Graphics Processing Unit-Accelerated Numerical Simulations and Theoretical Study of Qubit Dynamics in Realistic Systems

PhD Thesis

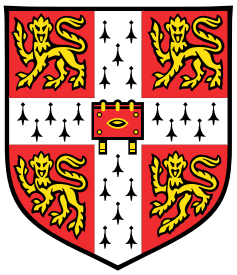

Aleksander Adam Lasek

Department of Physics
University of Cambridge

This thesis is submitted for the degree of
Doctor of Philosophy

Girton College

September 2021

Declaration

The work presented in this thesis was carried out at the Thin Film Magnetism / Quantum Information group at the Cavendish Laboratory, University of Cambridge, between October 2017 and July 2021. I state that this dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specified in the text. I further state that no part of my dissertation has already been submitted, or, is being concurrently submitted for any degree, diploma or other qualification at the University of Cambridge or any other University or similar academic institution. The thesis does not exceed the prescribed word limit for the Physics and Chemistry Degree Committee of 60,000 words, including summary/abstract, tables, footnotes and appendices, but excluding table of contents, photographs, diagrams, figure captions, list of figures/diagrams, list of abbreviations/acronyms, bibliography and acknowledgments.

Aleksander Adam Lasek
September 2021

Acknowledgements

I would like to acknowledge my supervisor, Professor Crispin Barnes. The work on Root-of-SWAP with Surface Acoustic Waves described in this thesis is suggested by him and based on Crispin's substantial research on the subject, dating back over two decades.
I would like to acknowledge Dr Aleksey Andreev for his suggestion and hints on the momentum space eigensolver, and Dr Hugo Lepage, Dr David Arvidsson-Shukur and Dr Thierry Ferrus for collaboration.
I acknowledge Dr Jacek Mosakowski, as his PhD project served as a seed and an inspiration for the single qubit control part of this thesis.
The main part of this work, described in Chapters 3 and 4 was done in collaboration with Hugo Lepage and Crispin Barnes, with input from David Arvidsson-Shukur and Echo Kexin Zhang. The work described in Chapter 5 was done in collaboration with Thierry Ferrus, who has conducted the experiments. The work described in Chapter 8 was done in collaboration with Nicole Yunger-Halpern, David Arvidsson-Shukur, Hugo Lepage and Seth Lloyd. The results presented here are published under our names [140, 13], or in preparation for publication.
This work was supported financially by the Project for Developing Innovation Systems of the Ministry of Education, Culture, Sports, Science and Technology (MEXT), Engineering and Physical Sciences Research Council (EPSRC) and Hitachi via CASE studentships RG 94632.
I would also like to thank all people mentioned above for support and invaluable discussion related and unrelated to this work.

Abstract

Quantum computers are thought to be the future of computation, using the properties of quantum mechanics to solve problems intractable to classical computers. Quantum computing leverages non-classical properties, such as entanglement, to achieve an exponential improvement in computational power. A quantum computer would enable us to address many real-world problems, such as how to synthesize fertilizers more efficiently; how to combat global warming; or to simulate protein folding in biological systems. Although much work has been done to describe the use and implementation of entanglement generation theoretically, it is still a challenge to develop such protocols experimentally. The bulk of this work is focused on creating Graphics Processing Unit (GPU)-accelerated computer simulations of quantum systems with advanced numerical and analytical techniques. Simulations can guide experiments attempting to create building blocks of quantum computers - qubits and their control devices. However, simulation of more realistic device setups in two dimensional systems has been facing problems owing to the space and time domain scaling associated with the solutions of the many-particle time dependent Schrodinger equation (TDSE). Nevertheless, recent advances in computer hardware performance has made previously intractable two-particle problems readily solvable. I have developed custom GPU-accelerated software based on a staggered-leapfrog algorithm that opens up new possibilities of simulating two-dimensional two-particle systems accurately.

I focus on three research projects. Firstly, optimally defining a charge-based solid state qubit, and controlling it in a simple and experimentally achievable way, while accounting for imperfections of the waveform generators. I simulate the physical qubit on a fine-grained lattice, and propose an innovative control scheme that accounts for finite rise/fall time of the experimental apparatus, while being relatively fast and resulting in very high operation fidelity. An optimal pulsing scheme with rise time-dependent parameters is found, and shown to be able to achieve an arbitrary qubit rotation. Since the proposed pulse sequence reduces to sine waves to minimize total pulse duration, it is straightforward to implement experimentally, and easily generalisable to different systems. I also show how the fidelity remains sufficiently high independently of the initial qubit state. The proposed sequence can

even reduce errors caused by charge noise under certain conditions. Readout techniques are discussed as well, and found to not present significant issues.

Secondly, I aid the effort to create a Surface Acoustic Wave quantum computer prototype by describing how to produce an universal quantum gate set with a Root-of-SWAP operation used as a physical two-qubit gate. Using realistic parameters, it is shown how this operation can be performed with high fidelity. Previous work has been done to simulate a proposed Root-of-SWAP method in one dimension [165] - this work focuses on extending this to two dimensions.
We find that the method of generating Root-of-SWAP mentioned above breaks down in two dimensions- unwanted excitations are introduced in the extra dimension, causing a phase difference to appear, and thus ruining coherence of the state.
I propose to implement the Root-of-SWAP operation via a tunneling interaction across the effective double dot instead. This was previously considered, however was thought to be unstable against variations in tunnel barrier height, which has exponential impact on the speed of the quantum operation. Using newly available computing power, we were able to run detailed two dimensional simulations investigating this method and its robustness against variations in the double dot potential. We find that the method produces high fidelity Root-of-SWAP states, and is robust against small variations in the tunnel barrier. Additionally, we find a relation between the tunnel barrier height and spin measurement probability, providing a way for experimentalists to estimate an actual device barrier indirectly.

Finally, I theoretically model and simulate transport through a single electron transistor (SET) device. It is shown that a single donor structure can reliably be engineered from doped quantum dots by taking advantage of the tunability of the electron tunneling rates as well as the interplay, at low temperatures, between disorder conferred by randomness in dopant distribution and electron-electron interaction originating from the high doping concentration. It is possible to electrostatically isolate a single donor from the large ensemble of dopants. I investigate how such a complex system is expected to conduct, and verify a hypothesis that two donors take part in the transport by numerically reproducing the experimental measurements. Finally, it is shown that this device can be used as a single atom detector of the charge occupancy of a nearby capacitively coupled double quantum dot. While this final part does not make use of the GPU-accelerated software, it is still closely related to the rest of this work, and the theme of modeling realistic quantum devices.

