag{2.5}$$

where δB is the smallest magnetic field (i.e., RMS magnetic-field noise) that can be measured with

$\mathrm{SNR} = 1$ after T_{acq} acquisition time. For a camera operating at an external frame rate F_s, a series of magnetic noise images are collected. Then the magnetic noise of each pixel, σ_{pxl}, is obtained from the standard deviation of the magnetometry data from that pixel across all recorded frames. The per-pixel magnetic sensitivity can then be written as

$$\eta = \frac{\sigma_{pxl}}{\sqrt{F_s}}. \tag{2.6}$$

The results shown in Figure 2.5 are obtained using 500 frames collected at $F_s \approx 1.4$ kHz with spin-bath driving and $F_s \approx 1.7$ kHz without spin-bath driving. The uncertainty reported for the measured median per-pixel sensitivity η with spin-bath driving is computed as the standard deviation of η across 15 repeated experiments, each consisting of 500 frames. The experiment without spin-bath driving is only repeated once. As an estimate of the uncertainty for η without spin-bath driving, the 500 frames are divided into ten 50-frame sets and the standard deviation of η across the ten frame sets is reported. Experiments for per-pixel magnetic sensitivity measurements with spin-bath driving are also performed at different values of F_s ($\approx$ 3.1 kHz, 2.2 kHz, 1.0 kHz and 0.5 kHz) to investigate potential technical noise at different imaging speeds. Across all these measurements, the difference between the maximal and minimal median magnetic sensitivity values is $\lesssim 5$ %.

For photon shot-noise-limited Ramsey magnetometry, the magnetic sensitivity scales as [1]

$$\eta_{\mathrm{shot}} \propto \frac{\sqrt{t_D + \tau}}{\tau C exp[-(\tau/T_2^*)^p]}, \tag{2.7}$$

where C is the NV spin-state readout contrast, τ is the Ramsey free evolution interval, p is a parameter used to describe the Ramsey envelope decay shape, and τ_D is the overhead time for NV initialization

and readout. The present QDM operates in a regime where the sensitivity nominally improves linearly with increased dephasing time [1, 57], as $\tau_D = 7.04\,\mu\text{s}$ is longer than the sensing interval $\tau \approx T_2^*$, yielding

$$\eta_{\text{shot}} \propto \frac{\sqrt{t_D}}{\tau C}. \tag{2.8}$$

The measured sensitivity enhancement is slightly greater than that expected from the T_2^* increase alone. This additional improvement is attributed to an approximately $1.2\times$ increase in the spin-state readout contrast C when spin-bath driving is applied, together with the discrete choices of τ imposed by Ramsey-fringe beating from the hyperfine splitting of the NV spin resonances [13].

The per-pixel magnetic sensitivities reported here describe the expected magnetic noise after one second of active signal acquisition, $T_{\text{acq}} = 1$ s, and exclude the camera-buffer transfer and data-processing overhead included in the corresponding wall-clock time T_{wall}.

2.5 Spatial magnetic noise floor

In magnetic imaging, the ability to resolve a signal of interest depends not only on the magnetic sensitivity of individual pixels, but also on the spatial magnetic noise floor across the full image [44]. Per-pixel magnetic sensitivity describes the temporal noise measured at a fixed location. However, even when the noise of individual pixels averages down with acquisition time, pixel-to-pixel offsets or spatially correlated fluctuations can remain in the time-averaged image. Such residual spatial structure can ultimately limit the smallest magnetic-field pattern that can be resolved. Following Ref. [44], the spatial magnetic noise floor, σ_{spatial}, is quantified after a given acquisition time, by calculating the standard deviation of magnetic-field measurements across the time-averaged image without any external signal sources present. Uncorrelated spatial noise can be time-averaged and reduced. However,

correlated spatial noise across an image, particularly when similar to the length scale of the signal of interest, inhibits the utility of temporal averaging and limits the minimum resolvable magnetic field. Specifically in the QDM, vibration and temperature changes can induce spatial and temporal variations in the bias magnetic field and optical illumination intensity, thereby introducing correlated noise into QDM magnetic images.

The differential measurement protocol described in Section 2.3.3 is used to suppress these slowly varying spatial artifacts. As shown in Figure 2.6, the spatial distribution of noise in QDM images, at different acquisition times T_{acq} and with no current flowing in the phantom, have insignificant correlation. The spatial magnetic noise floor, $\sigma_{spatial}$, can be averaged down as reflected in the observed scaling

$$\sigma_{\text{spatial}} \propto T_{\text{acq}}^{-1/2}, \tag{2.9}$$

which persists up to about $T_{acq} = 4000\,\text{s}$ ($T_{wall} \approx 8\,\text{h}$). This behavior indicates that the remaining image noise is dominated by uncorrelated contributions over these acquisition times.

Spatial averaging through pixel binning further reduces the spatial magnetic noise floor. When $n_{\text{bin}} \times n_{\text{bin}}$ camera pixels are combined, the measured noise floor follows approximately

$$\sigma_{\text{spatial}} \propto \frac{1}{n_{\text{bin}}}, \tag{2.10}$$

allowing spatial resolution to be exchanged for improved magnetic sensitivity and reduced acquisition time. Additional measurements of the spatial magnetic noise floor for different pixel-binning conditions, as well as measurements performed without the differential protocol, are presented in Fig. 2.7.

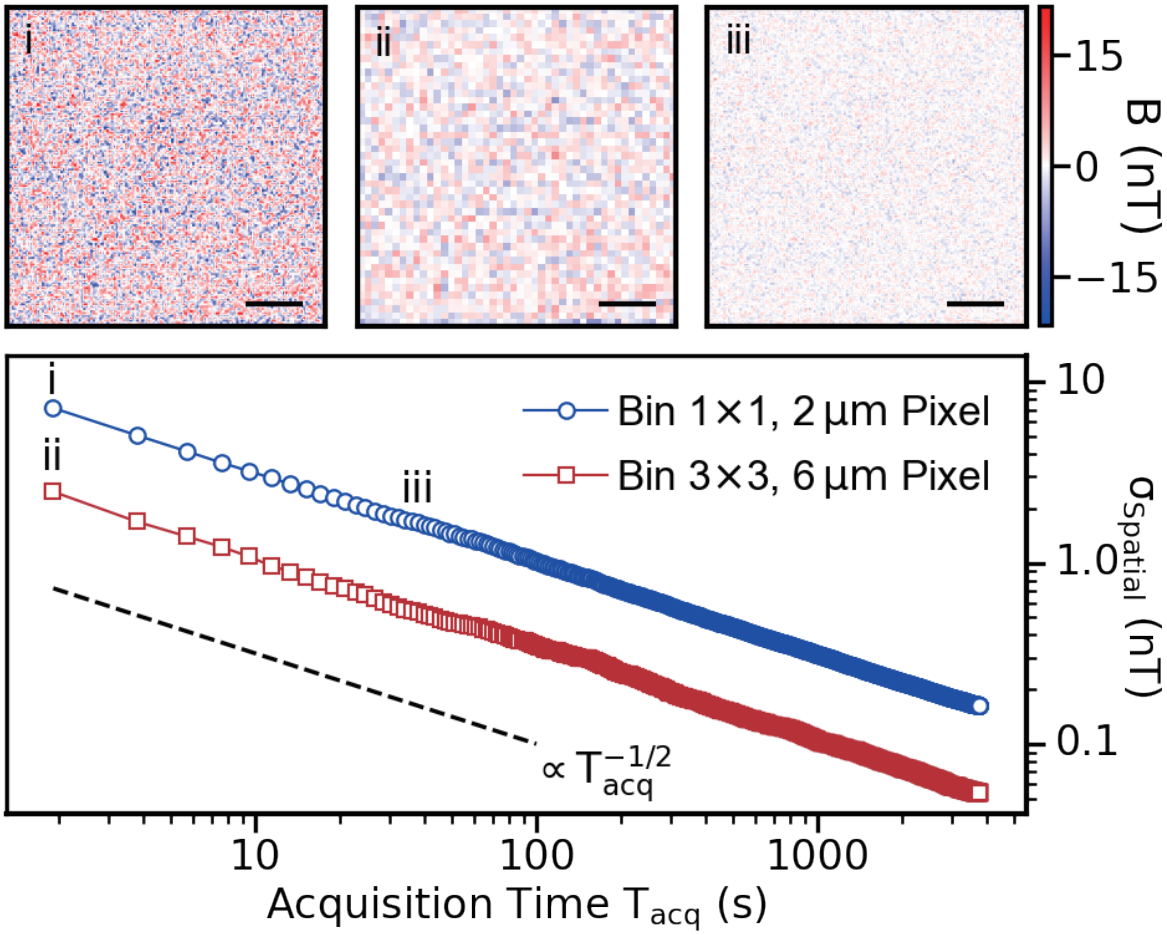


Figure 2.6: Spatial magnetic noise floor. Measurement data are obtained using the differential measurement protocol. Top: Images of magnetic noise with different acquisition times and pixel binning (see the bottom plot). Here, the acquisition time accounts for the duration of time spent acquiring measurements, but excludes the implementation-dependent time required to transfer and store data from the camera to a host computer. The magnetic noise is distributed randomly without significant spatial correlation. Field of view: $270 \times 270\,\mu m^2$. Scale bar: 50 µm. Bottom: Spatial magnetic noise floor $\sigma_{spatial}$ as a function of acquisition time T_{acq} for no pixel binning and 3×3 pixel binning. $\sigma_{spatial}$ is computed as the standard deviation of magnetic-field measurements using the entire image after any pixel binning. A dashed black line depicts power law scaling behavior $\propto T_{acq}^{-1/2}$ as a guide to the eye.

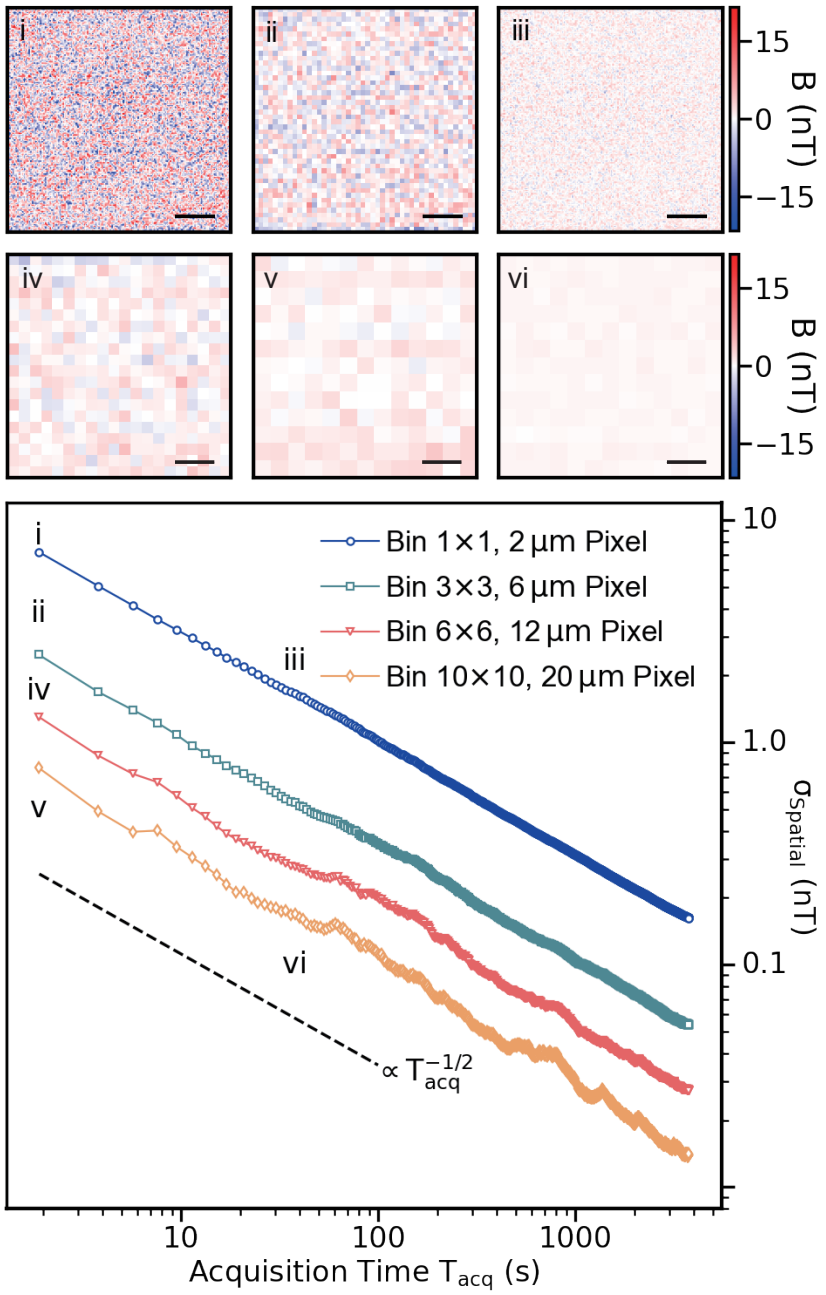


Figure 2.7: Additional pixel binning results of spatial magnetic noise floor in QDM images. Measurement data are obtained using the differential measurement protocol. Upper six panels: Images of the background magnetic noise with different acquisition times and pixel binning (values indicated in lower panel). Scale bar: 50 µm. Lower panel: Spatial magnetic noise floor $\sigma_{spatial}$ as a function of acquisition time T_{acq}. A dashed black line depicts power law scaling behavior $\propto T_{acq}^{-1/2}$ as a guide to the eye.

2.6 Static magnetic-field imaging

We characterize the capability of the Ramsey-based QDM to image static magnetic fields by applying steady currents to the fabricated wire phantom. The polarity of applied signals is reversed between acquisitions of successive frame sets to implement the differential measurement protocol described in Section 2.3.3. As shown in Figure 2.8(a), an imaged magnetic-field pattern with a peak-to-peak amplitude difference of about 14 nT is obtained by supplying a fixed voltage of 25 mV, corresponding to a phantom current of $436\ \mathrm{nA}$. The image is acquired with a signal-acquisition time of $T_{\mathrm{acq}} = 200\ \mathrm{s}$, corresponding to a wall-clock time of approximately $T_{\mathrm{wall}} \approx 21\ \mathrm{min}$.

The expected phantom magnetic-field pattern is simulated using finite-element software [61] to calculate the current-density distribution and resulting magnetic field for the measured phantom geometry and applied current. Additional details of the simulation are provided in Ref. [62]. The phantom-to-NV standoff distance d_{so} is the only free parameter in the simulation and is adjusted to obtain the best agreement with the measured magnetic-field pattern [48]. This procedure yields an effective standoff distance of $d_{\mathrm{so}} \approx 5\ \mu\mathrm{m}$. Additional simulated phantom magnetic-field patterns projected along the other NV crystallographic axes are provided in [62].

We further demonstrate QDM magnetic imaging of a picotesla-scale signal by applying a $500\ \mu\mathrm{V}$ voltage, corresponding to a phantom current of $9\ \mathrm{nA}$. After averaging for about $T_{acq} = 3600\ \mathrm{s}$ ($T_{wall} \approx 7\ \mathrm{h}$), a phantom magnetic-field pattern with a peak-to-peak amplitude difference of about 300 pT is resolved with $\mathrm{SNR} \approx 3$ as shown in Figure 2.8, left panel. Because the residual magnetic noise exhibits negligible pixel-to-pixel spatial correlation after differential processing (Fig. 2.6), image-denoising methods can be applied to improve the visibility of the measured field pattern.

A non-local mean (NLM) denoising algorithm [63] is applied to the picotesla-scale magnetic

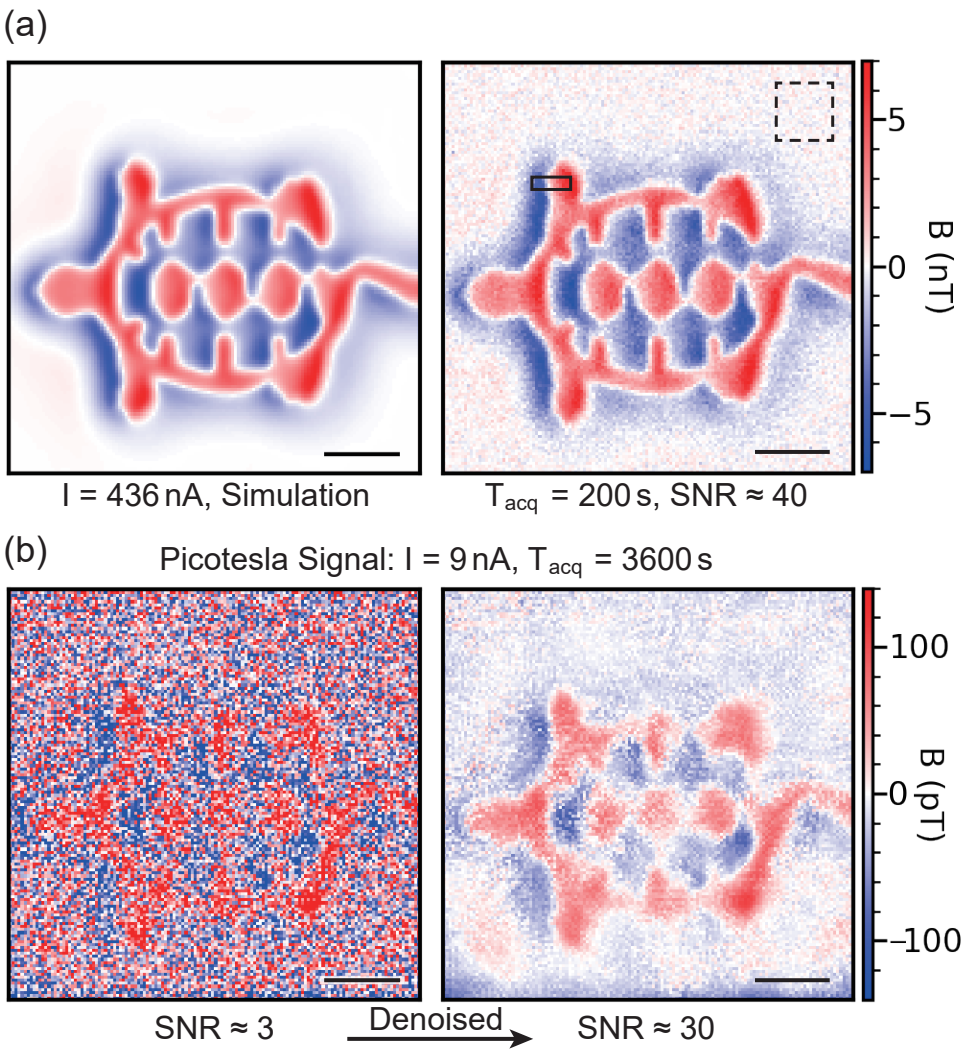


Figure 2.8: (a) Simulated (Left) and measured (Right) phantom static magnetic-field patterns using the Ramsey-based QDM and a differential measurement protocol. The measurement is conducted with a 25 mV voltage (436 nA current) applied to the phantom and time-averaged for about 200 s. To compute SNR, the peak-to-peak amplitude difference of the magnetic-field pattern (solid box) is divided by the spatial magnetic noise obtained from a source-free region (dashed box). The simulated phantom magnetic-field pattern is obtained using the applied current and geometry of the phantom. To mimic the parameters of the imaging system, the simulated magnetic fields are further averaged over a depth of 10 µm (the NV-layer thickness) and binned with a pixel area of $1.9 \times 1.9\,\mu\text{m}^2$. NV-phantom standoff distance is the only free parameter, and is tuned to a value of 5 µm to best match the measured magnetic-field pattern. (b) Demonstration of a picotesla-scale QDM magnetic image. Left: After about $T_{acq} = 3600\,\text{s}$ ($T_{wall} \approx 7\,\text{h}$) time-averaging, a magnetic-field pattern with a peak-to-peak amplitude difference of about 300 pT is resolved with $\text{SNR} \approx 3$. Right: Applying a non-local mean denoising technique enhances image SNR by about $10\times$. All images are not binned and scale bars are 50 µm for (a, b).

image shown in Figure 2.8(b), right panel, to improve SNR. The main feature of this denoising algorithm is the ability to restore both the large-scale geometries and fine structures (e.g., contours, patterns, and edges where the Laplacian is non-zero) from a noisy image; whereas local-mean filtering, such as Gaussian smoothing, can blur fine structures [64]. Briefly, to denoise a pixel p, the NLM algorithm will compute the weighted averages of all other pixels q within a search window centered at pixel p. The weight assigned to a pixel q is computed using the Euclidean L^2 distance between two template windows centered at q and p, respectively — i.e., summation over squared pixel-pair value differences with each pair consisting of two pixels from the template windows around q and p, respectively — and normalized using a 1D Gaussian function. If the template window around q has a similar intensity profile compared to that of p, a higher weight will be assigned to q; this process is then repeated for all q within the search window centered at p. Thus the denoised pixel p is predominantly the averaged value of similar q pixels, rather than its immediate neighbors in a local-mean denoising method.

This property is particularly useful for magnetic images of the fabricated wire phantom, in which similar portions of the field pattern can occur at spatially separated locations. We therefore expect NLM denoising to improve the image signal-to-noise ratio while introducing less distortion of the phantom magnetic-field pattern than local-mean filtering methods such as Gaussian smoothing. A comparison with Gaussian smoothing is provided in Appendix A.3. The NLM search and template windows are chosen to be 21×21 pixels and 7×7 pixels, respectively, following the recommended values in Ref. [63]. The NLM filtering parameter is set to approximately $44\ \mathrm{pT}$, corresponding to about $0.5 \times \sigma_{\mathrm{spatial}}$ for the picotesla-scale image after $T_{\mathrm{acq}} = 3600\ \mathrm{s}$. This parameter was manually optimized to maximize the reconstructed image signal-to-noise ratio, yielding $\mathrm{SNR} \approx 30$.

2.7 Dynamic magnetic-field imaging

Access to individual external frames from the QDM's lock-in camera permits imaging of dynamic magnetic signals. As an initial demonstration, a 1 s-long broadband voltage waveform that mimics a human cardiac signal is applied to the wire phantom, and the associated temporal variation of the magnetic-field pattern is imaged. The camera external frame rate F_s is set to about 0.5 kHz, corresponding to a temporal sampling interval of approximately 2 ms. This frame rate provides a balance between temporal sampling and the 500-frame camera-buffer limit during continuous acquisition. The signals applied to the phantom are gated between on and off across successive acquisitions of frame sets for differential measurement. The full experiment is repeated 400 times to allow signal averaging.

Figure 2.9 displays the imaged dynamic magnetic fields with 3×3 pixel binning (spatial resolution $\lesssim 10$ µm). The applied voltage trace is shown in Figure 2.9(a) and overlaid with the measured peak magnetic-field amplitude. To demonstrate the capability of high-fidelity waveform-reconstruction, the entire voltage trace is multiplied with a voltage-to-magnetic-field scaling factor obtained from the applied voltage and measured peak magnetic-field amplitude data in the static imaging experiment (Fig. 2.8). The good agreement demonstrates reconstruction of the temporal variation of a spatially structured magnetic-field pattern. This capability could support future comparisons between magnetic imaging and electrical measurements such as patch-clamp recordings. Magnetic images acquired at selected time points near the local extrema of the synthetic cardiac waveform are shown in Fig. 2.9(b). At time point (vi), the phantom magnetic-field pattern becomes indistinguishable from the background, yielding a spatial magnetic noise floor of approximately $\sigma_{\text{spatial}} \approx 3$ nT.

The present QDM can also operate with sub-millisecond external-frame sampling. To demonstrate this capability, we operate the Heliotis camera at an external frame rate of $F_s \approx 1.7$ kHz. A

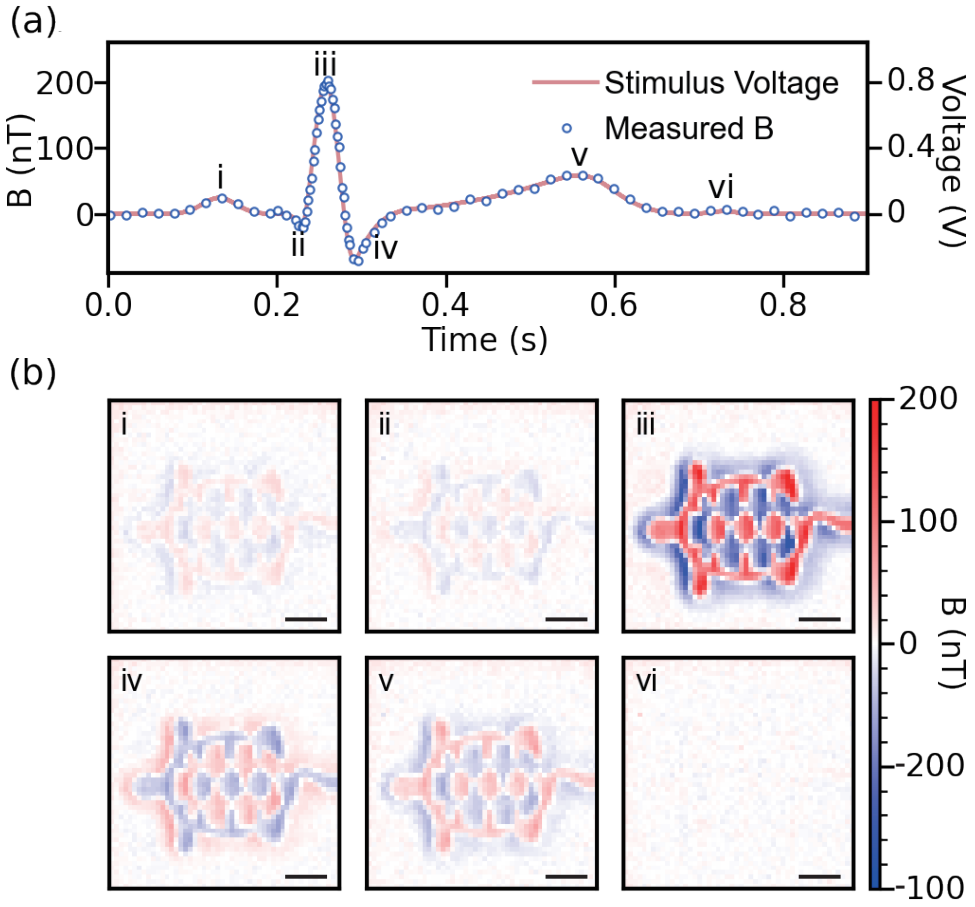


Figure 2.9: Time-resolved imaging of dynamic magnetic fields using the Ramsey-based QDM and a differential measurement protocol. A voltage trace mimicking a synthetic, broadband human cardiac signal is applied to the phantom. Sets of QDM magnetic images are collected with temporal resolution of about 2 ms and binned with 3×3 pixels. The experiments are repeated 400 times to allow signal averaging. (a) Temporal variation of the peak magnetic field from the signal-averaged images is overlaid with the applied stimulus. A voltage-to-magnetic-field scaling factor of 274.26 nT/V is applied to the entire voltage trace (see main text). A subset of peak magnetic-field amplitudes from the acquired images are displayed, indicating good agreement between the applied voltages and magnetic measurements. (b) Selected magnetic-field images at time points labeled in (a). Each image has been signal averaged. The spatial magnetic noise at time (vi) is about 3 nT.

single-tone AC signal voltage of fixed amplitude is then applied to the phantom; and 500 dynamic magnetic images are acquired (over $\sim$290 ms), with a time step between images $= 1/F_s \approx 0.6$ ms. The experiment is conducted for three different runs with signal frequencies of 200 Hz, 500 Hz, and 750 Hz; and each run is repeated once. A single pixel (no binning), selected to have significant amplitude variation over a signal cycle, is used to characterize the time-domain oscillations of the magnetic-field measurements (see Figure 2.10(a,b)); the time-domain data is then Fourier-transformed to identify the signal frequency components, as shown in Figure 2.10(c). The results are in good agreement with the signal frequencies.

The spectral peak amplitude is reduced for signal frequencies approaching the Nyquist frequency, $F_s/2$. This reduction is due to the effective low-pass filtering in generating a camera external frame, as each external frame contains the accumulated photoelectrons from multiple internal exposures for magnetic-signal measurements using the DQ 4-Ramsey sequence. In all the magnetic imaging data acquired in this sub-millisecond demonstration, the internal exposure (i.e., readout of NV fluorescence after a DQ Ramsey sequence) rate is 122 kHz; and a single external frame is produced by the camera after digitizing the cumulative photoeletrons from 72 internal exposures, giving a time step between external frames of about 0.6 ms irrespective of the AC signal frequency. A simulation incorporating the internal-exposure sampling rate, accumulation over the camera demodulation cycles, and subsequent Fourier transformation reproduces the measured frequency-dependent attenuation, as shown in Fig. 2.10(c). Reducing the number of accumulated internal exposures would permit faster external-frame sampling, ultimately limited by the maximum external camera frame rate of approximately 3.8 kHz.

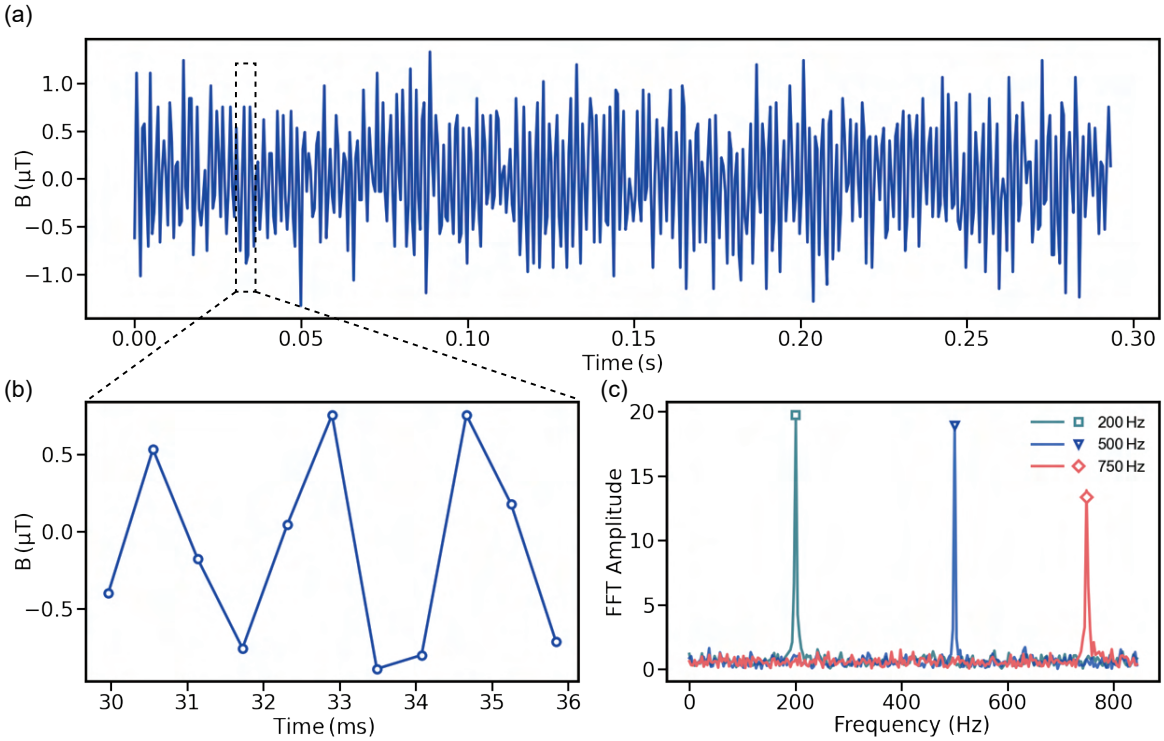


Figure 2.10: QDM magnetic imaging with sub-millisecond time resolution. Camera external frame rate is set to $F_s \approx 1.7\,\text{kHz}$ (i.e., time step between frames of about 0.6 ms). Magnetic signals are produced by applying single-tone AC signal voltages to the phantom. AC frequencies of 200 Hz, 500 Hz, and 750 Hz are used for different runs of the demonstration experiment. (a, b) Example time-domain magnetic-measurement data (500 Hz signal) from a single pixel in the QDM image, selected to have significant amplitude variation over a signal cycle. (c) Time-domain data from the single pixel is Fourier transformed and displayed, here for all three AC frequencies used in the demonstration experiments. The resulting spectrum is in good agreement with the applied signal frequencies; and with a simulation of the QDM measurement protocol that leads to a reduction in spectral peak amplitude near the Nyquist frequency ($F_s/2$).

2.8 Discussion and outlook

This chapter demonstrated a pulsed quantum diamond microscope for wide-field imaging of static and dynamic magnetic fields. By combining DQ 4-Ramsey magnetometry, spin-bath driving, lock-in camera acquisition, and differential measurement, the QDM achieved a median per-pixel magnetic sensitivity of $4.1(1)\ \mathrm{nT}/\sqrt{\mathrm{Hz}}$ without pixel binning. Spin-bath driving extended the median NV spin dephasing time to $T_2^* = 2.2(1)\ \mu\mathrm{s}$, approaching the estimated NV-NV interaction limit of the diamond sensor. The resulting spatial magnetic noise floor averaged down to the picotesla scale during long signal-acquisition times, enabling the imaging of static phantom magnetic-field patterns with micron-scale structure. In addition, the QDM reconstructed the temporal variation of a broadband synthetic cardiac waveform with approximately $2\ \mathrm{ms}$ sampling, while separate single-frequency measurements demonstrated sub-millisecond external-frame sampling capability.

The Ramsey-based QDM presented in this chapter is a key step toward applications in the physical and life sciences. Potential applications include magnetic imaging of iron-containing biomineralization products [65], high-throughput screening for biogenic magnetic particles [66–68], tip-growing plant cells, including pollen tubes [69–72], and spatially resolved measurements of current-generated magnetic fields from electrically active cells and tissues [73–77]. In the latter case, magnetic imaging could complement electrical measurements such as patch-clamp recording by non-invasively mapping a spatial distribution of current flow. These biological opportunities, together with expected signal magnitudes and experimental requirements, are considered further in Chapter 5.

The present QDM nevertheless remains subject to several technical limitations that affect the types of magnetic signals that can be measured efficiently. The following discussion identifies these limitations and potential routes for addressing them.

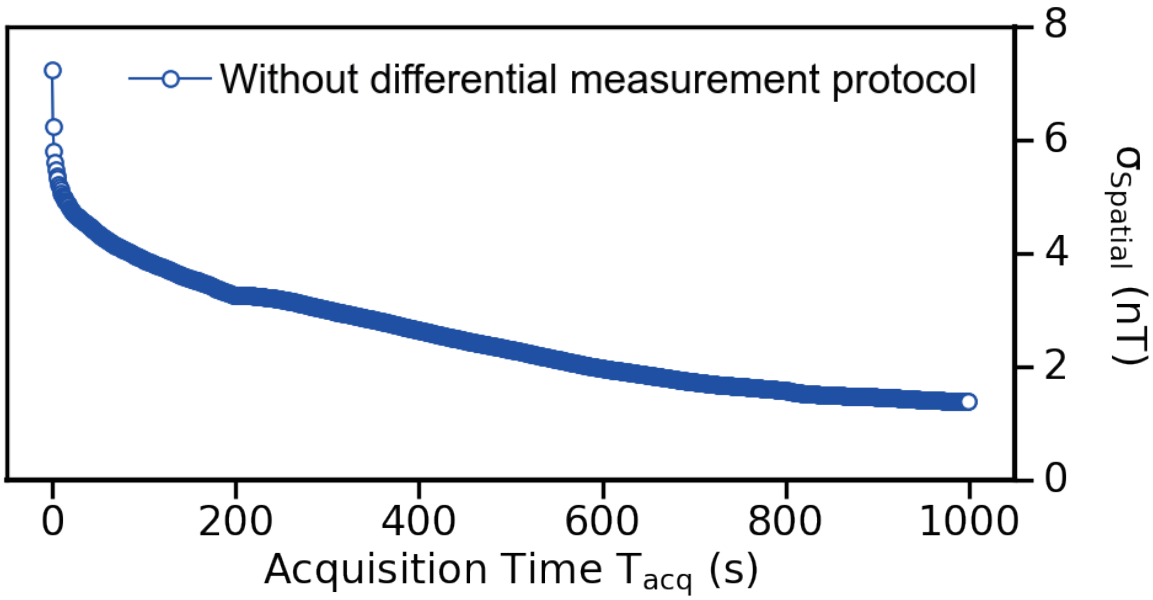


Figure 2.11: Spatial magnetic noise floor $\sigma_{spatial}$ as a function of acquisition time T_{acq}. The differential measurement protocol is not employed and thus the utility of time averaging is limited. No pixel binning is applied.

2.8.1 Technical limitations and paths forward

The first limitation is that the present imaging implementation relies on differential measurements to suppress slowly varying pixel-to-pixel offsets and bias-magnetic-field drift. The differential protocol is most effective for controllable magnetic sources, such as the wire phantom used here, for which the source field can be reversed or gated between successive acquisitions. However, this approach is less straightforward for magnetic objects whose signal cannot be easily switched on and off, such as magnetic particles attached directly to the diamond surface. Spatial magnetic-noise measurements acquired without differential processing are presented in Fig. 2.11. Under the present operating conditions, the spatial magnetic noise floor plateaus at approximately $2\ \mathrm{nT}$.

Bias magnetic field drift primarily produces smooth background gradients with spatial length scales of hundreds of micrometers to millimeters. Because the magnetic fields of interest in many

applications vary on micrometer length scales, these broad background gradients can be effectively suppressed using high-pass spatial filtering, provided that the filter scale is chosen to preserve the relevant signal structure. The more challenging contribution is the slowly varying change in the pixel-to-pixel baseline pattern. This effect is not fundamental to NV magnetometry, but instead reflects a technical challenge in maintaining stable camera-pixel baselines over long acquisition times. Improved stabilization of the optical illumination and bias magnetic field, together with improved correction of camera-pixel offsets, may reduce the need for source modulation for background rejection.

Alternatively, modulation of the measurement geometry may be useful for magnetic sources that cannot be electrically driven. For example, a magnetic object could be translated laterally or vertically relative to the diamond surface and field of view, thereby modulating the detected field through a controlled change in its position or sensor-sample separation. Such approaches could enable differential measurements of magnetic particles and other sources whose intrinsic magnetic state cannot be readily reversed or gated.

The second limitation is the difference between the active signal acquisition time and the total wall-clock time required to collect and process magnetic images. Each continuous acquisition of the lock-in camera is limited by the 500-frame camera buffer. In the present implementation, transfer of a full camera buffer to the control computer requires approximately $5\ \mathrm{s}$. For dynamic imaging, retaining and processing the individual camera frames introduces an additional implementation-specific overhead of approximately $17\ \mathrm{s}$. No magnetic sensing can be performed during these transfer and processing intervals. Consequently, the wall-clock time required for long averaging can substantially exceed the corresponding signal acquisition time.

This camera-related overhead is again not intrinsic to the NV sensing protocol and could be reduced through faster data interfaces, larger buffer, continuous streaming from the camera, real-time

averaging, or more efficient data-conversion and storage pipelines. Such improvements would be particularly valuable for applications that require averaging over many repetitions of a weak transient signal. For example, the projected magnetic measurements of neuronal action potential using the present QDM would require a large number of repeated excitations and, under the current workflow, could require wall-clock times of the order of 2 days. Reducing data-transfer and storage overhead could improve the experimental time to within several hours, even without changing the underlying per-pixel magnetic sensitivity. A more detailed discussion of these projected measurements is provided in Chapter 5.

The final limitation is the current magnetic sensitivity (volume-normalized sensitivity $25\,\mathrm{nT}\cdot\mathrm{Hz}^{-1/2}\cdot\mu\mathrm{m}^{3/2}$). One direct route toward improved sensitivity would be to increase the detected NV photoluminescence by increasing the optical excitation intensity. However, in the present diamond, higher optical intensity leads to NV charge-state conversion and reduced spin-state readout contrast. The experimentally optimal excitation intensity is $\sim 0.014\,\mathrm{mW/\mu m^2}$ substantially below the nominal NV saturation intensity [78] ($1-3\,\mathrm{mW/\mu m^2}$). This result indicates that improved charge stability in the NV-diamond material is an important requirement for further sensitivity improvement.

In addition, spin-bath driving extends the NV dephasing time close to the estimated limit imposed by NV-NV dipolar interactions. Further extension of T_2^* will require suppression of interactions within the NV ensemble itself. Homonuclear decoupling techniques [79] provide one possible route toward mitigating NV-NV interactions while preserving Sensitivity may also be improved by increasing the measurement duty cycle, namely the fraction of each pulsed measurement spent in the Ramsey free-evolution interval. Reducing the relative initialization and readout overhead could increase this duty cycle and may strengthen the practical advantage of alternative advanced readout schemes [1, 80]. These considerations partly motivate the material-engineering study presented in Chapter 4.

Chapter 3: NV-Based Characterization of Single-Ion Damage Tracks in Diamond

Chapter provenance and contributions. This chapter is adapted from Daniel G. Ang, Jiashen Tang, Maximilian Shen, Michael Titze, Gavishta Liyanage, Jordan Chapman, Chinmay Bharathulwar, Andrew T. Gilpin, Tanguy Terlier, Edward S. Bielejec, Vsevolod Ivanov, and Ronald L. Walsworth, "Multi-scale reconstruction of single-ion damage tracks in diamond via nitrogen-vacancy centers" [81].

I co-led this project with Daniel G. Ang. We conceived and designed the study, and performed the confocal measurements of NV formation and experimental yield analysis. I developed the SIIMPL-to-SPPARKS implantation-to-annealing simulation pipeline, performed the simulated NV-yield analysis, and implemented the correction to the simulated yield informed by first-principles DFT/NEB calculations carried out by Vsevolod Ivanov and Jordan Chapman. I also conducted the spin-echo measurements of the implantation-generated NV ensembles.

I worked with Chinmay Bharathulwar on the analysis of directional information retained in the simulated NV tracks. This includes asymmetry-based classification and simulation-based machine learning inference. Chinmay led the development of the machine learning methodology. I prepared the initial complete manuscript draft; and contributed to subsequent manuscript revision and editing with Daniel G. Ang and Ronald L. Walsworth.

3.1 Motivation and chapter overview

As discussed in Chapter 1, the formation and evolution of particle-induced damage tracks in solid states are relevant to a broad range of applications, including geoscience [82], particle detection

[83], radiation effects in materials [84], carbon-ion radiotherapy [85], and swift-heavy-ion studies [86–89]. More recently, damage-track studies have also attracted interest in emerging application including paleodetection of rare events [24, 39, 40, 90] and quantum defect engineering [41, 91–93]. A central motivation for the present work is the possible use of quantum diamond as a novel directional detector for rare-event, such as dark matter or neutrinos [24, 27, 30–32]. Compared to conventional nondirectional detectors, directionality of nuclear recoils may help distinguish a dark-matter signal from the solar-neutrino background in the WIMP neutrino-fog regime [25, 26].

As introduced in Chapter 1, realizing such a detector requires a multiscale readout pipeline [31, 32]. Candidate events must first be identified at the millimeter, single-diamond-chip scale, then localized to a smaller micrometer-scale region of the detector, and finally examined at the nanoscale to recover the morphology of the underlying damage track. We are primarily interested in sensing at the micrometer-to-nanometer scales, building on conventional phonon, charge, and/or photon readout [33–35] at the millimeter scale. At the micrometer scale, a candidate track could be localized through strain imaging using pre-existing NV ensembles or through optical detection of NV centers formed locally after annealing. At the nanoscale, the track morphology may be probed using approaches such as magnetic gradient-based imaging of individual NV centers [22], particularly newly formed NV centers, or high-resolution strain-imaging methods such as X-ray diffraction [36]. The combination of these possibilities constitutes a multiscale readout pipeline, but the appropriate diamond material choice and sensing modality at each stage remain open questions. Addressing them requires a quantitative understanding of damage-track formation, defect evolution during annealing, and the resulting optical, spin, and strain signatures.

This chapter investigates one specific route within this broader framework: ion-induced NV formation in nitrogen-rich high-pressure, high-temperature (HPHT) diamond. The high concentration of

substitutional nitrogen in HPHT diamond can provide nearby capture sites for implantation-created vacancies during annealing, potentially preserving a spatial record of the original damage track. The low density of pre-existing NV centers can allow newly formed NV photoluminescence to serve as an optical marker of an individual event. However, the viability of this approach requires systematic assessment. Previous ion-implantation studies have explored several regimes that do not directly address the present problem, including shallow low-energy irradiation in HPHT diamond [94–96], implantation into ultra-pure low-nitrogen CVD diamond [91, 95, 97], and hundreds-of-MeV ion irradiation in HPHT diamond [88, 98, 99]. Shallow implantation studies in HPHT diamond have reported low NV-formation efficiencies of the order of 0.01-0.1%, but these measurements are complicated by surface-related vacancy outdiffusion loss during annealing [100]. In contrast to ion implantation, rare-event interactions are not restricted to the diamond surface, and the resulting recoil tracks can occur in the bulk of the material. Consequently, it remain unclear what NV-formation yield could be expected for recoil damage inside bulk nitrogen-rich HPHT diamond, and whether that would produce a stable and larger-than-background NV signal.

Moreover, several additional questions need to be addressed even if implantation-induced NV formation provides a reliable optical marker. The annealing process needs to be understood quantitatively, because the resulting NV yield and spatial distribution influence how well the recoil direction can be reconstructed. The degree to which annealing smears out the directional information contained in the original damage track, and whether there are efficient methods to recover that information, need to be investigated. Finally, if future nanoscale reconstruction utilizes magnetic gradient-based quantum sensing of the resulting NV centers, these implantation-created NV ensembles must retain sufficient spin coherence to support such measurements.

In this chapter, I describe an experimental and computational study of individual $800\ \mathrm{keV}$

carbon-ion implantation events in nitrogen-rich HPHT diamond. First, spatially localized ensemble NV photoluminescence is shown to be capable of identifying individual implantation events after annealing and is used to measure the resulting NV yield. I then develop an implantation-and-annealing simulation framework to model track formation, defect evolution, and NV creation. The measured yield is used to benchmark the model, and a missing defect-reaction radius physics is identified as necessary for quantitative agreement with experiment. Using the resulting framework, I evaluate the directional information retained in the simulated post-annealing NV distributions using both a counting-based asymmetry metric and a simulation-based machine-learning approach. Finally, I measure the spin-echo coherence of the implantation-generated NV ensembles to assess their compatibility with future nanoscale reconstruction of individual damage tracks.

3.2 Experimental platform and methods

3.2.1 Ion implantation

A high-pressure, high-temperature (HPHT) type Ib single-crystal diamond sample (Element Six) with a specified nitrogen concentration of $[\mathrm{N_s^0}] \sim 200\,\mathrm{ppm}$ and $\{100\}$ surface orientation was used. The sample surface was polished to a roughness of $R_a < 1\,\mathrm{nm}$, and fiducial markers were fabricated to guide implantation and subsequent confocal optical characterization.

Carbon ion implantation at 800 keV was performed using the 6 MV Tandem accelerator at the MicroOne endstation at Sandia National Laboratories [101, 102]. This energy places the resulting damage track $\sim$600 nm deep into the surface, suppressing surface-related noise and defect outdiffusion [103] to produce a sufficient number of NV centers for optical detection while being near to the $<$500 keV energy range relevant to nuclear recoils induced by WIMP dark matter and neutrinos. Im-

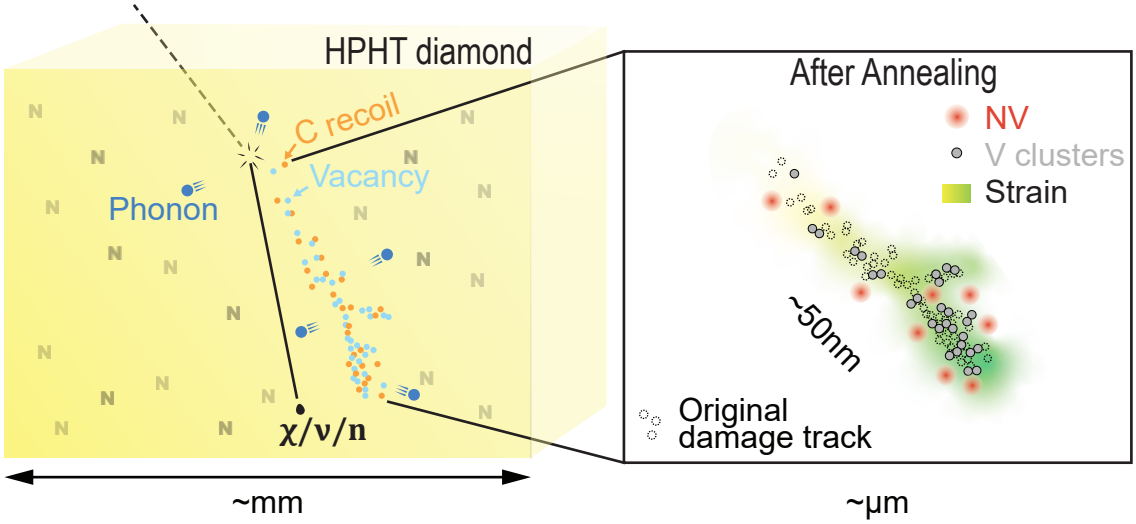


Figure 3.1: Schematic of damage track formation and multi-scale readout in nitrogen-rich HPHT diamond. An incoming particle induces a carbon nuclear recoil, generating a cascade of vacancies and interstitials. The recoil also produces phonons that enable identification of candidate interaction events at the millimeter scale (single diamond chip). During annealing, vacancies diffuse and either recombine with interstitials, form vacancy clusters, or combine with substitutional nitrogen to create optically-active NV centers near the original track. The resulting NV photoluminescence (PL) enables localization at the micrometer scale, while the spatial distribution of NV centers and residual defect-induced strain enables nanoscale reconstruction of track morphology and inference of the incoming particle direction.

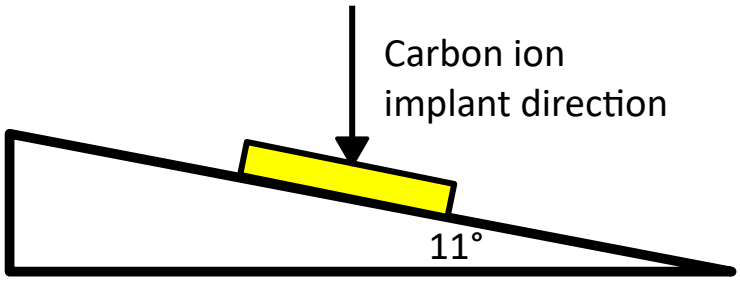


Figure 3.2: Schematic of ion implantation geometry. The diamond surface is tilted by $\sim$11° relative to the incident carbon ion beam to suppress channeling effects.

plantation was performed in square grids with the surface tilted by 11 degrees to suppress channeling effects [104] (Fig. 3.2). The ion beam had a full width at half maximum (FWHM) spot size of 2.62 μm (X) $\times$ 2.90 μm (Y), and implantation was carried out on a grid of sites with spacings of 6 μm or 10 μm. The beam current and dwell time were adjusted to achieve an average dose of about one ion per site, maximizing the probability of single-ion implantation events.

Following implantation, high temperature annealing was performed to induce the migration of vacancies into nearby nitrogen atoms. Annealing was performed in a vacuum tube furnace (MTI OTF-1200X) at 800°C and 10^{-6} torr for 2 hours, following standard recipes [97, 105]. The concentration of substitutional nitrogen in the HPHT diamond is measured to be $\sim$253 ppm (see Appendix B.1). This not only results in NV center formation near the orignal damage track, but also provides an electron-donor-rich environment that favors the negatively charged NV^- state [6]; in similarly nitrogen-dense diamond, NV^- has been reported to dominate over NV^0 by approximately two orders of magnitude [106]. NV center formation was confirmed by laser-induced PL using a confocal setup (See confocal setup in Section 3.2.2 and optical characterization in Section 3.3). To verify that annealing had proceeded to completion, a second identical cycle was performed; the absence of any detectable change in PL confirmed that the initial conditions were sufficient.

3.2.2 Confocal microscopy and NV measurement

A home-built confocal microscope was used to characterize the diamond samples. NV centers were optically excited using a 532 nm laser (Coherent Compass 315M) with a typical power of $\sim$1 mW at the sample. The excitation beam was circularly polarized to uniformly excite NV centers of different crystallographic orientations. NV PL in the 647-800 nm range was collected onto a single-photon counting module (Excelitas SPCM-AQRH-14).

The collimated excitation beam ($d \approx 3.5\,\mathrm{mm}$) was focused through a 100×, 0.8 NA objective (Nikon CFI60 TU Plan Epi ELWD), and collected PL was spatially filtered using a $100\,\mu\mathrm{m}$ pinhole to reject out-of-focus light. The microscope operated in a sample-scanning configuration using a nanopositioning stage (Mad City Labs Nano-3D200). The point spread function (PSF) was characterized by three-dimensional photoluminescence scans of a single NV center, yielding Gaussian widths of $\sigma_x = 189 \pm 12\,\mathrm{nm}$, $\sigma_y = 180 \pm 11\,\mathrm{nm}$, and $\sigma_z = 546 \pm 26\,\mathrm{nm}$.

For spin measurements, microwave (MW) control was provided by a continuous-wave source (SRS SG384), gated using a fast electronic switch driven by a pulse generator (Swabian Pulse Streamer 8/2). The same pulse generator synchronized MW pulses with optical excitation via an acousto-optic modulator (Brimrose TEM-85-10) to implement pulsed sequences such as Hahn echo. NV photoluminescence counts were recorded using a data acquisition system (NI USB-6363) for spin-state analysis.

3.2.3 Implantation and annealing simulation for NV formation

Ion implantation was modeled using the SIIMPL (Simulation of ion IMPLantation) binary collision approximation code [107, 108]. SIIMPL was chosen instead of the more commonly used SRIM/TRIM code [109] because it incorporates the crystalline structure of diamond and associated

ion-channeling effects, whereas SRIM treats the target as amorphous. The simulation was run at 800 keV carbon ion energy and 11° polar angle, matching the Sandia implantation parameters. The lattice displacement threshold energy and thermal vibration amplitude were set to the standard values of 43 eV [110] and 0.044 Å [111] respectively. For each implanted ion, the simulation produced the spatial coordinates of the resulting vacancies and interstitials.

Subsequent annealing was modeled using a kinetic Monte Carlo (KMC) simulation implemented in SPPARKS [112], which tracks defect diffusion and reactions on the diamond lattice. Transition rates were described by the Arrhenius expression

$$\gamma = \nu_0 \exp\left(-\frac{E_a}{k_\mathrm{B} T}\right), \tag{3.1}$$

using an annealing temperature of $800\ ^\circ\mathrm{C}$. Initial defect configurations of vacancies and interstitials were obtained from SIIMPL implantation simulations. To improve computational efficiency, each damage track was partitioned into spatially separated defect clusters that could be treated independently. The minimum separation between clusters was $\gtrsim 154$ Å, for which the probability of inter-cluster interaction was estimated to be $\lesssim 1\%$ (see Appendix B.2). For each cluster, the initial defect configuration was embedded within a diamond lattice padded by 100 unit cells in each direction, such that vacancy loss through diffusion to the simulation boundary was negligible. Nitrogen atoms were randomly distributed throughout the lattice at a concentration of 253 ppm, consistent with experimental measurements. Results from individual clusters were combined to obtain statistics for full damage tracks. Typical simulations contained on the order of $10^8 - 10^9$ atoms per cluster.

Defect diffusion and reactions were modeled as follows. Vacancies were allowed to hop with an activation energy of $E_a = 2.59\,\mathrm{eV}$ [6, 113], while carbon interstitials were assigned activation

energy of $E_a = 1.5\,\text{eV}$ [114]. All other defects, including substitutional nitrogen atoms, were assumed to be immobile. Reactions included vacancy recombination with interstitials (restoring the lattice), vacancy aggregation into divacancies and higher-order clusters ($\text{V} + \text{V} \rightarrow \text{VV}$, $\text{V} + \text{VV} \rightarrow \text{VVV}$), and NV center formation via vacancy capture by nitrogen ($\text{V} + \text{N} \rightarrow \text{NV}$). These reactions were initially assumed to occur upon first-nearest-neighbor contact, corresponding to effectively barrierless interactions.

3.2.4 Head-tail directionality and machine learning

The full directionality of a recoil damage track includes both the orientation of the recoil axis and the head-tail information. In this study, we evaluate only head-tail directionality, while reserving full vectorial reconstruction for future work.

Head-tail directionality was first evaluated using an asymmetry metric A, defined as the ratio of defect counts between the tail region, corresponding to the upper one-third of the track, and the head region, corresponding to the lower one-third of the track. Exchanging the head and tail labels maps $A \rightarrow 1/A$. For the simulated benchmark tracks considered here, we define $+z$ to be along the recoil direction, such that the tail is more positive in z than the head. Therefore, given a threshold $x \geq 1$, tracks with $A > x$ were classified as correctly identified, tracks with $A < 1/x$ were classified as incorrectly identified, and tracks with $1/x \leq A \leq x$ were treated as indeterminate.

To extract directional information beyond counting-based metrics, we employ simulation-based inference (SBI) [115], a Bayesian framework for inverse problems in which the likelihood $P(X|\theta)$ is not available in closed form. In the present case, the inverse problem maps a defect distribution X to the initial recoil parameters θ, including the implantation energy and head-tail parity. Due to

the stochastic nature of the SIIMPL collision cascade and NV formation during annealing, direct evaluation of the likelihood is intractable. Instead, the model learns an approximation to the posterior $P(\theta|X)$.

We implement a Mixed Neural Posterior Estimation (MNPE) architecture, which enables joint inference over continuous parameters, such as implantation energy, and discrete parameters, such as head-tail parity. The model combines a Neural Spline Flow [116] for the continuous parameter with a categorical softmax head for parity classification, implemented using the `sbi` toolkit [115]. Energy parameters are included when learning the general inversion from simulated damage tracks across different ion implantation energies. However, for prediction of the head-tail parity of the annealed $800\ \mathrm{keV}$ NV tracks considered here, the implantation energy is effectively fixed. Parity classification is binary, corresponding to motion along either the $+z$ or $-z$ direction in the global coordinate system, since the injected carbon ion is aligned along the z axis in the SIIMPL simulations.

Rather than using raw defect coordinates, the model learns from a set of summary features derived from defect spatial distributions. Each track is encoded as a 13-dimensional feature vector, listed in Table 3.1, designed to capture directional information contained in the track morphology. The model is trained on SIIMPL-simulated vacancy tracks with ground-truth parity labels using the Adam optimizer, with a learning rate of 5×10^{-4} and a batch size of 128. A $90/10$ train-validation split is used, with early stopping applied when the validation loss does not improve for 20 epochs. For inference, the trained model is directly applied to post-annealing NV distributions using the same feature representation. For each track, 200 samples are drawn from $P(\theta|X)$ to estimate the probability P of the $+z$ direction, with $1 - P$ corresponding to the $-z$ direction.

Under this convention, $P \to 1$ corresponds to motion along $+z$, while $P \to 0$ corresponds to motion along $-z$. For test tracks with ground-truth recoil direction along $+z$,a threshold $y \geq 0.5$

Table 3.1: The 13 summary features computed for each track and passed to the MNPE model. Here N is the number of defects and z_i is the position of the i-th defect along the track axis, measured relative to the track's centroid (each track is re-centered to $z = 0$). $N_\pm$ and $z_\pm$ denote the counts and positions of defects with $z_i > 0$ and $z_i < 0$, respectively.

#	Feature	Definition	Interpretation
1	Mean depth	$\frac{1}{N}\sum_i \lvert z_i \rvert$	The average extent of the track. Longer mean depth usually means a higher-energy ion.
2	Max depth	$\max_i \lvert z_i \rvert$	Distance of the farthest defect from track center.
3	Vacancy count	N	The total number of defects.
4 – 9	Depth-profile fractions ($k = 1, \ldots, 6$)	Fraction of defects in each of six log-spaced bins along the normalized track depth	Defect distribution binned along z axis.
10	Centered asymmetry	$(N_+ - N_-)/N$	Whether more defects sit on one side of the centroid than the other.
11	Asymmetry ratio defined in Ref. [24]	$N_{\text{last third}}/N_{\text{first third}}$	The asymmetry ratio from Ref. [24].
12	Longitudinal skewness	Skewness of $\{z_i\}$ along the track axis	Whether the defects lean more toward the head or the tail of the track.
13	Log-variance difference	$\log(1 + \mathrm{Var}(z_+)) - \log(1 + \mathrm{Var}(z_-))$	Compares how spread out the defects are on each side of the centroid.

is applied: tracks with $P > y$ were classified as correct, tracks with $P < 1 - y$ were classified as incorrect, and tracks with $1 - y \leq P \leq y$ were treated as indeterminate. For both the asymmetry and ML analyses, the threshold was chosen to evaluate the classification efficiency at a fixed false positive rate.

3.3 Optical detection of individual ion-implantation events

A coarse PL scan was first performed using a confocal setup, revealing arrays of bright PL spots with spacing matching the implantation pitch (6 or 10 μm). Small deviations from the ideal grid positions reflect the uncertainty in the finite ion-beam size (see Section 3.2.1), as shown in Fig. 3.3(a). These results demonstrate that individual carbon ion impacts generate spatially localized ensembles of NV centers that can be optically resolved.

To quantify the NV yield per ion, candidate implantation sites were identified using an automated grid-matching procedure, which also rejects backgrounds from surface polishing damage (see Appendix B.3.1). High-resolution three-dimensional confocal scans were then acquired around each identified site. For sites that received multiple ions, we observed spatially separated PL spots (Fig. 3.3(b)). The occurrence of such multi-spot events is consistent with a zero-truncated Poisson distribution,supporting the assignment of each separated PL spot within a multi-ion grid location to an individual ion impact (see Appendix B.3.2). The fitted Poisson rate of $\lambda = 0.99$ indicates that approximately one ion was delivered per nominal grid point, consistent with the experimental ion beam settings.

The number of NV centers associated with each implanted ion, N_{NV}, was then determined by comparing the integrated PL intensity of each spot to that of calibrated single NV centers under identical imaging conditions. The resulting distribution of N_{NV} across 222 implantation sites is shown in Fig. 3.3(c), with a mean value of 17.5 ± 0.4, where the uncertainty is the standard error of the mean. The site-to-site standard deviation is 5.7.

3.4 Experimentally benchmarked model of implantation, annealing, and NV formation

3.4.1 Implantation and annealing simulation

As outlined in Section 3.2.3, the simulation pipeline combines SIIMPL, which produces the spatial distributions of vacancies and interstitials that constitute the initial damage track, with kinetic Monte Carlo (KMC) annealing simulations performed in SPPARKS [112]. The KMC simulations use the SIIMPL-generated defect configurations as input and evolve the defect distribution through

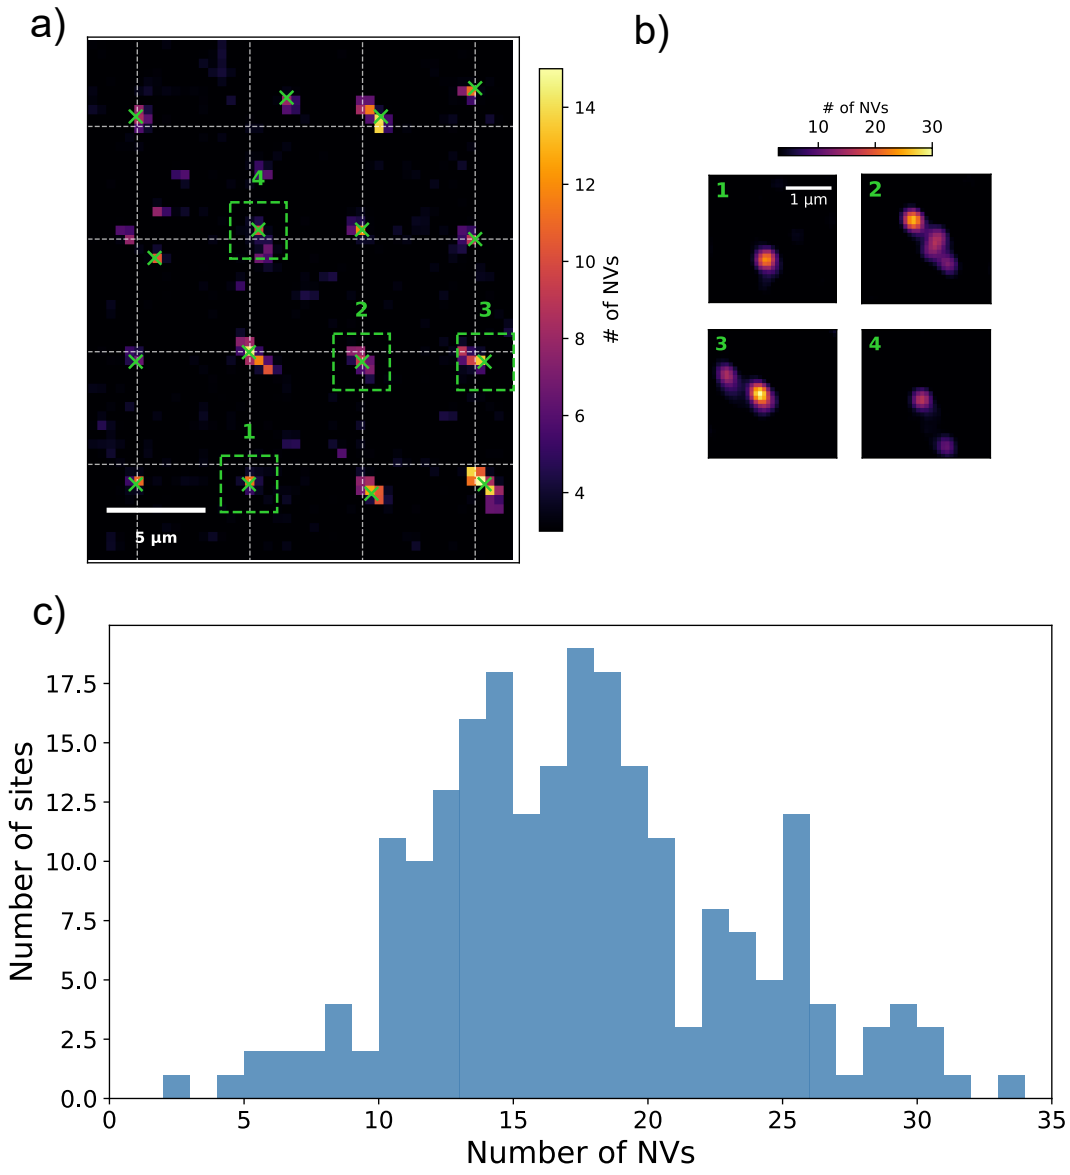


Figure 3.3: Optical characterization of NV formation. a) Wide-field confocal NV PL map acquired at coarse spatial resolution (∼500 nm). Bright PL spots correspond to NV ensembles formed at individual implantation sites, arranged in a regular grid consistent with the implantation pitch. Candidate sites selected for further high-resolution 3D confocal scans are marked with green crosses. The NV number is estimated from the PL intensity relative to calibrated single NV. Background NV counts in the HPHT sample are $\lesssim 2$. b) Representative high-resolution confocal scans of four implantation sites (fixed Z slice). Spatially separated PL spots are observed and are attributed to multi-ion implantation events, consistent with statistical analysis. c) Distribution of the number of NV centers per implanted ion (N_{NV}), extracted from integrated PL intensity (see main text). Data are collected from 222 sites, yielding a mean NV count of (17.5 ± 0.38), where the uncertainty denotes one standard error of the mean; the standard deviation is 5.7.

diffusion and reaction events during annealing. This framework predicts both the number and spatial distribution of NV centers following annealing (Fig. 3.4(a)). The KMC simulations do not model the charge state of the NV centers after formation. Because the experiment primarily measures NV^- PL, direct comparison with the simulation assumes that the NV^0 population does not substantially affect the inferred NV count. This assumption is expected to be reasonable for the nitrogen-rich HPHT diamond used here [6, 106].

The distribution of NV counts per simulated implantation event is shown in Fig. 3.4(b). Across 112 simulated damage tracks, approximately 328 vacancies are produced per track on average. After annealing, a mean of N_{NV} = 23.9 NV centers is predicted (standard error of the mean: 1.1; standard deviation: 11.5), corresponding to a conversion yield of $\sim$7%. The predicted mean yield is higher than the experimentally measured value of $N_{\mathrm{NV}} = 17.5$ by approximately 37%, indicating that the initial implantation-and-annealing model overestimates NV formation.

3.4.2 Testing sensitivity of NV-yield uncertainty to displacement energy and nitrogen concentration

We first consider two sources of uncertainty that may affect the predicted NV yield: the total vacancy production in the implantation simulation and the measured nitrogen concentration. As discussed below, neither is sufficient to account for the discrepancy between the modeled and experimentally measured NV yields.

Displacement energy. The threshold displacement energy E_d used in the SIIMPL simulation controls the number of vacancies generated in the damage track. In diamond, E_d is known to be directionally dependent [110], with reported values ranging from $\sim$37.5 eV for [100] to $\sim$47.6 eV for

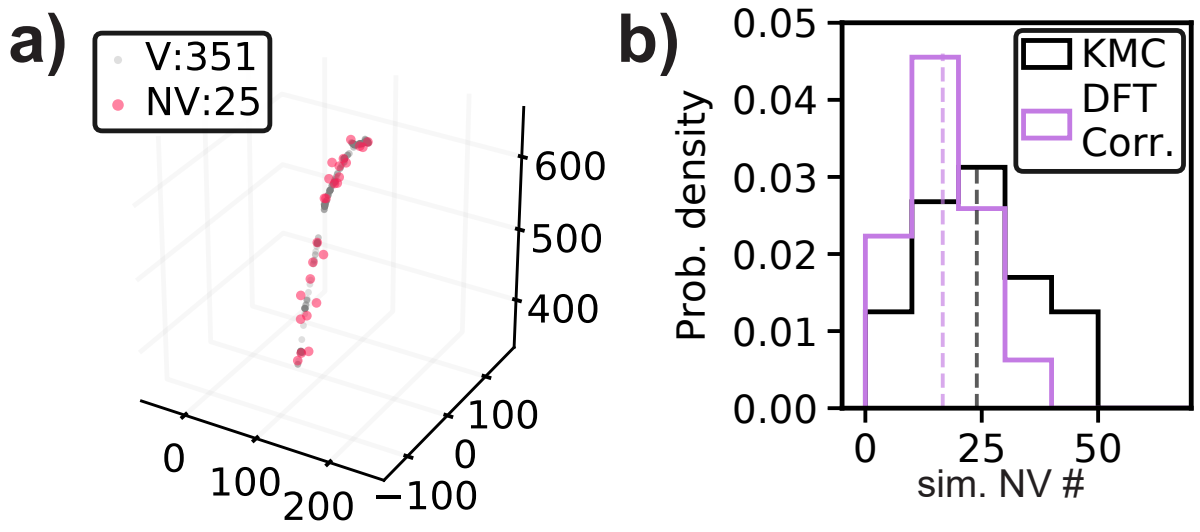


Figure 3.4: Modeling NV formation. a) Example damage track before and after annealing. Prior to annealing, vacancies (V) generated by an 800 keV carbon ion implantation are shown. During annealing, vacancies are consumed through recombination with interstitials, aggregation into vacancy clusters, or capture by substitutional nitrogen (not shown) to form NV centers. The resulting NV distribution after annealing is shown. b) Histogram of the number of NV centers per simulated implantation event. A mean NV number of 23.9 ± 1.1 (standard error of mean; standard deviation: 11.5) is obtained from 112 KMC-annealed tracks. A correction accounting for second-nearest-neighbor (2NN) defect interactions (DFT Corr.) reduces the predicted yield to 16.6 ± 0.7 NVs per event, in close agreement with experiment (Fig. 3.3(c)).

[110]. Similar to SRIM, SIIMPL employs a single scalar value of E_d, effectively averaging over crystallographic directions during the recoil cascade. In our analysis, we use $E_d = 43\,\text{eV}$. To evaluate the sensitivity of vacancy production to this parameter, we vary E_d over a broad range. As shown in Fig. 3.5, the mean vacancy number decreases monotonically with increasing E_d. Within the experimentally relevant range of $\sim$37–48 eV, this corresponds to a $\sim$10–20% uncertainty in the number of resulting vacancies. Assuming a fixed vacancy-to-NV conversion efficiency, this rescales the KMC-predicted NV yield from $\sim$23 to $\sim$21 at the upper end of E_d, and therefore does not fully account for the experimentally measured value of $\sim$17 NV centers.

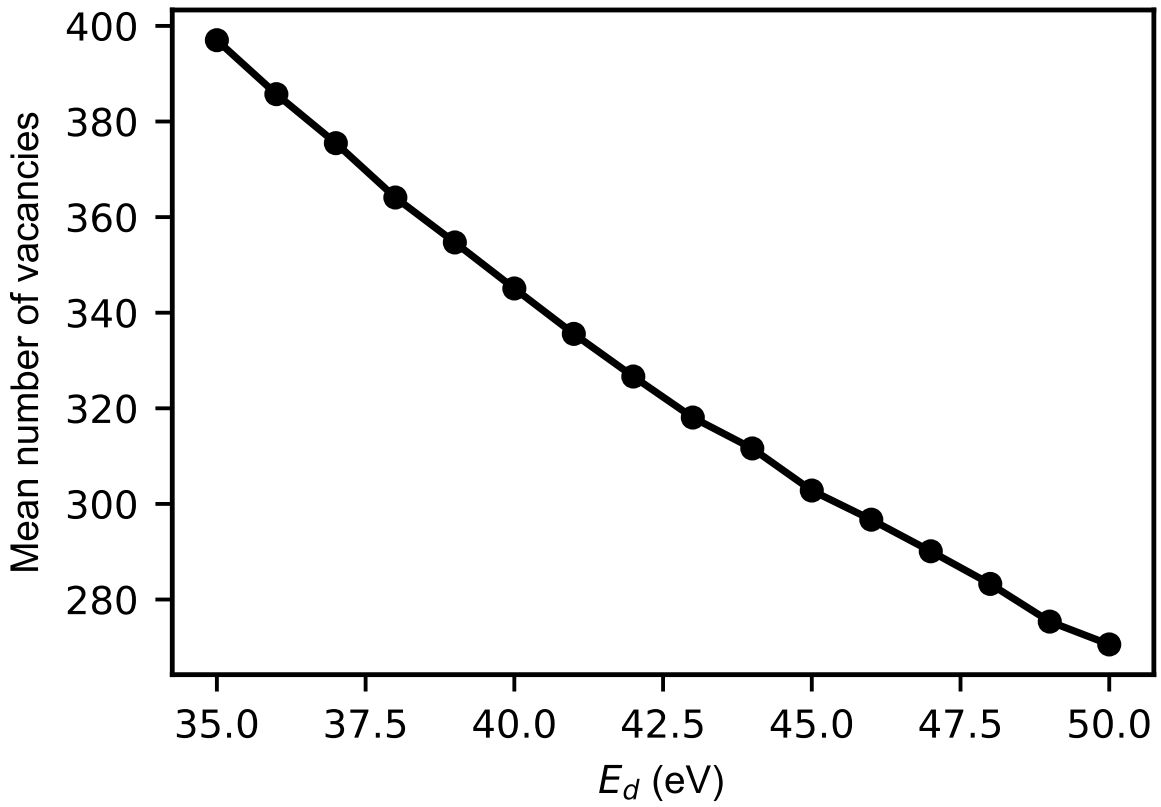


Figure 3.5: Sensitivity of vacancy production to displacement energy E_d. Mean number of vacancies generated per 800 keV carbon recoil as a function of the threshold displacement energy used in SIIMPL simulations. Each data point represents the average over 5000 simulated damage tracks.

Nitrogen concentration. The measured nitrogen concentration carries a relatively large uncertainty, ranging from ∼180–320 ppm as described above. To assess its impact on NV yield, we perform annealing simulations across a range of nitrogen concentrations, as shown in Fig. 3.6(a). For each concentration, a single damage track is simulated, and the nitrogen lattice configuration is randomized five times to estimate the mean NV yield. We find that the mean NV yield changes only weakly over the range from 150 ppm to 5000 ppm. Increasing nitrogen concentration primarily leads to NV formation occurring closer to the original damage track, as illustrated in Fig. 3.6b. This weak dependence can be understood from the relevant spatial scales. Vacancy and interstitial defects are typically separated by less than ∼0.5 nm, whereas the average nitrogen-nitrogen spacing is approximately 3 nm at 200 ppm and scales with the inverse cubic root of nitrogen concentration. As a result, vacancies predominantly react with nearby interstitials and other vacancies before diffusing sufficiently far from the damage track to encounter substitutional nitrogen. The overall NV yield is therefore much less sensitive to nitrogen concentration than might be expected from vacancy capture alone.

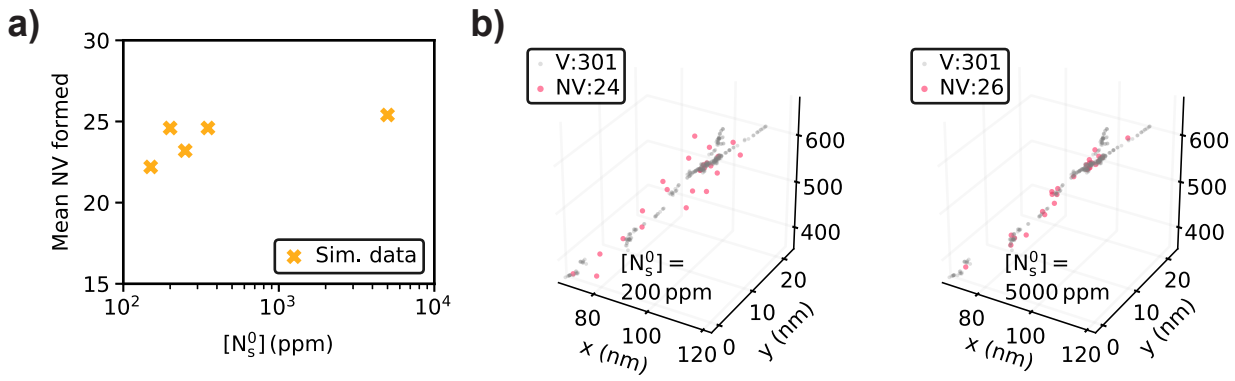


Figure 3.6: Impact of nitrogen concentration on annealing simulations. (a) Mean NV yield as a function of nitrogen concentration $[N_s^0]$, showing weak dependence over a broad range (150–5000 ppm). (b) Representative spatial distribution of vacancies (gray) and resulting NV centers (red) at $[N_s^0] = 200\,\mathrm{ppm}$ and 5000 ppm. Higher nitrogen concentration leads to NV formation closer to the original damage track, while the overall NV yield remains largely unchanged.

3.4.3 Short-range defect reactions beyond the nearest-neighbor approximation

we next investigated whether certain assumptions in the KMC simulation could account for the difference between the predicted and measured NV yields. In the present KMC simulation, vacancy-related interactions, including vacancy-interstitial recombination, divacancy formation, and NV creation, are restricted to first-nearest-neighbor (1NN) separation. In this treatment, defects initially separated by second-nearest-neighbor (2NN) distances must first undergo isolated defect hopping to reach 1NN separation before reacting. To assess whether this restriction is physically justified, nudged elastic band (NEB) calculations within density functional theory were performed to more precisely determine the barriers for relevant defect reactions [117] (see Appendix B.5). These calculations indicate that vacancy-related defect interactions can occur directly from 2NN configurations with either negligible or strongly reduced activation barriers. Thus, vacancies with 2NN defect partners are expected to react effectively immediately during annealing, rather than first undergoing isolated hopping to reach 1NN contact. This missing 2NN interaction range can affect the predicted NV yield in two ways. First, it reduces the initial population of free vacancies available for diffusion and later NV formation by removing vacancies that would react immediately with nearby interstitial or vacancy partners. Second, it may modify the subsequent free-vacancy-to-NV conversion efficiency during annealing. Although vacancies with a nitrogen atom at 2NN separation could also form NV centers directly, the number of such configurations is negligible (see Appendix B.5).

A full treatment of these effects would require implementing extended interaction ranges directly within the KMC framework. Doing so natively in SPPARKS is computationally expensive and non-trivial, as it requires replacing the lattice-local nearest-neighbor event catalog with configuration-dependent reaction rates. We therefore do not attempt a full 2NN-resolved KMC simulation here.

Instead, as a first-order correction estimate, we assume that the dominant effect is the reduction of the initial free-vacancy population, and correct this population by excluding vacancies that have either 1NN or 2NN defect partners in the initial configuration. Rescaling this reduced population using the free-vacancy-to-NV conversion factor obtained from the original KMC simulation reduces the predicted mean NV yield to 16.6 ± 0.7 per simulated damage track, bringing the model into close agreement with experiment (Fig. 3.4(b)). We note that this estimate is approximate. Nonetheless, the analysis identifies short-range defect reactions beyond the 1NN assumption as the dominant missing physics.

3.5 Directional information retained after annealing

3.5.1 Counting-based head-tail asymmetry

Using the simulated spatial distributions of defects before and after annealing, we investigated how directional information is modified during annealing. Unlike vacancies, NV centers provide a viable optical probe of the damage track because they are bright and can be detected down to the single-defect level. We evaluated head-tail directional information using a counting-based asymmetry metric following Ref. [24]. The asymmetry A, defined as the ratio of defect counts between the tail and head regions of a track, encodes the initial recoil direction because the tail region typically exhibits higher energy deposition, analogous to Bragg-peak-like behavior.

To benchmark classification performance, we choose an asymmetry threshold $x \geq 1$. Tracks with $A > x$ are classified in the correct recoil direction, whereas tracks with $0 < A < 1/x$ are classified in the opposite direction and counted as false positives. Tracks with $1/x \leq A \leq x$ are treated as indeterminate. We define the classification efficiency ϵ as the fraction of all tracks that are correctly

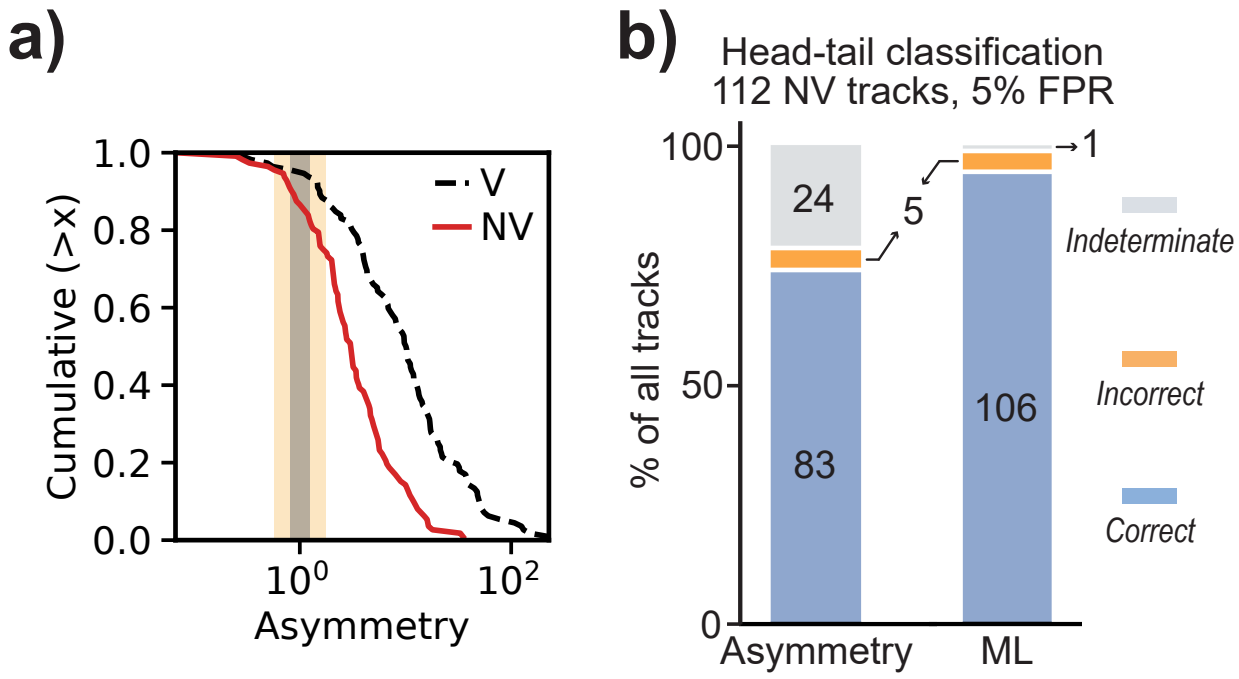


Figure 3.7: Head-tail directional information from simulated damage tracks. a) Cumulative distribution of the asymmetry metric for vacancy (before annealing) and NV (after annealing) tracks. The asymmetry is defined as the ratio of defect counts in the tail and head regions of the track. For a threshold $x \geq 1$, the shaded region between $1/x$ and x indicates the indeterminate range. Tracks to the left of the shaded region are incorrectly classified and counted as false positives, while tracks to the right are correctly classified and counted toward the classification efficiency, following Ref. [24]. For both vacancy and NV tracks, the thresholds are chosen to give a $\lesssim$5% false positive rate. The resulting classification efficiency decreases from 93% for vacancy tracks to 74% for NV tracks, due to annealing-induced diffusion and reduced NV yield. b) Comparison of head-tail classification performance using NV tracks with the counting-based asymmetry metric and a machine learning (ML) approach. The ML model outputs a posterior probability over track direction, allowing tracks to be similarly classified as correct, incorrect, or indeterminate (see main text). At the same $\lesssim$5% false positive rate, the ML method improves classification efficiency from 74% (asymmetry metric) to 95%, demonstrating enhanced recovery of directional information from NV spatial distributions.

classified, and the false-positive rate (FPR) as the fraction of all tracks that are incorrectly classified. An effective directional-detection scheme should therefore achieve high classification efficiency at low FPR. As shown in Fig. 3.7(a), choosing a threshold that gives a $\lesssim 5\%$ FPR for simulated $800\ \mathrm{keV}$ tracks yields a classification efficiency of 93% for vacancy tracks. After annealing, diffusion-induced broadening of the spatial distribution and the reduced number of NV centers lower the classification efficiency computed from NV tracks to 74%.

3.5.2 Simulation-based inference of track direction

We next tested whether a machine learning (ML) approach could recover directional information from simulated NV tracks beyond the counting-based asymmetry metric. The model is based on simulation-based inference (SBI) [115], in which a neural network learns an approximate posterior over recoil parameters given an observed defect distribution. This approach can exploit spatial features beyond simple count asymmetry. For example, the tail region of a track typically exhibits higher defect density, leading to reduced spacing between NV centers that is not captured by counting-based metrics. To enable efficient training and improve generalization, each track was encoded as a low-dimensional vector of summary features describing the longitudinal defect distribution, including binned depth profiles, asymmetry measures, and the spread of defects on either side of the track centroid (see Section 3.2.4). The neural network was trained on simulated vacancy tracks, which can be efficiently generated using SIIMPL, and then applied to NV distributions obtained after KMC annealing.

The model outputs a probability (P) for the recoil direction. Analogous to the asymmetry-threshold analysis, we choose a probability threshold $y \geq 0.5$, where tracks with $P > y$ are classified as correct, tracks with $P < 1 - y$ are counted as false positives, and tracks with $1 - y \leq P \leq y$

are treated as indeterminate. At a comparable $\lesssim$5% FPR, the ML approach yields a classification efficiency of 95%, a $\sim$20 percentage-point increase over the counting-based asymmetry metric (Fig. 3.7(b)). A full cumulative comparison between the count-based asymmetry metric and the machine learning approach is shown in Fig. 3.8. Additional improvement may be possible with models trained directly on post-annealing NV distributions. However, this would require substantially larger KMC annealing datasets than are presently available. Overall, these results show that, within the experimentally benchmarked simulation framework, data-driven inference can recover directional information beyond simple counting-based metrics. Extending this approach to full vectorial direction reconstruction remains an important next step.