Table of contents

List of figures

Part I

Main results

The main part of this thesis presents the work on Graphics Processing Unit-accelerated numerical simulations and theoretical study of qubit dynamics in realistic systems. I describe why simulations of realistic quantum systems are important for the future of quantum computing, explain the theory and methods used in the GPU-accelerated staggered-leapfrog software, and present research projects that utilise it. My work on quantum metrology, that, while interesting on its own, does not fit into the above theme, is presented in Part II of this thesis.

Chapter 1

Introduction

I Quantum computation

Since the dawn of civilization, humanity has always strived for a faster and better way to calculate. From the antiquity's geared mechanical calculation devices like the Antikythera mechanism that calculated the date and tracked the positions of planets and the Moon [77], through more complex and advanced machines like Charles Babbage's differential engine that would solve differential equations [44], and finally to modern general computers based on semiconductors, computers have been getting faster and more accessible throughout the ages. This has never been more obvious than now, with computational power roughly doubling every 2 years, according to Moore's law [28]. However, we seem to be close to hitting the quantum limit with semiconductor technology -as the device scale becomes smaller and smaller to fit in more transistors in a chip, quantum effects start to come into play [124], [222], [197]. As distances become as small as 10 nanometers or even less, phenomena like quantum tunneling start being a problem. After all, you cannot have a reliable computer if the state of the machine can uncontrollably and probabilistically change due to an electron tunneling out to where it's not supposed to be, and once you get down to the smallest building block, the atom, the architecture cannot be shrunk down any more.

Efforts are made to find a technology to overcome these limitations and keep improving classical computers, however another way would be to embrace the quantumness and exploit quantum mechanics to create a general fully programmable quantum computer. A quantum computer would use effects such as superposition and entanglement to perform calculations much faster than a classical computer ever could [74]. While we are still rather far away from a general quantum computer, many useful quantum algorithms have been found that solve important problems exponentially faster than any classical algorithm. For example, the Grover algorithm [98] can search an unsorted database, the Deutsch–Jozsa algorithm

can determine whether a function is constant or balanced [49], and perhaps the most famous Shor algorithm can perform integer factorisation [203] - a problem whose difficulty virtually all modern encryption relies on.

Building a robust quantum computer could significantly change the world we live in, solving previously untractable problems. However, while the theory behind the idea is sound, implementing it in practice is incredibly difficult. Building a reliable single qubit (a quantum bit) is expensive and not necessary scalable to thousands or millions of qubits that would be required. Indeed, quantum computers in operation today only have tens of qubits, an IBM holding the current record with 54 superconducting transmon qubits [9]. While that number is technically sufficient to achieve so called "quantum supremacy" - the point at which the quantum computer is able to solve a problem that no classical counterpart is able to solve - this is only true if you ignore errors that occur during the computation. Despite the best experimental efforts, the final fidelity achieved by IBM is still sub- 1 % due to issues with scalability. While a single operation can be made with very high fidelity, as the computer grows in size, the errors multiply and get out of control. With the physical size of the system growing, it becomes harder to have all the qubits connected while keeping errors low. Therefore, if we wish to make quantum supremacy practically achievable, great care must be taken to avoid decoherence - the loss of entanglement. The whole system must be cooled to extremely low temperatures close to absolute zero, and separated from the environment. The coherence times are typically up to milliseconds in most commonly used qubits [127, 33], and the computation must be performed before the qubits decohere and lose entanglement. While error correction techniques can be used to reduce errors [85], they require a much greater amount of qubits to implement, making them currently not viable. This places very significant difficulties on actually building a functioning quantum computer.

There is a search to find new feasible physical implementations. Examples of possible implementations are linear optics [126], [187], trapped ions [60], [112], nuclear spin [118], and many others [198], [226], [234]. Electrons carried by Surface Acoustic Waves (SAWs) [19] are one such idea that we specifically focus on in this work.

DiVincenzo [52] has laid out the following 5 criteria that a practical quantum computer must achieve in his seminal work:

1. *Scalable physical system with well characterized qubits*
   A qubit is a quantum 2 state system, like spin states of a spin-$\frac{1}{2}$ particle or ground and excited energy states of an atom. The system must contain a collection of qubits. "Well characterized" means that the exact qubit dynamics and couplings to other qubits should be known, and the occupation probability of any higher energy states must be small enough to bring any errors under a required threshold. Scalability means that it's

possible to build a device with a large enough number of qubits, and that these qubits should be able to entangle with one another so that the general quantum state spans the whole complex Hilbert space.

2. *The ability to initialize the state of the qubits to a simple fiducial state, such as* $|000\rangle$
   To perform any calculation, the initial state of the system must be well known, therefore there must be a way of reliably initialising the system to some reference state. This could be achieved, depending on the qubit nature, by cooling the system to put it in the ground state for energy level qubits, or by applying a strong magnetic field for spin qubits.

3. *Long relevant coherence times, much longer than the gate operation time*
   Coherence times characterize the dynamics of a qubit in contact with its environment. Decoherence means a loss of quantumness, and needs to be avoided. This can be generally done by cooling the system to extremely low temperatures to prevent thermal excitations, and by keeping interaction with the environment to a minimum (until the final measurement). The coherence times do not need to be of the same order as the entire calculation thanks to error correction [184], but they are still a major limiting factor.

4. *A universal set of quantum gates*
   To perform a general computation, one needs a set of physical unitary operations (gates) that can be used to express a general unitary transformation. There are many sets of gates that satisfy this requirement, typically containing an arbitrary single qubit SU(2) rotation and a two-qubit gate, for example CNOT or Root-of-SWAP.

5. *A qubit-specific measurement capability*
   At the end of any quantum computation, a measurement must be performed on each qubit, which will yield a classical 1 or 0 result. Therefore, there must be a reliable way of measuring the state of each qubit with low enough error rates. However, perfect efficiency is not required, as multiple computations can be performed.