3.6 Spin coherence of implantation-created NV centers

Having shown that the simulated post-annealing NV distributions retain directional information, we next assessed whether the implantation-created NV ensembles have sufficient spin coherence for prospective nanoscale readout. Magnetic-gradient-based imaging has emerged as a promising super-resolution approach, enabling few-nanometer spatial resolution of multiple NV centers [23]. The achievable resolution in such techniques depends on both the applied magnetic-field gradient and the spin coherence properties of the NV centers. We therefore measured the ensemble NV spin coherence time across implantation tracks using a Hahn-echo pulse sequence Figure 3.9(a-c). From these measurements, we extracted the coherence decay time T_2 and the decay exponent p, which provides information about the local spin environment [118]. As shown in Fig. 3.9(d), we observe an average $T_2 \approx 1.47\,\mu$s, consistent with a nitrogen spin-bath-limited regime at the measured nitrogen impurity concentration ($[\mathrm{N_s^0}] \approx 253\,\mathrm{ppm}$) [57]. Following Ref. [118], the extracted decay exponent of $p \approx 1.55$

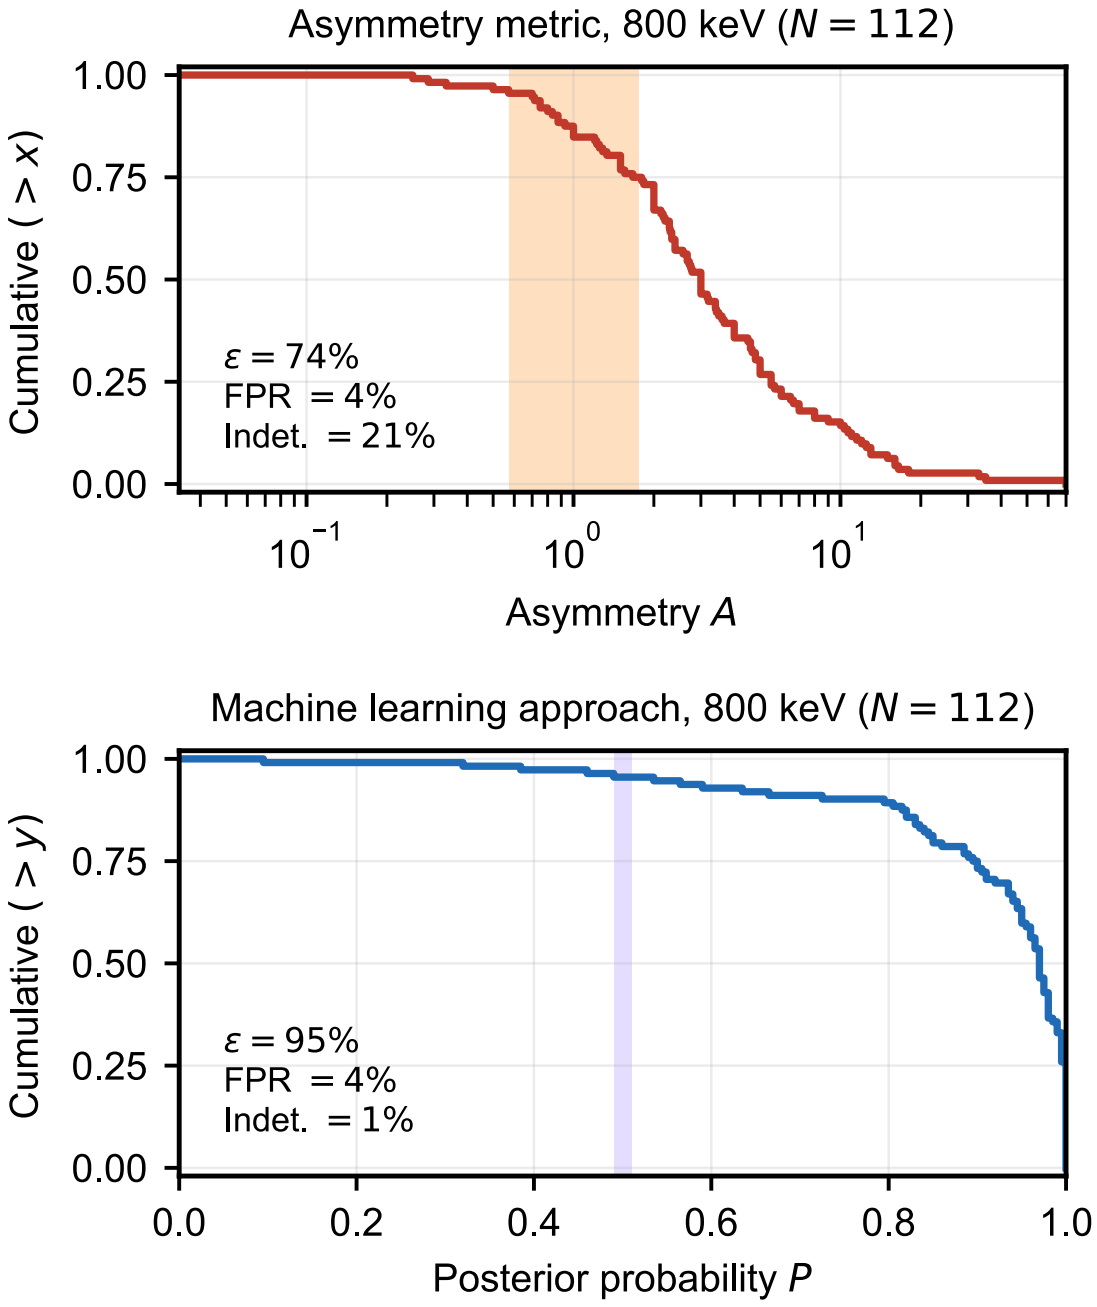


Figure 3.8: Cumulative plots of NV track classification using the counting-based asymmetry metric and machine learning approach. Top row: cumulative distributions of the asymmetry metric for 800 keV NV tracks. Bottom row: cumulative distributions of the machine learning posterior probability for the same tracks. Shaded regions indicate indeterminate ranges chosen to give a $\lesssim$5% false positive rate. The machine learning approach increases the classification efficiency relative to the asymmetry metric.

is also consistent with ensemble-NV decoherence behavior. Additional NV spin-echo measurements from 14 implantation sites are shown in Appendix B.6.

These measured coherence properties allow an estimate of the spatial resolution that may be achievable with gradient-based imaging of the implantation-created NV distributions. Using a modest magnetic-field gradient of $0.002\ \mathrm{G/nm}$, as demonstrated in Ref. [23], the measured coherence time corresponds to an estimated spatial resolution of $\sim 30\ \mathrm{nm}$. With state-of-the-art gradient strengths exceeding $0.01\ \mathrm{G/nm}$ [22], this estimate can be improved to below $6\ \mathrm{nm}$. For the $800\ \mathrm{keV}$ implantation energy considered here, the average simulated NV–NV separation is $\sim 12\ \mathrm{nm}$. These results do not yet demonstrate nanoscale imaging of an individual damage track. However, they show that the implantation-created NV ensembles retain spin coherence compatible with prospective gradient-based readout of the post-annealing track morphology in HPHT diamond.

3.7 Summary and outlook

This chapter establishes an experimental and computational framework for studying individual ion-induced damage tracks in nitrogen-rich HPHT diamond through implantation-created NV centers. Spatially localized NV photoluminescence was used to identify individual $800\ \mathrm{keV}$ carbon-ion implantation events after annealing and to directly measure their NV yield. The measured mean yield of NV created per implanted ion is $N_{\mathrm{NV}} = 17.5 \pm 0.4$, which further provides an experimental benchmark for the implantation-and-annealing model. The initial SIIMPL-KMC framework overestimated this yield, predicting $N_{\mathrm{NV}} = 23.9 \pm 1.1$. The discrepancy could not be simply accounted for by uncertainty in either the threshold displacement energy or nitrogen concentration, instead, DFT-informed analysis identified extended defect reaction radius beyond the first-nearest-neighbor approximation as

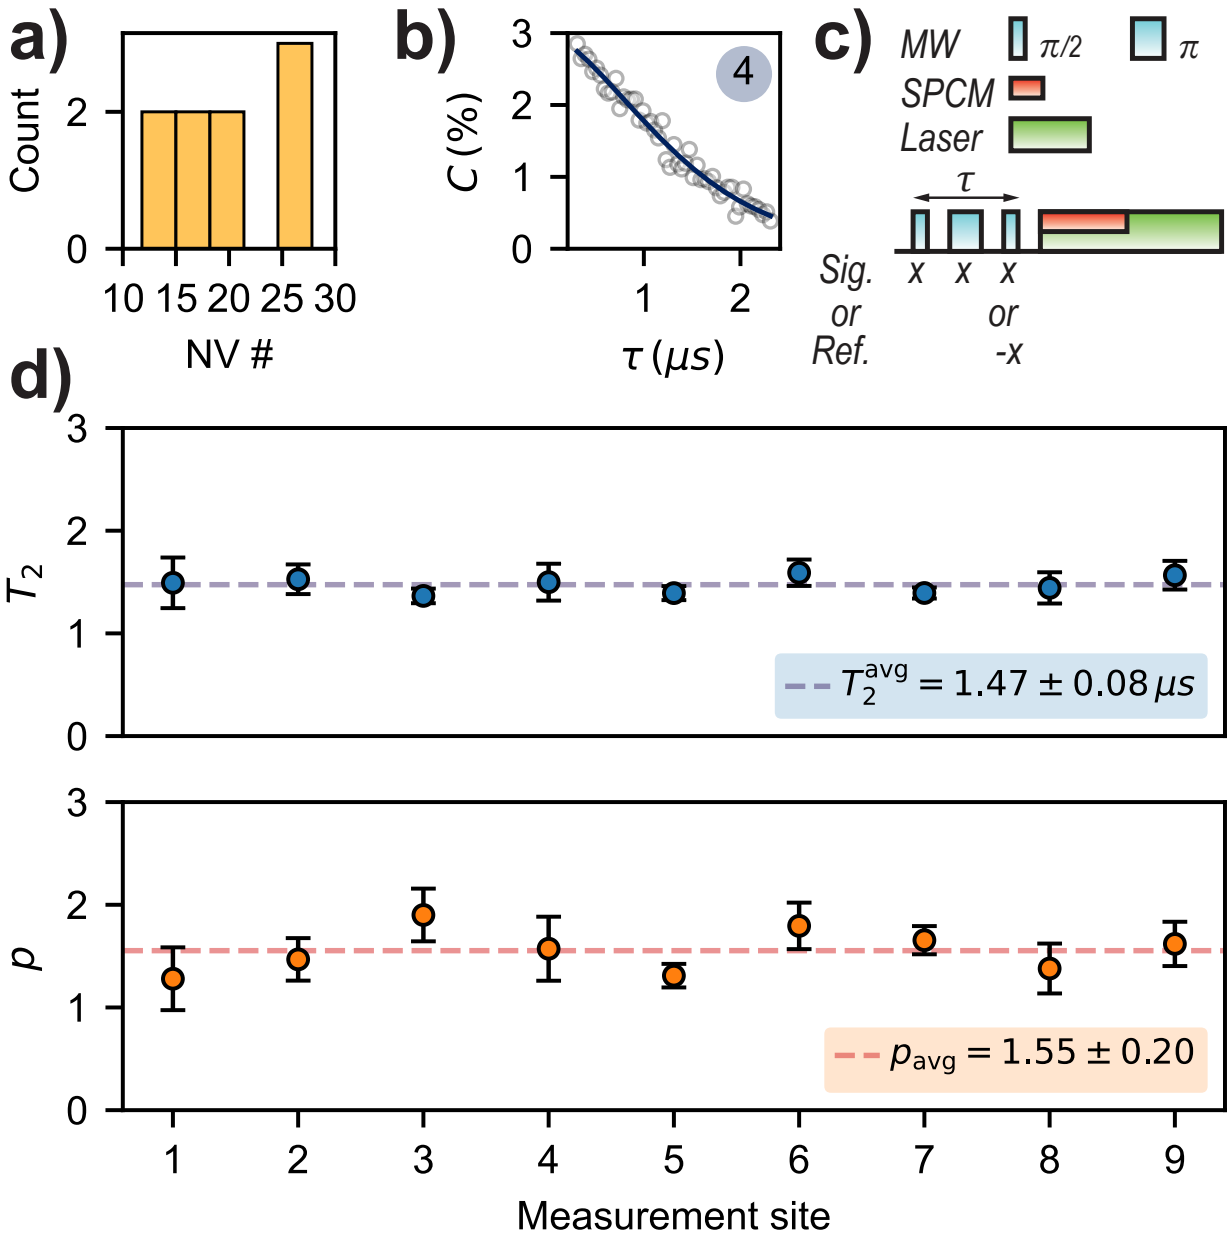


Figure 3.9: Spin coherence of ion-implanted NV centers. a) Distribution of NV numbers in the implantation sites selected for coherence measurements. b) Representative Hahn echo measurement (site 4). The signal is fitted to $C_0 \exp[-(\tau/T_2)^p]$ to extract the coherence time T_2 and decay exponent p. c) Hahn echo pulse sequence used for the measurements. The phase of the final microwave $\pi/2$ pulse is alternated between signal and reference shots to obtain the NV PL contrast C. d) Extracted coherence time T_2 (top) and decay exponent p (bottom) for nine single ion implantation sites. Error bars represent the standard deviation from fitting. The measured average values ($T_2^{\text{avg}} \approx 1.47\,\mu$s, $p_{\text{avg}} \approx 1.55$) are consistent with nitrogen spin-bath-limited ensemble decoherence.

the dominant missing physics. Applying a first-order correction for these effects reduced the predicted mean yield to $N_{\mathrm{NV}} = 16.6 \pm 0.7$, in close agreement with experiment.

The experimentally benchmarked simulation framework was then used to evaluate the directional information retained in the post-annealing NV distribution. Although annealing broadens the original damage track and produces limited NV yield, a counting-based head-tail asymmetry metric retained a 74% head-tail direction classification efficiency at a false-positive rate below approximately 5% for simulated $800\ \mathrm{keV}$ tracks. Simulation-based inference further improved the corresponding classification efficiency to 95% by using spatial information beyond the total defect count in the head and tail regions.

Hahn-echo measurements showed that the implantation-created NV ensembles possess an average coherence time of $T_2 \approx 1.47\ \mu\mathrm{s}$ in the nitrogen-rich HPHT material. This coherence is compatible with prospective magnetic gradient-based imaging approaches that could resolve NV distributions with spatial resolution comparable to, or smaller than, the simulated NV-NV separation in the present tracks.

In terms of limitations, these directionality results remain model-based projections rather than experimental measurements. The nanoscale positions of individual NV centers within a single post-annealing track have not yet been directly resolved, and their measurement represents an important next step for testing the model predictions of defect spatial distributions and the inferred directional performance. In addition, the second-nearest-neighbor correction introduced here is partially incorporated as a first-order estimate. Fully incorporating extended-range defect-defect interactions into the KMC treatment is warranted, given the initial encouraging agreement between the corrected model and experiment.

This framework can be extended to assess the applicability of NV-based directional detection

in HPHT diamond at lower recoil energies. The resulting analysis, together with its implications for rare-event detection, is considered further in Chapter 5.

Chapter 4: Low-Nitrogen Quantum Diamond: Engineering and Characterization

Chapter provenance and contributions. This chapter is adapted from Jiashen Tang, Connor A. Roncaioli, Andrew M. Edmonds, Atli Davidsson, Connor A. Hart, Matthew L. Markham, and Ronald L. Walsworth, "Characterization of low-nitrogen quantum diamond for pulsed magnetometry applications" [119]. This work was carried out as part of the longstanding collaboration between the Walsworth group and Element Six, with the primary aim to develop, characterize, and improve NV diamond materials for quantum sensing. I contributed to the conception and development of the experimental study with Andrew M. Edmonds and Matthew L. Markham. Element Six synthesized the low-nitrogen CVD diamond samples and performed birefringence strain measurements. I performed the NV optical, spin, charge-state measurements, and led the corresponding data analysis. Connor A. Roncaioli conducted the wide-field NV-based strain measurements and analysis. I prepared the complete initial manuscript draft and contributed to subsequent revision and editing with Ronald L. Walsworth and the other authors,. This work illustrates an iterative material-development process.

4.1 Motivation and chapter overview

As discussed in Chapter 1, the magnetic sensitivity of an NV-diamond sensor depends on a combination of spin, material, measurement-protocol, and experimental parameters. Relevant spin properties include the NV spin dephasing time T_2^* and the spin-readout photoluminescence contrast. At the material level, the substitutional nitrogen concentration, NV concentration, NV charge-state fraction, residual $^{13}\mathrm{C}$ concentration, strain homogeneity, and other defect populations can also affect magnetic sensing performance. Regarding the measurement protocol, as discussed in Chapter 2,

pulsed-magnetometry protocols such as Ramsey sensing typically offer better sensitivity than CW-ODMR, at the cost of more demanding quantum control. Finally, experimental parameters such as optical excitation intensity can influence both the spin-readout contrast and the charge state of the NV ensemble.

Control and optimization of these parameters are not independent. For example, increasing the substitutional nitrogen concentration can increase the density of NV centers formed after irradiation and annealing of a nitrogen-doped as-grown diamond plate. Substitutional nitrogen atoms also function as electron donors that help stabilize the negatively charged NV^- state used for sensing. However, nitrogen also forms a dominant paramagnetic spin bath that shortens the NV spin dephasing time. As another example, increasing the optical excitation intensity can reduce the time required for spin initialization and increase the detected PL rate. However, it can also ionize the sensing NV^- state into NV^0, reducing the effective NV^- density through a change in the NV charge-state fraction and reducing the observed spin-readout contrast. The interrogated sensing volume also changes the strain inhomogeneity experienced by the selected NV ensemble. Thus, a material that exhibits low strain variation in a confocal-scale measurement may show greater inhomogeneous broadening when used over a much larger sensing area. Therefore, engineering a diamond sample requires a broader understanding of the causal relationships among the relevant NV sensing properties.

Over time, diamond material engineering has identified several broadly beneficial techniques as well as key material-engineering parameters. For the former, $^{12}\mathrm{C}$ isotopic enrichment can suppress dephasing from the residual $^{13}\mathrm{C}$ nuclear spin bath, while controlled CVD growth and careful substrate selection can reduce strain inhomogeneity. Careful choice of irradiation and annealing parameters can also increase the NV concentration without introducing deleterious charge-stability issues. Among the latter, the concentration of neutral substitutional nitrogen, $[\mathrm{N}_\mathrm{s}^0]$, plays a particularly central role

because it strongly affects NV density, NV^- charge-state fraction, and the nitrogen spin bath that often limits ensemble T_2^*.

Previous material studies identified a nitrogen concentration of roughly $1\text{–}10\ \mathrm{ppm}$ as a favorable regime, balancing NV density and coherence time for ensemble magnetic sensing [57, 58]. This guidance was largely based on a simplified sensitivity metric, expressed as

$$\eta \propto \frac{1}{\sqrt{[\mathrm{N_s^0}]T_2^*}}. \tag{4.1}$$

Here, $[\mathrm{N_s^0}]$ is used rather than the NV concentration because the achievable NV density scales with the available substitutional nitrogen concentration. However, this metric does not adequately capture the experimental overhead associated with NV spin initialization and readout. Such overhead is particularly important for pulsed magnetometry protocols, where optical initialization, microwave spin manipulation, and optical readout are required and temporally separated from the free-precession interval during which active sensing occurs. Including a finite initialization-and-readout overhead in the sensitivity estimate shifts the preferred nitrogen concentration toward lower values. For example, as shown in Figure 4.1, assuming nitrogen-spin-bath-limited dephasing times and comparable values of other sensing parameters, including contrast and effective NV^- density, a low-nitrogen diamond with $[\mathrm{N_s^0}] \approx 0.8\ \mathrm{ppm}$ is predicted to provide a nearly threefold sensitivity improvement over a previously characterized higher-nitrogen diamond with $[\mathrm{N_s^0}] \approx 14\ \mathrm{ppm}$ when the overhead time is approximately $10\ \mu\mathrm{s}$. The development of lower-nitrogen diamond samples therefore enables an experimental test of whether the long coherence, increased sensing duty cycle advantage predicted can be realized in practice once material strain, NV charge-state behavior, spin-readout contrast, and optical initialization overhead are taken into account.

In this chapter, I describe the synthesis and characterization of low-$[N_s^0]$ CVD diamond designed for pulsed NV ensemble magnetometry, with emphasis on Ramsey-based dc and broadband magnetic-field sensing. Low-strain CVD growth, ^{12}C isotopic enrichment, and controlled electron irradiation and annealing are used to produce NV-enriched diamonds with $[N_s^0] \approx 0.8$ ppm. I first present the characterization of nitrogen concentration, birefringence, and post-treatment NV concentration for a cohort of diamonds produced in a common growth run. I then examine the achieved NV spin-dephasing time at the confocal scale, together with the strain variation as a function of sensor size and its implications for bulk magnetic-sensing performance. The NV spin-state contrast, initialization time, and charge-state fraction are then measured as functions of optical excitation intensity.

Table 4.1 summarizes the principal parameters characterized in this chapter and their relevance to practical pulsed magnetometry. This overview is intended to make the central material and protocol tradeoffs visible before the detailed characterization, since diamond material engineering can otherwise become highly specialized and obscure the connection between individual measurements and magnetic sensitivity.

4.2 Methods for sample synthesis and characterization

4.2.1 Key equations and parameters relevant to sensitivity

The photon shot-noise-limited magnetic sensitivity of an NV ensemble Ramsey-based DC magnetometry protocol is given by:

$$\eta_{\text{Ramsey}}^{\text{photon-shot}} = \frac{1}{\Delta m_s \gamma_e} \frac{1}{\sqrt{N\tau}} \frac{1}{e^{-(\tau/T_2^*)^p}} \sqrt{1 + \frac{1}{C^2 n_{\text{avg}}}} \sqrt{\frac{\tau + t_O}{\tau}}, \tag{4.2}$$

Table 4.1: Principal material and measurement parameters characterized in this chapter, together with their relevance to practical pulsed Ramsey magnetometry.

Parameter characterized	**Why it matters for pulsed Ramsey sensitivity**	**How it is treated in this chapter**
Substitutional nitrogen concentration, $[\mathrm{N_s^0}]$	Sets a central tradeoff: higher nitrogen can support greater NV creation and stabilize $\mathrm{NV^-}$, but also creates a stronger paramagnetic spin bath that shortens T_2^*.	Measured across a cohort of low-nitrogen CVD diamonds to establish the material regime and growth reproducibility.
Post-treatment NV concentration and charge fraction, $[\mathrm{NV}]$ and ψ	The effective density of sensing centers is $[\mathrm{NV^-}] = \psi[\mathrm{NV}]$. Both total NV density and the fraction remaining in the $\mathrm{NV^-}$ charge state affect PL signal and contrast.	Total NV concentration is estimated after electron irradiation and annealing. The charge fraction ψ is measured as a function of optical excitation intensity.
$^{12}\mathrm{C}$ enrichment / residual $^{13}\mathrm{C}$	Suppresses dephasing from the $^{13}\mathrm{C}$ nuclear-spin bath, which becomes increasingly relevant when the nitrogen spin bath is reduced.	Controlled through isotopically enriched CVD growth and considered when interpreting the measured T_2^*.
Strain inhomogeneity, Δ	Spatial variation in strain broadens the ensemble resonance and can limit T_2^*, particularly in large-area sensing.	Evaluated through birefringence imaging and wide-field ODMR strain mapping.
Sensing area, L_{sensor}^2	A larger sensing area contains more NV centers and can improve sensitivity, but may also sample greater strain variation and reduce the effective T_2^*.	Examined by evaluating Δ over progressively larger square sensor regions with lateral dimension L_{sensor}.
NV spin-dephasing time, T_2^*	Sets the useful Ramsey free-precession interval and therefore the fraction of each measurement cycle spent actively sensing the magnetic field.	Measured using confocal single-quantum and double-quantum Ramsey protocols.
Optical excitation intensity, I	Influences PL rate, spin initialization, spin-readout contrast, and charge-state conversion. Its optimum is therefore not determined by PL brightness alone.	Varied systematically in measurements of t_{I}, C, and ψ.
Spin-readout contrast and initialization time, C and t_{I}	Contrast determines the magnetic signal extracted from spin-dependent PL, while initialization time contributes to pulsed measurement overhead and reduces sensing duty cycle.	Measured as functions of I.

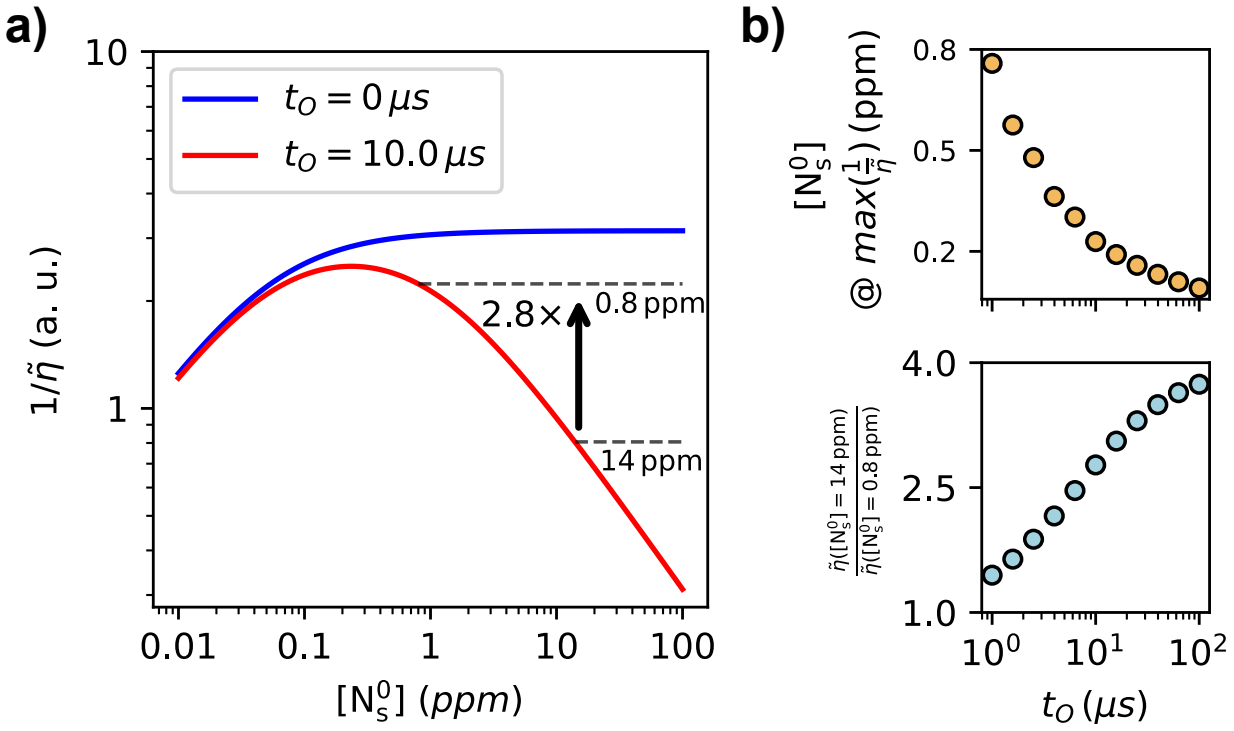


Figure 4.1: (a) NV ensemble DC magnetic field sensitivity estimated using the simplified metric $1/\tilde{\eta} = \sqrt{[\mathrm{N}_\mathrm{s}^0] \times T_2^*}$, assuming spin-bath-limited T_2^* and 99.995% isotopic enrichment of ^{12}C. Corrections accounting for experimental overhead time (t_O) are included by multiplying with $1/\sqrt{(T_2^* + t_O)/T_2^*}$ (cf. Eq. 4.2). With zero overhead (blue line), sensitivity plateaus for $[\mathrm{N}_\mathrm{s}^0] \gtrsim 1$ ppm. However, incorporating an overhead of 10μ s, typical in pulsed Ramsey protocols, shifts the optimal $[\mathrm{N}_\mathrm{s}^0]$ towards lower values (red line). A preliminary comparison between the current low-$[\mathrm{N}_\mathrm{s}^0]$ samples (~ 0.8 ppm) and a previously characterized high-$[\mathrm{N}_\mathrm{s}^0]$ sample (~ 14 ppm suggests a nearly threefold sensitivity improvement. (b) Estimated optimal nitrogen concentration, and sensitivity improvement between 0.8 ppm and 14 ppm samples as functions of overhead time (t_O).

where Δm_s denotes the spin-state transition (1 for single-quantum and 2 for double-quantum transitions [19]), γ_e is the electron gyromagnetic ratio, N is the total number of NV^- contributing to the measurement, τ is the spin free-precession time during sensing, T_2^* and p characterize the NV spin dephasing time and Ramsey decay envelope shape, C is the NV spin-state-dependent photoluminescence (PL) measurement contrast, n_{avg} is the average PL photon number detected per NV^- per measurement, and t_O is the total experimental overhead time including initialization and readout.

T_2^* for an NV ensemble can be expressed as the inverse sum of several common dephasing mechanisms [1, 57]:

$$\frac{1}{T_2^*} = \frac{1}{T^*_{2,[\mathrm{N_s^0}]}} + \frac{1}{T^*_{2,[^{13}\mathrm{C}]}} + \frac{1}{T^*_{2,\text{NV-NV}}} + \frac{1}{T^*_{2,\text{ strain}}} + \frac{1}{T^*_{2,\text{bias}}}. \tag{4.3}$$

The first two terms represent the spin-bath noise from substitutional nitrogen and residual ^{13}C in the diamond lattice, with their contributions scaling directly with respective concentrations. The third term accounts for NV–NV dipolar interactions, which can be significant in high NV density samples created by irradiation and annealing. The last two terms account for spatial variation (inhomogeneity) in lattice strain and the applied bias magnetic field. Quantification of [N_s^0], [NV], ^{13}C content, and strain is therefore essential for confirming whether a given diamond sample achieves spin-bath-limited dephasing.

The NV optical excitation intensity further affects the photon-shot-noise-limited sensitivity by modifying the NV^- charge fraction, and hence the effective number of sensors N. Since NV^- and NV^0 have overlapping emission spectra and different PL brightness, both the measurement contrast (C) and detected photon number per NV^- (n_{avg}) vary with excitation laser intensity. Additionally, the experimental overhead time t_O, which includes NV spin initialization and readout, also depends on

optical NV excitation intensity. Therefore, in addition to measuring T_2^*, it is important to characterize these parameters as functions of excitation intensity to enable a complete sensitivity assessment.

4.2.2 Diamond synthesis and treatment

The diamond samples studied here were grown by Element Six on diamond substrates using microwave-plasma-assisted chemical vapor deposition (CVD). Nitrogen-doped CVD diamond can exhibit an increased fraction of positively charged substitutional nitrogen, $[\mathrm{N_s^+}]$, which is associated with undesirable brown coloration. This coloration increases absorption of both the excitation light and NV PL, and is correlated with elevated concentrations of vacancy-related paramagnetic defects [120–123] and hydrogen-related defects [124, 125]. To mitigate these effects while achieving a homogeneous nitrogen distribution with $[\mathrm{N_s^0}] \approx 1\ \mathrm{ppm}$, key synthesis parameters—including substrate temperature, the $\mathrm{CH_4/H_2}$ ratio, and the $\mathrm{N_2}$ concentration—were carefully selected according to methods described elsewhere [58]. Methane sources isotopically enriched with $^{12}\mathrm{C}$ ($> 99.99\%$) were also used. Deposition conditions were maintained throughout the growth run to avoid the formation of non-epitaxial crystallites on the primary growth surface. Twinning near the sample edges was likewise minimized to preserve usable area in the final processed plates.

Its worth noting that the crucial physical parameters that governs CVD diamond synthesis include the density of surface radical sites, the rate of CH_x addition relative to etching by H atoms, and the near-surface NH_x (or CN)/CH_x ratio [58]. The mapping of CVD process conditions (such as power, gas flows, and surface temperature) to these parameters is dependent on the specific CVD reactor design [14, 58, 126]. For the diamonds studied here, Element Six's proprietary and non-disclosed reactor design is used; and therefore it is not possible to generalize for users of other reactor platforms.

Crystal strain in CVD-grown diamond layers is strongly influenced by the dislocation density, and dislocations typically propagate along the growth direction [16, 127]. Key sources of dislocations include those nucleated at the substrate–growth interface due to substrate polishing damage, and dislocations threading from the substrate into newly grown layers. Additionally, mismatches in nitrogen concentration between the substrate and growth layers can induce bulk stress due to lattice mismatch [128]. Thus, careful selection and preparation of substrates is critical.

The substrates were low-nitrogen CVD single-crystal diamonds with $[N_s^0] \leq 1$ ppm, closely matching the nitrogen concentration of the intended growth layer (~ 0.8 ppm). They were selected for low birefringence, indicating a low density of dislocations likely to propagate into the new growth layer. The substrates had typical dimensions of approximately $3.5 \times 3.5 \times 0.3\ \mathrm{mm}^3$, with mechanically polished $\{100\}$ surfaces prepared using standard scaife techniques [129, 130]. After approximately 1.2 mm of growth, the CVD layers were laser sliced from the substrates. Only the central $3 \times 3\ \mathrm{mm}^2$ region was retained to avoid dislocations propagating from the substrate edges.

Following processing, electron irradiation and annealing were performed to convert grown-in $[N_s^0]$ defects into NV centers. Samples were irradiated using a 1.5 MeV electron beam at an estimated dose of $\sim 4.8 \times 10^{17}\ \mathrm{cm}^{-2}$, targeting approximately 50% N_s^0-to-NV conversion rate. This target was selected to maximize NV PL without introducing an excessive population of residual vacancies that are not converted into NV centers and could negatively affect NV^- charge stability and coherence time [131]. Annealing followed a previously established temperature-ramp profile, with a final hold at 1200 °C for 2 h [58].

4.2.3 NV-based wide-field strain mapping

Strain variations across each NV-diamond sample are mapped using CW-ODMR in a wide-field imaging configuration, as previously described in Ref. [132]. Briefly, a 532 nm laser beam (~ 1 W) illuminates the diamond at a shallow incidence angle. Before reaching the diamond, a 30 mm focal-length cylindrical lens shapes the laser beam into an elliptical profile, measuring $\sim$ 10 mm by $\sim$ 700 μm, ensuring uniform illumination across the diamond. NV PL is collected using a 0.1NA, 4$\times$ objective and imaged onto a camera.

A bias magnetic field $\sim$ 50 G is applied to the NV-diamond sample under study using a pair of permanent magnets, with the field oriented such that spin resonances from all four NV orientations in the diamond lattice are resolved. This configuration allows extraction of the NV spin-strain coupling terms (M_z) following the method described in [11]. Microwaves are delivered to the sample using a loop-gap resonator [133]. Pixel-by-pixel fitting of the NV Hamiltonian to the CW-ODMR spectra is performed using GPU-accelerated fitting [134], generating full-diamond strain maps.

4.2.4 Confocal-level characterization of NV spin and charge properties

NV spin-dephasing times (T_2^*), charge-state fractions, PL contrast, and experimental overhead associated with spin initialization and readout are measured using a custom-built confocal setup. These measurements are performed as functions of the optical NV excitation intensity.

To achieve relatively uniform optical excitation over the selected NV ensemble, the 532 nm laser beam is focused to a waist diameter of $\sim$ 15 μm by underfilling a 0.75-NA, 20$\times$ objective. A pinhole in the detection path restricts PL collection to a smaller central region, with an approximate diameter of 6 μm and depth of 10 μm. The optical intensity at the diamond surface is determined from the average

intensity within this restricted collection volume and is varied controllably from approximately 10^{-3} to $10^{1}\ \mathrm{mW}/\mu\mathrm{m}^2$.

Microwave pulses for NV spin-state manipulation are delivered using a coaxial loop positioned near the diamond surface, producing Rabi frequencies of approximately $5\ \mathrm{MHz}$. A bias magnetic field of approximately $20\ \mathrm{G}$ is aligned along a single NV crystallographic axis and optimized to minimize magnetic-field inhomogeneity within the probed region, similar to the approach in Ref. [57].

For NV charge-state determination, PL emission is collected through a multimode optical fiber coupled to a spectrometer. Recorded emission spectra are decomposed into NV^- and NV^0 spectral profiles [135]. NV^- charge fractions are then quantified using the intensity ratio between NV^- and NV^0 PL [113].

4.3 Birefringence and nitrogen concentration of synthesized diamonds

We first characterize ten low-$[\mathrm{N}_\mathrm{s}^0]$ diamonds grown in a single synthesis run, prior to electron irradiation and annealing. Measurements of $[\mathrm{N}_\mathrm{s}^0]$ obtained from UV–Vis absorption [58] yield values between 0.73 and $0.88\ \mathrm{ppm}$. Optical birefringence imaging at a wavelength of $590\ \mathrm{nm}$, performed using the methods described in [128, 136], gives average birefringence values, Δn, ranging from 2.6×10^{-6} to 5.2×10^{-6} across the sample cohort [Fig. 4.2]. These values indicate generally good strain homogeneity among the synthesized diamonds [128]. Following irradiation and annealing, the NV concentration in each sample increases to approximately $0.39(2)\ \mathrm{ppm}$, as measured by UV–Vis absorption [58].

For subsequent characterizations (described below and in the Appendix C.1) we focus on sample 10 This sample contains spatially restricted regions with relatively higher birefringence compared

with the other diamonds from this growth run, which therefore allows examination of both the typical behavior of the low-$[N_s^0]$ material, measured away from the higher-birefringence regions, and the more unfavorable influence of strain on NV spin dephasing and sensing performance within those regions. The high-strain pattern is indicative of dislocation bundles, most likely originating from residual damage due to mechanical polishing of the CVD substrates, rather than from the underlying bulk dislocation content of the substrates. If further strain reduction is required, alternative polishing techniques and enhanced substrate characterization for subsurface damage can be implemented [137]. The main factors governing NV ensemble Ramsey dephasing time (T_2^*) – namely nitrogen and ^{13}C concentrations and strain uniformity – are well established [1, 57, 58], allowing reliable estimation of T_2^* from these parameters. The diamonds grown in the same synthesis run exhibit closely matched birefringence and nitrogen concentrations, indicating good sample-to-sample reproducibility.

4.4 NV spin dephasing time and strain contribution

Confocal Ramsey measurements of the single-quantum (SQ) NV ensemble spin dephasing time at three representative low-birefringence locations within sample 10 (Fig. 4.3(a,b)) yield an average value $T_{2,\mathrm{SQ}}^* = 17.5 \pm 1.0\,\mu$s (an additional high-birefringence location measurement is shown in the Appendix C.1). These results are only slightly smaller than the theoretical spin-bath-limited value of approximately 20 μs, consistent with NV ensemble spin dephasing in the low-$[N_s^0]$ diamond samples being primarily limited by intrinsic nitrogen spin-bath interactions (see Appendix C.2). Achieving $T_{2,\mathrm{SQ}}^*$ values near the spin-bath limit – without using double-quantum techniques as in [57] – is enabled by improved strain control during CVD growth, including the use of low-dislocation CVD substrates instead of HPHT substrates and the growth optimizations described in Section 4.2.2. Similar measure-

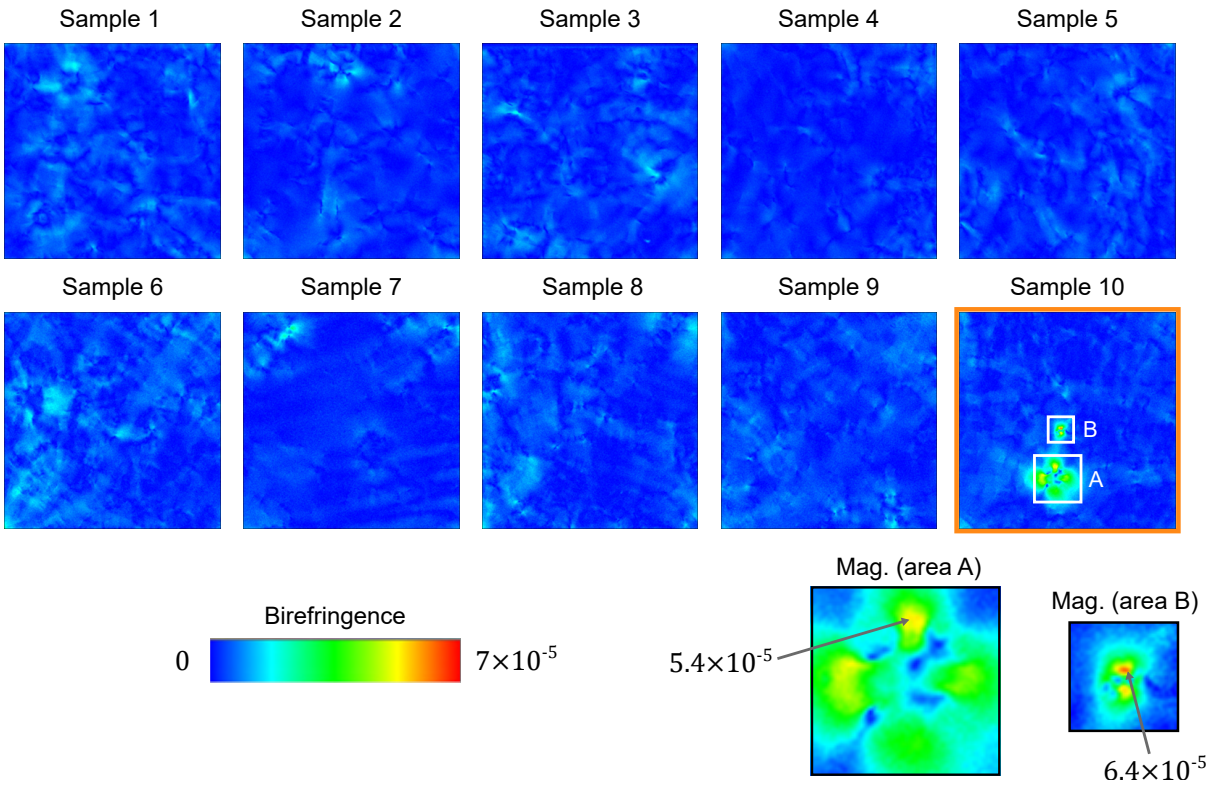


Figure 4.2: Optical birefringence images of ten low-[N_s^0] diamond samples from a single CVD synthesis run, prior to irradiation and annealing. Sample 10, exhibiting the strongest birefringence features, is selected for subsequent NV-based measurements after irradiation and annealing. Magnified views are shown of two regions of sample 10 with higher levels of birefringence (labeled A and B), with example values of high birefringence indicated.

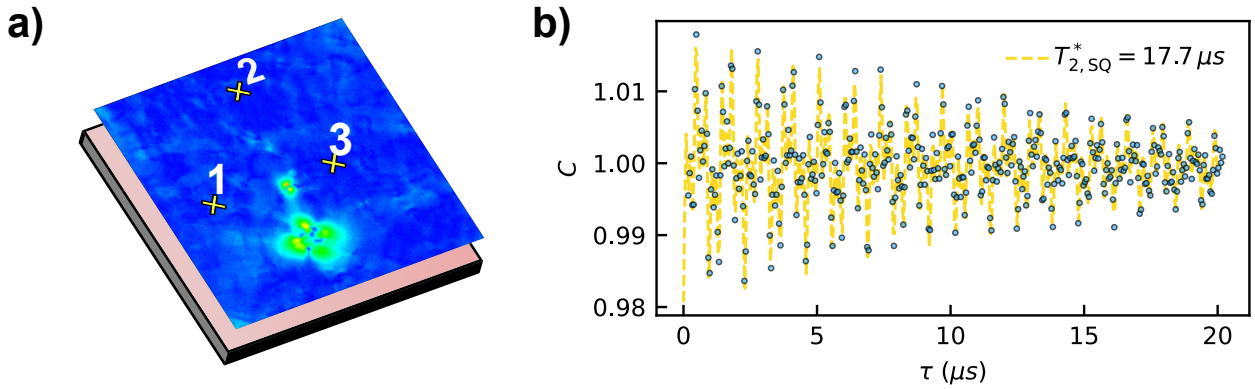


Figure 4.3: Single-quantum (SQ) Ramsey measurements on diamond sample 10, far from the high birefringence regions. (a) Ramsey measurements performed at three representative locations of low strain, overlaid on the birefringence image from Fig. 4.2. (b) Ensemble NV spin dephasing at location 2. Here the microwave frequency is detuned from the center of the hyperfine-split NV spin transitions, allowing the decay envelope to be separated from hyperfine-induced beating. The extracted ensemble NV dephasing time is $T^*_{2,\mathrm{SQ}} = 17.7(4)\,\mu$s.

ments on a previously studied higher-[N_s^0] (~ 14 ppm) diamond sample [58] yield $T^*_{2,\mathrm{SQ}} \approx 1\,\mu$s.

Figure 4.4(a) shows the strain-induced frequency shifts, M_z, of the NV spin resonances measured across sample 10 for one NV crystallographic orientation. Corresponding maps for all four NV orientations are provided in Appendix C.3. The distribution of these frequency shifts provides a quantitative estimate of the strain-limited NV electronic spin-dephasing time [1],

$$T^*_{2,\mathrm{strain}} = (\pi\Delta)^{-1}\,, \tag{4.4}$$

where Δ is the full width at half maximum (FWHM) of the strain-shift distribution. Histogram analysis of strain across the full diamond [Fig. 4.4(c)] yields $\Delta \approx 31$ kHz, corresponding to $T^*_{2,\mathrm{strain}} \approx 10\,\mu$s. These wide-field measurements include strain variation through the full sample thickness, approximately $500\,\mu$m for the diamonds studied here, and across the full lateral extent of the diamond. This thickness is roughly five times greater than that of the strain-imaged higher- [N_s^0] diamond studied in

Ref. [58]. The full-diamond measurement therefore overestimates the strain contribution relevant to smaller sensing volumes. For example, histogram analysis within a smaller low-birefringence region of sample 10, with an approximate diameter of $500\ \mu\mathrm{m}$, yields a narrower linewidth of $\Delta \approx 15\ \mathrm{kHz}$ and therefore $T^*_{2,\mathrm{strain}} \approx 20\ \mu\mathrm{s}$ [Fig. 4.4(b,d)]. This result indicates substantially reduced strain inhomogeneity in a smaller sensing region. Extrapolating to the much smaller volume probed by the confocal measurements suggests a further order-of-magnitude reduction in strain inhomogeneity, such that strain makes only a modest contribution to the measured confocal $T^*_{2,\mathrm{SQ}}$. At lower nitrogen concentration, the spin-bath limit occurs on a longer timescale, making good strain homogeneity essential for reaching that limit. The present results confirm that this condition is satisfied in the low-$[\mathrm{N}^0_\mathrm{s}]$ diamond studied here.

To evaluate the feasibility of scaling sensor size for bulk NV magnetometry, we analyze how strain-induced linewidth broadening changes with the lateral sensing dimension, L_{sensor}. We partition the wide-field strain map shown in Fig. 4.4(a) into square sub-regions with areas ranging from approximately $30\ \mu\mathrm{m}^2$ to $3000\ \mu\mathrm{m}^2$, and compute the strain-induced frequency-shift distribution within each sub-region, as illustrated in Fig. 4.4(c,d). We find that the median FWHM of these distributions, Δ, varies by only approximately a factor of three over this range (Fig. 4.5(a)), indicating modest strain heterogeneity over large length scales within the diamond sample.

To illustrate how size-dependent strain-induced dephasing affects NV ensemble sensor performance, we compute an effective Ramsey DC magnetometry sensitivity metric, $(T^*_{2,\mathrm{eff}} \times L_{\mathrm{sensor}})^{-1}$, where $T^*_{2,\mathrm{eff}}$ is calculated from Eq. 4.3 using the median Δ associated with each sensor size. The $\propto L^{-1}_{\mathrm{sensor}}$ comes from the following: assuming a constant sensor thickness across the diamond sample, the sensing volume, and therefore the number of interrogated NV centers, scales as L^2_{sensor}. In the ideal case where increasing L_{sensor} does not introduce additional ensemble dephasing, the effective sensitiv-

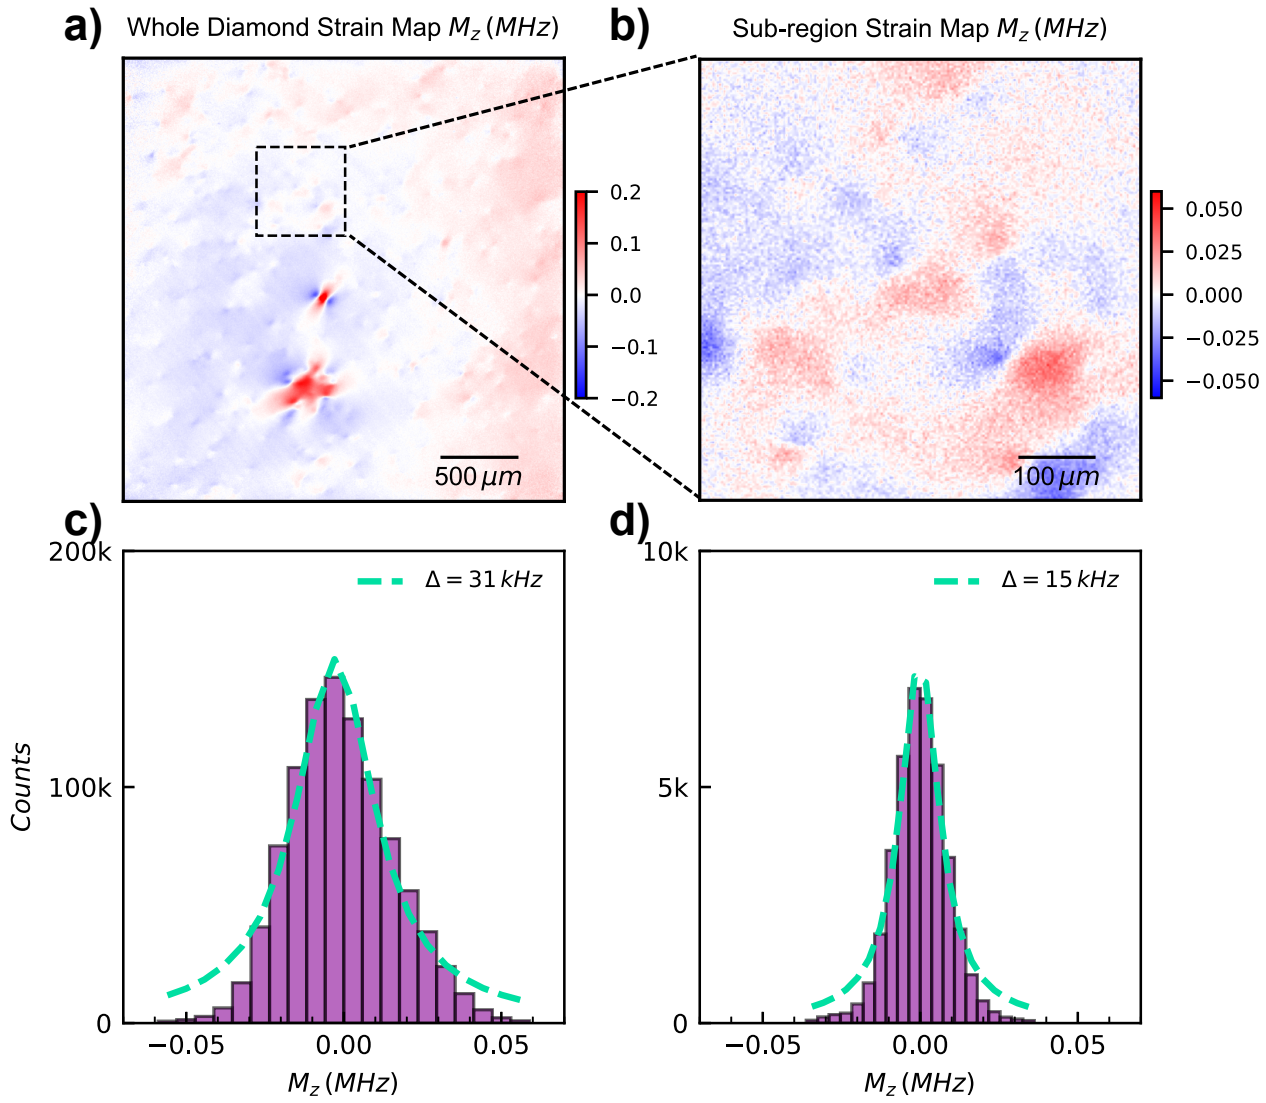


Figure 4.4: Wide-field imaging of diamond strain for sample 10. (a, b) Spatial maps of strain-induced frequency shifts in NV spin transitions across the entire diamond sample and within a selected sub-area. Strain fields shown are projections onto one NV crystallographic orientation. Each spatial map is mean-subtracted so that the global average is zero, emphasizing relative strain variations relevant to NV spin dephasing. (c, d) Histogram distributions of the strain-induced frequency shifts extracted from the maps in (a) and (b), respectively. Lorentzian fits to the histograms are also shown, from which linewidths (FWHM) Δ are determined.

ity metric should therefore scale inversely with L_{sensor}. Regarding the inverse linear relationship with dephasing time T_2^*: for photon-shot-noise-limited Ramsey-based DC magnetic sensing for an NV ensemble, optimal sensitivity typically occurs when the NV spin free-precession time (τ) is close to the spin dephasing time T_2^* (i.e., $\tau \approx T_2^*$) [1]. Consequently, the relationship between sensitivity (η) and spin dephasing time can be approximated as:

$$\eta \propto \frac{\sqrt{T_2^* + t_O}}{T_2^*}, \tag{4.5}$$

In the limit of large experimental overhead time ($t_O \gg T_2^*$), sensitivity scales inversely with dephasing time ($\eta \propto 1/T_2^*$). Conversely, when overhead time is negligible ($t_O \ll T_2^*$), sensitivity scales with the inverse square root of dephasing time ($\eta \propto 1/\sqrt{T_2^*}$). Thus, when evaluating the impact of spatially varying diamond strain on sensitivity through its effect on spin dephasing time, the most conservative scenario is considered.

The observed near-inverse-linear scaling of this metric for sample 10 (Fig. 4.5(b)) indicates that Ramsey T_2^* is only weakly reduced as the sensing region is increased over the measured length scales. This result suggests that increasing sensor size remains a viable strategy for improving bulk NV magnetometry using this low-$[\mathrm{N_s^0}]$ diamond material.

4.5 Photoluminescence contrast and overhead time

We characterize the NV photoluminescence (PL) spin-state contrast, C, and spin-initialization time, t_I, as functions of optical excitation intensity. Figure 4.6(a) illustrates the protocol used to measure the PL contrast and NV spin-polarization dynamics. After initializing the NV spins into either $m_s = 0$ or $m_s = -1$, the delay time, t_{delay}, before spin-state readout is varied. The PL contrast

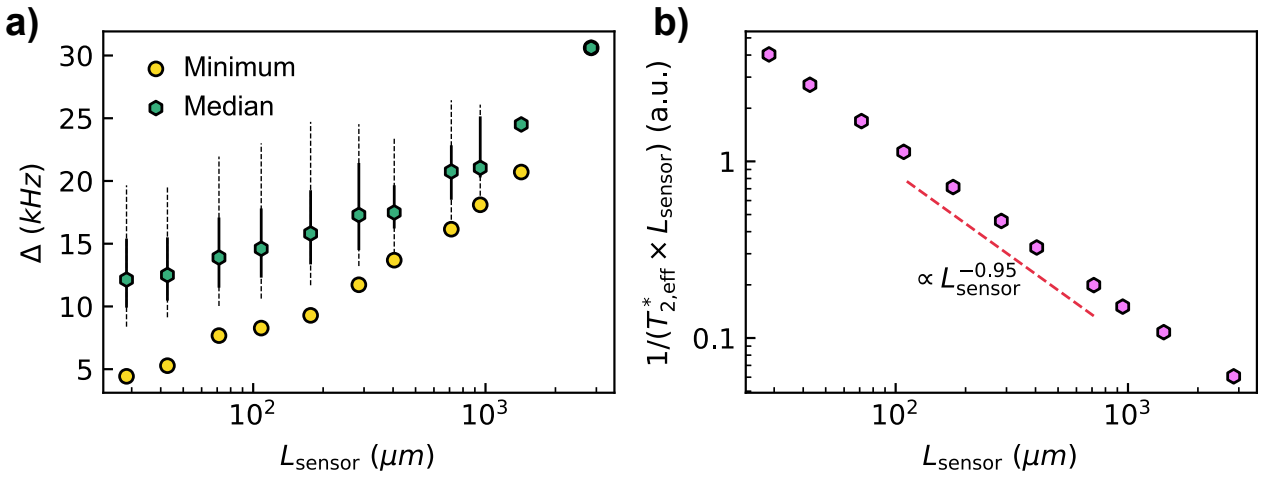


Figure 4.5: (a) Lorentzian-fitted FWHM (Δ) of strain-induced frequency shifts as a function of sensor size L_{sensor} for sample 10. Sensor size is systematically varied by partitioning the full-diamond strain map (Fig. 4.4(a)) into smaller square sub-regions. Statistical analysis of Δ across sub-regions is presented, including minimum and median values. When more than five sub-regions are available, interquartile ranges are included: 25-75% (vertical solid lines) and 10-90% (vertical dotted lines). (b) Scaling of effective Ramsey DC magnetometry sensitivity metric with increasing sensor size. In the ideal scenario, where increasing the sensing volume introduces no additional NV ensemble dephasing – and thus no reduction in Ramsey T_2^* – the sensitivity metric scales inversely with sensor size. Here, fitting a power-law function to the effective sensitivity metric vs. L_{sensor} yields nearly inverse linear scaling behavior, indicating minimal strain limitation on the Ramsey dephasing time for bulk NV ensemble sensing using the present low-[N_s^0] diamonds. The overall low-strain effect of these diamonds applies even to sample 10 with its isolated regions of relatively high strain.

initially reaches a maximum and then decreases as the NV spins repolarize toward $m_s = 0$. This behavior allows both C and t_I to be extracted from a single measurement (Fig. 4.6(b,c)).

To validate the experimental calibration of laser power to incident optical excitation intensity for the measured NV ensemble, we simulate NV spin polarization dynamics using a five-level photophysics model [43]. The NV electronic and spin energy levels included in the simulation are illustrated schematically in Fig. 4.7. The spin-conserving radiative decay rate is set to $\Gamma = 0.67\,\mu\text{s}^{-1}$ [43]. Transitions between excited and singlet states, as well as subsequent transitions to the ground state, are characterized by decay rates given by $\kappa_{ij}\Gamma$, where i, j denote the initial and final states, respectively. The optical excitation rate from the ground to the excited state is modeled as $s\Gamma$, with saturation parameter $s = I/I_{sat}$, where I is the laser intensity and I_{sat} is the NV saturation intensity. The saturation intensity is taken to be approximately $1 - 3\,\text{mW}/\mu\text{m}^2$ [1], providing bounds for our simulations. The rate-constant prefactors κ used in the model are averaged from experimental values summarized in Ref. [43]: $\kappa_{45} = 1, \kappa_{35} = 1/7, \kappa_{52} = 1/50, \kappa_{51} = 1/25$.

For the experimental measurement protocol described in the Figure 4.6, we simulate NV spin population dynamics starting from the initial states $m_s = 0$ and $m_s = \pm 1$. The following system of coupled first-order differential equations describes the population n_i of each of the five modeled NV

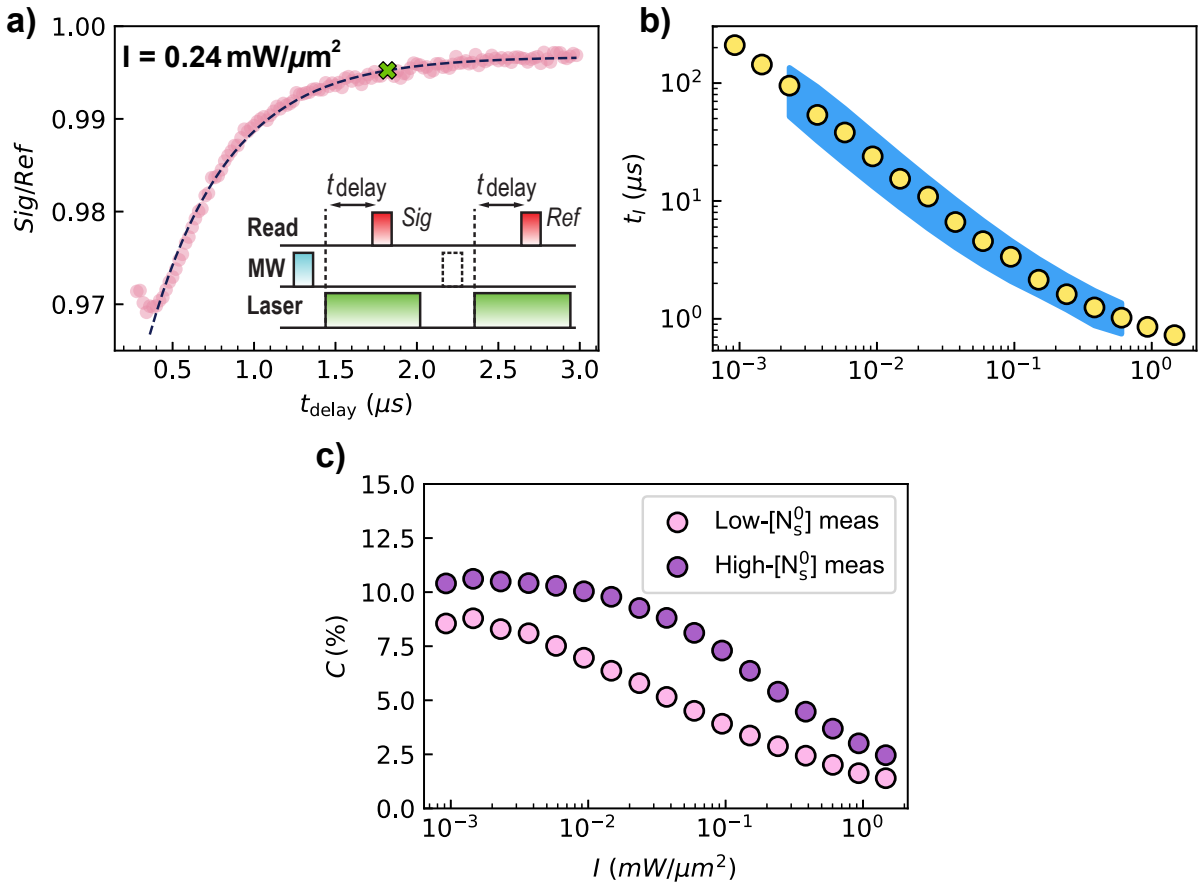


Figure 4.6: NV spin polarization dynamics, initialization time, and PL contrast as functions of laser excitation intensity. (a) Measured NV spin polarization dynamics and schematic measurement protocol (inset). A 1-ms laser pulse (not shown) initializes the NV ensemble into the $m_s = 0$ state prior to a microwave (MW) π pulse that initializes the NV ensemble into $m_s = -1$. A laser pulse then begins to repolarize the NV spins, with PL readout (Sig) acquired after a delay time t_{delay}. The measurement protocol is then repeated without the MW π pulse, and a reference PL measurement (Ref) is acquired. Example PL contrast measurement (Sig/Ref) at a laser intensity of 0.24 mW/μm^2, shown as a function of t_{delay}. Spin initialization time t_I (green cross) is extracted by fitting an exponential decay to the PL contrast data following its peak; we set t_I as the delay at which contrast decays to $1/e^3$ of its peak value, corresponding to repolarization of $\sim 95\%$ of the ensemble NV spins. With this measurement protocol, the readout time is included in the reported initialization time. (b) Measured NV initialization time (t_I) as a function of laser intensity. The shaded blue region represents the range of initialization times predicted based on a five-level model of NV photophysics, with saturation intensities of 1 mW/μm^2 and 3 mW/μm^2 used to produce the lower and upper bounds, respectively. (c) Measured peak PL contrast for low- and high-[N_s^0] diamonds as a function of laser intensity.

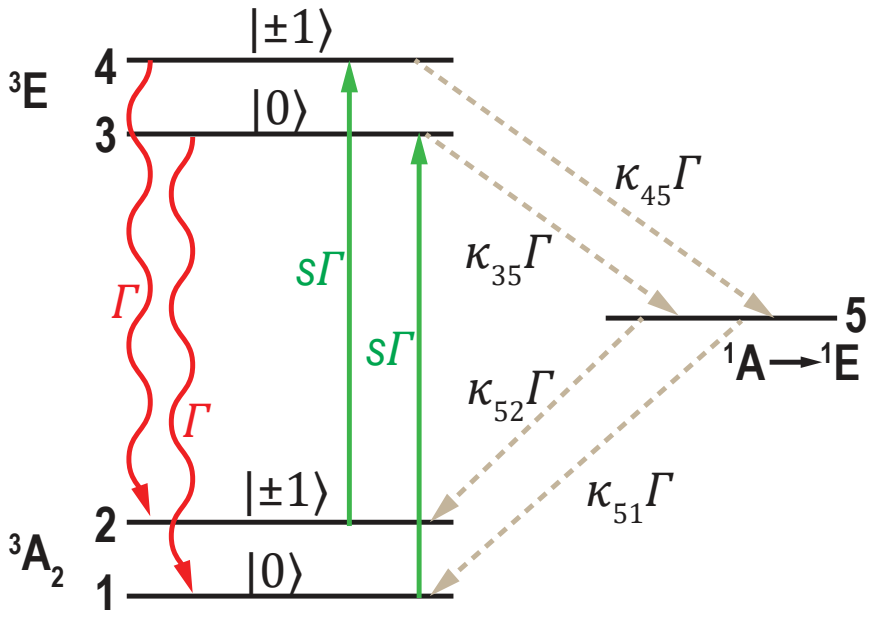


Figure 4.7: Schematic of the simplified five-level NV photophysics model used in spin polarization dynamics simulations.

states:

$$
\begin{aligned}
dn_1/dt &= \Gamma(-sn_1(t) + n_3(t) + \kappa_{51}n_5(t)) \\
dn_2/dt &= \Gamma(-sn_2(t) + n_4(t) + \kappa_{52}n_5(t)) \\
dn_3/dt &= \Gamma(sn_1(t) - (1+\kappa_{35})n_3(t)) \\
dn_4/dt &= \Gamma(sn_2(t) - (1+\kappa_{45})n_4(t)) \\
dn_5/dt &= \Gamma(\kappa_{35}n_3(t) + \kappa_{45}n_4(t) - (\kappa_{51}+\kappa_{52})n_5(t)).
\end{aligned}
\tag{4.6}
$$

The simulated PL emission rate is calculated as:

$$
R(t) = \Gamma(n_3(t) + n_4(t)), \tag{4.7}
$$

and is further processed by a low-pass filter (4th-order Butterworth with -3dB cutoff at 1.7 MHz) to mimic the effect of data acquisition hardware used experimentally (NI DAQ USB-6363).

The stimulated PL contrast is determined by normalizing the PL trace for spins initialized in the $m_s = \pm 1$ state by the PL for spins initialized in the $m_s = 0$ state. An exponential fit is applied to the PL contrast decay, with the initialization time (t_I) defined as the delay time (t_{delay}) at which the contrast decays to $1/e^3$ of its peak value, consistent with the employed experimental approach.

Simulated initialization times (t_I) closely match experimental data across a wide range of optical intensities (Fig. 4.6(b)), confirming accurate calibration of illumination conditions within the probed region.

4.6 NV charge state fraction

The observed degradation of PL contrast at higher optical excitation intensities (Fig. 4.6(c)) is primarily attributed to increased ionization of NV^- centers [138, 139]. This process increases the fraction of NV^0 and its associated background PL. A similar reduction in contrast at high optical intensity is observed in the higher-$[N_s^0]$ (~ 14 ppm) diamond sample. However, this effect is partially mitigated in the higher-nitrogen material because the increased concentration of nitrogen donors helps stabilize the NV^- charge state. Direct measurements of the NV^- charge-state fraction, ϕ, are performed using PL emission-spectrum decomposition [113, 135]. As shown in Fig. 4.8, the measured NV^- fraction decreases with increasing optical excitation intensity. The charge-state measurement is limited to intensities below approximately 10^{-1} mW/μm^2, where the relative PL contributions from NV^0 and NV^- can be reliably determined below optical saturation conditions [113]. The measured charge fraction also shows that the low-$[N_s^0]$ diamond has a substantially lower NV^- fraction than the higher-$[N_s^0]$

reference diamond (Fig. 4.8). However, the NV^0 PL emission is approximately $2.5\times$ dimmer than that of NV^- under typical optical excitation conditions [113]. In addition, the spectral filtering used in the present NV measurements, which transmits PL above $647\ \mathrm{nm}$, further suppresses the detected contribution from NV^0. The reduced NV^- fraction therefore affects the effective density of sensing centers more strongly than it contributes background PL to the observed NV spin-state contrast.

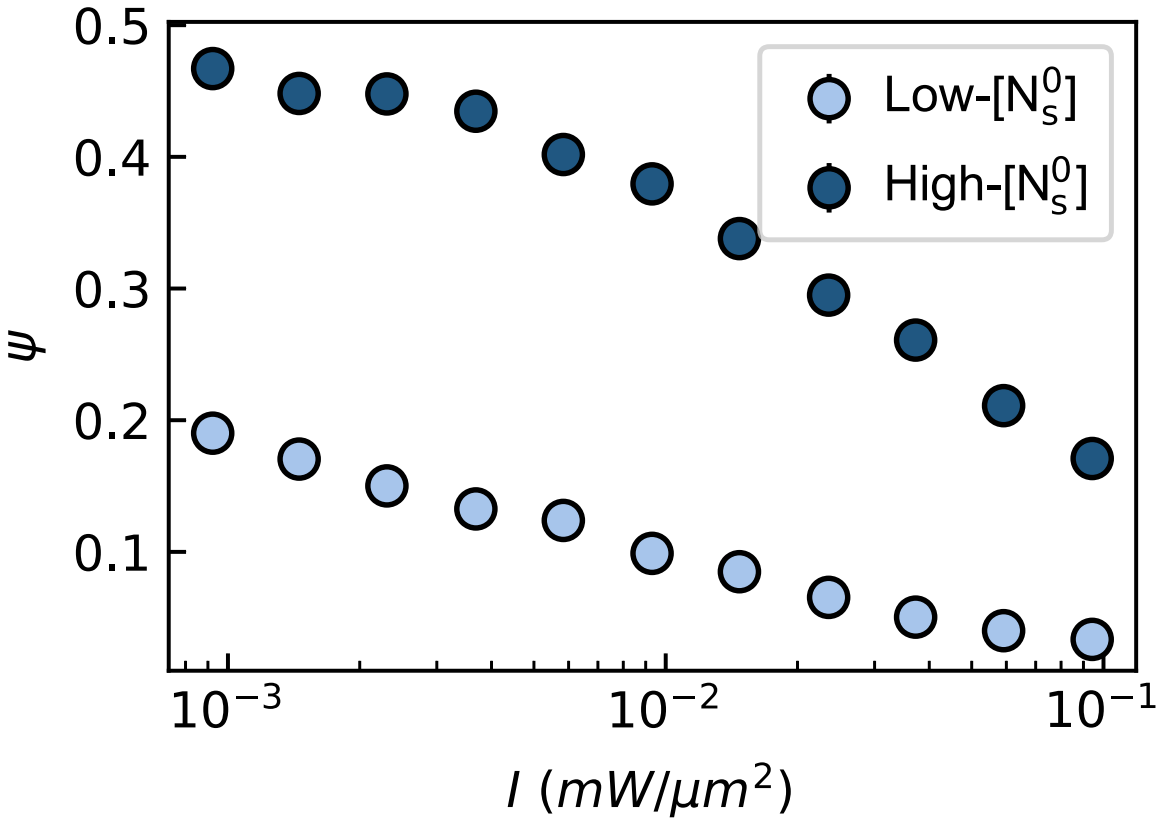


Figure 4.8: Experimentally-determined NV charge-state fraction ψ, defined as $[NV^-]/([NV^-]+[NV^0])$, as a function of laser intensity for both low- and high-$[N_s^0]$ diamond samples. NV PL spectra used for decomposition analysis (see the main text) are collected using a 550 nm long-pass filter.

4.7 Discussion

The availability of low-[N_s^0], NV-enriched, and ^{12}C-enriched diamond material with good strain homogeneity expands the material options available for NV ensemble sensing. In particular, our systematic characterization of NV spin dephasing, optical contrast, initialization dynamics, and charge-state behavior in both the low- and high- $[N_s^0]$ samples makes it possible to compare their expected photon-shot-noise-limited DC magnetic-field sensitivities over a range of optical excitation intensities.

Figure 4.9 summarizes volume-normalized DC sensitivity estimates for the two diamond types. In this comparison, the total number of sensing NV^- centers, N in Eq. 4.2, is expressed through the corresponding NV^- concentration, $[NV^-]$. At moderate and low optical excitation intensities, below $\sim 0.03\,\mathrm{mW}/\mu\mathrm{m}^2$, the low-[$N_s^0$] sample exhibits better (i.e., smaller) estimated sensitivity compared to the high-[N_s^0] sample, despite its lower NV density and reduced NV^- charge fraction. The main advantage comes from the substantially longer NV spin-dephasing time, T_2^*, which increases the fraction of each pulsed measurement spent in the active sensing interval when initialization and readout overhead are appreciable. At higher optical excitation intensities, faster initialization reduces this duty-cycle advantage, making differences in the NV^- charge fraction and effective sensing-spin density then become increasingly important in determining the estimated sensitivity.

Double-quantum (DQ) Ramsey protocols can typically further improve DC magnetic-field sensitivity by doubling the effective gyromagnetic ratio ($\Delta m_s = 2$ in Eq. 4.2) and suppressing common-mode strain-induced dephasing [13]. However, in diamond samples with good strain homogeneity,a such as the low-$[N_s^0]$ diamond studied here, or in sufficiently small sensing volumes where strain variation is limited, the advantages of DQ sensing strongly depend on the relationship between the spin

dephasing time (T_2^*) and experimental overhead time (t_O). Due to increased sensitivity to spin-bath noise, DQ dephasing times are typically about half those of single-quantum (SQ) protocols. The DQ protocol is therefore most favorable when the coherence time approaches or exceeds the optical overhead time, $T_2^* \gtrsim t_{\mathrm{O}}$, where sensitivity depends sublinearly on T_2^* while retaining a linear benefit from the doubled magnetic response. This criterion can be readily satisfied by low-[$\mathrm{N_s^0}$] diamonds, whereas the shorter dephasing times ($\lesssim 1\,\mu$s) characteristic of higher-[$\mathrm{N_s^0}$] diamonds typically result in linear scaling of sensitivity with T_2^*.

Our comprehensive assessment of NV PL contrast, charge fraction, spin dephasing time, and photon-shot-noise-limited sensitivity provides guidance to the community in selecting suitable NV-diamond materials for specific sensing applications and operational regimes. In particular, the projected sensitivity advantage of low-$[\mathrm{N_s^0}]$ diamond at moderate-to-low optical intensity makes it attractive for sensing applications constrained by size, weight, power, and cost (SWaP-C), as well as for biological measurements where reducing intense green illumination can help limit phototoxicity [140]. Additionally, the current limitations of low-[$\mathrm{N_s^0}$] diamonds under high-intensity optical illumination motivate the continued need to optimize $\mathrm{NV^-}$ charge-state stability in future diamond material engineering efforts.

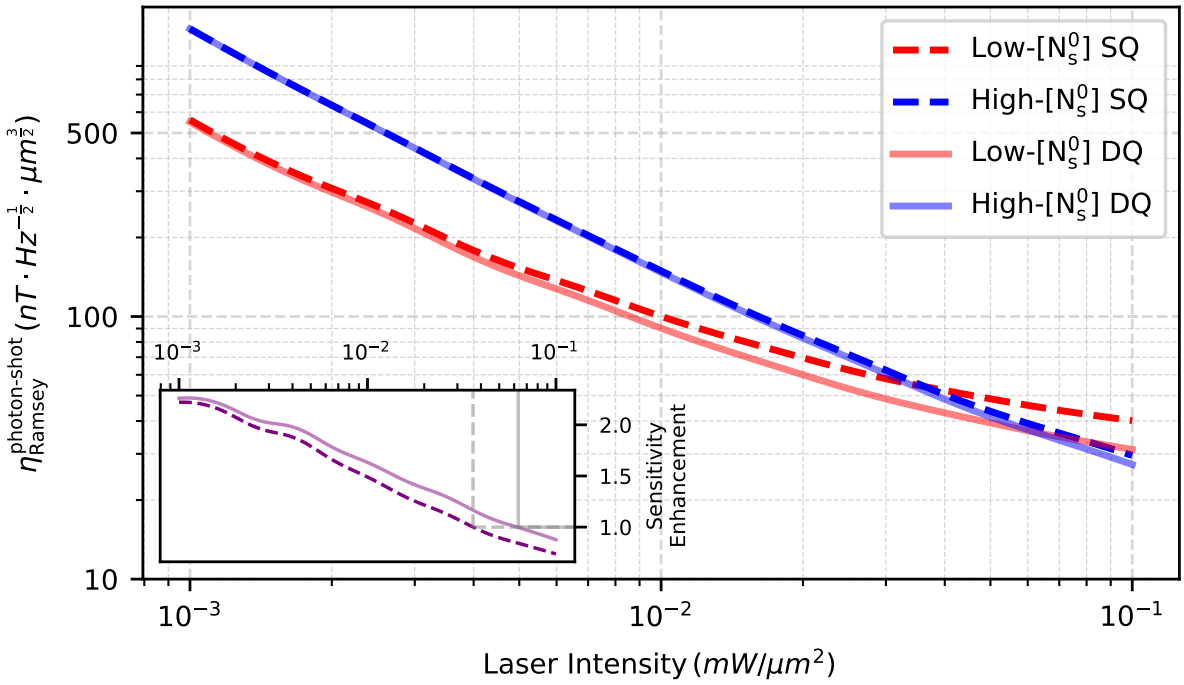


Figure 4.9: Estimated volume-normalized photon shot-noise limited DC magnetic field sensitivity for low- and high-[N_s^0] diamond samples as functions of optical excitation intensity. We assume a photon emission rate of 30 kcps per single NV at 1 mW/μm^2 illumination intensity. Sensitivity-relevant parameters (see Eq. 4.2), including the Ramsey dephasing times (T_2^* for both single- and double-quantum measurements), spin-state-dependent PL contrast (C), and experimental overhead time (t_O), are directly measured as functions of laser intensity. Nitrogen concentration and NV charge-state fraction, experimentally determined using the UV-Vis absorption method (Section 4.2.2), are used to estimate the number of NV$^-$ centers contributing to sensing at a given laser intensity. The inset shows the ratio of estimated sensitivity between high-[N_s^0] and low-[N_s^0] samples as a function of optical excitation intensity, for SQ (solid) and DQ (dashed) Ramsey protocols, highlighting the relative advantage of low-nitrogen material at moderate to low optical excitation intensities.

Chapter 5: Conclusions and Outlook

5.1 Summary of dissertation contributions

This dissertation explored methods to advance nitrogen-vacancy (NV) quantum sensors in diamond in order to unlock more microscopic information across different applications. The three main projects span high-speed and sensitive magnetic imaging, quantum readout of particle-induced damage track, and diamond material engineering for improved magnetometry. Through these projects, I used a combination of experimental demonstrations and modeling techniques to benchmark the performance of NV systems and identify the limitations that must be further addressed to enable applications to weaker or more demanding signals.

Chapter 2 developed the pulsed-Ramsey quantum diamond microscope (QDM) for dynamic magnetic-field imaging. This project was primarily motivated by the need to combine advanced pulse-sequence techniques with fast imaging capability, in order to benchmark the best magnetic sensitivity performance using NV ensembles while enabling measurements of transient magnetic fields, such as those from electrically active cells. I addressed the central challenge of integrating DQ 4-Ramsey sensing, spin-bath driving, lock-in camera acquisition, and differential measurement methods into a stable QDM platform capable of long-term averaging. The achieved volume-normalized sensitivity $\approx 25\ \mathrm{nT} \cdot \mathrm{Hz}^{-1/2} \cdot \mu\mathrm{m}^{3/2}$ represents state-of-the-art performance for NV magnetic imaging [141]. In addition, the demonstration of sub-millisecond temporal resolution and picotesla-scale magnetic-field imaging represents a major advance for QDM development, moving beyond previous regimes of seconds-scale temporal resolution and nanotesla-scale resolvable fields.

Chapter 3 used NV centers in a role different from magnetic sensing. Here, NV centers func-

tioned as optical reporters for localizing particle-induced damage tracks in diamond, as model benchmarks for defect evolution during annealing, and potentially as nanoscale probes of the final track morphology. The key motivation was to experimentally advance a previous proposal for using NV quantum diamond as a directional dark-matter and rare-event detector [24]. I co-led the experimental study of single carbon-ion implantation into HPHT diamond. I performed the NV yield measurements and analysis, and developed an implantation-to-annealing simulation pipeline for NV yield modeling. I also interpreted the physical implications of DFT results from co-authors and developed a correction model that brought the simulation into close agreement with experiment, crucially without free-parameter fitting. Using this experimentally grounded simulation pipeline, I showed that head-tail directional information is retained after annealing and co-developed machine-learning methods to improve directional-information recovery. I further showed that the implantation-created NV centers have spin coherence compatible with future nanoscale imaging. These results are useful not only for understanding which techniques could constitute a multiscale directional detector readout pipeline, but also for providing atomic-scale defect distributions as a starting point for modeling other damage-track properties, such as crystal strain.

Chapter 4 examined a new quantum diamond material developed in collaboration with Element Six. The motivation was to improve our understanding of how to synthesize quantum diamond with favorable spin and material properties, and how different diamond materials compare in practical sensing conditions. Magnetic sensitivity is determined not only by strain, NV density, and spin coherence time, but also by factors such as optical excitation intensity, which affects PL rate, spin contrast, NV charge state, and initialization time. Low-nitrogen, $^{12}\mathrm{C}$-enriched CVD diamonds were synthesized and characterized, showing long ensemble spin-dephasing times T_2^* near the nitrogen-spin-bath limit, even for single-quantum Ramsey measurements. This result indicates good control of strain homogeneity in

the material. Measurements of charge-state fraction, PL contrast, and initialization overhead as functions of optical intensity were then combined to estimate practical photon-shot-noise-limited Ramsey sensitivity for both the new low-nitrogen diamond and previously developed higher-nitrogen diamond samples. This study showed that low-nitrogen diamond can be advantageous for pulsed magnetometry under moderate-to-low optical excitation, as expected from the increased sensing duty cycle. At the same time, charge-state stability at high optical intensity remains an important practical limitation, where higher-nitrogen diamond retains an advantage.

In the following outlook sections, I consider the implications of the results achieved in this dissertation for biological magnetic-field imaging and for NV-based reconstruction of damage tracks toward directional rare-event detection.

5.2 Outlook for biomagnetic-field imaging

5.2.1 Prospects for biomineralization and electrically active cell imaging

There are several canonical biological magnetic targets that motivate the development of broadband DC magnetic-field imagers. These sources are often magnetic particles, either synthetic particles that bind to biological targets or magnetic structures that are directly biosynthesized by cells. In addition, action potentials that propagate through electrically active cells are also of interest. In particular, the ability of a QDM to noninvasively resolve current flow in a coupled cellular network has been a long-standing motivation for developing more sensitive diamond magnetic imagers [51]. Below, I consider how the dynamic QDM developed in Chapter 2 meets the requirements for these classes of biological measurements.

A first example of potential interest is the study of iron-loaded bacterial encapsulin compart-

ments expressed in genetically engineered mammalian cells [65]. Magnetization in such $30\ \mathrm{nm}$-diameter shell structures has been shown to produce useful T_2 contrast in MRI, and further cellular engineering to enhance iron loading and screen for superparamagnetism could support optimization of such MRI contrast agents. The primary attractive feature of the QDM developed in this dissertation is its ability to image picotesla-scale magnetic-field patterns. Assuming a volumetric magnetization similar to that of commercial superparamagnetic iron-oxide nanoparticles [50], an engineered $30\ \mathrm{nm}$-diameter compartment is expected to produce a magnetic field of approximately $600\ \mathrm{pT}$ at a $10\ \mu\mathrm{m}$ standoff, which could be resolved by the Ramsey-based QDM within a reasonable averaging time, approximately $100\ \mathrm{s}$ of active signal acquisition. A remaining technical challenge is that the present QDM relies on differential measurements to suppress slowly varying camera-pixel baseline drift. This protocol is naturally suited to controllable current sources, where the magnetic signal can be reversed or gated, but it is less directly applicable to static magnetic objects such as intracellular magnetic particles. This is a technical limitation of the present imaging implementation rather than a fundamental physics limitation. Possible routes forward include improved stabilization and correction of camera-pixel baselines, which is primarily an engineering issue specific to heliCam, and modulation of the measurement geometry. For the latter approach, magnetic particle-containing cells could be embedded on a sample holder that can be translated vertically relative to the diamond surface, thereby converting a static magnetic field into a differential signal. I have performed initial fabrication validation of such sample holders designed to interface with the diamond sensor; additional details are provided in Appendix D.

A second possible application idea for the Ramsey QDM is a high-throughput biomagnetism screening tool, such as searches for biogenic magnetite particles related to vertebrate magnetoreception [66] or population-based measurements. The localization of cells containing magnetite has been chal-

lenging, as organism-scale volumes of candidate tissues need to be screened. Previous identification of magnetite in salmon using magnetic force microscopy [67, 68] suggests an order-of-magnitude dipole moment of 10^{-16} A m^2. Assuming magnetite in other model animals has similar dipole moments, the magnetic field magnitude immediately outside a 10 µm cell is expected to be tens of nanotesla, which can be imaged by the present QDM with SNR $\approx$ 10 in a few minutes of signal averaging. We estimate a tissue area magnetite screening rate of a few mm^2 per hour using the present QDM, assuming a 10 µm-thick tissue. Although these samples are again not naturally amenable to signal modulation, the expected magnetic-field amplitude at a conservative standoff is already above the measured non-differential spatial noise floor of about 2 nT. The QDM may therefore provide a useful tool for screening candidate tissues or cell populations for rare, strongly magnetic biological inclusions before more targeted nanoscale or single-cell characterization.

Beyond imaging of static magnetic fields, the Ramsey QDM's capability to measure temporal dynamics can benefit studies of electrically-active cells. For example, spatially resolved measurements of electrical currents in cardiac tissue – via imaging of induced dynamic magnetic fields – could inform biophysics modeling at the cellular level [142–144] or aid in pharmaceutical studies [145]. Previous measurements of the heart have recorded nanotesla-scale magnetic fields at $>$ 100 µm standoff distances [54, 146]. As the NV sensing layer can be brought to few micrometer standoff distances from live bio-samples [51], and externally controlled currents ($\gtrsim$ 1 µA) can be applied to cardiac tissue [147], we expect the Ramsey-based QDM may find applications in mapping the microscopic electrical properties of such tissue [73, 74].

The sub-millisecond time resolution of the present QDM is sufficient for mapping the dynamic magnetic fields produced by neuronal action potential (AP) currents [75]. However, resolving the expected $\lesssim$ 1 nT AP magnetic signal [76] with a $\sim$ 0.5 ms temporal resolution requires $\gtrsim$ 4 k averages

($\gtrsim 2$ day of wall-clock time) to reduce the spatial magnetic noise floor $\sigma_{spatial}$ to $\lesssim 0.8$ nT using 5×5 pixel binning ($\sim 10 \times 10 \times 10\,\mu m^3$ sensing volume) with the present sensitivity and technical overhead time. While repeated AP excitation of cultured neurons $in\ vitro$ is feasible (continuous stimulation at 4 Hz over days has been reported [77]), the long experimental time poses a significant challenge. For the QDM to become an attractive tool for this application, at least an order-of-magnitude improvement in volume-normalized sensitivity is likely required.

5.2.2 Prospects for pollen tube studies

After the development of the dynamic QDM, I spent some time looking for biological systems where magnetic imaging could provide information that is otherwise difficult to obtain with existing methods. Prof. José Feijó of UMD introduced to me and my colleagues the idea of applying magnetic imaging to growing pollen tubes. This remains an early-stage idea for the NV research community, but it is a particularly interesting one because pollen tubes are robust plant cells and has large ionic fluxes during rapid tip growth. I worked with Custódio de Oliveira Nunes from Prof. Feijó's group to test whether pollen grains could be integrated with the Ramsey QDM platform and whether germination could proceed under experimental conditions involving laser illumination and microwave spin control. These system-integration results are described in Appendix E. Here, I discuss why magnetic imaging could be useful from a pollen-tube perspective, mostly written through the lens of a physicist trying to understand the biological problem.

When a pollen grain germinates, it forms a thin tube that grows directionally to deliver sperm cells to the female reproductive tissue; and pollen tubes are widely used as a model system for studying apical growth and cell polarity. The dimensions and growth rate depend strongly on species. For exam-

ple, Arabidopsis pollen tubes have diameters of approximately $5\ \mu\mathrm{m}$, about two to three times smaller than tobacco or lily pollen tubes [70]. On the other hand, lily pollen tubes can grow at $6-10\,\mu\mathrm{m/min}$ [72], while Arabidopsis pollen tubes are typically smaller and slower [70]. From a bioelectric perspective, a key feature of pollen tubes is that ion channels and pumps are distributed in a spatially polarized way along the tube, organizing large ionic fluxes [69]. A hypothesis of interest based on this picture is that there could exist a cytosolic standing electric field along the tube.