II Quantum operations

Any calculation on a quantum computer can be expressed using the universal unitary gates set [18]. There are many such sets, typically expressed in terms of single qubit operations and a two qubit operation that is the physical gate specific to the implementation.

To define what all those gates do, first let us define what a qubit is. In general, a qubit is a system that can take 2 discrete values, $|0\rangle$ and $|1\rangle$. In this, the qubit is exactly like a classical bit. It would typically be physically represented by a two-level system, with one state assigned to $|0\rangle$ and the other to $|1\rangle$. What makes the qubit different from a bit is the ability to exist in a superposition state

$$|\psi\rangle = a|0\rangle + b|1\rangle \tag{1.1}$$

where a and b are complex numbers and $a^2 + b^2 = 1$ due to normalisation. This means that there is a a^2 probability to measure the qubit in $|0\rangle$ state and b^2 probability to measure it in a $|1\rangle$ state.

Such a general single qubit state can be geometrically represented on a Bloch sphere, as per Fig. 1.1 below. The North and South poles of the sphere correspond to $|0\rangle$ and $|1\rangle$ states respectively, while any superposition of the two is specified by the polar and azimuthal angles, θ and ϕ (Spherical polar coordinates). θ takes values between 0 and π, while ϕ is between 0 and 2π The state can then be written in terms of those angles as:

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle \tag{1.2}$$

where i is the imaginary number.

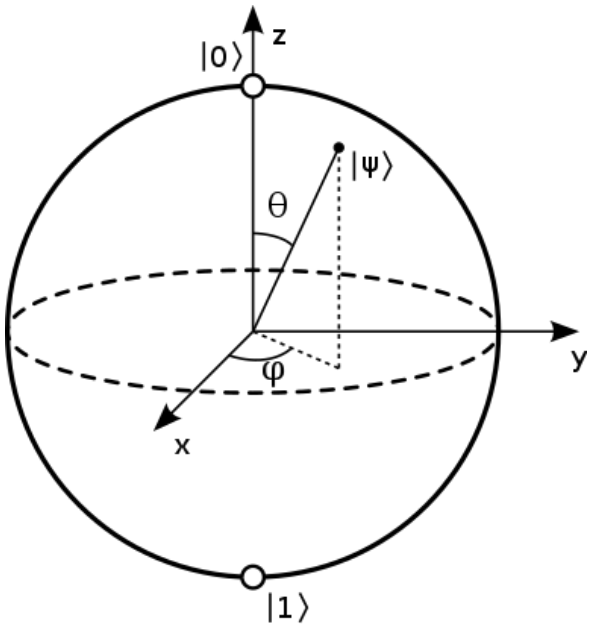

Fig. 1.1 Bloch sphere. The angles θ and ϕ that define the qubit position on the sphere are shown. The logical $|0\rangle$ state is located on the "north pole" of the sphere, while the $|1\rangle$ is on the "south pole". Image taken from [237].

An advantage of this representation is that it maps the qubit operations, described by the somewhat abstract group of unitary and hermitian SU(2) matrices, to 3D rotations on a 2-sphere described by the SO(3) group, which is isomorphic to the former. A single qubit operation then corresponds to rotating the state vector on a Bloch sphere, and can take any superposition of $|0\rangle$ and $|1\rangle$ to any other one. Such single qubit rotations are necessary in any quantum computer implementation.

However, to perform meaningful calculations, there must be a way for the qubits to interact and entangle with one another. This is where two qubit operations come in. A pure non-entangled two qubit state can be represented in braket notation as $|n_1 n_2\rangle$, where n_1 is the state of the first qubit and n_2 is the state of the second qubit. A general two qubit state can then be written as:

$$|\psi\rangle = a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle \tag{1.3}$$

where $a^2 + b^2 + c^2 + d^2 = 1$ due to normalisation. A two qubit operation, or gate, transforms between such two qubit states. It can be shown [51] that a two qubit gate such as CNOT or (Square) Root-of-SWAP forms a universal gate set together with single qubit rotations. Therefore, it is sufficient to only have a physical implementation of those gates to achieve general quantum computation.

The CNOT, or controlled NOT gate, has the following matrix representation in basis mentioned above:

$$CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \tag{1.4}$$

This can be interpreted as : if the first (control) qubit is in a $|1\rangle$ state, then flip the second qubit state. Otherwise, do nothing.

The other universal two qubit gate, Root-of-SWAP, has the following representation:

$$\sqrt{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{2}(1+i) & \frac{1}{2}(1-i) & 0 \\ 0 & \frac{1}{2}(1-i) & \frac{1}{2}(1+i) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{1.5}$$

It can be interpreted as a halfway *SWAP*, where *SWAP* is a gate that swaps qubits 1 and 2. Therefore, Root-of-SWAP applied twice will yield a SWAP. The focus of this work is to investigate the Root-of-SWAP operation as a basic physical to qubit gate.

An N qubit non-entangled state can be represented as $|n_1 n_2 ... n_N\rangle$, where n_N is the state of the Nth qubit. A general quantum computation would use gates from the universal set to take an initial state, typically initialised to $|00...0\rangle$, to some final state that is to be measured for a result:

$$|\psi\rangle = U_1 U_2 ... U_M |00...0\rangle \tag{1.6}$$

where U_M represents some operation from the universal gate set. It is worth noting that the final result is classical - each qubit will be measured as either $|0\rangle$ or $|1\rangle$. The final state probability distribution can be inferred from repeated measurements.

III Outline of work

III.I Single qubit control

Firstly, we investigate control of a single solid-state qubit, trying to find a simple and experimentally realistic control scheme that takes rise/fall time of the waveform generator into account. .