The cytosolic electric field can be written as

$$E_{\mathrm{cyt},x}(x) = -\frac{\partial \phi_{\mathrm{cyt}}(x)}{\partial x}, \tag{5.1}$$

where ϕ_{cyt} is the electric potential inside the cell and x is the coordinate along the pollen tube. Directly probing E_{cyt} or ϕ_{cyt} is challenging. For example, a microelectrode array requires penetrating the cell wall, which can cause catastrophic bursting in growing pollen tubes because of their high internal pressure. Extracellular bioelectric fields and ion currents around growing pollen tubes have been widely observed, and recent work using voltage-sensitive dyes further measured a gradient of membrane potential from the tip to the shank of the pollen tube [148]. However, a complete picture of the internal electric-potential landscape and the associated cytosolic electric field has not yet been established.

Magnetic-based measurement could serve as a complementary approach by probing the ionic currents associated with this bioelectric state. Earlier vibrating probe measurements demonstrated that growing lily pollen tubes sustain a steady transcellular current – current leaves the pollen grain and tube base and enters much of the growing tube of order a few hundred picoamperes [72]. These measurements also found that the current can contain both a steady component and periodic pulses superimposed on that baseline. However, the path of the internal return current was inferred from ex-

tracellular measurements, thus not directly imaged. This leaves an important unresolved measurement problem. In the previous picture [72], current may return through the cytoplasm over an extended region of the pollen tube.However, this interpretation has been questioned, since the spatial distribution of ion transport activity may instead favor a more localized circuit in which net current enters near the tip and returns or dissipates over the apical or subapical region, rather than along most of the shank [149]. A spatially resolved magnetic image could directly test these hypotheses by mapping the net ionic current distribution along the pollen tube. Ohm's law,

$$\mathbf{J}_{\mathrm{cyt}} = \sigma_{\mathrm{cyt}} \mathbf{E}_{\mathrm{cyt}}, \tag{5.2}$$

then connects the measured current density to the underlying cytosolic electric field, once an effective cytosolic conductivity is assumed. From an electrophysiological modeling perspective, the conductivity σ_{cyt} in the apical region and along the shank is itself an important missing quantity. However, this parameter cannot be directly measured from magnetic-field imaging alone. High-spatial-resolution magnetic measurements could also constrain the current distribution and potentially enable cross-sectional current profiling, similar in spirit to NV imaging of current flow in graphene [46, 150].

These are the specific utilities I see that the QDM can bring to in this problem. In terms of experimental feasibility, however, the measurement remains beyond the demonstrated sensitivity of the Ramsey QDM developed in this thesis. The total transcellular current inferred in lily pollen tubes is approximately $200\ \mathrm{pA}$, whereas the smallest current resolved in Chapter 3 is approximately $9\ \mathrm{nA}$. This leaves a gap of roughly a factor of 50 in current sensitivity. In a recent discussion, Prof. Feijó noted that such current estimates can depend sensitively on the measurement geometry: the use of a membrane in the previous study introduces a finite standoff from the pollen-tube surface, and the membrane itself

may also perturb the current path near the tip and shank [149]. These effects could change the inferred current amplitude by an order-unity factor, possibly reducing the sensitivity gap by roughly a factor of three. This highlights the value of direct experimental probing, though QDM imaging of pollen-tube currents would still remain challenging. In the near term, mapping out the fuller opportunities of NV magnetic imaging for pollen tube studies could itself be a useful vision statement for the quantum and plant biology communities. In parallel, QDM development along the lines of improved magnetic sensitivity, better low-frequency camera stability, and identifying biological conditions that enhance the current signal would all be needed before the QDM can become a practical tool for pollen tube bioelectricity.

5.3 Outlook for NV-based damage-track reconstruction

The simulation framework developed in Chapter 3 can be extended to assess the applicability of NV-based directional detection in HPHT diamond at lower recoil energies. At 100 keV, we performed annealing simulation for 84 tracks. The model predicts a mean of 13.8 ± 0.7 NV centers per track, formed from an average of 217 initial vacancies, corresponding to a yield of $\sim$6%, similar to the 800 keV case. However, the reduced number of NV defects lowers head-tail discrimination efficiency to 37% using the counting-based asymmetry metrics at $\lesssim$5% FPR. Applying the same ML analysis improves the efficiency to 96% at the same FPR (see Fig. 5.1). At recoil energies of $\sim$10 keV, assuming a similar yield ratio, an average of $\sim$4 NV centers is expected to form from $\sim$50 vacancies, making directionality reconstruction based solely on NV positions challenging. Nevertheless, the NV PL signal remains above background ($\lesssim$ 2 NVs per confocal spot), enabling event localization at the micron scale.

These results suggest that a hybrid detection strategy will likely be required at lower recoil energies. While NV centers can provide an optical marker for event localization, a significant fraction of vacancies remains after annealing as divacancies or higher-order vacancy clusters, while another large fraction recombines with interstitials. In the simulations discussed in Chapter 3, approximately 40% of vacancies remain in vacancy-related defect clusters and approximately 50% recombine with interstitials. The remaining vacancy-related defects can generate local strain fields that may be detected using high-resolution X-ray diffraction techniques [151]. Based on a previous study [36], it is estimated that for a typical 10 keV recoil, strain fields on the order of 2×10^{-4} can be generated within $\sim$30 nm of the damage track, which is within experimentally accessible limits. Accounting for vacancy recombination during annealing reduces this estimate by approximately a factor of two, but the resulting strain fields remain measurable. More detailed modeling of strain fields from annealed defect configurations will be required for quantitative predictions. Overall, these considerations indicate that combining NV-based readout with nanoscale strain-based imaging provides a promising pathway for investigating directional detection of low-energy recoil events.

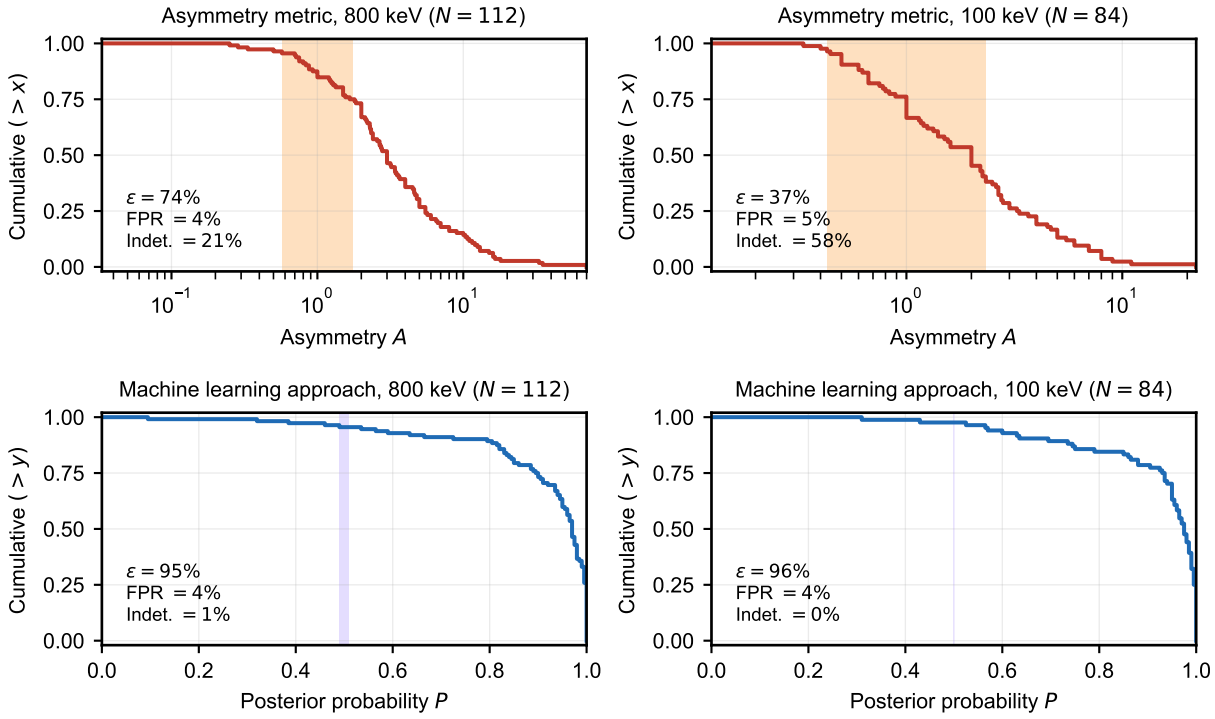


Figure 5.1: Cumulative plots of NV track classification using the counting-based asymmetry metric and machine learning approach. Top row: cumulative distributions of the asymmetry metric for 100 keV NV tracks. Bottom row: cumulative distributions of the machine learning posterior probability for the same tracks. Shaded regions indicate indeterminate ranges chosen to give a $\lesssim$5% false positive rate. The machine learning model trained on 20,000 100 keV vacancy tracks..

Appendix A: Appendix for Dynamic QDM

This appendix is adapted from J. Tang, Z. Yin, C. A. Hart, J. W. Blanchard, J. T. Oon, S. Bhalerao, J. M. Schloss, M. J. Turner, and R. L. Walsworth, "Quantum Diamond Microscope for Dynamic Imaging of Magnetic Fields" [42].

A.1 Estimate of NV-NV interaction limited Ramsey dephasing time

Efficient spin-bath driving should result in the per-pixel DQ T_2^* approaching the NV-NV interaction limit due to the high NV density in the employed diamond. Ignoring insignificant per-pixel T_2^* limitations from the 0.005% ^{13}C, $T_2^*\{^{13}\mathrm{C}\} = 200\,\mu\mathrm{s}$, and the $\sim 0.013\,\mu\mathrm{T/\mu m}$ bias magnetic field gradient, $T_2^*\{\mathrm{bias}\} = 850\,\mu\mathrm{s}$, we estimate the NV-NV interaction limit as [1]

$$\begin{aligned}\frac{1}{T_2^*} \approx & \frac{1}{T_2^*\{\mathrm{NV}\}_{||}} + \frac{1}{T_2^*\{\mathrm{NV}\}_{\nparallel}} \\ \approx & A_{\mathrm{NV}_{||}}[\mathrm{NV}_{||}] + \zeta_{\nparallel} A_{\mathrm{NV}_{\nparallel}}[\mathrm{NV}_{\nparallel}].\end{aligned} \tag{A.1}$$

Here $[\mathrm{NV}]_{||}$ is the concentration of NV spins in the same crystal axis group used for sensing and $[\mathrm{NV}]_{\nparallel}$ is the concentration of NV spins in other groups. Since the bias magnetic field is aligned to one of the four NV crystal axes, $[\mathrm{NV}]_{||} = \frac{1}{3}[\mathrm{NV}]_{\nparallel} = 0.59\,\mathrm{ppm}$ [58]. The constants $A_{\mathrm{NV}_{||}} = 0.247\,\mu\mathrm{s}^{-1}\cdot\mathrm{ppm}^{-1}$ and $A_{\mathrm{NV}_{\nparallel}} = 0.165\,\mu\mathrm{s}^{-1}\cdot\mathrm{ppm}^{-1}$ [1] characterize the dipolar interaction strength for pairs of NV spins in the same and different groups, respectively. $\zeta_{\nparallel}$ is a dimensionless factor of order unity accounting for imperfect initialization of $\mathrm{NV}_{\nparallel}$ into the $|m_s = 0\rangle$ state. Experimental determination of $\zeta_{\nparallel}$ is complicated by unequal laser polarization projected among the three $\mathrm{NV}_{\nparallel}$ axes

in a total internal reflection excitation scheme. As an estimation for the NV-NV interaction limited T_2^*, we provide the results for perfect ($\zeta_{\nparallel} = 0$) and partial ($\zeta_{\nparallel} = 0.5$) initialization cases, thus $T_2^*(\zeta_{\nparallel} = 0) = 6.6\,\mu s$ and $T_2^*(\zeta_{\nparallel} = 0.5) = 3.4\,\mu s$. In addition, accounting for the $2\times$ sensitivity to the magnetic noise in a DQ scheme ($T_2^*(\mathrm{DQ}) = \frac{1}{2}T_2^*$) [57], we finally expect the NV-NV interaction limited T_2^* using DQ measurement to be approximately 2.5 µs, the average of $T_2^*(\zeta_{\nparallel} = 0, \mathrm{DQ})$ and $T_2^*(\zeta_{\nparallel} = 0.5, \mathrm{DQ})$, which is reflected in the results shown in main text (Sec. 2.4). Experimental per-pixel T_2^* with and without spin-bath driving are extracted by fitting the measured Ramsey fringes to a sum of oscillations with a common Lorentzian exponential decay [13]. The uncertainties reported in the main text (Sec. 2.4) for the median per-pixel T_2^* are estimated from fitting uncertainties using the covariance matrices returned by the PYTHON `scipy.optimize.curve_fit` function.

A.2 Estimate of shot noise and quantization noise

The photon shot-noise-limited magnetic sensitivity for NV Ramsey magnetometry (either SQ or DQ) is given by [1]

$$\eta_{shot} = \frac{\hbar}{\Delta m_s g_e \mu_B} \frac{1}{C e^{(-\tau/T_2^*)^p} \sqrt{N}} \frac{\sqrt{t_{I,R} + \tau}}{\tau}, \tag{A.2}$$

where $g_e \approx 2.003$ is the NV center's electronic g factor, μ_B is the Bohr magneton, $\hbar$ is the reduced Planck constant, Δm accounts for the difference of m_s states used for sensing, C is the optical measurement contrast at Ramsey free evolution time $\tau = 0$, T_2^* is the NV spin dephasing time, p is a parameter used to describe the Ramsey envelope decay shape, N is the average number of photons collected per measurement, and $t_{I,R}$ is the combined initialization and readout time during the Ramsey sequence. We use the following parameters to give an estimate of η_{shot} for a single QDM pixel using

the present NV-diamond sensor. In a DQ sensing scheme, $\Delta m_s = 2$. By focusing the total NV fluorescence within the field of view onto a photodiode, we experimentally determine the ensemble-averaged Ramsey oscillation parameters to be about $C = 3\,\%$, $T_2^* = 2.3\,\mu s$, and $p = 1$. We use $t_{I,R} = 7.04\,\mu s$ and $\tau = 2.29\,\mu s$, typical for a QDM magnetic imaging experiment, and estimate $N = 4 \times 10^4$ photons are collected in a single camera pixel (detailed in the next paragraph) during the $2.15\,\mu s$ camera readout window. Using these values, the estimated photon shot-noise-limited per-pixel magnetic sensitivity is $1.7\,\mathrm{nT}/\sqrt{\mathrm{Hz}}$.

As discussed in the main text (Sec. 2.3.3) camera differential measurement section, the camera accumulates a differential photon number result ΔN between two successive internal exposures. When outputting an external frame, ΔN is additionally converted to the camera device unit (DU) via a conversion factor (CF), giving $\mathrm{DU} = \mathrm{CF}\Delta N$. To calibrate CF and infer the photon number N collected by a single pixel during the $2.15\,\mu s$ camera exposure window, we perform two sets of measurements to extract the photon noise and signal scaling behavior as a function of the camera readout duration t. For the noise measurement, we vary t while ensuring the readout times for two successive camera internal exposures are identical, thus the mean $\Delta N = 0$. However, the characteristic standard deviation (std) in the device unit for each pixel scales as

$$(\mathrm{std})^2 = A_{noise} t = (\mathrm{CF})^2 2N, \tag{A.3}$$

where A_{noise} is the per-pixel noise scaling slope factor to be determined. For the signal measurement, we vary the exposure time difference Δt between two successive camera internal exposures, thus the

signal S in terms of the camera device unit is

$$S = A_{signal}\Delta t = \mathrm{CF}\Delta N, \tag{A.4}$$

where A_{signal} is the per-pixel signal scaling slope factor to be determined. Further recognizing that $\Delta N/N = \Delta t/t$, the CF can be deduced by

$$\mathrm{CF} = \frac{A_{noise}}{2A_{signal}}. \tag{A.5}$$

For the camera operating at unity gain and an internal exposure rate matching the Ramsey sequence readout intervals, we use the measured photon noise and signal scaling behavior to find an average $\mathrm{CF} \approx \frac{1}{446}$ across the field of view. By further recording the DU as a function of Δt, an estimated photon number of $N = 4 \times 10^4$ is collected in a single pixel after a 2.15 µs readout duration, for typical QDM operating parameters (detailed above). We note that N is technically the photoelectron number, which differs by a multiplicative factor from the actual photon number collected by a camera pixel (about 70 % optical fill factor and 90 % quantum efficiency around the NV fluorescence wavelengths for the Heliotis camera with micro-lenses).

The quantization noise of the Heliotis camera includes electronic and digitization noise. For each external frame, the noise is dependent on the cycles (N_c) spent accumulating differences between successive internal exposures, and can be estimated in DU as [152]

$$\sqrt{0.81 + 0.16N_c}. \tag{A.6}$$

Converting the estimate in DU to equivalent magnetic noise and including the contributions from

both the I and Q readout channels (see main text camera acquisition method section, Sec. 2.3.1), the estimated quantization-noise-limited per-pixel magnetic sensitivity is $2.7\,\mathrm{nT}/\sqrt{\mathrm{Hz}}$. After adding photon shot-noise and camera quantization noise in quadrature, the combined estimate of per-pixel magnetic sensitivity is $3.2\,\mathrm{nT}/\sqrt{\mathrm{Hz}}$. As discussed in the main text (Sec. 2.4), the measured median per-pixel magnetic sensitivity is $4.1(1)\,\mathrm{nT}/\sqrt{\mathrm{Hz}}$.

In future work, the magnetic sensitivity limit due to quantization noise can be reduced for a QDM using the Heliotis camera by increasing the NV fluorescence intensity, NV spin-state readout contrast, and NV spin dephasing time. The dynamic range of the Heliotis camera is fixed (DU from 0 to 1023, 10-bit), and the camera sensor is optimized for unity gain [152]. For a given accumulation cycle N_c of internal exposures, the quantization noise is a constant in DU (see Equation A.6). Thus, if the maximally achievable slope of the magnetometry curve, which is used for converting the camera external frame output from DU to a magnetic-field value, can be increased, then the magnetic-noise contribution from quantization noise will decrease. The maximized slope of a magnetometry curve for a given N_c is proportional to both NV fluorescence intensity and spin-state readout contrast (which affect the accumulated differential photon number ΔN in a camera external frame, see above); and also to NV spin dephasing time T_2^* (which affects the optimal sensing interval). Hence, further diamond engineering — e.g., permitting the use of increased laser illumination intensity while mitigating the contrast degradation due to NV charge-state conversion [58], and also improving the NV spin dephasing time — is beneficial for mitigating the quantization-noise limit on magnetic sensitivity.

A.3 Denoising of Picotesla-Scale Magnetic Image

The NLM denoised results shown in the main text (Sec. 2.6) are reproduced in Figure A.1. For Gaussian smoothing in comparison, the measured picotesla-scale magnetic image is convolved with a Gaussian filter and its kernel size h is set to about 2.6 pixels to increase SNR of the denoised image to ~ 30. For pure additive Gaussian white noise in an image, the standard deviation of the noise reduces as $\frac{1}{h\sqrt{8\pi}}$ [153]; and we thus expect $h \approx 2.9$ pixel to reach the desired SNR — similar to the results found for NLM denoising. As shown in Figure A.1(b), the contours of Gaussian-smoothed magnetic-field patterns are qualitatively more distorted than the NLM results. A quantitative perspective to evaluate the performance of the denoising algorithms is comparing the mean square errors (MSE) [63], obtained using the Euclidean L^2 distance between the original and the denoised images. The NLM denoising method produces about 17 % smaller MSE than the Gaussian smoothing result.

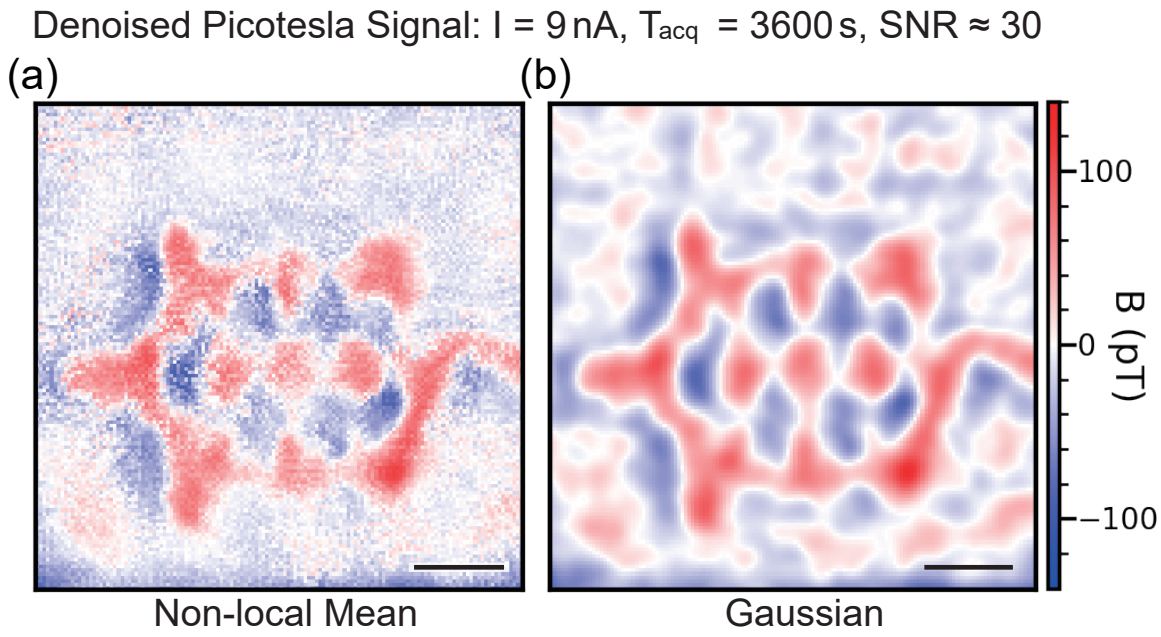


Figure A.1: Denoised QDM picotesla-scale magnetic-field images after about 3600 s signal acquisition time. Original measurement data is presented in the main text (Sec. 2.6). Here, two denoising methods, Non-local mean (NLM) and Gaussian smoothing, are applied. Both algorithms are successful at improving SNR. However, NLM denoising, as shown in (a), additionally excels at restoring fine structures (e.g., contours, patterns, and edges where the Laplacian is non-zero) from a noisy image that can be otherwise distorted by a Gaussian smoothing, as shown in (b). Scale bar: 50 µm

Appendix B: Appendix for damage track study

This appendix is adapted from Daniel G. Ang, Jiashen Tang, Maximilian Shen, Michael Titze, Gavishta Liyanage, Jordan Chapman, Chinmay Bharathulwar, Andrew T. Gilpin, Tanguy Terlier, Edward S. Bielejec, Vsevolod Ivanov, and Ronald L. Walsworth, "Multi-scale reconstruction of single-ion damage tracks in diamond via nitrogen-vacancy centers" [81].

B.1 SIMS measurement of nitrogen concentration

The kinetic Monte Carlo (KMC) annealing simulations described in the main text (Sec. 3.2.3) take the substitutional nitrogen concentration $[\mathrm{N_s^0}]$ as an input. We determined this value for the implanted sample ("Sandia 1") using two independent methods: Time-of-Flight Secondary Ion Mass Spectrometry (TOF-SIMS) and an estimate based on the NV background PL of the annealed sample. Both methods are anchored to a commercial SIMS measurement of a second HPHT diamond sample ("Sandia 2") from the same manufacturer (Element Six). We use the average of the two estimates as the simulation input, treating the spread between them as a conservative uncertainty range; as shown in main text (Sec. 3.4.2), the predicted NV yield is only weakly dependent on $[\mathrm{N_s^0}]$, so the residual calibration uncertainty does not affect the conclusions of the main text.

B.1.1 TOF-SIMS depth profiling

Negative high mass-resolution depth profiles were acquired with a TOF-SIMS NCS instrument, which combines a TOF.SIMS5 spectrometer (ION-TOF GmbH, Münster, Germany) and an in-situ

scanning probe microscope (NanoScan, Switzerland) at the Shared Equipment Authority from Rice University. The analysis field of view was $80 \times 80\,\mu\mathrm{m}^2$ (Bi_3^+ at 30 keV, 0.3 pA) with a raster of 256 by 256 along the depth profile. Charge compensation was provided by an electron flood gun, with surface potential adjusted to 20 V. The cycle time was $70\,\mu$s, corresponding to a mass range of $m/z = 0$–911 amu. Sputtering was performed with Cs^+ at 2 keV, 105 nA, rastered over $100 \times 100\,\mu\mathrm{m}^2$. The two beams were operated in non-interlaced mode, alternating one analysis cycle with ten sputtering frames followed by a 5 s charge-compensation pause. Depth calibration was performed using the interface tool in SurfaceLab v7.3 (ION-TOF GmbH), with sputter rates referenced to crater depths measured ex-situ with a DekTak stylus profilometer. For each region of interest, the analyzed area was further cropped to remove crater edges and ensure homogeneity of the extracted signal.

Quantification of the CN^- ion signal into an absolute nitrogen concentration requires an external calibration. To this end, we sent a second HPHT diamond with similar properties (Sandia 2) to EAG for commercial SIMS analysis at multiple regions of interest, and then measured the same regions on Sandia 2 with the Rice TOF-SIMS. The resulting linear EAG-Rice calibration curve (Fig. B.1b) was applied to all subsequent Rice TOF-SIMS depth profiles acquired on Sandia 1. A representative depth profile is shown in Fig. B.1a. This SIMS procedure was performed across six different regions on Sandia 1 adjacent to implanted regions, yielding an average value of $[\mathrm{N}_\mathrm{s}^0]_\mathrm{SIMS} \approx 318$ ppm.

B.1.2 NV background PL as an independent estimate

As an independent cross-check, we exploit the correlation between substitutional nitrogen concentration and grown-in NV-related PL in HPHT type Ib diamond [154]. NV background PL intensities were measured in implanted regions of Sandia 1 (i.e., the regions surrounding the NV clusters

at the ion impact sites) and at multiple regions of the EAG-measured Sandia 2 sample, the latter providing calibration. Plotting confocal PL intensity against the corresponding EAG-reported $[N_s^0]$ values yields a linear correlation (Fig. B.1c) which confirms the validity of this proxy for our HPHT material. Applying this calibration to the averaged background PL in implanted regions of Sandia 1 gives $[N_s^0]_{PL} \approx 188$ ppm.

B.1.3 Combined estimate and propagation into the simulations

The two methods yield estimates that differ by approximately $\pm 25\%$ from their mean: ≈ 318 ppm from TOF-SIMS and ≈ 188 ppm from NV background PL. We adopt the average, $[N_s^0] = 253$ ppm, as the nominal input to the KMC annealing simulations, and treat the spread between the two methods as a conservative uncertainty range of $\sim$180–320 ppm. As shown in the main text (Sec. 3.4.2), the predicted NV yield varies only weakly with $[N_s^0]$ across this range and indeed across nearly two orders of magnitude in concentration. The dominant determinants of NV yield are instead the vacancy-interstitial and vacancy-vacancy reactions discussed in the main text (Sec. 3.4.3), and the calibration uncertainty on $[N_s^0]$ does not propagate appreciably into the conclusions of the main text.

B.2 Cluster decomposition and KMC annealing simulation details

The defect distribution – consisting of both vacancies and carbon interstitials – produced by ion implantation forms an extended but spatially discontinuous damage track. The defect–defect interaction probability decreases with initial separation r (assuming defects are mobile), and in a three-dimensional continuous limit, this interaction (or "hitting") probability is described as $P(r) = a_0/r$, where a_0 is the interaction radius. Based on this scaling, we partition each damage track into smaller

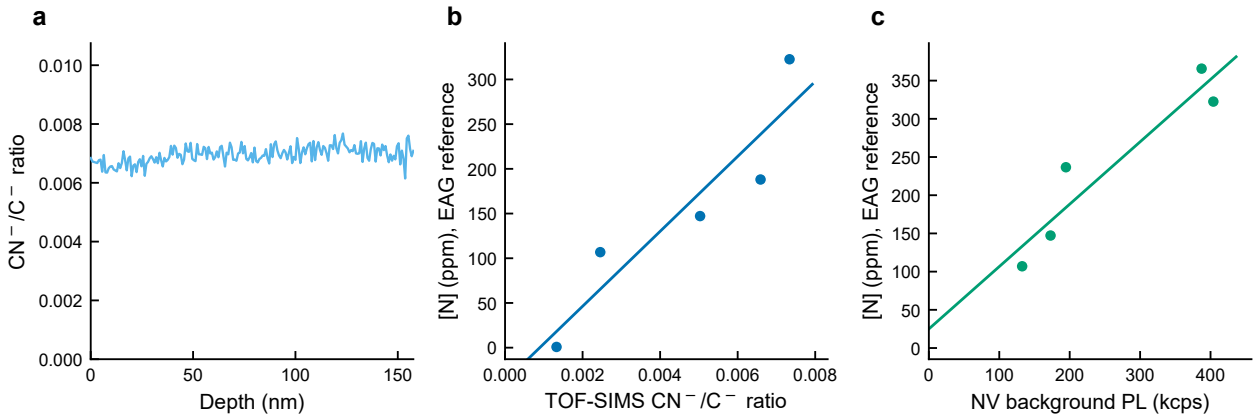


Figure B.1: **Nitrogen-concentration measurement. (a)** Representative TOF-SIMS depth profile of the CN^-/C^- ion ratio, showing a uniform signal over the profiled depth. **(b)** EAG–Rice TOF-SIMS calibration: CN^-/C^- ratio measured at Rice University plotted against the EAG reference nitrogen concentration for five regions of Sandia 2 (solid line: linear fit). **(c)** NV background PL calibration: confocal PL intensity plotted against the EAG reference nitrogen concentration for the same five regions (solid line: linear fit).

clusters and simulate each cluster independently, effectively assuming negligible interaction between clusters.

To ensure the validity of this approximation, defects belonging to different clusters are required to be separated by more than $154\,\text{Å}$. Under the assumption that defect reactions occur at nearest-neighbor (1NN) separation (i.e. $a_0 = 1.54\,\text{Å}$), this corresponds to an interaction probability of $\leq 1\%$ between any defect pairs from different clusters in the absence of other species. The presence of immobile substitutional nitrogen impurities can further suppress defect–defect interactions between clusters by providing an additional reaction channel. We quantify this effect for vacancy–vacancy interactions by directly simulating two-vacancy systems with varying initial separations r and nitrogen concentrations $[\text{N}_\text{s}^0]$, as shown in Fig. B.2(a). The interaction probability as a function of r is fitted to a power-law dependence $1.54\,\text{Å}/r^p$. For our experimentally measured nitrogen concentration of $[\text{N}_\text{s}^0] = 253\,\text{ppm}$, the expected vacancy-vacancy interaction probability at a separation of $154\,\text{Å}$ is reduced to $\sim$0.3%. In practice, vacancy–vacancy interactions and vacancy–interstitial recombination

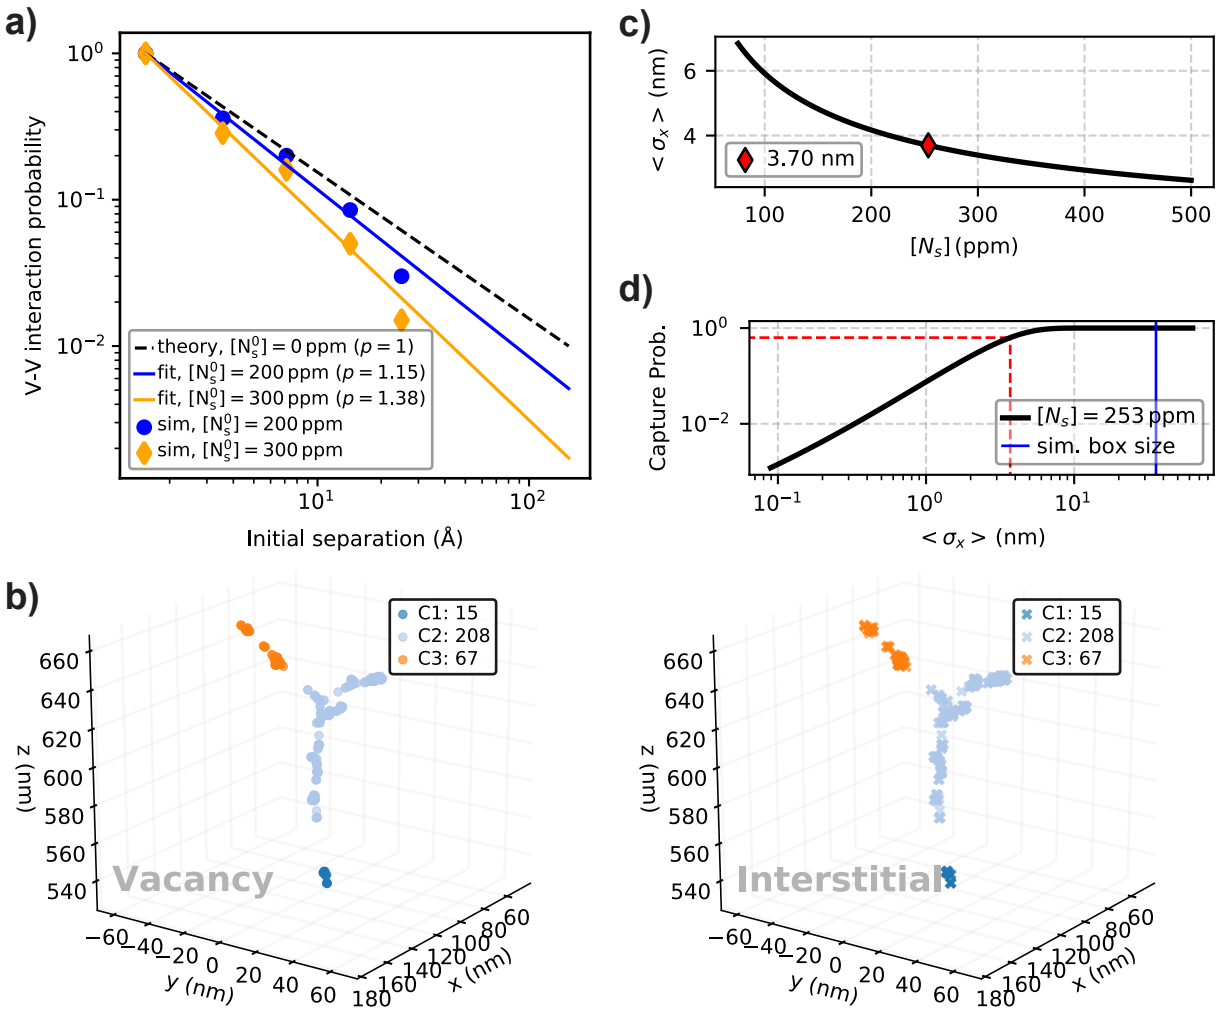


Figure B.2: Cluster decomposition and lattice domain size selection for annealing simulations. (a) Probability of vacancy–vacancy interaction as a function of initial separation r, obtained from two-vacancy simulations at different nitrogen concentrations. The dashed theoretical line is computed using the 3D hitting probability, assuming an effective interaction radius of $1.54\,\text{Å}$. Solid lines show power-law fits of the form $1.54\,\text{Å}/r^p$. For the experimental nitrogen concentration $[\text{N}_\text{s}^0] = 253\,\text{ppm}$, the interaction probability at $154\,\text{Å}$ is $\sim$0.3%, supporting negligible cross-cluster interactions. (b) Example decomposition of an 800 keV implantation damage track into spatially separated clusters. The minimum pairwise defect spacing between clusters exceeds $154\,\text{Å}$. Vacancy (left) and carbon interstitial (right) clusters are shown. The legend indicates the number of defects in each cluster. (c) Root-mean-square (RMS) displacement of a diffusing vacancy as a function of nitrogen concentration, evaluated at the point where the number of unique sites visited equals the number of available capture sites ($N_{uni} = N_{cap}$), corresponding to a capture probability of $\sim$64%. (d) Vacancy capture probability as a function of RMS displacement for a nitrogen concentration of $[\text{N}_\text{s}^0] = 253\,\text{ppm}$. For each cluster, the simulation domain is defined by padding 100 unit cells in each direction from all vacancy defects. This ensures that the probability of vacancy escape is $\lesssim 0.01\%$.

within each cluster further reduce the likelihood of cross-cluster interactions, supporting the validity of independent cluster simulations. An example of an 800 keV damage track and its decomposition into clusters is shown in Fig. B.2(b). The corresponding carbon interstitial distributions are also shown; due to their close spatial proximity to vacancies, interstitial clusters closely follow the vacancy distribution, with approximately equal numbers of vacancies and interstitials within each cluster.

After partitioning the damage track into clusters, a diamond lattice is constructed around each cluster. The simulation domain must be sufficiently large such that vacancies are unlikely to diffuse out of the box and instead undergo reactions within the domain. To ensure this condition, the box size is determined by requiring that trapping by substitutional nitrogen impurities alone is sufficient to capture diffusing vacancies.

For a vacancy undergoing a random walk on the diamond lattice, the number of unique lattice sites visited after N hops is given by [113, 155]

$$N_{\text{uni}} \approx 0.56N + 0.63\sqrt{N}. \tag{B.1}$$

For a given nitrogen concentration $[\mathrm{N_s^0}]$ (in ppm), the expected number of distinct lattice sites that must be visited before encountering a nitrogen capture site is

$$N_{\text{cap}} = \frac{10^6}{4[\mathrm{N_s^0}]}, \tag{B.2}$$

where the factor of 4 reflects the four nearest-neighbor sites adjacent to each substitutional nitrogen, each of which is a possible capture site for NV formation. The cumulative probability χ that a vacancy

is captured by nitrogen after N hops is

$$\chi = 1 - \exp\left(-\frac{N_{\text{uni}}}{N_{\text{cap}}}\right), \tag{B.3}$$

and the root-mean-square (RMS) displacement along any Cartesian direction is given by

$$\langle \sigma_{x/y/z} \rangle = \sqrt{\frac{N}{3}} \times 1.54\,\text{Å}. \tag{B.4}$$

Figure B.2(c) shows the RMS displacement as a function of nitrogen concentration at the point where $N_{uni} = N_{cap}$, corresponding to a capture probability of $\sim$64%. To ensure that vacancies remain within the simulation domain, the lattice is padded by at least 100 unit cells in each direction from any vacancy defects in the simulated cluster. At this scale, the probability of a vacancy reaching the boundary without being captured is $\lesssim 0.01\%$, as shown in Fig. B.2(d).

During the annealing simulation, the kinetic Monte Carlo (KMC) algorithm [112] samples allowed processes – including defect diffusion and reactions – according to Arrhenius rates of the form $\gamma = \nu_0 \exp(-\frac{E_a}{k_B T})$, which determine the relative frequency of each event. Here, ν_0 is taken as the Debye frequency of diamond [113], and the annealing temperature is set to $800\,^{\circ}$C. After each event is selected and applied, the global lattice configuration is updated and the simulation time advances, resulting in non-uniform time steps that depend on the selected processes. The simulation is terminated when no further events are available, ensuring that all diffusion and reaction processes are exhausted. Additional damage tracks before and after annealing are shown in Fig. B.3, with corresponding head-tail asymmetry values computed using the defect counting method. Diffusion-induced broadening during annealing reduces the asymmetry of the NV distribution relative to the initial vacancy track.

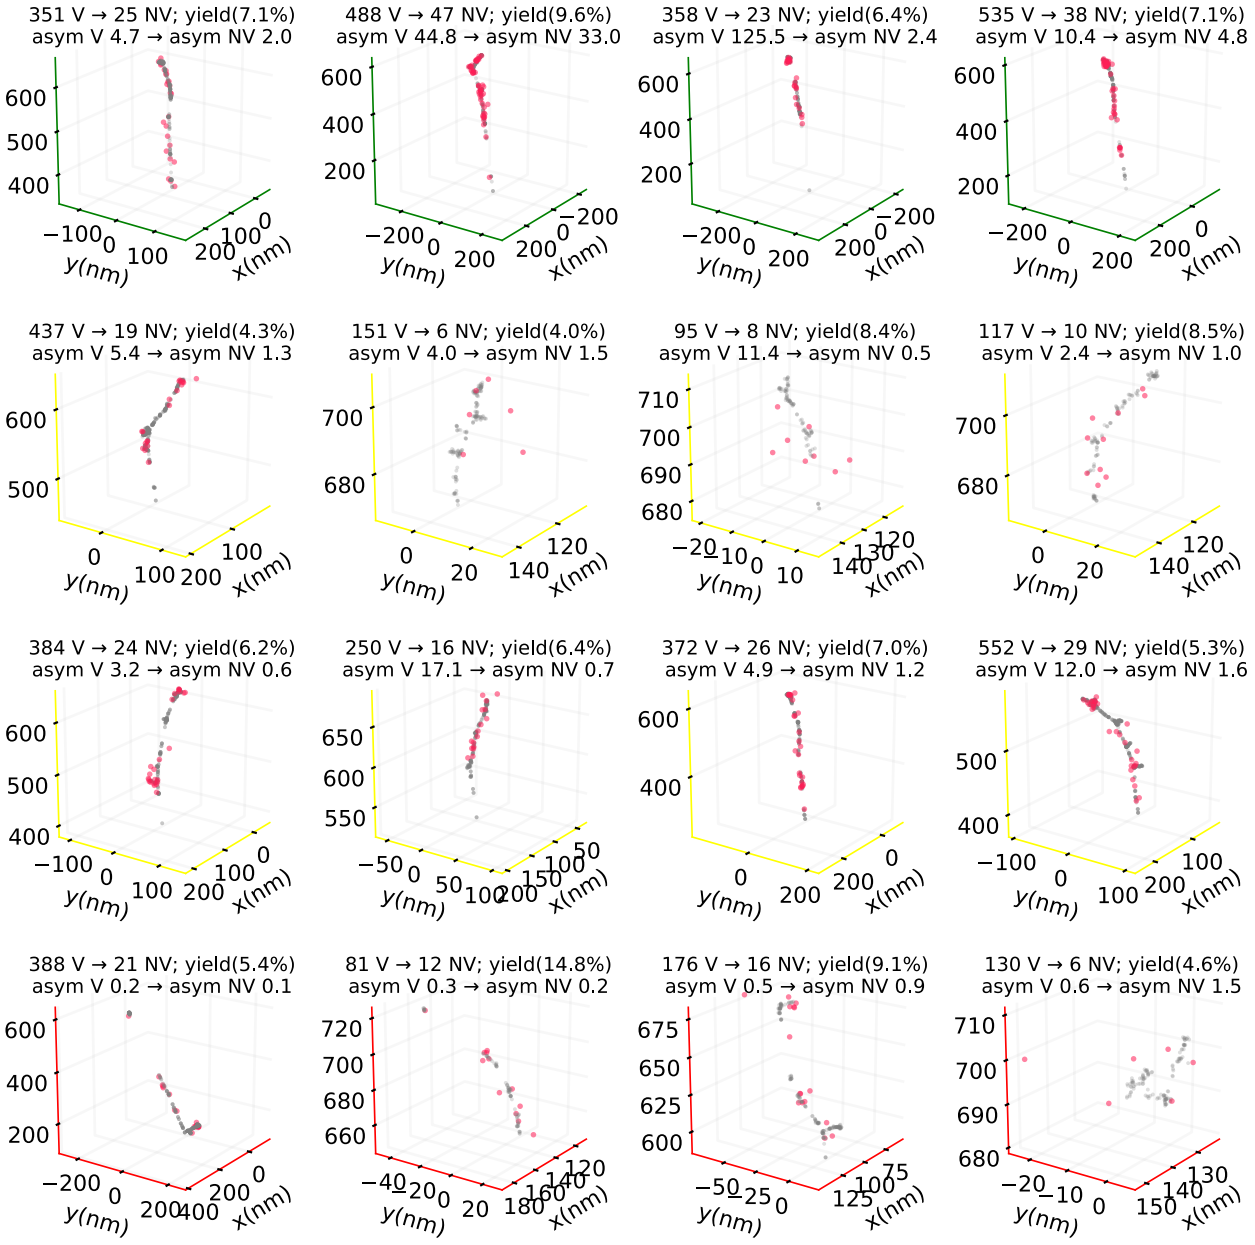


Figure B.3: Representative damage tracks before and after annealing simulations. Gray dots indicate the initial vacancy distribution, and red dots indicate the resulting NV distribution after annealing. The vacancy-to-NV conversion yield and head-tail asymmetry values, computed using a defect counting method, are shown for both vacancy and NV tracks.