Solid-state quantum devices can be manufactured using the mature semiconductor technology and infrastructure, thus allowing for low costs and availability. Quantum dots in silicon and III-V materials can be made with relatively long coherence times (10s of μ s), making them attractive candidates for qubits [102, 81, 173, 172, 231]. While not quite as good as superconducting transmon qubits with coherence times of up to 100 μ s, it is enough with short enough operation times. Double quantum dots (DQDs), dots with a potential barrier in the middle separating the two halves, can be used to realise a variety of qubits in a straightforward manner [145, 52]. Charge-based and spin-based qubits are commonly used. The charge qubit is formed by having a single electron in the DQD, with the charge distribution information being used to represent the qubit. An excess charge in the left dot represents one logical state, say $|0\rangle$, while charge in the right dot represents a $|1\rangle$ state. The charge can be read out by coupling a single electron transistor (SET) to one of the sides [73]. Alternatively, spin can be used to store the qubit information, although the readout would not be as straightforward, requiring a full state tomography. While this work focuses on the study of charge qubits, as they are easier to measure in practice [91], the results can be generalised to spin qubits as well.

In the case of the charge qubit, it is common to use the two-site localised state model, where we distinguish between the left and right charge states only. However, we must note that in realistic DQD potentials, such fully localised states are quite far from energy

eigenstates. Therefore, requiring perfect localisation will not result in a two-level system, making it unsuitable for a qubit- the higher energy contributions will introduce significant effective noise/decoherence [119]. This simple picture including only localised states will fail to define a high fidelity, high coherence time qubit. We look at an alternative definition here.

In both the charge and spin cases, single qubit operations, which move our qubit on the Bloch sphere, are achieved by using surface gates [173, 172], or nearby electron/hole reservoirs [102, 81], with external voltages applied. Applying voltage pulses on these gates alters the relative energies of the two wells, allowing the electron to move between them, and thus change state. Alternatively, one can think of the voltage gradient tilting the DQD potential landscape, which will ”push” the electron towards the lower side.

These pulses alter the relative energies of the two quantum wells as well as the tunnel barrier between them, allowing the qubit to oscillate between the basis states. Such a manipulation of charge qubits has been successfully implemented in GaAs/AlGaAs [102, 56, 119, 173], Si/SiGe [200] and Si:P [189] devices. Alternatively, photon-assisted tunneling can be used with AC pulses generally in the MHz-GHz range . But for defining a set of unitary gate operations, one should be able to address any quantum state on the Bloch sphere with high fidelity. The tunnel barrier and the quantum dot levels are not independently controllable, neither by a static gate voltage nor an AC signal. In gated or lithographically defined architectures, the applied electric field always contains some planar components that couple various parts of the device and capacitive coupling can be significant, making local control by an electric field challenging in nanoscale structures. Consequently, there is a need for defining a realistic pulsing scheme that can take into account the specificities of the architecture.

III.II Root-of-SWAP with Surface Acoustic Waves (SAWs)

Secondly, we investigate theoretically and by numerical simulation the possible ways of implementing a two qubit Root-of-SWAP operation in a SAW quantum computer implementation as proposed by Barnes [19]. Firstly, a scheme proposed by Owen and Barnes [165], where electrons collide in a harmonic confining potential, will be investigated. This method was shown to work in 1D. This will be confirmed, and then extended to 2D to verify that it is still feasible in a more realistic scenario. Additionally, a different method, where the electrons interact by tunneling across a double dot, will be investigated. The methods will then be compared. This work assumes that the SAW system is implemented in a Gallium Arsenide (GaAs) semiconductor structure, but is easily generalisable to other materials, and a significant part is in fact applicable to general double quantum dot systems.

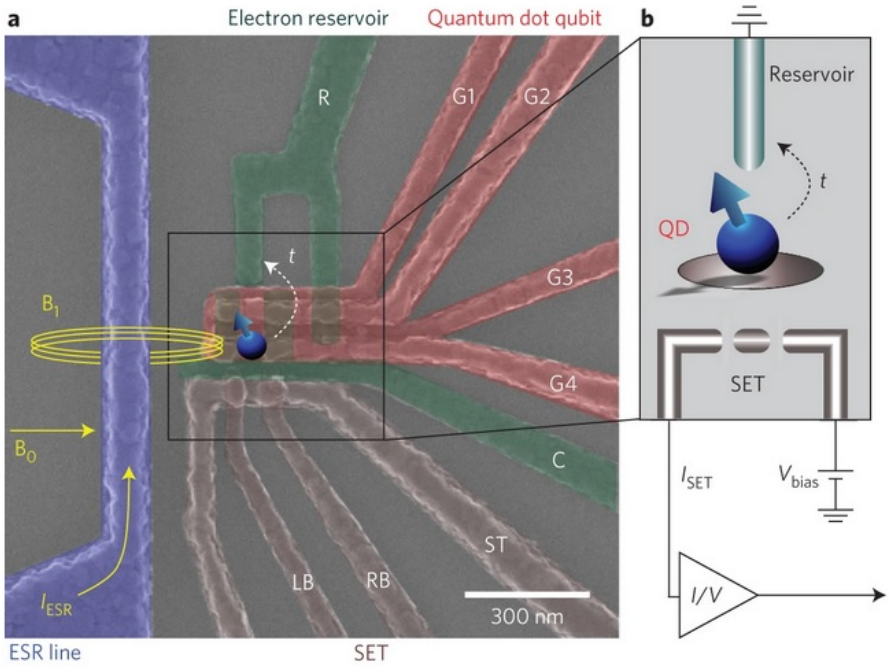

Fig. 1.2 **(a)** False-coloured SEM image of the qubit device. Gate electrodes G1 to G4 are used to control the qubit electrostatically. Gate C helps to confine the qubit, while R is the electron reservoir. This system is shown operating as a single qubit. The electron spin resonsnce (ESR) line seen on the left would not be used when operating as a charge-only qubit (no spin). **(b)** the inset shows a schematic of how an SET positioned under the dot is used as a detector. Image taken from [231].

A large part of this work focuses on investigating a quantum computer implementation proposed by Barnes et al [19]. This implementation would use Surface Acoustic Waves (SAWs) to carry single electrons, to be used as spin qubits (i.e. spin up/down corresponds to $|1\rangle$/ $|0\rangle$ qubit state. SAWs are coupled electromechanical waves that can be generated in piezoelectric materials. An interdigitated transducer consisting of two interlocking comb-shaped arrays of metallic electrodes is placed on such a material, e.g. quartz, and oscillating electric field is applied to every other electrode. This causes a sinusoidal electric field that generates a traveling mechanical wave in the material, which in turn generates a sinusoidal electric field via the piezoelectric effect. This effectively causes a sinusoidal electric field to travel along the material. Such a wave can then be used to capture single electrons to use as qubits from a 2 Dimensional Electron Gas (2DEG). Such a 2DEG can be created in a semiconductor structure - see Fig. 1.3 below.