B.3 NV photoluminescence (PL) analysis details

B.3.1 Data analysis procedure for confocal scans

The number of NV centers formed at each implantation site is extracted from the confocal PL data in several stages (Fig. B.4). First, 100 μm $\times$ 100 μm implantation squares on the sample containing a regular grid of sites at 6 or 10 μm pitch are imaged in a coarse, wide-field two-dimensional scan at 500 nm pixel pitch, focused at $\sim$600 nm below the surface (Fig. B.4(a)). The following procedure is then used to locate candidate implantation sites for finer scans. PL count rates are converted to a number of NV centers using a single-NV reference rate $I_1 = 25$ kcts/s, calibrated on isolated NV centers in a low-background region of the same diamond sample under identical excitation and collection conditions. Sub-regions within implantation squares containing high NV backgrounds from residual subsurface polishing damage are manually excluded, and an intensity window of 5–100 I_1 (125–2500 kcts/s) is applied to suppress the residual background ($\lesssim$ 2 NV per confocal spot; Fig. B.4(b)). See Appendix B.4 for results of additional diamond surface preparation experiments to suppress such backgrounds for future implantations.

Next, a rectangular grid of points spaced by the nominal implantation pitch (6 or 10 μm) is overlaid on the image, and an elliptical acceptance region with semi-axes equal to $\pm 3\sigma$ of the measured beam profile is placed at each grid node. Within each ellipse, the brightest pixel $\mathbf{r}_i^{\mathrm{peak}}$ is taken to mark the implantation site. The grid origin is then shifted by an offset (δ_x, δ_y) chosen to minimize the total distance between the grid nodes and their enclosed brightest pixels,

$$D(\delta_x, \delta_y) = \sum_i \left\| \mathbf{r}_i^{\mathrm{grid}}(\delta_x, \delta_y) - \mathbf{r}_i^{\mathrm{peak}} \right\|, \tag{B.5}$$

where $\mathbf{r}_i^{\mathrm{grid}}$ is the position of the i-th grid node, which shifts with the applied offset (δ_x, δ_y). Grid nodes whose brightest enclosed pixel exceeds $25\, I_1$ are excluded from the sum as anomalously bright artifacts (Fig. B.4(c)). This procedure is similar to that of Ref. [156].

From the procedure above, each located site is matched to a high-resolution three-dimensional confocal scan (∼200 nm pixel pitch) acquired at that position, from which the NV number is obtained as $N_{\mathrm{NV}} = (I_{\mathrm{max}} - I_{\mathrm{bg}})/I_1$, where I_{max} is the brightest pixel of the in-focus slice and I_{bg} the median of the scan's border pixels (Fig. B.4(d)). Sites with one to six resolved PL spots are retained, and per-site outliers are removed if the NV count deviates from the population median by more than three times the rescaled median absolute deviation ($1.4826 \times \mathrm{MAD}$, which estimates the standard deviation for normally distributed data). Across the 800 keV region the procedure yields $N = 222$ analyzed sites with a mean NV yield $\langle N_{\mathrm{NV}} \rangle = 17.5 \pm 0.38$ (standard error of the mean, $5.7/\sqrt{N}$) and a site-to-site standard deviation of 5.7 (see main text Sec. 3.3).

The reported mean NV count is robust to the analysis choices. Tests were performed varying the grid offset, elliptical acceptance area, outlier rejection brightness threshold, and the per-site intensity estimator (brightest pixel versus the mean of the ten brightest pixels), as well as restricting the data set to sites with a single instead of multiple resolved spots. None of these data analysis variations resulted in significant changes to $\langle N_{\mathrm{NV}} \rangle$.

B.3.2 Poisson analysis of multi-ion impacts

At sites that received more than one ion, the high-resolution scans resolve multiple, spatially separated PL spots (Fig. B.4(d)), each attributed to a single ion impact. The distribution of the number of resolved spots per site, pooled over all 222 analyzed sites, is shown in Fig. B.4(e). Because sites

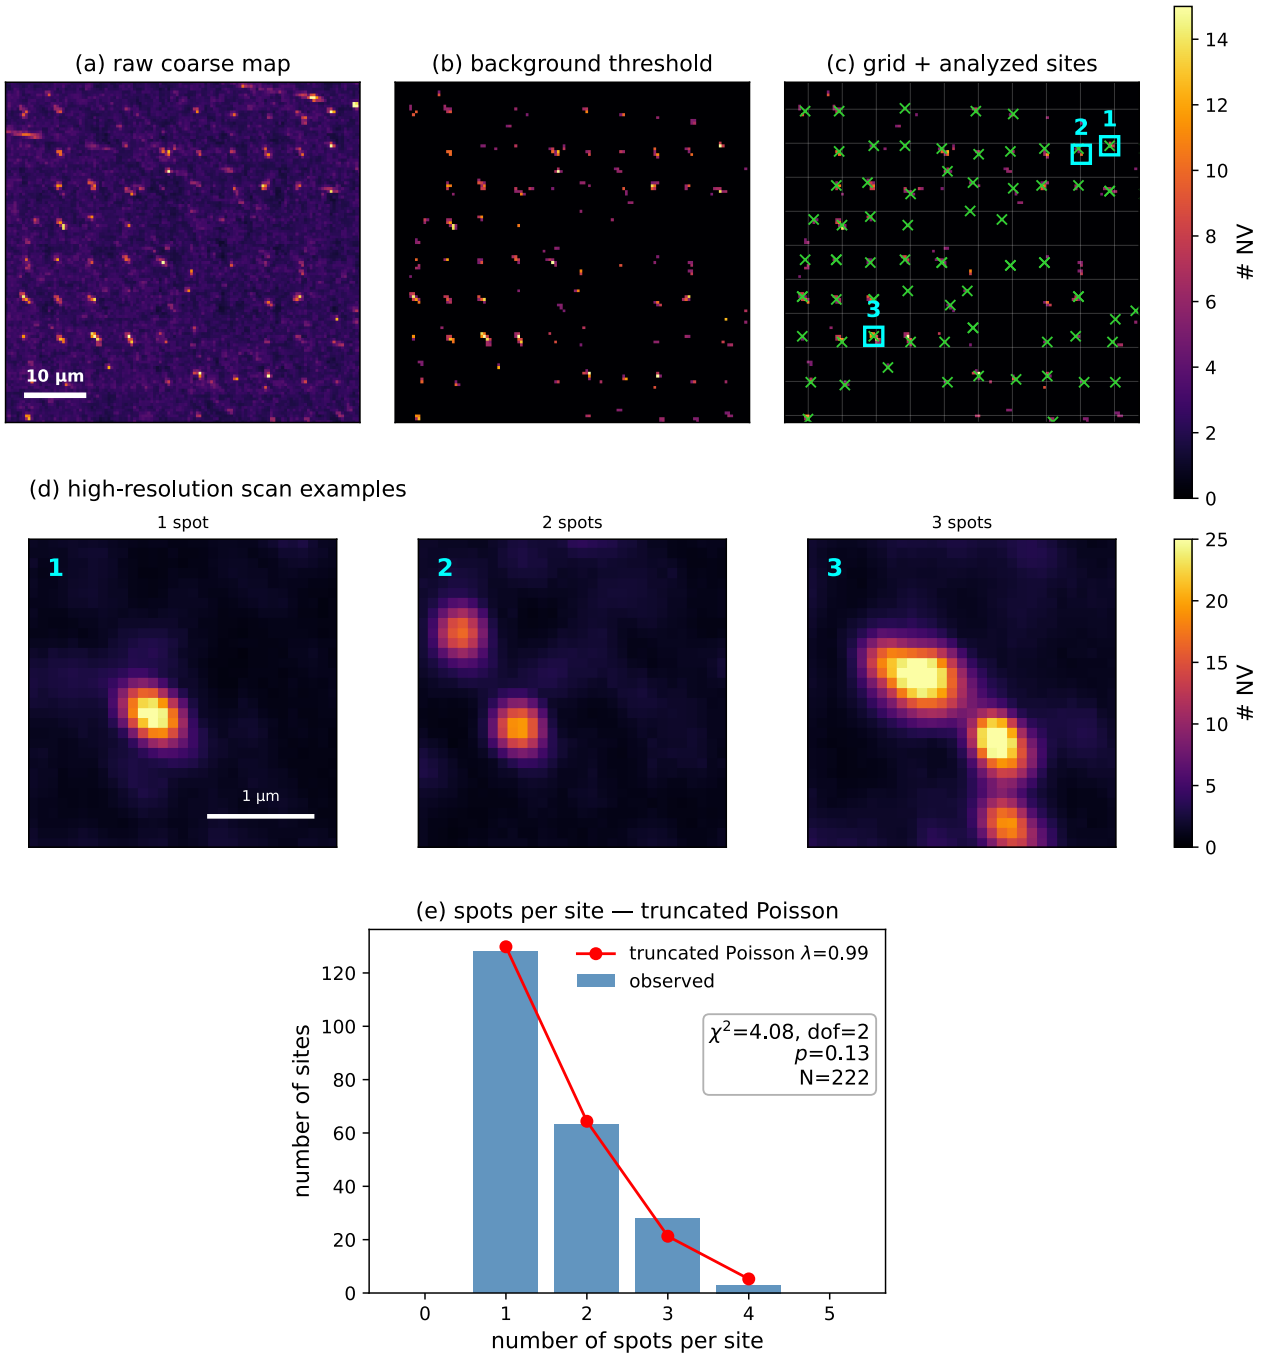


Figure B.4: **NV-count extraction pipeline and multi-ion statistics**, illustrated for a representative 800 keV implantation square (6 μm pitch). **(a)** Raw coarse confocal map (500 nm pixels), in units of calibrated single-NV PL. **(b)** After region-of-interest and scratch masking and intensity thresholding, isolating the implantation-site PL. **(c)** Elliptical grid mask ($\pm 3\sigma$ of the beam profile) overlaid on the thresholded map; green crosses mark the analyzed sites and cyan boxes the three example sites shown in (d). **(d)** High-resolution nanovolume scans of the three boxed sites, resolving one, two, and three PL spots. **(e)** Distribution of the number of resolved spots per site over all $N = 222$ analyzed 800 keV sites (bars), with a maximum-likelihood truncated-Poisson fit (red; $\lambda = 0.99$, $\chi^2 = 4.1$, 2 d.o.f., $p = 0.13$).

with zero spots are not observed, we model the counts with a Poisson distribution truncated to $n \geq 1$,

$$P(n) = \frac{\lambda^n e^{-\lambda}}{n!\left(1 - e^{-\lambda}\right)}, \qquad n = 1, 2, 3, \ldots, \tag{B.6}$$

whose single parameter λ is the mean number of ions delivered per grid node. A maximum-likelihood fit gives $\lambda = 0.99$. A Pearson χ^2 goodness-of-fit test, with the counts grouped into bins $n = 1, 2, 3$, and ≥ 4 so that every expected count exceeds five, yields $\chi^2 = 4.1$ for 2 degrees of freedom ($p = 0.13$); the data are thus statistically consistent with a truncated Poisson distribution. This confirms that each resolved PL spot corresponds to a single, independent ion impact, since the arrival of ions at a given grid node is expected to follow Poisson statistics.

B.4 Improving diamond surface preparation

The NV count analysis of Appendix B.3.1 required manually excluding sub-regions with high NV backgrounds from residual subsurface polishing damage. These backgrounds arise from the diamond polishing process. Synthetic HPHT diamond is typically laser-cut into individual plates and mechanically polished to achieve optically smooth surfaces. However, mechanical polishing is known to introduce subsurface damage [157], generating lattice vacancies that can form NV centers after annealing. In our ion implantation studies, where damage tracks are formed near the surface, such polishing-induced defects can obscure or interfere with the implantation signal. As shown in Fig. B.5(a), a scaife-polished diamond surface (with roughness $R_a \lesssim 1\,\mathrm{nm}$) exhibits strong and highly non-uniform photoluminescence after annealing. Although regions with low background can still be found after annealing, their spatial distribution is difficult to control. In contrast, prior to annealing, the same material typically exhibits fewer than $\sim$2 NV centers per confocal spot across the sample

surface, except for the non-$\{100\}$ sectors (e.g., the four "petal" features near the diamond center).

We performed several experiments to suppress this background through improved surface preparation. These procedures were developed after the implanted sample measured in this work and are intended for future studies. The primary approach we investigated was argon–chlorine (Ar/Cl_2) inductively coupled plasma reactive ion etching (ICP-RIE) [158–160], which removes polishing-induced damage from the near-surface region. While such etching processes have shown success in low-nitrogen CVD diamond, their application to high-nitrogen HPHT diamond ($[N_s^0] \sim 200\,\mathrm{ppm}$) requires further evaluation, as etching-induced damage may be more pronounced. We performed etching with varying RF power (Table B.1), followed by annealing and PL imaging at different depths to quantify the resulting NV density. In each case, more than $\sim 10\,\mu\mathrm{m}$ of material is removed to eliminate the original polishing-damaged layer. For high RF power (250 W), the NV density peaks at ~5–7 per confocal spot at a depth of $\sim 500\,\mathrm{nm}$ (Fig. B.5(b)), whereas reducing the power to 20 W lowers the density to ~1 per spot (Fig. B.5(c)), comparable to the as-received HPHT diamond prior to annealing. Further reduction in RF power yields diminishing improvement while reducing the etch rate. We note that similar surface NV density can also be achieved with shorter etching time by combining an initial high-power etch with a subsequent low-power step to remove 3–4 $\mu\mathrm{m}$ of fast-etch-induced damage.

Table B.1: Ar/Cl_2 ICP-RIE etching recipes and O_2 surface finishing step.

Step	**Gas**	**Flow (sccm)**	**RF Power (W)**	**ICP Power (W)**	**Etch Rate (nm/hr)**	**NVs / confocal spot**
Etch Recipe 1	Ar/Cl_2	25 / 40	250	400	3000	~5-7
Etch Recipe 2	Ar/Cl_2	25 / 40	20	400	480	$\lesssim 1$
Etch Recipe 3	Ar/Cl_2	25 / 40	8	400	117	$\lesssim 1$
O_2 finish	O_2	30	20	700	–	–

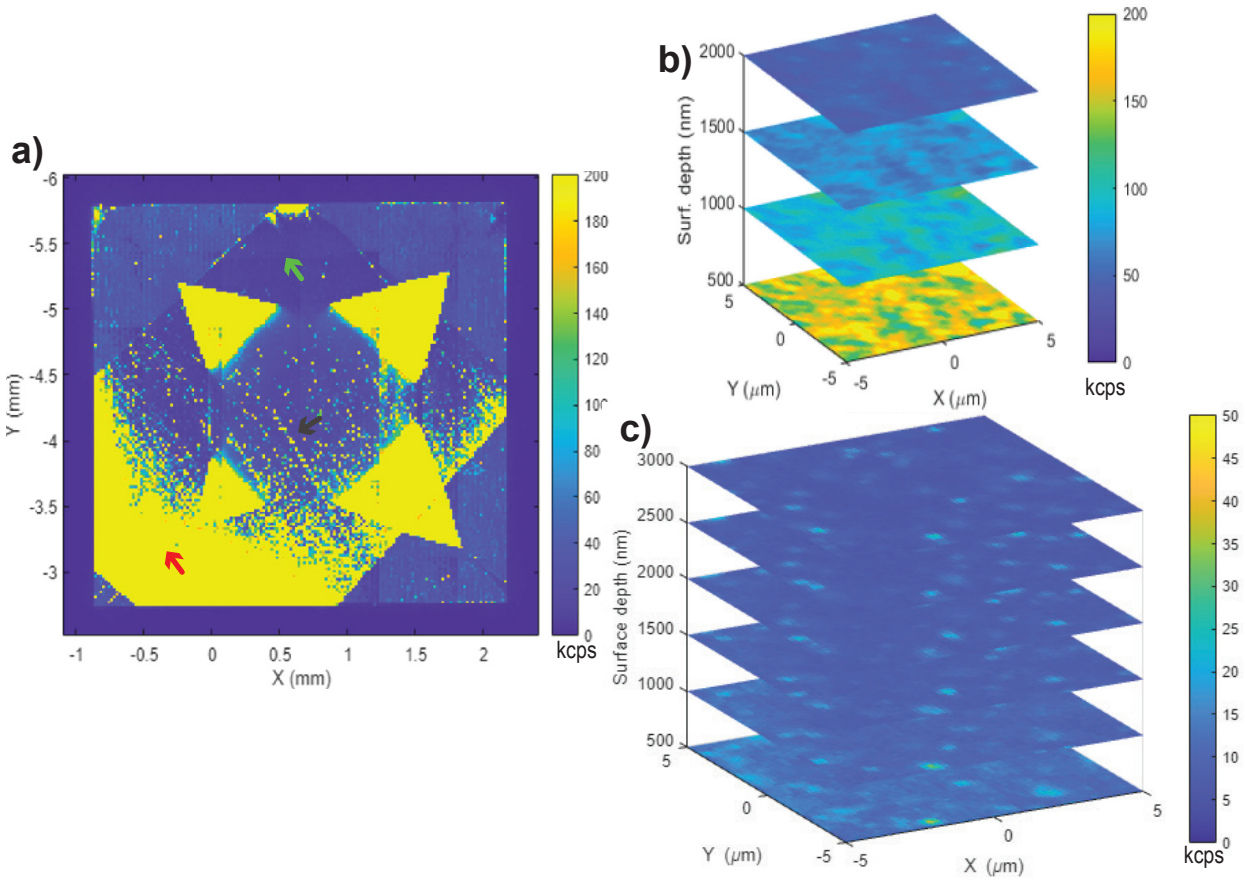


Figure B.5: Mitigation of polishing-induced NV background. (a) Confocal PL map of a scaife-polished HPHT diamond measured at a depth of $\sim 1\,\mu\mathrm{m}$ after annealing. The diamond is predominantly $\{100\}$-oriented, with four high-PL "petal" regions arising from growth sectors. Strong and highly non-uniform background PL is observed due to polishing-induced damage. Red arrows indicate regions of severe damage that can overwhelm implantation signals, black arrows highlight streak-like features consisting of isolated high-NV spots, and green arrows indicate regions with relatively low background. Three corners have been previously etched, showing indented features. (b) Confocal PL maps at increasing surface depths following high-power (250 W) Ar/Cl_2 ICP-RIE etching. The near-surface region ($\lesssim 1\,\mu\mathrm{m}$) exhibits elevated PL corresponding to $\sim$5-7 NV centers per confocal spot, indicating etch-induced damage. (c) Confocal PL maps at increasing depths following low-power (20 W) Ar/Cl_2 ICP-RIE etching. The near-surface PL is significantly reduced, with NV densities ~ 1 per confocal spot, comparable to as-received HPHT diamond without annealing.

B.5 DFT calculations of defect interactions

B.5.1 Methods

First principles calculations of defect formation in diamond are performed with the Vienna *Ab initio* Simulation Package (VASP) [161, 162], which implements plane-wave basis sets using a projector-augmented wave (PAW) scheme [163]. We examined two elementary formation mechanisms of vacancy defects in diamond:

$$\mathrm{C_C} \leftrightarrow \mathrm{V_C} + \mathrm{C_i} \quad \text{(Reaction 1)} \tag{B.7}$$

$$\mathrm{V_{2C}} \leftrightarrow 2\,\mathrm{V_C} \quad \text{(Reaction 2)} \tag{B.8}$$

where Reaction 1 describes the formation of a carbon monovacancy ($\mathrm{V_C}$) and carbon interstitial ($\mathrm{C_i}$) via ejection of a lattice-site carbon atom ($\mathrm{C_C}$), and Reaction 2 describes the ejection of carbon monovacancies from a divacancy site ($\mathrm{V_{2C}}$).

Defects were embedded within a $3 \times 3 \times 3$ (216 atoms) supercell. Ionic geometries of vacancy defects were relaxed at the Perdew-Burke-Ernzerhof (PBE) level of theory according to the generalized gradient approximation (GGA) [164, 165]. Sampling of the Brillouin zone for defect structures was done at the Γ point with a cutoff energy of 520 eV. Convergence criteria for electronic self-consistency and ionic relaxation were set to 10^{-10} eV and 10^{-3} eV Å^{-1}, respectively.

Defect formation energies (E_f) were calculated according to

$$E_f = E_{\mathrm{def}} - E_{\mathrm{pristine}} + \sum_i n_i \mu_i, \tag{B.9}$$

where E_{def} and E_{pristine} denote the energies of the defective and pristine lattices, respectively, and μ_i and n_i represent the chemical potential and number of atoms corresponding to species i. Because the kinetic Monte Carlo (KMC) simulations for annealing do not explicitly account for charge states of defects, only neutral configurations were considered, and no charge correction term is applied.

The energy barriers (E_B) associated with Reactions 1 and 2 were determined using a climbing image-modified nudged elastic band (NEB) approach [117]. The NEB calculations were carried out using the PBE functional with a spring constant of $5.0\,\text{eV}\,\text{Å}^{-2}$ between images. Here, E_f and E_B are distinct quantities, where E_B corresponds to the energy difference between the transition state and the initial configuration.

B.5.2 Results

We performed first principles calculations to qualitatively improve the prediction of NV center formation in KMC annealing simulations. Using the NEB method, we elucidate the transition state geometries and associated energy barriers for Reactions 1 and 2. The MEPs corresponding to V_C formation and V_C ejection from V_{2C} are shown in Figures B.6 and B.7, respectively; we have included images of the initial state, transition state, and final state of each reaction pathway, as well.

As shown in Figure B.6, we predict that formation of a V_C defect and accompanying C_i defect from a pristine diamond structure is achieved via the formation of a split interstitial located two lattice sites away from the vacancy position. By examining the reaction pathway, we can calculate the radial distance between the ejected C_i atom and the resulting V_C defect at both the transition state and final state geometries. We have found that at its stable endpoint, the C_i defect lies about $1.94\,\text{Å}$ away from V_C, whereas the C_i and V_C defects are separated by only $1.75\,\text{Å}$ at the transition state. The

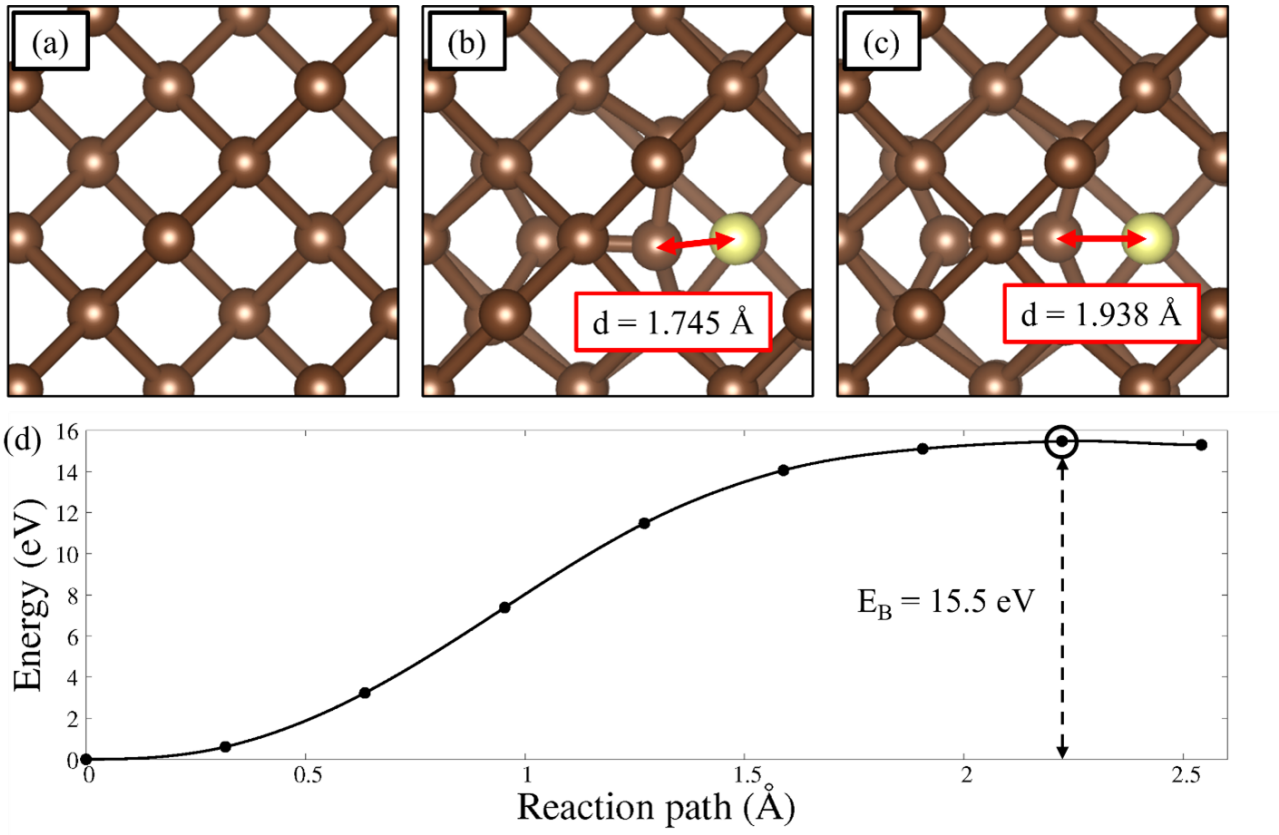


Figure B.6: NEB-DFT calculations of Reaction 1. Geometries of the (a) initial, (b) transition geometry, and (c) final states along the reaction path. V_C and C_i defects are formed by the ejection of a C_C atom from its lattice site. Carbon atoms are shown in brown, while the resulting vacancy is shown in yellow. (d) MEP of Reaction 1. The transition state occurs at a V_C-C_i spacing of 1.75 Å corresponding to E_B = 15.5 eV.

transition state is predicted to have a formation energy of approximately 15.5 eV, and the endpoint has a formation energy of about 15.2 eV.

The MEP for the dissociation of a V_C in a V_{2C} defect is shown in Figure B.7. Unlike Reaction 1, we find that the constituent V_C defects require a separation of three lattice sites ($\sim$3.88 Å) to stably form two distinct vacancies. Instead, the transition state occurs when the V_C defects are separated by two lattice sites ($\sim$2.53 Å). The dissociation of a V_C defect from a V_{2C} defect is predicted to have an energy barrier of 5.95 eV, which is in good agreement with previously reported values [166].

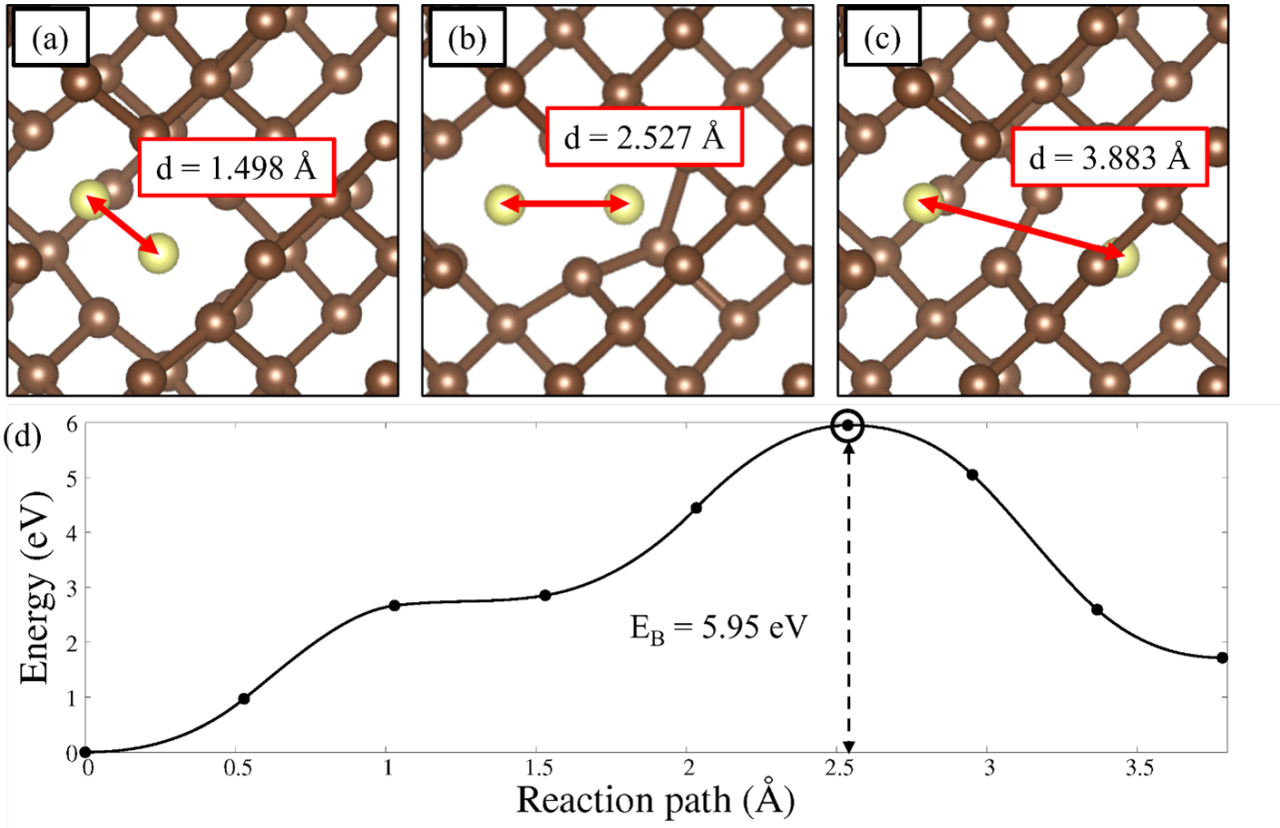


Figure B.7: NEB-DFT calculations of Reaction 2. Geometries of the (a) initial, (b) transition geometry, and (c) final states along the reaction path. Two V_C defects are formed by the dissociation of constituent V_C from the V_{2C} defect. Carbon atoms are shown in brown, while the resulting vacancy is shown in yellow. (d) MEP of Reaction 2. The transition state occurs at a V_C-V_C spacing of 2.53 Å corresponding to E_B = 5.95 eV.

These NEB results provide a basis for refining the defect–defect interaction distances used in the

KMC simulations, which are restricted to first-nearest neighbor (1NN) separation. Vacancy–vacancy interactions can occur spontaneously when defects are within 2NN separation. For vacancy–interstitial pairs, although the transition state occurs at a separation of $1.75\,\text{Å}$, the relaxed configuration lies only $\sim$0.2 eV lower in energy. The apparent atom-vacancy separation in this configuration is $1.94\,\text{Å}$; however, in the KMC lattice picture, the effective interstitial defect position is defined by the arithmetic mean of the two carbon atoms forming the split interstitial, corresponding to a 2NN separation from the vacancy. In the current KMC framework, defect migration barriers are treated as environment-independent, with isolated hopping barriers of 1.5 eV for interstitials and 2.59 eV for vacancies. In contrast, the NEB results indicate a barrier of only $\sim$0.2 eV for V–I recombination from the 2NN configuration. At the annealing temperature of $800\,^{\circ}\text{C}$, this reduced barrier leads to a rate enhancement on the order of $\exp([1.5-0.2]\,\text{eV}/\text{k}_\text{B}\text{T}) \sim 10^6$ compared with isolated defect hopping, indicating that 2NN V–I pairs will preferentially approach and recombine.

As a first-order correction to the 1NN assumption in the KMC model, we modify the initial population of free vacancies available for NV formation – here, "free" denotes vacancies that can diffuse rather than react immediately at the beginning of the annealing simulation. Specifically, in the current model, free vacancies are defined as those without 1NN defect neighbors. In the corrected estimate, we further exclude vacancies that have 2NN vacancy or interstitial neighbors, accounting for enhanced spontaneous divacancy formation and vacancy–interstitial recombination. The resulting reduced vacancy population is then rescaled using the same conversion factor (free-vacancy-to-NV yield) as in the uncorrected model to estimate the NV yield.

Finally, we note that NEB results indicate that vacancy–nitrogen interactions can also occur at approximately 2NN separation. However, on average, only $\sim$0.2% of the initial vacancies have a nitrogen partner within 1NN or 2NN distance; therefore, this effect is neglected in the first-order

correction.

B.6 Additional spin echo measurement data

Additional spin echo measurements are performed on 14 implantation sites in the implanted HPHT sample to investigate potential correlations among spin contrast, coherence properties, and the number of NV centers per site. Individual decay traces and the corresponding NV number distribution are shown in Fig. B.8. These data were acquired with reduced averaging compared to the main text (Sec. 3.6) and were fitted using a stretched exponential model with the exponent fixed at $p = 1.55$, informed by the main text results, i.e. $C_0 \exp[-(\tau/T_2)^{1.55}]$.

The extracted spin echo parameters and NV numbers are summarized in Fig. B.9. The coherence time T_2 remains relatively constant across measurement sites, with an average value of $\sim 1.4\,\mu$s. In contrast, the PL contrast exhibits a larger spread ($\sim$1–3%). Neither shows a clear dependence on NV count. This variation in contrast may arise from differences in the distribution of NV orientations within each site, as the measurement selectively probes a single NV axis while the remaining three crystal orientations contribute to background PL.

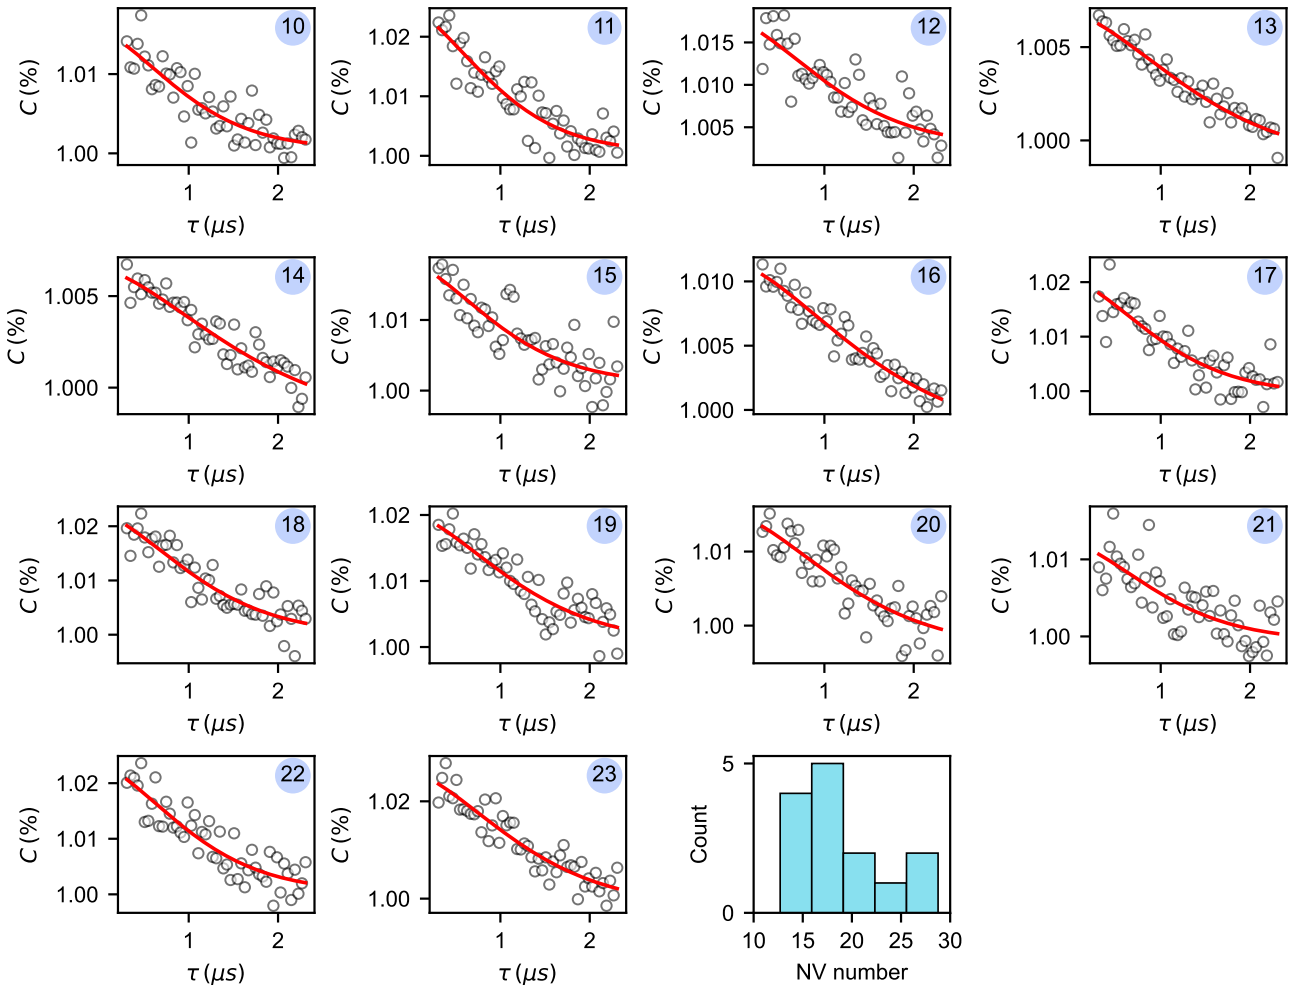


Figure B.8: Additional NV spin echo measurements. Spin echo measurements from 14 additional implantation sites in the HPHT sample. Each trace is fitted to a stretched exponential decay (red curve) with fixed exponent $p = 1.55$. The final panel shows the distribution of NV PL counts for the characterized sites.

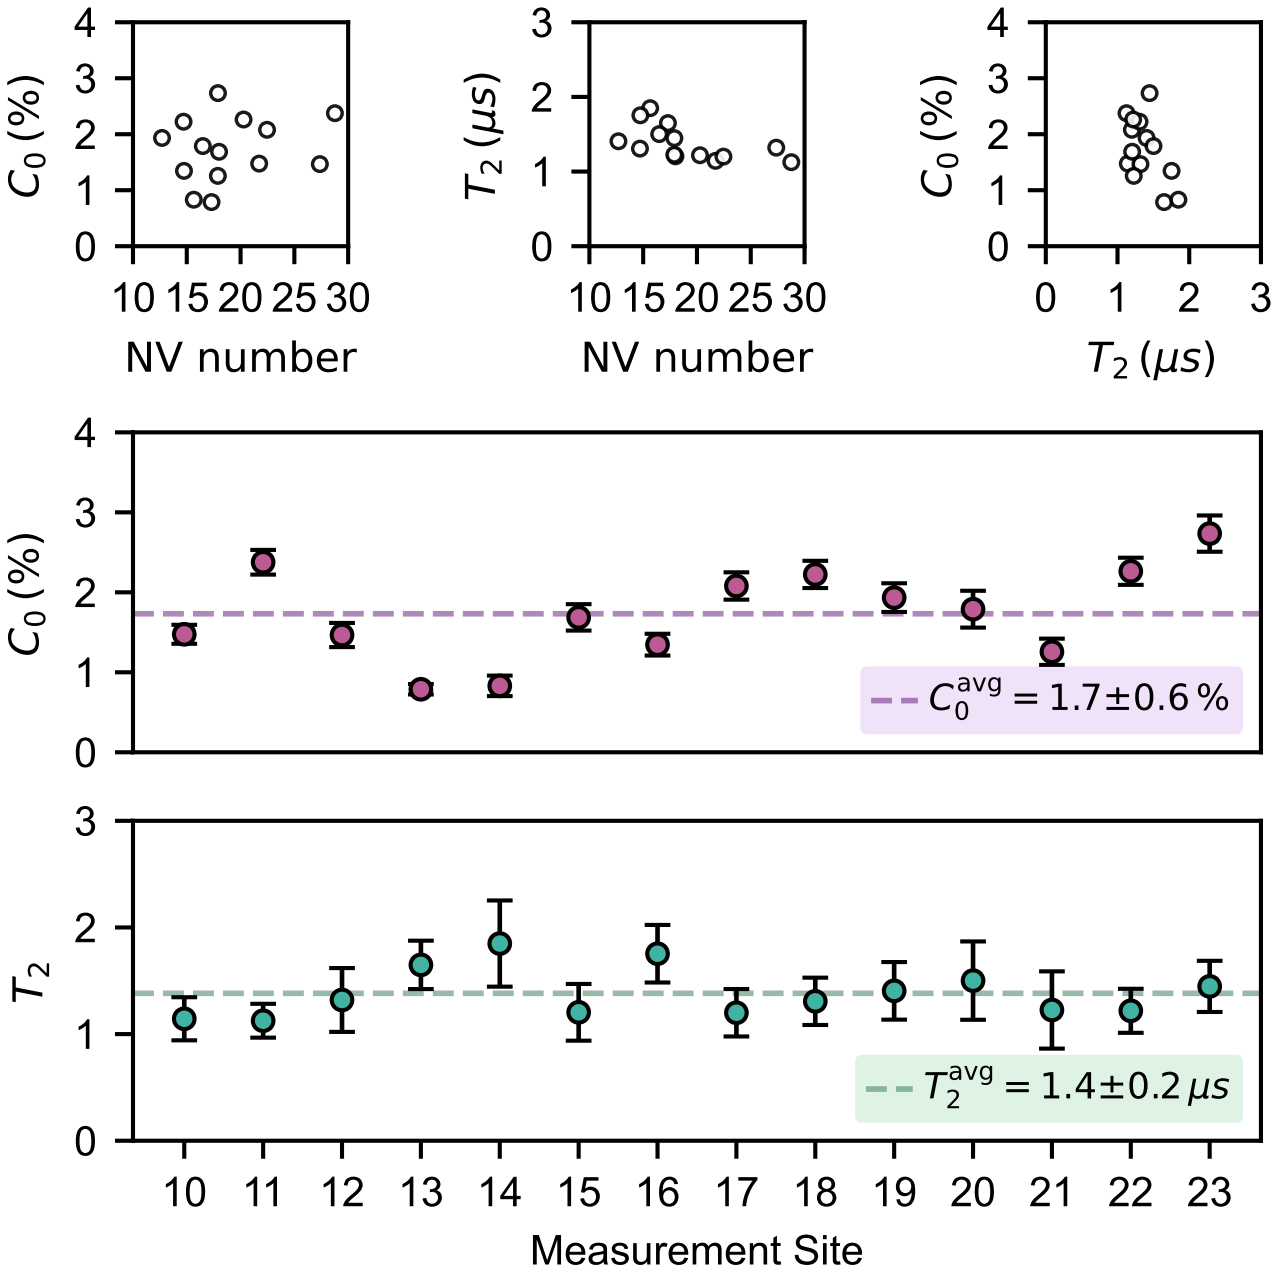


Figure B.9: Spin properties versus NV number. (a) Scatter plots of spin-state-dependent PL contrast C_0, coherence time T_2, and NV number per site, showing no clear correlation between spin properties and NV count. (b) Distribution of C_0 and T_2 across measurement sites. T_2 remains relatively constant around $1.4\,\mu$s, while C_0 exhibits a larger spread ($\sim$1-3%).

Appendix C: Appendix for low nitrogen quantum diamond

This appendix is adapted from Jiashen Tang, Connor A. Roncaioli, Andrew M. Edmonds, Atli Davidsson, Connor A. Hart, Matthew L. Markham, and Ronald L. Walsworth, "Characterization of low-nitrogen quantum diamond for pulsed magnetometry applications" [119].

C.1 Ramsey measurement at high-birefringence location

As shown in the main text (Sec. 4.3), sample 10 contains two regions with elevated birefringence, indicating higher strain variation. To illustrate the effect of such inhomogeneity on NV ensemble properties, we perform a single-quantum (SQ) Ramsey measurement at a spot within area B. The extracted dephasing time is $T^*_{2,\mathrm{SQ}} = 5.1(6)\,\mu$s, significantly shorter than the spin-bath-limited value. In contrast, a double-quantum (DQ) Ramsey measurement at this same spot, which is insensitive to longitudinal strain and temperature variations [13], yields $T^*_{2,\mathrm{DQ}} = 8.6(5)\,\mu$s. This value is approximately half of the SQ dephasing time observed in low-birefringence regions (see main text Sec. 4.4), consistent with the doubled gyromagnetic ratio and increased sensitivity to spin-bath noise of DQ measurements.