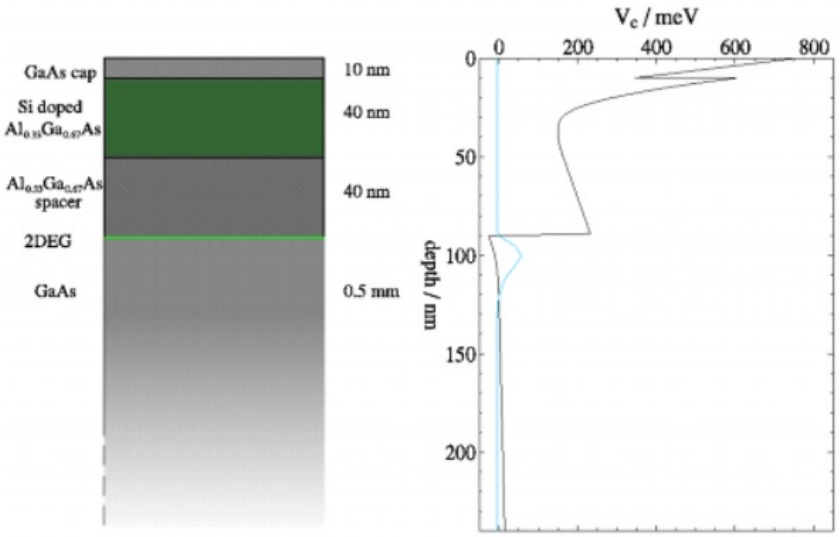

Fig. 1.3 The structure of a typical SAW wafer and its conduction band and electron wave-function is shown on the right. If the well is sufficiently narrow, and the electron density sufficiently low, then only the first sub-band in the confinement direction will be occupied. Image from [97].

Proposed device schematic

Single electrons carried by SAW potential would travel along one-dimensional channels, with single qubit rotations being realised by magnetic split gates, and two qubit Root-of-SWAP operation being realised by bringing the channels together for some time, so that the electrons can interact and entangle as per Eq. 1.5. While such an implementation has not

yet been fully demonstrated experimentally, single electron transport by SAWs already has [158], [105], [120], [31], [63]. A protocol for Fermionic qubit transport and spin control has been described by Arvidsson-Shukur, H.V Lepage et al [14]. Another important feature of such a system is the constant confinement of electrons by the SAW potential. Fermionic particles have the advantage over photons in that they interact and entangle much more easily. However, since they are massive, they would usually tend to spread ballistically, which is an unwanted effect that would cause increasing errors in the system. But with SAW confinement present, this potential drawback is completely negated. The system is also readily scalable - electrons can be reflected back at the end of their channels, while the single and two qubit gates are adjusted to continue the computation. Many parallel channels can be used, with the Root-of-SWAP operation enabling swapping and interaction between them. Overall, such a SAW system is a promising quantum computer implementation. While relative lack of maturity holds this technology back, for example issues with keeping the electrons cold enough to stay confined in the SAW potential, and reading them out with high fidelity, there has been a lot of progress towards low-error single- and two-qubit operations [217].

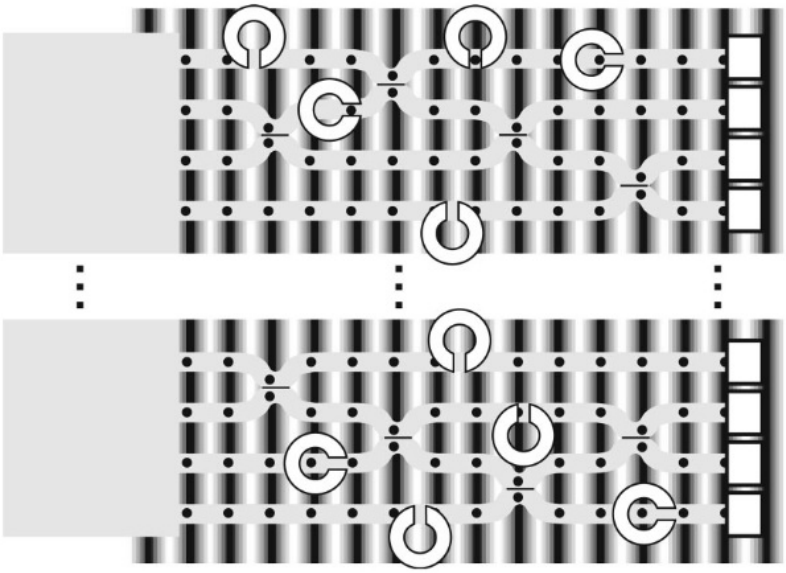

Fig. 1.4 Device schematic diagram showing the gate pattern layout for a SAW quantum gate network. Black and white vertical lines represent the SAW effective potential. The network of gray lines represents a set of Q1DC. Black dots represent qubits. White squares on the right hand side represent readout gates. Image from [19].

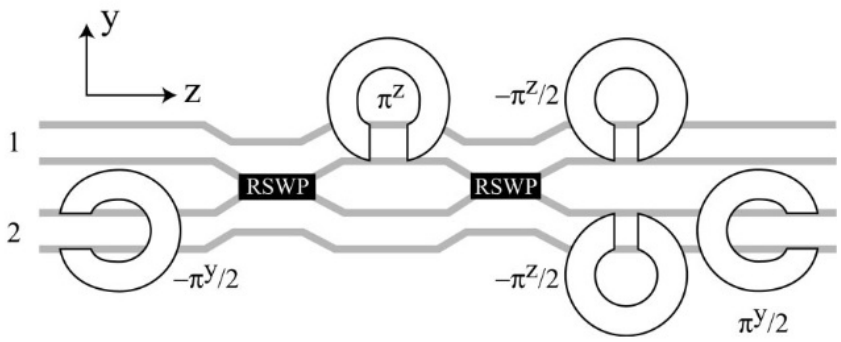

Fig. 1.5 CNOT in terms of Root-of-SWAP and single qubit rotations. The diagram illustrates the equivalency of the two operations, as one can be used to express the other. It also shows an example of how the universal gate set can be used to express an operation/gate from outside the set. Image from [19].