C.2 NV spin dephasing time at nitrogen spin-bath limit

The spin-bath-limited NV spin dephasing time ($T^*_{2,\mathrm{bath}}$) can be estimated using the following relation [1]:

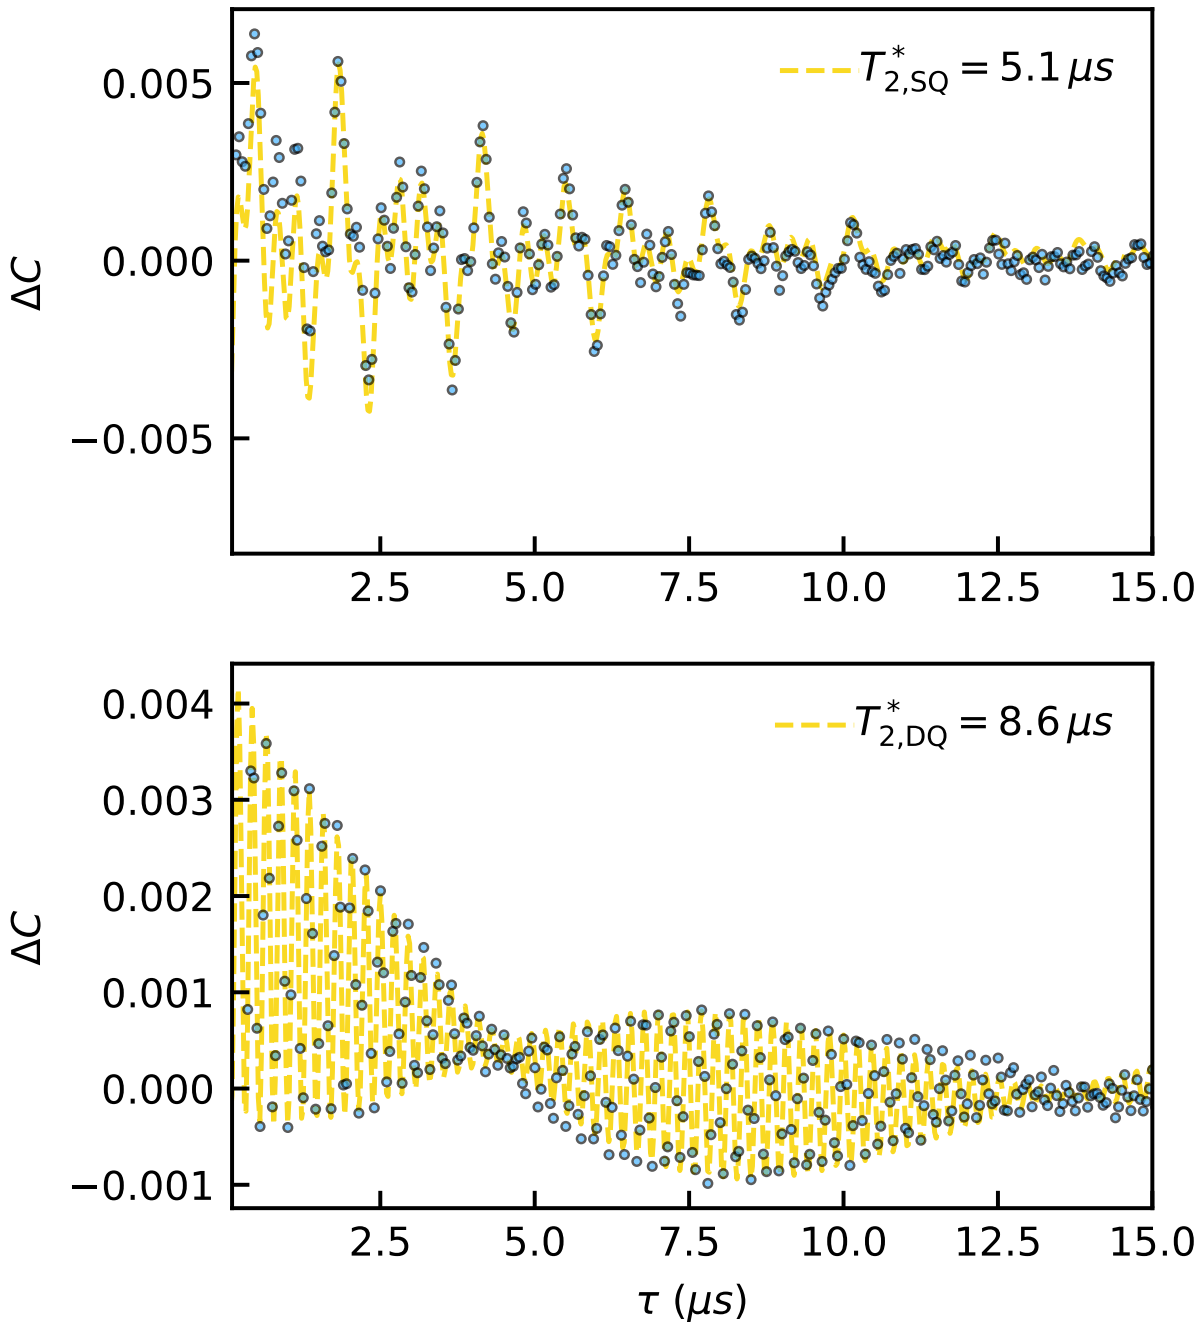


Figure C.1: Single-quantum (SQ) and double-quantum (DQ) Ramsey measurements and fits yield the ensemble NV spin dephasing times $T^*_{2,\mathrm{SQ}}$ and $T^*_{2,\mathrm{DQ}}$ at a high-birefringence region (area B of sample 10, see main text Fig. 2). The reduced $T^*_{2,\mathrm{SQ}}$ reflects the strong local strain variation, while the DQ protocol mitigates these effects. ΔC is the change in NV PL contrast.

$$\begin{aligned}\frac{1}{T^*_{2,\text{bath}}} &= \frac{1}{T^*_{2,[\text{N}_\text{s}^0]}} + \frac{1}{T^*_{2,[^{13}\text{C}]}} + \frac{1}{T^*_{2,\text{NV-NV}}} \\ &\approx A_{\text{N}_\text{s}^0}[\text{N}_\text{s}^0] + A_{^{13}\text{C}}[^{13}\text{C}] + \zeta_{||} A_{\text{NV}^-_{||}}[\text{NV}^-_{||}] + \zeta_{\nparallel} A_{\text{NV}^-_{\nparallel}}[\text{NV}^-_{\nparallel}],\end{aligned} \tag{C.1}$$

where the scaling constants are $A_{\text{N}_\text{s}^0} = 0.101\,\mu\text{s}^{-1}\text{ppm}^{-1}$, $A_{^{13}\text{C}} = 0.1\,\text{ms}^{-1}\text{ppm}^{-1}$, $A_{\text{NV}^-_{||}} = 0.247\,\mu\text{s}^{-1}\text{ppm}^{-1}$ and $A_{\text{NV}^-_{\nparallel}} = 0.165\,\mu\text{s}^{-1}\text{ppm}^{-1}$. Here, $[\text{NV}^-_{||}]$ represents the concentration of negatively-charged NV centers (NV^-) aligned along the sensing axis, while $[\text{NV}^-_{\nparallel}]$ corresponds to the concentration of NV centers oriented along other crystal axes. For our experimental study, the bias magnetic field aligns with a single NV crystal axis, resulting in $[\text{NV}^-_{||}] \approx \frac{1}{3}[\text{NV}^-_{\nparallel}]$. For the sensing axis and other axes, the dimensionless correction factor ζ accounts for the incomplete polarization of NV^- spins into the $m_s = 0$ state. We optimize laser polarization such that NV centers aligned along two axes are predominantly excited, with one orientation used for sensing. Thus, we approximate $\zeta_{[\text{NV}^-_{||}]} \approx 0$. Determining $\zeta_{[\text{NV}^-_{\nparallel}]}$ is complicated due to unequal laser excitation intensities and polarizations along the three non-sensing NV orientations under epi-illumination conditions; as a conservative estimate for the NV-NV interaction-limited $T^*_{2,\text{NV-NV}}$, we assume partial non-sensing NV spin polarization with $\zeta_{\nparallel} = 0.5$.

Before electron irradiation and annealing of the ten low-$[\text{N}_\text{s}^0]$ diamond samples, the substitutional nitrogen concentration is measured to be $[\text{N}_\text{s}^{0,\text{as-grown}}] \approx 0.8\,\text{ppm}$ via UV-Vis absorption [58]. After processing, the total NV concentration is 0.39(2) ppm determined by UV-Vis absorption [58], with approximately 20% of NVs in the negatively charged (NV^-) state (as determined from the PL emission spectrum decomposition technique, see main text Sec. 4.6). Formation of NV centers consumes substitutional nitrogen, with typically one nitrogen atom per NV^0 and two per NV^- center [1]. Therefore, the post-treatment substitutional nitrogen concentration $[\text{N}_\text{s}^0] \approx 0.35\,\text{ppm}$. Residual ^{13}C concentration

is 0.01% (equivalent to 108 ppm) given the isotopic purification used in the CVD process. Using these values, we estimate the spin-bath-limited NV spin dephasing time to be $T^*_{2,\text{bath}} \approx 20\,\mu$s, slightly larger than the experimentally measured value for a single quantum (SQ) Ramsey measurement (see main text Sec. 4.4).

C.3 Full strain and stress maps

The NV electronic spin Hamiltonian contains spin-strain coupling terms, including longitudinal (M_z) and transverse (M_x, M_y) components [167]. Under a strong bias magnetic field ($|\vec{B}| \gtrsim 10\,G$), the longitudinal component M_z dominates NV spin resonance frequency shifts, and transverse contributions can be neglected [11]. By resolving NV spin transitions corresponding to the four NV orientations in the diamond lattice, the spin stress terms $M_{z,i}$ for each orientation i are experimentally determined using continuous-wave optically detected magnetic resonance (CW-ODMR) measurements, as described in [11]. Representative spatial maps of these terms across diamond sample 10 are shown in Fig. C.2.

Additionally, from these strain measurements, the stress tensor components of the diamond – comprising diagonal terms ($\sigma_{diag} = \sigma_{xx} + \sigma_{yy} + \sigma_{zz}$) and off-diagonal shear terms ($\sigma_{xy}, \sigma_{yz}, \sigma_{xz}$) – can be uniquely determined. The calculation method is detailed in Ref. [11] and the results are also shown in Fig. C.2.

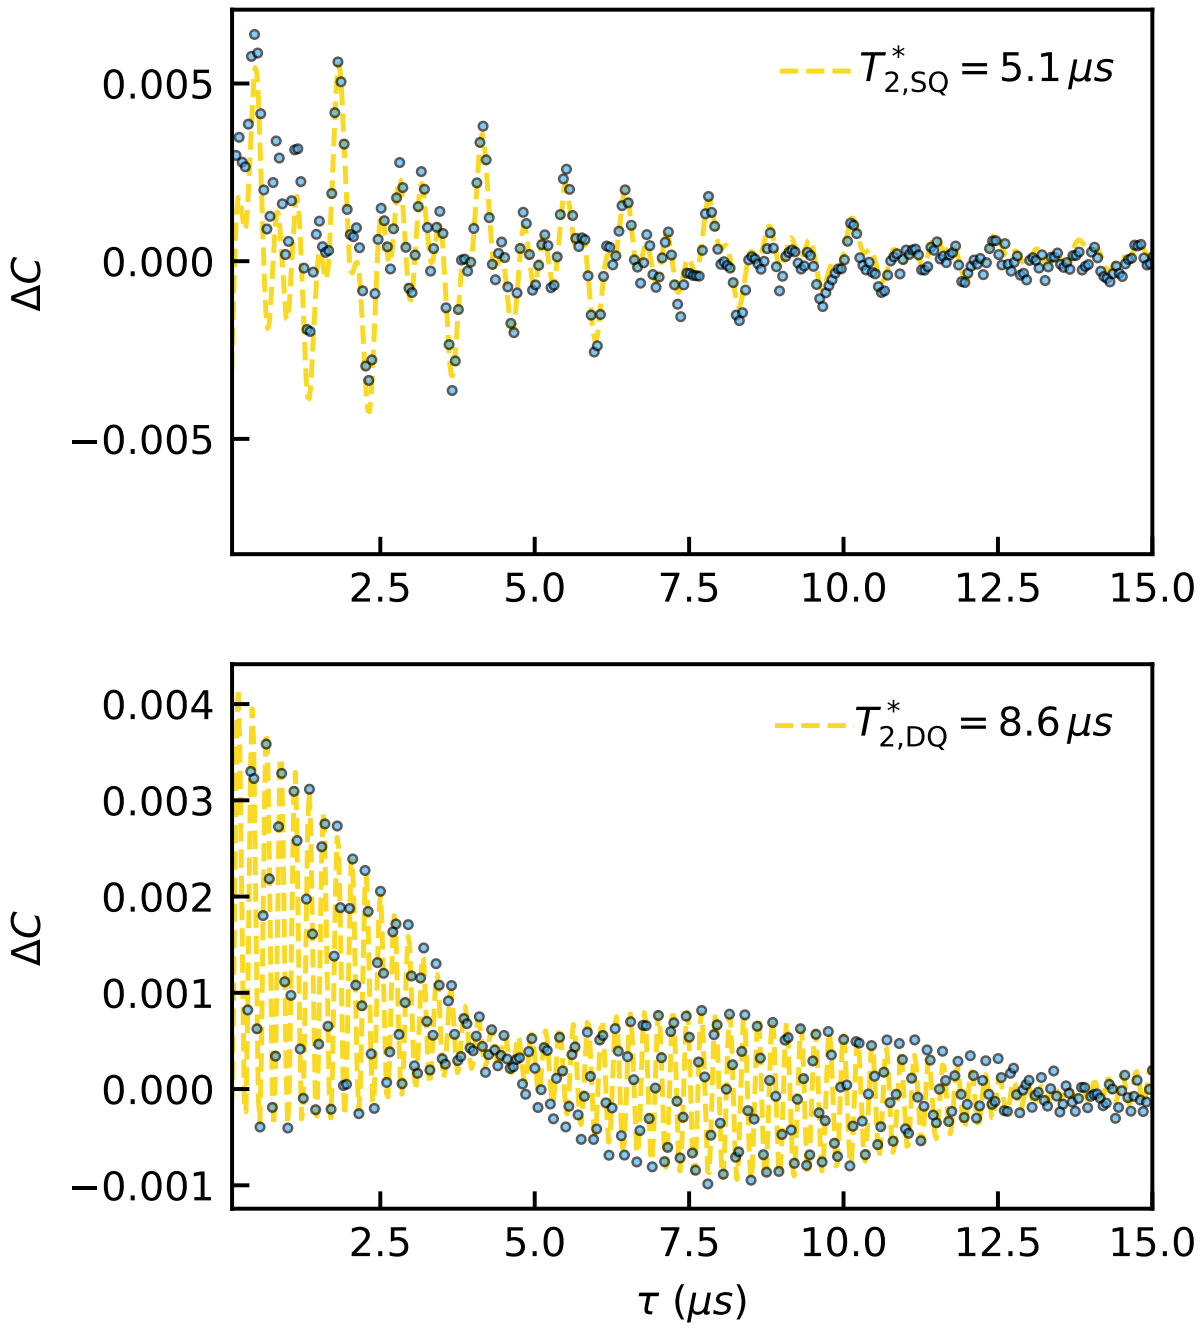


Figure C.2: Spatial maps of experimentally determined NV spin-strain coupling terms (M_z) for each of four NV orientations and derived diamond stress tensors components σ across the entire diamond sample 10. Scale bars: $500\,\mu$m

Appendix D: Interfacing Sample with Diamond in Ramsey QDM

In benchmarking the Ramsey-based QDM (Chapter 2), we used a wire phantom to provide a magnetic-field pattern for imaging. The wire phantom was lithographically defined together with the microwave loop used for NV spin control. Both electrical circuits were fabricated on a 2-inch sapphire or SiC wafer, and the diamond was manually glued onto the wafer to form close contact with the microwave and wire circuits Fig. D.1. Close contact with the microwave loop ensures that the near-surface NV layer receives strong near-field microwave irradiation, while close contact with the wire phantom minimizes the source-to-NV standoff, thereby improving signal strength and magnetic-field spatial resolution. However, in this configuration there is no direct access for interfacing an external sample with the diamond. A sample could in principle be placed above the wafer, but this would introduce a sample-sensor standoff of at least several hundred micrometers.

To address this issue and allow direct, non-circuit sample interfacing with the diamond while maintaining mechanical support and heat spreading, a slot opening can be created in the wafer by laser cutting. The microwave circuit can then be lithographically fabricated around the opening, after which the diamond is glued onto the wafer to complete the assembly. As shown schematically in Fig. D.2, a square slot is used here to allow insertion of a glass pillar carrying the sample. The motivation is to enable modulation of the pillar-diamond standoff for differential magnetic measurements. Laser cutting is versatile, and other opening geometries could be used.

For the sample-embedded glass pillar, I developed the fabrication process described below. An example final product is shown in Fig. D.3. The image is a top view of the pillar, where three magnetic

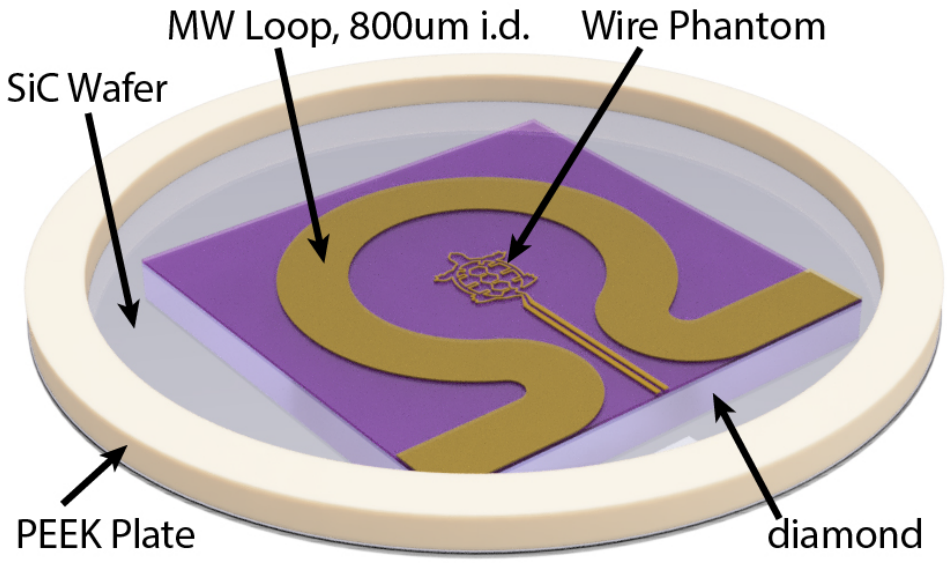


Figure D.1: Schematic of the Ramsey QDM sensor head assembly, including the wafer, microwave circuit, wire phantom, and diamond.

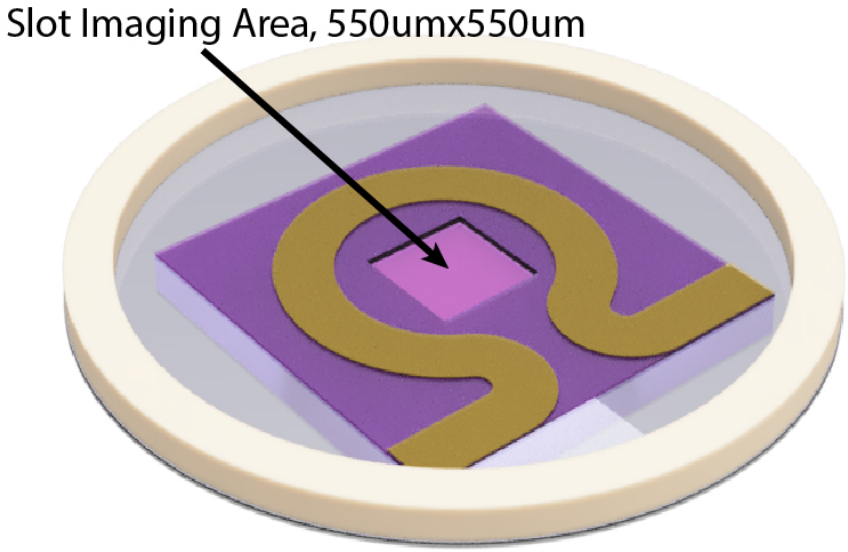


Figure D.2: Schematic of a wafer opening that enables direct sample interfacing with the diamond in the Ramsey QDM.

beads, each approximately 3 μm in diameter, are fixed near the top surface. The magnetic beads are embedded in S1813 photoresist so that they remain stably attached to the glass pillar.

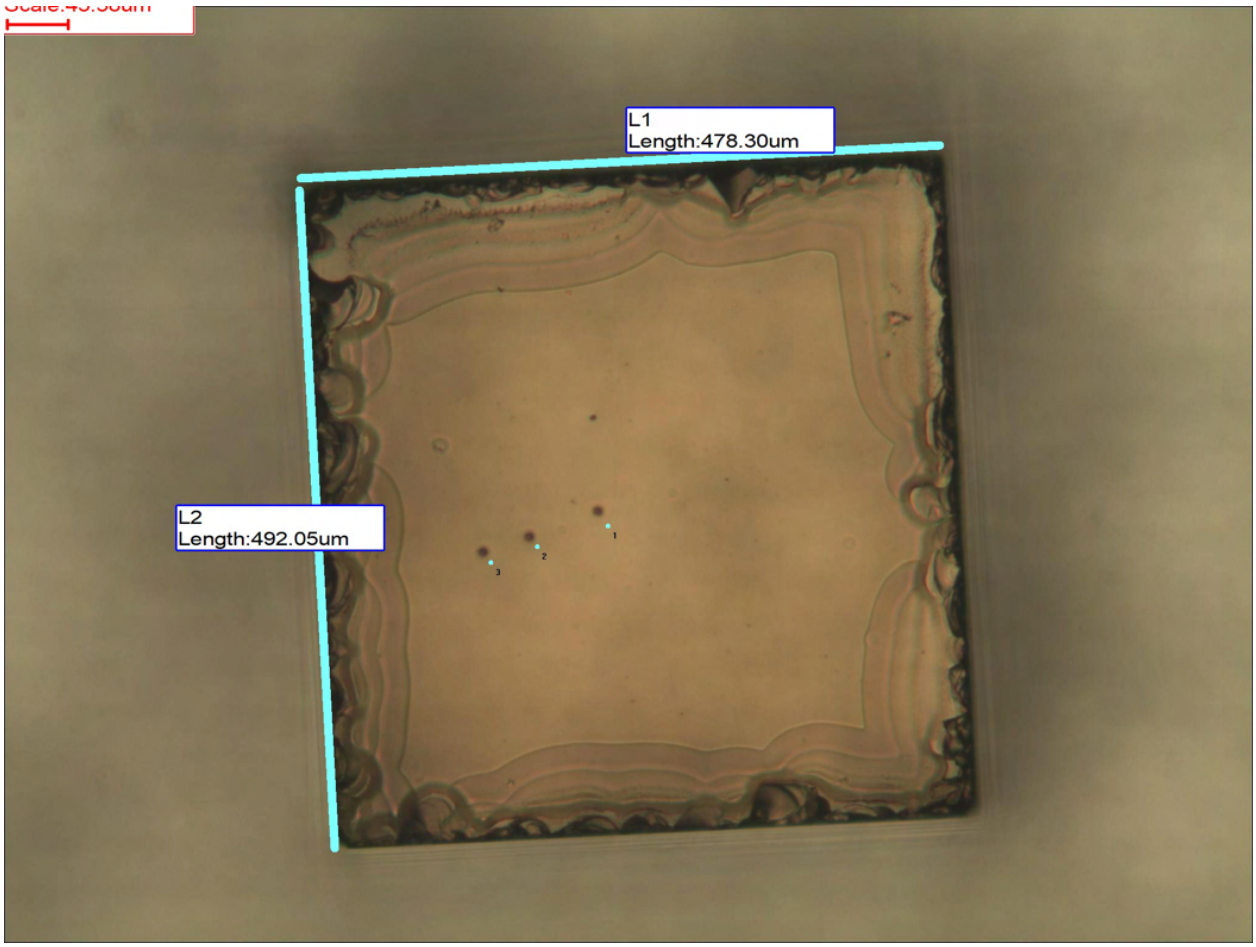


Figure D.3: Top-view image of a glass pillar carrying magnetic beads fixed with S1813 photoresist.

The fabrication process is as follows. First, a microscope glass slide is chosen and cleaned. The thickness of the slide can be varied depending on the desired pillar height. Magnetic beads are then attached to the glass either by directly drop-casting a bead containing solution or by mixing the bead solution with S1813 photoresist and spin-coating the mixture onto the glass slide. After spin coating, the S1813 is baked to harden the photoresist and stably fix the beads to the glass. Once the beads are attached, the glass slide is cut with a dicing saw to the desired pillar size. We found that beads directly

drop-cast onto diamond are typically less secure and can be washed away during subsequent handling, while S1813 fixation provides improved stability.

As a demonstration, we chose a square pillar size of approximately $500\ \mu\mathrm{m}$ so that the pillar could be carefully picked up with tweezers. The pillar height was approximately $1\ \mathrm{mm}$, thicker than the wafer and therefore easier to remove and exchange with a tweezers. The pillar can also be attached to a mechanical arm with translational degrees of freedom, allowing its stand-off from the diamond sensor surface to be adjusted. This capability could be used to modulate the measurement geometry for differential magnetic measurements, although this approach remains to be tested. An example setup image is shown in Fig. D.4, and an example magnetic image is shown in Fig. D.5.

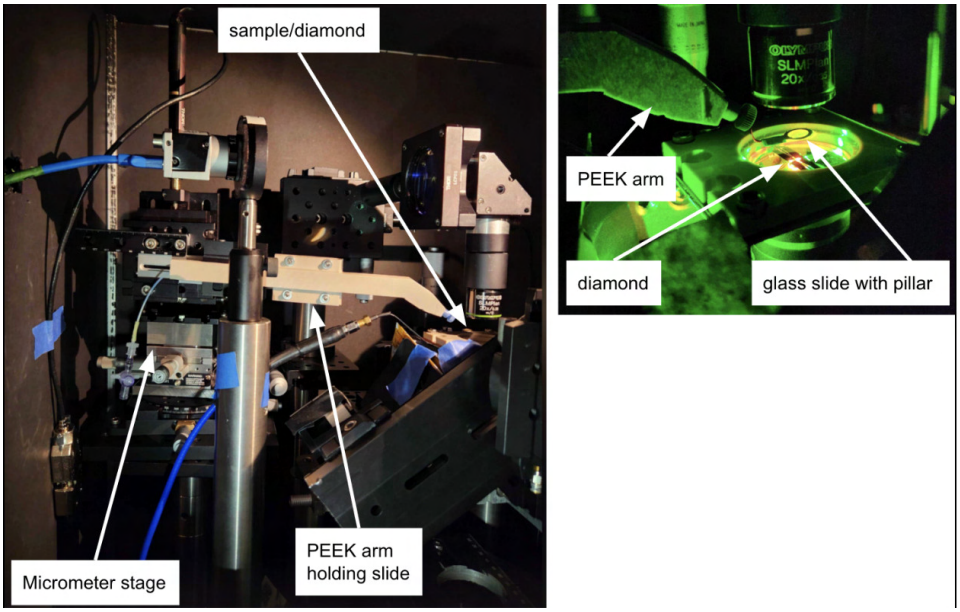


Figure D.4: Ramsey QDM setup with a mechanical arm holding the glass pillar.

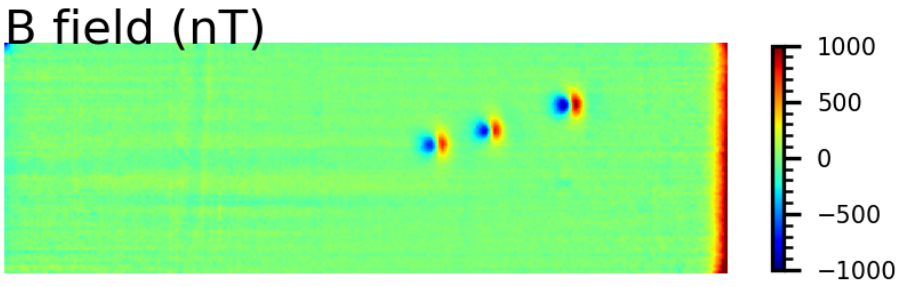


Figure D.5: Magnetic image of the pillar sample measured using the Ramsey-based QDM. The pillar is positioned within the square imaging area.

Appendix E: Germinating Pollen Grains in the Ramsey QDM

To germinate pollen grains inside the Ramsey QDM under normal operating conditions, including strong laser illumination and microwave irradiation, several engineering modifications were implemented. These efforts are documented below.

The pollen grains are placed inside the square imaging area described in Appendix D. The opening slot allows the pollen grains to sit directly above the diamond surface, in close proximity to the near-surface NV sensing layer. As shown in Fig. E.1, a bio-mount is constructed above the sensor head to hold the culture medium required for pollen germination.

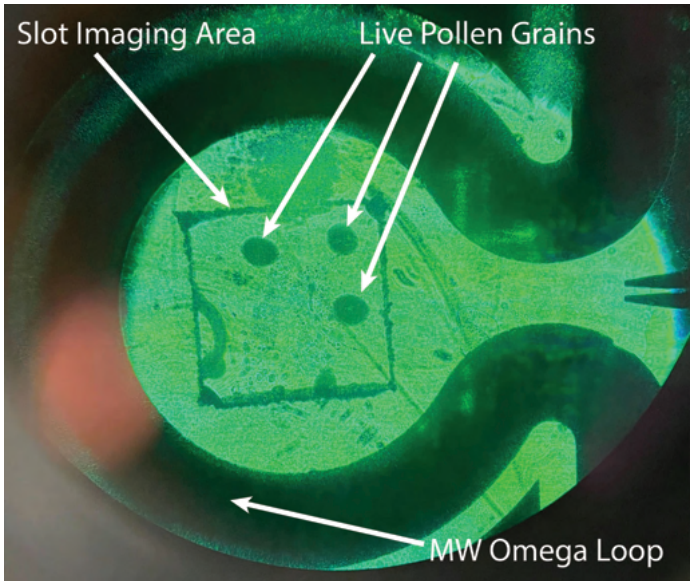


Figure E.1: Lily pollen grain inside the square imaging area above the diamond surface.

Images of the bio-mount are shown in Figs. E.2. The main component is a Warner Instruments RC-27LD perfusion chamber, which is tightly clamped onto the sapphire or SiC wafer containing the microwave loop and imaging diamond. Non-magnetic PEEK components and an EPDM gasket are

used to prevent culture-medium leakage. The inlet delivers fresh culture medium, while the suction outlet removes excess medium to maintain a relatively stable fluid level above the diamond.

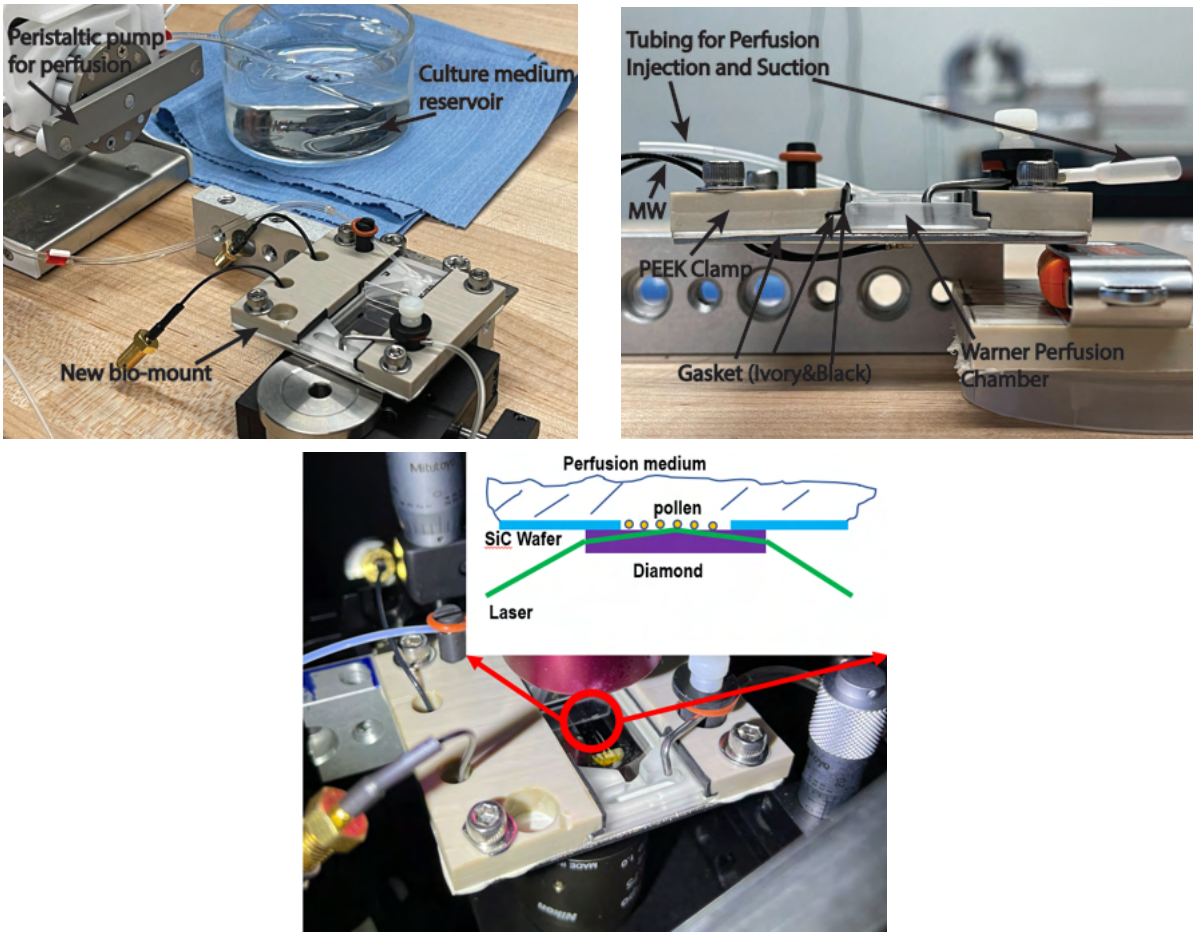


Figure E.2: (Top) Images of the bio-mount designed to integrate with the Ramsey QDM sensor head and create a perfusion environment for pollen-tube growth. A peristaltic pump is shown here, but in the final experiment it was replaced by an HPLC pump for improved stability when injecting culture medium into the bio-mount. (Bottom) Bio-mount inside the Ramsey QDM setup. The zoomed-in panel shows a schematic cross-sectional view through the central imaging area.

I also built a water bath system, shown in Fig. E.3, to stabilize the temperature of the culture medium. This was motivated by the fact that the laboratory temperature fluctuates and may not be ideal for pollen-grain germination. A typical turnkey water-bath system with both heating and cooling capabilities can cost more than $7k. However, for the pollen-tube experiment, we only needed to main-

tain the culture medium near room temperature, within approximately $\pm 5\ ^{\circ}\mathrm{C}$. I therefore developed a lower-cost alternative.

The final system (Fig. E.3) consists of a TEC laser chiller, such as a Coherent T257P, which circulates temperature-regulated coolant. Several heat exchange plates are placed in the circulation loop to regulate the temperature of a water bath containing flasks of culture medium. I measured the temperature of the culture medium directly and found that this TEC-based water-bath system provides stable temperature control, as shown in Fig. E.4.

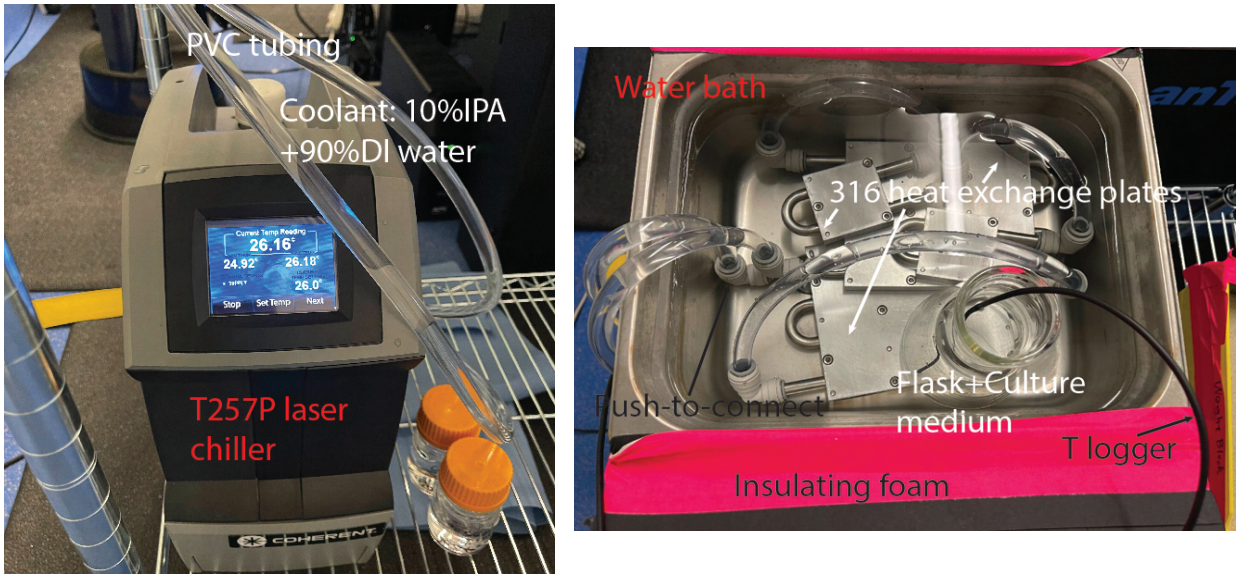


Figure E.3: TEC-based water bath system used to regulate the temperature of the culture medium.

With these engineering modifications, we performed germination experiments with tobacco pollen grains. During the experiment, the QDM operated in continuous magnetic-imaging mode, while a camera above the diamond periodically imaged the pollen grains located near the diamond surface. Successful germination of multiple pollen grains was observed, as shown in Fig. E.5.

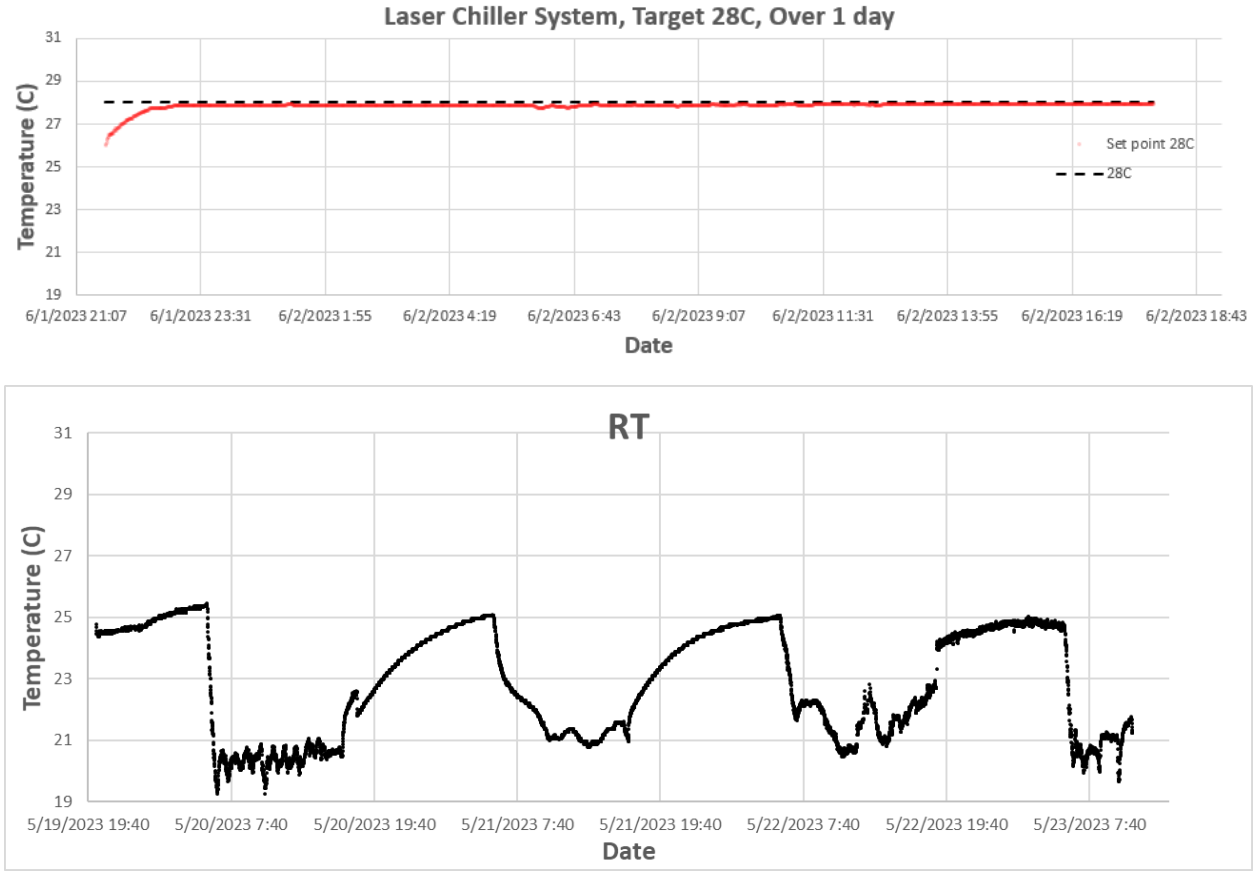


Figure E.4: Example temperature profile of culture medium submerged in the water bath (top) and an example of laboratory temperature variation (bottom).

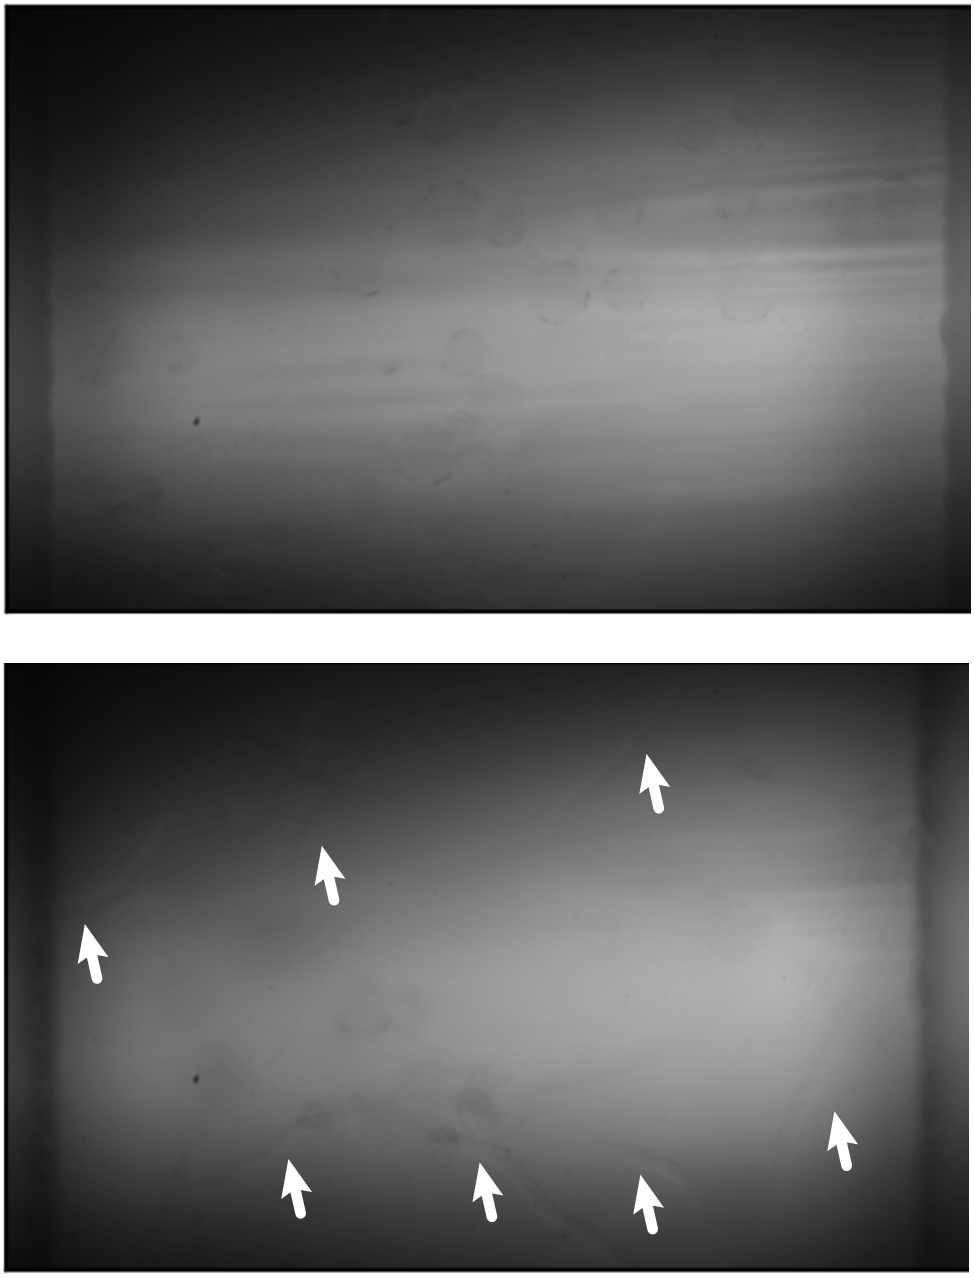


Figure E.5: Tobacco pollen grains germinating under normal NV magnetic-imaging operation. Top: before continuous magnetic imaging, pollen grains are loaded above the diamond surface. Bottom: after approximately 1 hour of continuous magnetic imaging, multiple pollen tubes are visible across the image. These images were acquired with a $20\times$ 0.35 NA objective for quick viewing of pollen grains from above the diamond surface (cf. Fig. D.4), rather than for NV photoluminescence imaging or optimized optical imaging of pollen tubes.

Bibliography

[1] J. F. Barry, J. M. Schloss, E. Bauch, M. J. Turner, C. A. Hart, L. M. Pham, and R. L. Walsworth, "Sensitivity optimization for nv-diamond magnetometry", Rev. Mod. Phys. **92**, 015004 (2020).

[2] M. W. Doherty, F. Dolde, H. Fedder, F. Jelezko, J. Wrachtrup, N. B. Manson, and L. C. L. Hollenberg, "Theory of the ground-state spin of the nv^- center in diamond", Phys. Rev. B **85**, 205203 (2012).

[3] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. Hollenberg, "The nitrogen-vacancy colour centre in diamond", Physics Reports **528**, The nitrogen-vacancy colour centre in diamond, 1 (2013).

[4] Á. Gali, "Ab initio theory of the nitrogen-vacancy center in diamond", Nanophotonics **8**, 1907 (2019).

[5] J. M. Schloss, J. F. Barry, M. J. Turner, and R. L. Walsworth, "Simultaneous broadband vector magnetometry using solid-state spins", Phys. Rev. Appl. **10**, 034044 (2018).

[6] P. Deák, B. Aradi, M. Kaviani, T. Frauenheim, and A. Gali, "Formation of NV centers in diamond: A theoretical study based on calculated transitions and migration of nitrogen and vacancy related defects", Physical Review B **89**, 075203 (2014).

[7] B. J. Shields, Q. P. Unterreithmeier, N. P. de Leon, H. Park, and M. D. Lukin, "Efficient readout of a single spin state in diamond via spin-to-charge conversion", Physical review letters **114**, 136402 (2015).

[8] D. McCloskey, N. Dontschuk, A. Stacey, C. Pattinson, A. Nadarajah, L. Hall, L. Hollenberg, S. Prawer, and D. Simpson, "A diamond voltage imaging microscope", Nature Photonics **16**, 730 (2022).

[9] J. T. Oon, J. Tang, C. A. Hart, K. S. Olsson, M. J. Turner, J. M. Schloss, and R. L. Walsworth, "Ramsey envelope modulation in nv diamond magnetometry", Physical Review B **106**, 054110 (2022).

[10] M. L. Goldman, M. Doherty, A. Sipahigil, N. Y. Yao, S. Bennett, N. Manson, A. Kubanek, and M. D. Lukin, "State-selective intersystem crossing in nitrogen-vacancy centers", Physical Review B **91**, 165201 (2015).

[11] P. Kehayias, M. J. Turner, R. Trubko, J. M. Schloss, C. A. Hart, M. Wesson, D. R. Glenn, and R. L. Walsworth, "Imaging crystal stress in diamond using ensembles of nitrogen-vacancy centers", Phys. Rev. B **100**, 174103 (2019).

[12] M. J. Turner, "Quantum diamond microscopes for biological systems and integrated circuits", phdthesis (Harvard University, 2020).

[13] C. A. Hart, J. M. Schloss, M. J. Turner, P. J. Scheidegger, E. Bauch, and R. L. Walsworth, "N-V–diamond magnetic microscopy using a double quantum 4-ramsey protocol", Phys. Rev. Appl. **15**, 044020 (2021).

[14] J. Gracio, Q. Fan, and J. Madaleno, "Diamond growth by chemical vapour deposition", Journal of Physics D: Applied Physics **43**, 374017 (2010).

[15] S. S. Dossa, I. Ponomarev, B. N. Feigelson, M. Hainke, C. Kranert, J. Friedrich, and J. J. Derby, "Analysis of the high-pressure high-temperature (hpht) growth of single crystal diamond", Journal of Crystal Growth **609**, 127150 (2023).

[16] M. Gaukroger, P. Martineau, M. Crowder, I. Friel, S. Williams, and D. Twitchen, "X-ray topography studies of dislocations in single crystal cvd diamond", Diamond and Related Materials **17**, 262 (2008).

[17] R. Schirhagl, K. Chang, M. Loretz, and C. L. Degen, "Nitrogen-vacancy centers in diamond: nanoscale sensors for physics and biology", Annual review of physical chemistry **65**, 83 (2014).

[18] A. Dréau, M. Lesik, L. Rondin, P. Spinicelli, O. Arcizet, J.-F. Roch, and V. Jacques, "Avoiding power broadening in optically detected magnetic resonance of single nv defects for enhanced dc magnetic field sensitivity", Phys. Rev. B **84**, 195204 (2011).

[19] H. J. Mamin, M. H. Sherwood, M. Kim, C. T. Rettner, K. Ohno, D. D. Awschalom, and D. Rugar, "Multipulse double-quantum magnetometry with near-surface nitrogen-vacancy centers", Phys. Rev. Lett. **113**, 030803 (2014).

[20] M. C. Marshall, R. Ebadi, C. Hart, M. J. Turner, M. J. Ku, D. F. Phillips, and R. L. Walsworth, "High-precision mapping of diamond crystal strain using quantum interferometry", Phys. Rev. Appl. **17**, 024041 (2022).

[21] K. Arai, C. Belthangady, H. Zhang, N. Bar-Gill, S. J. DeVience, P. Cappellaro, A. Yacoby, and R. L. Walsworth, "Fourier magnetic imaging with nanoscale resolution and compressed sensing speed-up using electronic spins in diamond", en, Nature Nanotechnology **10**, 859 (2015).

[22] H. Zhang, K. Arai, C. Belthangady, J.-C. Jaskula, and R. L. Walsworth, "Selective addressing of solid-state spins at the nanoscale via magnetic resonance frequency encoding", en, npj Quantum Information **3**, 1 (2017).

[23] M. T. Amawi, A. Trelin, Y. Huang, P. Weinbrenner, F. Poggiali, J. Leibold, M. Schalk, and F. Reinhard, "Three-dimensional magnetic resonance tomography with sub-10 nanometer resolution", en, npj Quantum Information **10**, 16 (2024).

[24] S. Rajendran, N. Zobrist, A. O. Sushkov, R. Walsworth, and M. Lukin, "A method for directional detection of dark matter using spectroscopy of crystal defects", Physical Review D **96**, 035009 (2017).

[25] XENON Collaboration, "First indication of solar $^{8}\mathrm{B}$ neutrinos via coherent elastic neutrino-nucleus scattering with XENONnT", Phys. Rev. Lett. **133**, 191002 (2024).

[26] C. A. J. O'Hare, "New Definition of the Neutrino Floor for Direct Dark Matter Searches", Physical Review Letters **127**, 251802 (2021).

[27] R. Ebadi, M. C. Marshall, D. F. Phillips, J. Cremer, T. Zhou, M. Titze, P. Kehayias, M. Saleh Ziabari, N. Delegan, S. Rajendran, A. O. Sushkov, F. J. Heremans, E. S. Bielejec, M. V. Holt, and R. L. Walsworth, "Directional detection of dark matter using solid-state quantum sensing", AVS Quantum Science **4**, 044701 (2022).

[28] S. E. Vahsen, C. A. O'Hare, and D. Loomba, "Directional Recoil Detection", Annual Review of Nuclear and Particle Science **71**, _eprint: https://doi.org/10.1146/annurev-nucl-020821-035016, 189 (2021).

[29] S. E. Vahsen, C. A. J. O'Hare, W. A. Lynch, N. J. C. Spooner, E. Baracchini, P. Barbeau, J. B. R. Battat, B. Crow, C. Deaconu, C. Eldridge, A. C. Ezeribe, M. Ghrear, D. Loomba, K. J. Mack, K. Miuchi, F. M. Mouton, N. S. Phan, K. Scholberg, and T. N. Thorpe, *CYGNUS: Feasibility of a nuclear recoil observatory with directional sensitivity to dark matter and neutrinos*, arXiv:2008.12587 [astro-ph, physics:hep-ex, physics:hep-ph, physics:physics], **december** 2020.

[30] M. C. Marshall, M. J. Turner, M. J. H. Ku, D. F. Phillips, and R. L. Walsworth, "Directional detection of dark matter with diamond", Quantum Science and Technology **6**, 024011 (2021).

[31] D. G. Ang, X. Liu, J. Tang, M. Shen, R. Ebadi, and R. Walsworth, "Progress toward a solid-state directional dark matter detector", **in**Mdvdm 2024 proceedings arXiv:2405.01626 (2024), **pages** 18–21.

[32] D. G. Ang, J. Tang, M. Shen, M. Camp, A. Gilpin, G. Liyanage, and R. Walsworth, "Progress toward a solid-state directional dark matter detector", **in**Mdvdm 2025 proceedings arXiv:2508.20482 (2025), **pages** 18–21.

[33] N. Kurinsky, T. C. Yu, Y. Hochberg, and B. Cabrera, "Diamond detectors for direct detection of sub-GeV dark matter", Physical Review D **99**, 123005 (2019).

[34] I. Kim, N. Kurinsky, H. Kagan, S. Boyd, and G. Kim, "Athermal phonon collection efficiency in diamond crystals for low mass dark matter detection", Physical Review D **111**, 072009 (2025).

[35] A. Umemoto, T. Iida, M. Yoshino, A. Yoshikawa, and S. Nomura, "Basic characteristics of synthetic-diamond scintillator", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1057**, 168789 (2023).

[36] M. C. Marshall, D. F. Phillips, M. J. Turner, M. J. H. Ku, T. Zhou, N. Delegan, F. J. Heremans, M. V. Holt, and R. L. Walsworth, "Scanning X-Ray Diffraction Microscopy for Diamond Quantum Sensing", Physical Review Applied **16**, 054032 (2021).

[37] J. Wierbik, H. Heimes, C. Notthoff, S. Dutt, T. Alwadi, A. Kiy, P. Mota-Santiago, N. Kirby, and P. Kluth, "Anisotropic fine structure of ion tracks in single crystals", Physical Review B **113**, 035306 (2026).

[38] M. Abdullah, D. Aristizabal Sierra, B. Dutta, and L. E. Strigari, "Coherent elastic neutrino-nucleus scattering with directional detectors", Physical Review D **102**, 015009 (2020).

[39] A. K. Drukier, S. Baum, K. Freese, M. Górski, and P. Stengel, "Paleo-detectors: searching for dark matter with ancient minerals", Physical Review D **99**, 043014 (2019).

[40] S. Baum, P. Stengel, N. Abe, J. F. Acevedo, G. R. Araujo, Y. Asahara, F. Avignone, L. Balogh, L. Baudis, Y. Boukhtouchen, J. Bramante, P. A. Breur, L. Caccianiga, F. Capozzi, J. I. Collar, R. Ebadi, T. Edwards, K. Eitel, A. Elykov, R. C. Ewing, K. Freese, A. Fung, C. Galelli, U. A. Glasmacher, A. Gleason, N. Hasebe, S. Hirose, S. Horiuchi, Y. Hoshino, P. Huber, Y. Ido, Y. Igami, N. Ishikawa, Y. Itow, T. Kamiyama, T. Kato, B. J. Kavanagh, Y. Kawamura, S. Kazama, C. J. Kenney, B. Kilminster, Y. Kouketsu, Y. Kozaka, N. A. Kurinsky, M. Leybourne, T. Lucas, W. F. McDonough, M. C. Marshall, J. M. Mateos, A. Mathur, K. Michibayashi, S. Mkhonto, K. Murase, T. Naka, K. Oguni, S. Rajendran, H. Sakane, P. Sala, K. Scholberg, I. Semenec, T. Shiraishi, J. Spitz, K. Sun, K. Suzuki, E. H. Tanin, A. Vincent, N. Vladimirov, R. L. Walsworth,

and H. Watanabe, "Mineral detection of neutrinos and dark matter. A whitepaper", Physics of the Dark Universe **41**, 101245 (2023).

[41] S. Kim, P. London, D. Yang, L. B. Hughes, J. Ahlers, S. Meynell, W. J. Mitchell, K. Mukherjee, and A. C. Bleszynski Jayich, "Scalable nanoscale positioning of highly coherent color centers in prefabricated diamond nanostructures", Nature Communications **16**, 9803 (2025).

[42] J. Tang, Z. Yin, C. A. Hart, J. W. Blanchard, J. T. Oon, S. Bhalerao, J. M. Schloss, M. J. Turner, and R. L. Walsworth, "Quantum diamond microscope for dynamic imaging of magnetic fields", AVS Quantum Science **5**, 044403 (2023).

[43] J. M. Schloss, "Optimizing nitrogen-vacancy diamond magnetic sensors and imagers for broadband sensitivity", Advisors: Ronald L. Walsworth and Isaac L. Chuang, Ph.D. thesis (Massachusetts Institute of Technology, Cambridge, MA, 2019).

[44] D. R. Glenn, R. R. Fu, P. Kehayias, D. Le Sage, E. A. Lima, B. P. Weiss, and R. L. Walsworth, "Micrometer-scale magnetic imaging of geological samples using a quantum diamond microscope", Geochemistry, Geophysics, Geosystems **18**, 3254 (2017).

[45] R. R. Fu, B. A. Maher, J. Nie, P. Gao, T. Berndt, E. Folsom, and T. Cavanaugh, "Pinpointing the mechanism of magnetic enhancement in modern soils using high-resolution magnetic field imaging", Geochem. Geophys. Geosyst. **24**, e2022GC010812 (2023).

[46] M. J. H. Ku, T. X. Zhou, Q. Li, Y. J. Shin, J. K. Shi, C. Burch, L. E. Anderson, A. T. Pierce, Y. Xie, A. Hamo, U. Vool, H. Zhang, F. Casola, T. Taniguchi, K. Watanabe, M. M. Fogler, P. Kim, A. Yacoby, and R. L. Walsworth, "Imaging viscous flow of the Dirac fluid in graphene", Nature **583**, 537 (2020).

[47] M. J. Turner, N. Langellier, R. Bainbridge, D. Walters, S. Meesala, T. M. Babinec, P. Kehayias, A. Yacoby, E. Hu, M. Lon čar, R. L. Walsworth, and E. V. Levine, "Magnetic field fingerprinting of integrated-circuit activity with a quantum diamond microscope", Phys. Rev. Appl. **14**, 014097 (2020).

[48] P. Kehayias, E. V. Levine, L. Basso, J. Henshaw, M. Saleh Ziabari, M. Titze, R. Haltli, J. Okoro, D. R. Tibbetts, D. M. Udoni, E. Bielejec, M. P. Lilly, T.-M. Lu, P. D. D. Schwindt, and A. M. Mounce, "Measurement and simulation of the magnetic fields from a 555 timer integrated circuit using a quantum diamond microscope and finite-element analysis", Phys. Rev. Appl. **17**, 014021 (2022).

[49] D. Le Sage, K. Arai, D. R. Glenn, S. J. DeVience, L. M. Pham, L. Rahn-Lee, M. D. Lukin, A. Yacoby, A. Komeili, and R. L. Walsworth, "Optical magnetic imaging of living cells", Nature **496**, 486 (2013).

[50] H. C. Davis, P. Ramesh, A. Bhatnagar, A. Lee-Gosselin, J. F. Barry, D. R. Glenn, R. L. Walsworth, and M. G. Shapiro, "Mapping the microscale origins of magnetic resonance image contrast with subcellular diamond magnetometry", Nat. Commun. **9**, 131 (2018).

[51] J. F. Barry, M. J. Turner, J. M. Schloss, D. R. Glenn, Y. Song, M. D. Lukin, H. Park, and R. L. Walsworth, "Optical magnetic detection of single-neuron action potentials using quantum defects in diamond", Proc. Natl. Acad. Sci. U.S.A. **113**, 14133 (2016).

[52] D. R. Glenn, K. Lee, H. Park, R. Weissleder, A. Yacoby, M. D. Lukin, H. Lee, R. L. Walsworth, and C. B. Connolly, "Single-cell magnetic imaging using a quantum diamond microscope", Nat. Methods **12**, 736 (2015).

[53] R. W. de Gille, J. M. McCoey, L. T. Hall, J.-P. Tetienne, E. P. Malkemper, D. A. Keays, L. C. Hollenberg, and D. A. Simpson, "Quantum magnetic imaging of iron organelles within the pigeon cochlea", Proc. Natl. Acad. Sci. U.S.A. **118**, e2112749118 (2021).

[54] K. Arai, A. Kuwahata, D. Nishitani, I. Fujisaki, R. Matsuki, Y. Nishio, Z. Xin, X. Cao, Y. Hatano, S. Onoda **andothers**, "Millimetre-scale magnetocardiography of living rats with thoracotomy", Commun. Phys. **5**, 200 (2022).

[55] L. Troise, N. W. Hansen, C. Olsson, J. L. Webb, L. Tomasevic, J. Achard, O. Brinza, R. Staacke, M. Kieschnick, J. Meijer **andothers**, "In vitro recording of muscle activity induced by high intensity laser optogenetic stimulation using a diamond quantum biosensor", AVS Quantum Sci. **4**, 044402 (2022).

[56] Z. Kazi, I. M. Shelby, H. Watanabe, K. M. Itoh, V. Shutthanandan, P. A. Wiggins, and K.-M. C. Fu, "Wide-field dynamic magnetic microscopy using double-double quantum driving of a diamond defect ensemble", Phys. Rev. Appl. **15**, 054032 (2021).

[57] E. Bauch, C. A. Hart, J. M. Schloss, M. J. Turner, J. F. Barry, P. Kehayias, S. Singh, and R. L. Walsworth, "Ultralong dephasing times in solid-state spin ensembles via quantum control", Phys. Rev. X **8**, 031025 (2018).

[58] A. M. Edmonds, C. A. Hart, M. J. Turner, P.-O. Colard, J. M. Schloss, K. S. Olsson, R. Trubko, M. L. Markham, A. Rathmill, B. Horne-Smith **andothers**, "Characterisation of CVD diamond with high concentrations of nitrogen for magnetic-field sensing applications", Mater. Quantum Technol. **1**, 025001 (2021).

[59] G. De Lange, T. Van Der Sar, M. Blok, Z.-H. Wang, V. Dobrovitski, and R. Hanson, "Controlling the quantum dynamics of a mesoscopic spin bath in diamond", Sci. Rep. **2**, 1 (2012).

[60] A. Cox, M. Newton, and J. Baker, "^{13}C, ^{14}N and ^{15}N endor measurements on the single substitutional nitrogen centre (P1) in diamond", Journal of Physics: Condensed Matter **6**, 551 (1994).

[61] *COMSOL Multiphysics® v. 6.0*, 2021.

[62] Z. Yin, "Advanced quantum magnetic microscopy with solid-state spin ensemble in diamond", phdthesis (University of Maryland, 2026).

[63] A. Buades, B. Coll, and J.-M. Morel, "Non-local means denoising", Image Process. Line **1**, 208 (2011).

[64] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising", **in**2005 ieee comput. soc. conf. comput. vis. pattern recognit. (cvpr) **volume** 2 (IEEE, 2005), **pages** 60–65.

[65] F. Sigmund, C. Massner, P. Erdmann, A. Stelzl, H. Rolbieski, M. Desai, S. Bricault, T. P. Wörner, J. Snijder, A. Geerlof **andothers**, "Bacterial encapsulins as orthogonal compartments for mammalian cell engineering", Nat. Commun. **9**, 1990 (2018).

[66] M. Winklhofer, and J. L. Kirschvink, "A quantitative assessment of torque-transducer models for magnetoreception", J. R. Soc. Interface **7**, S273 (2010).

[67] C. E. Diebel, R. Proksch, C. R. Green, P. Neilson, and M. M. Walker, "Magnetite defines a vertebrate magnetoreceptor", Nature **406**, 299 (2000).

[68] M. M. Walker, C. E. Diebel, C. V. Haugh, P. M. Pankhurst, J. C. Montgomery, and C. R. Green, "Structure and function of the vertebrate magnetic sense", Nature **390**, 371 (1997).

[69] J. A. Feijó, *Bioelectricity in plants: from so simple a beginning*, 2023.

[70] D. S. Damineli, M. T. Portes, and J. A. Feijó, "Oscillatory signatures underlie growth regimes in arabidopsis pollen tubes: computational methods to estimate tip location, periodicity, and synchronization in growing cells", Journal of experimental botany **68**, 3267 (2017).

[71] D. S. Damineli, M. T. Portes, and J. A. Feijó, "Electrifying rhythms in plant cells", Current opinion in cell biology **77**, 102113 (2022).

[72] M. H. Weisenseel, R. Nuccitelli, and L. F. Jaffe, "Large electrical currents traverse growing pollen tubes.", The Journal of cell biology **66**, 556 (1975).

[73] N. G. Sepulveda, B. J. Roth, and J. Wikswo, "Current injection into a two-dimensional anisotropic bidomain", Biophys. J. **55**, 987 (1989).

[74] S. Rossi, and B. E. Griffith, "Incorporating inductances in tissue-scale models of cardiac electrophysiology", Chaos **27** (2017).

[75] D. R. Hochbaum, Y. Zhao, S. L. Farhi, N. Klapoetke, C. A. Werley, V. Kapoor, P. Zou, J. M. Kralj, D. Maclaurin, N. Smedemark-Margulies, and et al., "All-optical electrophysiology in mammalian neurons using engineered microbial rhodopsins", Nat. Methods **11**, 825 (2014).

[76] P. Sundaram, A. Nummenmaa, W. Wells, D. Orbach, D. Orringer, R. Mulkern, and Y. Okada, "Direct neural current imaging in an intact cerebellum with magnetic resonance imaging", Neuroimage **132**, 477 (2016).

[77] S. J. Ihle, S. Girardin, T. Felder, T. Ruff, J. Hengsteler, J. Duru, S. Weaver, C. Forró, and J. Vörös, "An experimental paradigm to investigate stimulation dependent activity in topologically constrained neuronal networks", Biosens. Bioelectron. **201**, 113896 (2022).

[78] T.-L. Wee, Y.-K. Tzeng, C.-C. Han, H.-C. Chang, W. Fann, J.-H. Hsu, K.-M. Chen, and Y.-C. Yu, "Two-photon excited fluorescence of nitrogen-vacancy centers in proton-irradiated type ib diamond", The Journal of Physical Chemistry A **111**, PMID: 17705460, 9379 (2007).

[79] P. Balasubramanian, C. Osterkamp, Y. Chen, X. Chen, T. Teraji, E. Wu, B. Naydenov, and F. Jelezko, "Dc magnetometry with engineered nitrogen-vacancy spin ensembles in diamond", Nano Lett. **19**, 6681 (2019).

[80] N. Arunkumar, K. S. Olsson, J. T. Oon, C. A. Hart, D. B. Bucher, D. R. Glenn, M. D. Lukin, H. Park, D. Ham, and R. L. Walsworth, "Quantum logic enhanced sensing in solid-state spin ensembles", Physical Review Letters **131**, 100801 (2023).

[81] D. G. Ang, J. Tang, M. Shen, M. Titze, G. Liyanage, J. Chapman, C. Bharathulwar, A. T. Gilpin, T. Terlier, E. S. Bielejec **andothers**, "Multi-scale reconstruction of single-ion damage tracks in diamond via nitrogen-vacancy centers", arXiv preprint arXiv:2606.23621 (2026).

[82] T. Tagami, and P. B. O'Sullivan, "Fundamentals of fission-track thermochronology", Reviews in Mineralogy and Geochemistry **58**, 19 (2005).

[83] R. Fleischer, P. Price, and R. Walker, "Solid-state track detectors: applications to nuclear science and geophysics", Annual Review of Nuclear Science **15**, 1 (1965).

[84] K. Nordlund, S. J. Zinkle, A. E. Sand, F. Granberg, R. S. Averback, R. E. Stoller, T. Suzudo, L. Malerba, F. Banhart, W. J. Weber, F. Willaime, S. L. Dudarev, and D. Simeone, "Primary radiation damage: a review of current understanding and models", Journal of Nuclear Materials **512**, 450 (2018).

[85] C. P. Karger, and P. Peschke, "RBE and related modeling in carbon-ion therapy", Physics in Medicine & Biology **63**, 01TR02 (2017).

[86] M. Lang, F. Djurabekova, N. Medvedev, M. Toulemonde, and C. Trautmann, "Fundamental phenomena and applications of swift heavy ion irradiations", Comprehensive Nuclear Materials **1**, 10.1016/b978-0-12-803581-8.11644-3 (2020).

[87] H. Amekura, A. Chettah, K. Narumi, A. Chiba, Y. Hirano, K. Yamada, S. Yamamoto, A. A. Leino, F. Djurabekova, K. Nordlund, N. Ishikawa, N. Okubo, and Y. Saitoh, "Latent ion tracks were finally observed in diamond", en, Nature Communications **15**, 1786 (2024).

[88] R. E. Lake, A. Persaud, C. Christian, E. S. Barnard, E. M. Chan, A. A. Bettiol, M. Tomut, C. Trautmann, and T. Schenkel, "Direct formation of nitrogen-vacancy centers in nitrogen doped diamond along the trajectories of swift heavy ions", Applied Physics Letters **118**, 084002 (2021).

[89] W. Liu, A. A. M. Leino, A. Persaud, Q. Ji, K. Jhuria, E. S. Barnard, S. Aloni, C. Trautmann, M. Tomut, R. Wunderlich, C. Nozais, S. Mogan, H. Ocker, N. Anand, Z. Hao, F. Djurabekova, and T. Schenkel, "Optical and spin properties of nitrogen vacancy centers in diamond formed along high-energy heavy ion tracks", en, Communications Materials **6**, Publisher: Nature Publishing Group, 242 (2025).

[90] G. A. Araujo, L. Baudis, N. Bowden, J. Chapman, A. Erickson, M. G. Perez, A. A. Hecht, S. C. Hedges, P. Huber, V. Ivanov, I. Jovanovic, G. A. Khodaparast, B. A. Magill, J. M. Mateos, M. Morrison, N. W. G. Smith, P. Stengel, S. Surani, N. Vladimirov, K. Walkup, C. Wittweg, and X. Zhang, *Nuclear recoil detection with color centers in bulk lithium fluoride*, arXiv:2503.20732 [nucl-ex], **march** 2025.

[91] F. Fávaro de Oliveira, D. Antonov, Y. Wang, P. Neumann, S. A. Momenzadeh, T. Häußermann, A. Pasquarelli, A. Denisenko, and J. Wrachtrup, "Tailoring spin defects in diamond by lattice charging", en, Nature Communications **8**, 15409 (2017).

[92] T. Lühmann, R. John, R. Wunderlich, J. Meijer, and S. Pezzagna, "Coulomb-driven single defect engineering for scalable qubits and spin sensors in diamond", Nature Communications **10**, 4956 (2019).

[93] N. Delegan, S. J. Whiteley, T. Zhou, S. L. Bayliss, M. Titze, E. Bielejec, M. V. Holt, D. D. Awschalom, and F. J. Heremans, "Deterministic nanoscale quantum spin-defect implantation and diffraction strain imaging", Nanotechnology **34**, 385001 (2023).

[94] P. Räcke, R. Wunderlich, J. W. Gerlach, J. Meijer, and D. Spemann, "Nanoscale ion implantation using focussed highly charged ions", en, New Journal of Physics **22**, Publisher: IOP Publishing, 083028 (2020).

[95] A. Healey, A. Stacey, B. Johnson, D. Broadway, T. Teraji, D. Simpson, J.-P. Tetienne, and L. Hollenberg, "Comparison of different methods of nitrogen-vacancy layer formation in diamond for wide-field quantum microscopy", Physical Review Materials **4**, 104605 (2020).

[96] C. Santori, P. E. Barclay, K.-M. C. Fu, and R. G. Beausoleil, "Vertical distribution of nitrogen-vacancy centers in diamond formed by ion implantation and annealing", Physical Review B **79**, 125313 (2009).

[97] S. Pezzagna, B. Naydenov, F. Jelezko, J. Wrachtrup, and J. Meijer, "Creation efficiency of nitrogen-vacancy centres in diamond", en, New Journal of Physics **12**, 065017 (2010).

[98] S. Onoda, M. Haruyama, T. Teraji, J. Isoya, W. Kada, O. Hanaizumi, and T. Ohshima, "New application of nv centers in cvd diamonds as a fluorescent nuclear track detector", Physica Status Solidi (a) **212**, 2641 (2015).

[99] S. Onoda, K. Tatsumi, M. Haruyama, T. Teraji, J. Isoya, W. Kada, T. Ohshima, and O. Hanaizumi, "Diffusion of vacancies created by high-energy heavy ion strike into diamond", Physica Status Solidi (a) **214**, 1700160 (2017).

[100] P. Räcke, L. Pietzonka, J. Meijer, D. Spemann, and R. Wunderlich, "Vacancy diffusion and nitrogen-vacancy center formation near the diamond surface", Applied Physics Letters **118** (2021).

[101] P. Chu, D. Hansen, B. Doyle, K. Jobe, R. Lopez-Aguado, M. Shoga, and D. Walsh, "Ion-microbeam probe of high-speed shift registers for SEE analysis-part i: sige", IEEE Transactions on Nuclear Science **53**, 1574 (2006).

[102] F. Martin, "Electrical rotation of quadrupole lenses", Nuclear Instruments and Methods in Physics Research **189**, 93 (1981).

[103] P. Räcke, L. Pietzonka, J. Meijer, D. Spemann, and R. Wunderlich, "Vacancy diffusion and nitrogen-vacancy center formation near the diamond surface", Applied Physics Letters **118**, 204003 (2021).

[104] K. Nordlund, F. Djurabekova, and G. Hobler, "Large fraction of crystal directions leads to ion channeling", Physical Review B **94**, 214109 (2016).

[105] S. Sangtawesin, B. L. Dwyer, S. Srinivasan, J. J. Allred, L. V. H. Rodgers, K. De Greve, A. Stacey, N. Dontschuk, K. M. O'Donnell, D. Hu, D. A. Evans, C. Jaye, D. A. Fischer, M. L. Markham, D. J. Twitchen, H. Park, M. D. Lukin, and N. P. de Leon, "Origins of diamond surface noise probed by correlating single-spin measurements with surface spectroscopy", Physical Review X **9**, 031052 (2019).

[106] N. B. Manson, M. Hedges, M. S. J. Barson, R. Ahlefeldt, M. W. Doherty, H. Abe, T. Ohshima, and M. J. Sellars, "NV^-–N^+ pair centre in 1b diamond", New Journal of Physics **20**, 113037 (2018).

[107] M. S. Janson, "Hydrogen diffusion and ion implantation in silicon carbide", Ph.D. thesis (KTH-Royal Institute of Technology, 2003).

[108] M. S. Janson, *SIIMPL: Simulation of Ion IMPLantation*, https://github.com/msjanson01/siimpl, Accessed: 2025-07-23, 2022.

[109] J. F. Ziegler, M. Ziegler, and J. Biersack, "SRIM – the stopping and range of ions in matter (2010)", Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **268**, 19th International Conference on Ion Beam Analysis, 1818 (2010).

[110] J. Koike, D. Parkin, and T. Mitchell, "Displacement threshold energy for type iia diamond", Applied Physics Letters **60**, 1450 (1992).

[111] J. Beyer, T. B. E. Grønbech, J. Zhang, K. Kato, and B. Brummerstedt Iversen, "Electron density and thermal motion of diamond at elevated temperatures", en, Acta Crystallographica Section A: Foundations and Advances **79**, 41 (2023).

[112] J. A. Mitchell, F. Abdeljawad, C. Battaile, C. Garcia-Cardona, E. A. Holm, E. R. Homer, J. Madison, T. M. Rodgers, A. P. Thompson, V. Tikare, E. Webb, and S. J. Plimpton, "Parallel simulation via SPPARKS of on-lattice kinetic and Metropolis Monte Carlo models for materials processing", Modelling and Simulation in Materials Science and Engineering **31**, 055001 (2023).

[113] S. T. Alsid, J. F. Barry, L. M. Pham, J. M. Schloss, M. F. O'Keeffe, P. Cappellaro, and D. A. Braje, "Photoluminescence decomposition analysis: a technique to characterize N-V creation in diamond", Physical Review Applied **12**, 044003 (2019).

[114] S. Breuer, and P. Briddon, "Ab initio investigation of the native defects in diamond and self-diffusion", Physical Review B **51**, 6984 (1995).

[115] A. Tejero-Cantero, J. Boelts, M. Deistler, J.-M. Lueckmann, C. Durkan, P. J. Gonçalves, D. S. Greenberg, and J. H. Macke, *SBI – A toolkit for simulation-based inference*, arXiv:2007.09114 [cs.LG], **july** 2020.

[116] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Neural spline flows", Advances in Neural Information Processing Systems **32** (2019).

[117] D. Sheppard, P. Xiao, W. Chemelewski, D. D. Johnson, and G. Henkelman, "A generalized solid-state nudged elastic band method", The Journal of Chemical Physics **136** (2012).

[118] E. Bauch, S. Singh, J. Lee, C. A. Hart, J. M. Schloss, M. J. Turner, J. F. Barry, L. M. Pham, N. Bar-Gill, S. F. Yelin, and R. L. Walsworth, "Decoherence of ensembles of nitrogen-vacancy centers in diamond", Phys. Rev. B **102**, 134210 (2020).

[119] J. Tang, C. A. Roncaioli, A. M. Edmonds, A. Davidsson, C. A. Hart, M. L. Markham, and R. L. Walsworth, "Characterization of low-nitrogen quantum diamond for pulsed magnetometry applications", Frontiers in Quantum Science and Technology **4**, 1701548 (2025).

[120] R. Khan, P. Martineau, B. Cann, M. Newton, H. Dhillon, and D. Twitchen, "Color alterations in cvd synthetic diamond with heat and uv exposure: implications for color grading and identification", Gems & Gemology **46**, 18 (2010).

[121] R. U. A. Khan, P. M. Martineau, B. L. Cann, M. E. Newton, and D. J. Twitchen, "Charge transfer effects, thermo- and photochromism in single crystal cvd synthetic diamond", Journal of Physics: Condensed Matter **21**, 364214 (2009).

[122] L. S. Hounsome, R. Jones, P. M. Martineau, D. Fisher, M. J. Shaw, P. R. Briddon, and S. Oberg, "Origin of brown coloration in diamond", Phys. Rev. B **73**, 125203 (2006).

[123] N. Fujita, R. Jones, S. Öberg, and P. Briddon, "Large spherical vacancy clusters in diamond – origin of the brown colouration?", Diamond and Related Materials **18**, Proceedings of Diamond 2008, the 19th European Conference on Diamond, Diamond-Like Materials, Carbon Nanotubes, Nitrides and Silicon Carbide, 843 (2009).

[124] C. Glover, M. E. Newton, P. Martineau, D. J. Twitchen, and J. M. Baker, "Hydrogen incorporation in diamond: The nitrogen-vacancy-hydrogen complex", Phys. Rev. Lett. **90**, 185507 (2003).

[125] C. Glover, M. E. Newton, P. M. Martineau, S. Quinn, and D. J. Twitchen, "Hydrogen incorporation in diamond: The vacancy-hydrogen complex", Phys. Rev. Lett. **92**, 135502 (2004).

[126] J. A. Cuenca, S. Mandal, E. L. Thomas, and O. A. Williams, "Microwave plasma modelling in clamshell chemical vapour deposition diamond reactors", Diamond and Related Materials **124**, 108917 (2022).

[127] P. M. Martineau, S. C. Lawson, A. J. Taylor, S. J. Quinn, D. J. F. Evans, and M. J. Crowder, "Identification of synthetic diamond grown using chemical vapor deposition", Gems. Gemmol. **40**, 2 (2004).

[128] I. Friel, S. Clewes, H. Dhillon, N. Perkins, D. Twitchen, and G. Scarsbrook, "Control of surface and bulk crystalline quality in single crystal diamond grown by chemical vapour deposition", Diamond and Related Materials **18**, Proceedings of Diamond 2008, the 19th European Conference on Diamond, Diamond-Like Materials, Carbon Nanotubes, Nitrides and Silicon Carbide, 808 (2009).

[129] T. Schuelke, and T. A. Grotjohn, "Diamond polishing", Diamond and Related Materials **32**, 17 (2013).

[130] J. R. Hird, and J. E. Field, "Diamond polishing", Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **460**, 3547 (2004).

[131] C. B. Hartland, "A study of point defects in cvd diamond using electron paramagnetic resonance and optical spectroscopy", phdthesis (University of Warwick, 2014).

[132] C. Roncaioli, C. A. Hart, R. Walsworth, and D. P. Fahey, "All-in-one quantum diamond microscope for sensor characterization", AVS Quantum Science **6** (2024).

[133] E. Eisenach, J. Barry, L. Pham, R. Rojas, D. R. Englund, and D. Braje, "Broadband loop gap resonator for nitrogen vacancy centers in diamond", Review of Scientific Instruments **89**, https://doi.org/10.1063/1.5037465 (2018).

[134] A. Przybylski, B. Thiel, J. Keller-Findeisen, B. Stock, and M. Bates, "Gpufit: an open-source toolkit for gpu-accelerated curve fitting", Scientific reports **7**, 15722 (2017).

[135] D. Aude Craik, P. Kehayias, A. Greenspon, X. Zhang, M. Turner, J. Schloss, E. Bauch, C. Hart, E. Hu, and R. Walsworth, "Microwave-assisted spectroscopy technique for studying charge state in nitrogen-vacancy ensembles in diamond", Physical Review Applied **14**, 014009 (2020).

[136] A. Glazer, J. Lewis, and W. Kaminsky, "An automatic optical imaging system for birefringent media", Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **452**, 2751 (1996).

[137] T. Graziosi, A. Rathmill, R. L. Patel, N. Palmer, A. M. Edmonds, M. L. Markham, and D. J. Twitchen, "Control of strain in thin diamond nitrogen-vacancy layers for magnetometry applications", **in**Quantum sensing, imaging, and precision metrology **volume** 12447, edited by J. Scheuer, and S. M. Shahriar (International Society for Optics **and** Photonics, 2023), 124470F.

[138] N. Aslam, G. Waldherr, P. Neumann, F. Jelezko, and J. Wrachtrup, "Photo-induced ionization dynamics of the nitrogen vacancy defect in diamond investigated by single-shot charge state detection", New Journal of Physics **15**, 013064 (2013).

[139] N. Manson, and J. Harrison, "Photo-ionization of the nitrogen-vacancy center in diamond", Diamond and related materials **14**, 1705 (2005).

[140] P. P. Laissue, R. A. Alghamdi, P. Tomancak, E. G. Reynaud, and H. Shroff, "Assessing phototoxicity in live fluorescence imaging", en, Nature Methods **14**, 657 (2017).

[141] J. Rovny, S. Gopalakrishnan, A. C. B. Jayich, P. Maletinsky, E. Demler, and N. P. de Leon, "Nanoscale diamond quantum sensors for many-body physics", Nature Reviews Physics **6**, 753 (2024).

[142] K. H. Jæger, and A. Tveito, "Deriving the bidomain model of cardiac electrophysiology from a cell-based model; properties and comparisons", Front. Physiol. **12**, 2439 (2022).

[143] S. Rohr, "Role of gap junctions in the propagation of the cardiac action potential", Cardiovasc. Res. **62**, 309 (2004).

[144] S. Dhein, and A. Salameh, "Remodeling of cardiac gap junctional cell–cell coupling", Cells **10**, 2422 (2021).

[145] S. Barral, and M. A. Kurian, "Utility of induced pluripotent stem cells for the study and treatment of genetic diseases: focus on childhood neurological disorders", Front. Mol. Neurosci. **9**, 78 (2016).

[146] J. R. Holzer, L. E. Fong, V. Y. Sidorov, J. P. Wikswo, and F. Baudenbacher, "High resolution magnetic images of planar wave fronts reveal bidomain properties of cardiac tissue", Biophys. J. **87**, 4326 (2004).

[147] S. Rossi, A. Buccarello, P. R. Ershler, R. L. Lux, S. Callegari, D. Corradi, L. Carnevali, A. Sgoifo, M. Miragoli, E. Musso **andothers**, "Effect of anisotropy on ventricular vulnerability to unidirectional block and reentry by single premature stimulation during normal sinus rhythm in rat heart", Am. J. Physiol. Heart Circ. Physiol. **312**, H584 (2017).

[148] C. d. Oliveira Nunes, "Molecular and biophysical bases of intracellular electric fields in pollen tubes", Ph.D. dissertation (University of Maryland, College Park, 2022).

[149] J. A. Feijó, *Private communication on pollen tube bioelectric current models*, Private communication, Discussion with Jiashen Tang, **july** 2026.

[150] A. Jenkins, S. Baumann, H. Zhou, S. A. Meynell, Y. Daipeng, K. Watanabe, T. Taniguchi, A. Lucas, A. F. Young, and A. C. Bleszynski Jayich, "Imaging the breakdown of ohmic transport in graphene", Physical Review Letters **129**, 087701 (2022).

[151] M. Holt, R. Harder, R. Winarski, and V. Rose, "Nanoscale Hard X-Ray Microscopy Methods for Materials Studies", Annual Review of Materials Research **43**, _eprint: https://doi.org/10.1146/annurev-matsci-071312-121654, 183 (2013).

[152] *Heliotis AG heliCam™ C3 Lock-in Camera*, 2023.

[153] A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one", Multiscale Model. Simul. **4**, 490 (2005).

[154] W. J. P. van Enckevort, and H. G. M. Lochs, "Photoluminescence tomography as a method to image point-defect distributions in crystals: nitrogen-vacancy pairs in synthetic diamonds", Journal of Applied Physics **64**, 434 (1988).

[155] R. Fastenau **andothers**, *Diffusion limited reactions in crystalline solids*, techreport (Technische Hogeschool Delft (Netherlands), 1982).

[156] T. Yamamoto, S. Onoda, T. Ohshima, T. Teraji, K. Watanabe, S. Koizumi, T. Umeda, L. P. McGuinness, C. Müller, B. Naydenov, F. Dolde, H. Fedder, J. Honert, M. L. Markham, D. J. Twitchen, J. Wrachtrup, F. Jelezko, and J. Isoya, "Isotopic identification of engineered nitrogen-vacancy spin qubits in ultrapure diamond", Physical Review B **90**, 081117 (2014).

[157] T. Schuelke, and T. A. Grotjohn, "Diamond polishing", Diamond and Related Materials **32**, 17 (2013).

[158] P. Appel, E. Neu, M. Ganzhorn, A. Barfuss, M. Batzer, M. Gratz, A. Tschöpe, and P. Maletinsky, "Fabrication of all diamond scanning probes for nanoscale magnetometry", Review of Scientific Instruments **87** (2016).

[159] T. X. Zhou, R. J. Stöhr, and A. Yacoby, "Scanning diamond NV center probes compatible with conventional AFM technology", Applied Physics Letters **111** (2017).

[160] C. Lee, E. Gu, M. Dawson, I. Friel, and G. Scarsbrook, "Etching and micro-optics fabrication in diamond using chlorine-based inductively-coupled plasma", Diamond and Related Materials **17**, 1292 (2008).

[161] G. Kresse, and J. Hafner, "Ab initio molecular dynamics for liquid metals", Physical Review B **47**, 558 (1993).

[162] G. Kresse, and J. Furthmüller, "Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set", Physical Review B **54**, 11169 (1996).

[163] P. E. Blöchl, "Projector augmented-wave method", Physical Review B **50**, 17953 (1994).

[164] J. P. Perdew, "Generalized gradient approximations for exchange and correlation: a look backward and forward", Physica B: Condensed Matter **172**, 1 (1991).

[165] J. P. Perdew, K. Burke, and M. Ernzerhof, "Generalized gradient approximation made simple", Physical Review Letters **77**, 3865 (1996).

[166] B. Slepetz, and M. Kertesz, "Divacancies in diamond: a stepwise formation mechanism", Physical Chemistry Chemical Physics **16**, 1515 (2014).

[167] A. Barfuss, M. Kasperczyk, J. Kölbl, and P. Maletinsky, "Spin-stress and spin-strain coupling in diamond-based hybrid spin oscillator systems", Phys. Rev. B **99**, 174102 (2019).