III.III Mapping the charge states of semi-isolated DQDs using a single donor device

We demonstrate that single donor structure can reliably be engineered from doped quantum dots by taking advantage of the tunability of the electron tunneling rates as well as the interplay, at low temperatures, between disorder conferred by randomness in dopant distribution and electron-electron interaction originating from the high doping concentration [171, 170, 22, 202]. For this purpose, phosphorous-doped silicon quantum dots offers ideal conditions for experimentation. This work on mapping the charge states of semi-isolated double quantum dots using a single donor device consists of an experimental part, conducted by Dr Thierry Ferrus, and a simulation and theory part, conducted by me. The goal is to show that single donors can be isolated in a single electron transistor (SET) device, which affects its characteristics. This device can then be used as a detector to map the charge states of a nearby DQD. I show that isolation of single donors explains the observed transport characteristics of the device, and that one can indeed use it to read out the charge state of a capacitively coupled DQD. Such a detector is an important part in the toolbox of many quantum computing architectures.

Chapter 2

Methods and Theory

I Numerical methods

Simulating the dynamics of a general nonrelativistic quantum system requires solving the Schrödinger equation, by the 6^{th} postulate of quantum mechanics. The time evolution is described by the Time-Dependent Schrödinger Equation (TDSE) [26]:

$$i\hbar\frac{\partial}{\partial t}|\psi(\boldsymbol{r},t)\rangle = \hat{H}|\psi(\boldsymbol{r},t)\rangle\,, \tag{2.1}$$

where $\hat{H}$ is the Hamiltonian operator.

The stationary states (eigenstates) of the system are found by solving the Time-Independent Schrödinger Equation (TISE) [26]:

$$\hat{H}|\psi(\boldsymbol{r},t)\rangle = E|\psi(\boldsymbol{r},t)\rangle\,, \tag{2.2}$$

where $\hat{H}$ is the Hamiltonian operator and E is the eigenenergy. Finding the eigenstates (performing the eigendecomposition) is necessary to find the ground/excited state of the system that could be assigned to a $|0\rangle/|1\rangle$ qubit state. Knowing the eigenstates and their energies also helps to understand, visualise and characterise the system dynamics.

Analytical solutions to the Schrödinger equation exist only in a limited number of simple cases. General solutions to the equation require the use of numerical methods. To solve the Schrödinger equation numerically, a space/time discretisation needs to be applied as a computer only has finite memory, and has to be able to index the space/time coordinate in a discrete way. Solving the problem numerically also poses many challenges - sophisticated algorithms must be used to avoid compounding errors and unfeasible computation times.

In this work, a custom Graphical Processing Unit (GPU)-accelerated solver written in C++ and CUDA [164], and rewritten later in Python, was used for majority of results. This program was already being used in the group at the time my work began. We have rewritten it, together with Dr Hugo Lepage, from C# to C++ and for the GPU kernel from OpenCL to CUDA, as well as extended its functionality from 1 particle systems to N particles in any number dimensions. The code was later again rewritten in Python instead of C++, to make it easier to develop and avoid issues with explicit memory management. Various other minor improvements and tweaks have been added as well. Work on this method was done in the group before by Jacek Mosakowski, Matthew Dean and Edmund Owen, and while we acknowledge their contributions, there is no direct link between their work and the program used here. I have also implemented the eigensolver method in momentum space, described in this chapter, for finding the initial condition for the time-dependent solver.
CUDA library was used to perform matrix diagonalisation on the GPU, and FFTW libray was used to calculate the Discrete Fourier Transform. On top of the main program, MATLAB (with its in-build functions) was also used for smaller calculations and testing. The relevant code can be found in the Appendix.
As for the hardware, majority of the calculations were done on a server jointly build by me and Hugo Lepage for this project specifically. The server specifications are:

- GPUs: 7 x GTX 1080 Ti 11GB GDDR5X
- Memory: 128 GB of 2133 MHz DDR4 ECC RAM
- CPUs: 2 x Intel 8 Core Xeon E5-2620
- Storage: 512 GB SDD + 2TB HDD

The GPUs were later upgraded to the faster GTX 2080 Ti, while the specifications of other components did not change significantly. This GPU upgrade yielded an approximately 20% speedup in code execution.

I.I Spatial Discretisation - Eigensolver

The nonrelativistic 1D Hamiltonian $\hat{H}$ of a spinless particle in a potential is given by [26]:

$$\hat{H} = -\frac{\hbar^2}{2m}\frac{d^2}{dx^2} + V(x) \tag{2.3}$$

where m is the particle mass and $\hbar$ is the reduced Planck constant. Note that in condensed matter systems considered here, the mass m must be replaced with an effective mass m_{eff}.

The effective mass is the mass the particle seems to have in the bulk material, owing to the influence of the periodic atomic potential. For our purposes, it can be considered a constant of the material [125]. Its value in GaAs is $0.067m_e$ in terms of the free electron rest mass $m_e = 9.10938356 \cdot 10^{-31}$ kg. To discretise this equation, the position is discretised as an array of some length N. The potential V and wave function $\psi(x)$ will therefore also be arrays of length N, while the Hamiltonian matrix H will have N^2 elements. The choice of N will depend on the particular problem - the greater it is, the closer the solution will be to the continuum limit. However, increasing the array size will also increase memory and time required to solve the problem, proportionally to the Hamiltonian size. In practice, N is increased until an acceptably small error in wave function normalisation is reached, while the simulation run time is reasonable.

The treatment of second derivative operator $\frac{d^2}{dx^2}$ appearing in the equation requires extra attention - it needs to be expressed as a finite difference. By Taylor-expanding the wave function in forward and backward directions to second order we get:

$$|\psi(x+dx)\rangle = |\psi(x)\rangle + \frac{d|\psi\rangle}{dx}dx + \frac{1}{2}\frac{d^2|\psi\rangle}{dx^2}dx^2 + \mathscr{O}(dx^3), \tag{2.4}$$

$$|\psi(x-dx)\rangle = |\psi(x)\rangle - \frac{d|\psi\rangle}{dx}dx + \frac{1}{2}\frac{d^2|\psi\rangle}{dx^2}dx^2 + \mathscr{O}(dx^3). \tag{2.5}$$

By adding these equations and solving for the second order differential operator we get:

$$\frac{d^2|\psi\rangle}{dx^2} = \frac{|\psi(x+dx)\rangle - 2|\psi(x)\rangle + |\psi(x-dx)\rangle}{dx^2} + \mathscr{O}(dx^3). \tag{2.6}$$

This equation allows us to express the Hamiltonian as a discretised matrix, since $|\psi(x+dx)\rangle$ is just the next array element of $|\psi(x)\rangle$ (and $|\psi(x-dx)\rangle$ is the previous one). The elements that extend beyond the array to the left and right are handled by applying a Dirichlet (infinite hard wall) boundary condition - these elements are effectively set to be 0. This requires that the solution (wave function $|\psi(x)\rangle$) goes to 0 sufficiently fast on the side boundaries. It is enforced by choosing a potential such that the wave function is localised toward the center of region of interest. This is not a stringent condition, as the same is also true analytically for bound states.

Writing down the Hamiltonian matrix explicitly gives:

$$\hat{H}|\psi(x)\rangle = \begin{pmatrix} V(x_1)+H_k & -2H_k & \dots & 0 \\ -2H_k & V(x_2)+H_k & -2H_k & \dots \\ \vdots & \vdots & \ddots & \dots \\ 0 & \dots & -2H_k & V(x_N)+H_k \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \dots \\ x_N \end{pmatrix} \tag{2.7}$$

where $H_k = -\frac{\hbar^2}{2m}\frac{d^2}{dx^2}$ is the kinetic term.

This tridiagonal matrix can then be numerically diagonalised using standard packages like Eigen [66] to find the eigenvectors (wave functions) and their eigenenergies. This method can be easily extended to higher dimensions and more particles. in 3D, the second derivative operator $\frac{d^2}{dx^2}$ is replaced by ∇^2 operator. The wave function can still be represented by a vector of size N^3, with some indexing scheme to assign 3D position to the array elements. A row-major order is used here, where the last index is varying the fastest. In general, a 3D Hamiltonian for M particles is:

$$\hat{H} = \sum_{n=1}^{M} -\frac{\hbar^2}{2m_n}\nabla_n^2 + V(x_1, x_2, ...x_M), \tag{2.8}$$

where n is the particle index. The wave function size scales as N^{dM}, where N is the number of points in each dimension, d is the number of dimensions and M is the number of particles. The Hamiltonian matrix is naturally the square of this size. This exponential growth of required memory causes many particle problems to become intractable very fast. While the above method works well for 1 particle in 2D or 2 particles in 1D, for 2 particles in 2D, which is the interest of this work, the problem is not directly solvable on currently existing hardware. For about 100 points which are required for stability of the time dependent solver, the Hamiltonian would take thousands of terabytes of memory, well beyond the capabilities of any existing GPU hardware.

Since the matrix is large but sparse - tridiagonal - sparse solver packages were tried to solve this problem. However, these methods suffer from some limitations. They are often only able to find a single eigenstate, closest to an initial guess, and the available routines do not benefit from GPU acceleration. Moreover, it was found that they can have issues with converging for such large matrix sizes. In general, we are interested in finding multiple eigenstates, and wish to leverage the GPU hardware available. Therefore, an alternative was found - to solve the Schrödinger equation in momentum space.

I.II Eigensolver in Momentum Space

Eigenfunctions that are well localised converge rather fast in momentum space. It is found that around 10 points in each dimension gives a very good result that is stable when used as an initial condition for the time-dependent solver. Therefore, by solving a Fourier-transformed Schrödinger equation with reduced number of points for discretisation, and then inverse Fourier-transforming the result, padding it beforehand to about 100 real-space points, accurate eigenfunctions can be found for 2 particle 2D systems. A real space wave function in 1D can

be written as a sum (integral in continuous case) of the momentum eigenfunctions:

$$\psi(\boldsymbol{r}) = \int_{\boldsymbol{k}-space} \phi(\boldsymbol{k})\psi_k(\boldsymbol{r})d(k), \tag{2.9}$$

where $\psi_k(\boldsymbol{r})$ are the momentum eigenfunctions.
Due to the canonical relation between momentum and position operators:

$$\boldsymbol{p} = -i\hbar\frac{\partial}{\partial \boldsymbol{r}}, \tag{2.10}$$

the eigenfunctions are:

$$\psi_k(\boldsymbol{r}) = \frac{1}{\sqrt{2\pi}^d} e^{i\boldsymbol{k}\cdot\boldsymbol{r}}, \tag{2.11}$$

where d is the dimension index. Therefore, the real space wave function can be written as:

$$\psi(\boldsymbol{r}) = \frac{1}{\sqrt{2\pi}^d} \int_{\boldsymbol{k}-space} \phi(\boldsymbol{k}) e^{i\boldsymbol{k}\cdot\boldsymbol{r}} d(k). \tag{2.12}$$

This is just an inverse Fourier transform:

$$\psi(\boldsymbol{r}) = \mathscr{F}^{-1}\big(\phi(\boldsymbol{k})\big), \tag{2.13}$$

where $\mathscr{F}^{-1}$ denotes the inverse Fourier transform.

By similar logic, the inverse is also true:

$$\phi(\boldsymbol{k}) = \mathscr{F}\big(\psi(\boldsymbol{r})\big), \tag{2.14}$$

where $\mathscr{F}$ denotes the Fourier transform.

For a numerical solution, momentum space can be discretised in multiples of $\frac{2\pi}{L}$, where L is a real space extent over which the wave function should have decayed close to 0 near the edges. Defining some cutoff N, momentum will take the values:

$$\boldsymbol{k}_n = n\frac{2\pi}{L}, \tag{2.15}$$

where $n = [-N, -N+1, ..., 0, 1, ..., n]$.

When we Fourier-transform the Schrödinger equation into momentum space, the kinetic term $-\frac{\hbar^2}{2m}\nabla^2$ will become $\frac{\hbar^2}{mL^2}(\cosh(kL) - 1)$ owing to discreteness and the canonical momentum-position relation (Eq. 2.10). The Fourier-transformed potential matrix element becomes

$$\hat{V}_{kq} = (2\pi)^{\frac{d}{2}} \sum_{r}^{N} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}} V(\boldsymbol{r}) e^{i\boldsymbol{k}\cdot\boldsymbol{r}} = (2\pi)^{\frac{d}{2}} \sum_{r}^{N} V(\boldsymbol{r}) e^{i(\boldsymbol{k}-\boldsymbol{q})\cdot\boldsymbol{r}} = \mathscr{F}\left(V(\boldsymbol{r})\right)(k-q). \tag{2.16}$$

The elements of the Hamiltonian matrix to be solved by diagonalisation in this method take the form

$$H_{kq} = \frac{\hbar^2}{mL^2}\left[\cosh(kL) - 1\right]\delta_{kq} + \hat{V}_{kq}. \tag{2.17}$$

With a fine enough resolution,, the kinetic term will become indistinguishable from a simpler form of $\frac{k^2}{2m}$. The Fourier-transformed potential matrix element will read:

$$H_{kq} = \frac{k^2}{2m}\delta_{kq} + \hat{V}_{kq}. \tag{2.18}$$

For the discretised case, the Fourier integral is evaluated using a Fast Fourier Transform (FFT) algorithm. The Hamiltonian matrix equation to be solved by diagonalisation in this method looks like:

$$\hat{H}\left|\phi(\boldsymbol{k})\right\rangle = \begin{pmatrix} \frac{k_1^2}{2m} + \mathscr{F}\big(V(\boldsymbol{r})\big)(0) & \mathscr{F}\big(V(\boldsymbol{r})\big)(k_1-k_2) & \dots & \mathscr{F}\big(V(\boldsymbol{r})\big)(k_1-k_N) \\ \mathscr{F}\big(V(\boldsymbol{r})\big)(k_2-k_1) & \frac{k_2^2}{2m} + \mathscr{F}\big(V(\boldsymbol{r})\big)(0) & \dots & \mathscr{F}\big(V(\boldsymbol{r})\big)(k_2-k_N) \\ \vdots & \vdots & \ddots & \vdots \\ \mathscr{F}\big(V(\boldsymbol{r})\big)(k_N-k_1) & \mathscr{F}\big(V(\boldsymbol{r})\big)(k_N-k_2) & \dots & \frac{k_N^2}{2m} + \mathscr{F}\big(V(\boldsymbol{r})\big)(0) \end{pmatrix} \begin{pmatrix} k_1 \\ k_2 \\ \dots \\ k_N \end{pmatrix} \tag{2.19}$$

Note that this matrix is not sparse as was the case in real space. This could be thought of as reducing the memory space (smaller matrix) at the cost of computational complexity - every matrix element has to be calculated (this is actually reduced by almost half owing to the matrix being Hermitian :$\hat{H} = \hat{H}^\dagger$. However due to the efficiency of FFT algorithms, the problem is still relatively fast to solve. This method extends readily to two or more particles by indexing the wave function accordingly as described in the previous subsection. A two-particle 2D problem with 10 points in each dimension can be solved in tens of minutes on a modern desktop computer, and gives an accurate enough result to be used as a starting point for the time-dependent solver and keep the errors small (<1%). It is found that the above method of solving the Hamiltonian is very competitive compared to sparse matrix diagonalisation in a wide variety of realistic potentials, while being able to take advantage of GPU acceleration readily, and being able to find a large number of eigenstates (up to the

momentum mode cutoff imposed). It is however required that the underlying potential be well-approximated by a limited number of momentum eigenfunctions.

I.III Temporal Discretisation - Staggered Leapfrog Algorithm

To solve the TDSE numerically (Eq. 2.1), the time derivative $\frac{\partial}{\partial t}|\psi(\boldsymbol{r},t)\rangle$ must be discretised. Some standard approaches include the Euler or Crank-Nicholson methods. However for the case of Eq. 2.1, these methods either diverge fast, or are very computationally intensive [16]. Here a Staggered Leapfrog method is used instead, as seen in a paper by Maestri and Landau [153] on 2 particle wave packet simulation/animation. This algorithm is self correcting - the normalisation (wave function modulus) error will oscillate about zero. The amplitude of this oscillation can be brought to acceptable level (<1% of modulus, which is equal to 1 for wave functions) with small enough step in time and space discretisation. The derivation of this algorithm for 2 spinless particles in 1D based on [153] is as follows. Natural units are used here for simplicity ($\hbar = c = 1$). The TDSE equation to be solved (Eq. 2.1) with explicit Hamiltonian is:

$$i\hbar\frac{\partial}{\partial t}|\psi(x_1,x_2,t)\rangle = \hat{H}|\psi(x_1,x_2,t)\rangle, \tag{2.20}$$

$$\hat{H} = -\frac{1}{2m_1}\frac{\partial^2}{\partial {x_1}^2} - \frac{1}{2m_2}\frac{\partial^2}{\partial {x_2}^2} + V(x_1,x_2). \tag{2.21}$$

The wave function can be discretised in time and space by using an evenly spaced grid :

$$\psi(x_1,x_2,t) = \psi(x_1 = l\Delta x_1, x_2 = m\Delta x_2, t = n\Delta t) = \psi^n_{l,m}, \tag{2.22}$$

where n,l,m are integers and Δ denotes step size in space/time. In this discrete notation, and by using the finite difference expression for the second derivative operator as in subsection2.1.1, the right hand side of Eq. 2.20 becomes:

$$\hat{H}\psi = -\frac{\psi_{l+1,m} - 2\psi_{l,m} + \psi_{l-1,m}}{2m_1\Delta x_1^2} - \frac{\psi_{l,m+1} - 2\psi_{l,m} + \psi_{l,m-1}}{2m_2\Delta x_2^2} + V_{l,m}\psi_{l,m}. \tag{2.23}$$

The time derivative in Eq. 2.20 can be written by approximating the formal time evolution operator:

$$\psi^{n+1}_{l,m} = e^{-i\Delta t\hat{H}}\psi^n_{l,m} \approx (1 - it\Delta\hat{H})\psi^n_{l,m}. \tag{2.24}$$

However, this simplistic approximation is unstable and will diverge, as the eigenvalue equation has modulus $\sqrt{1+E^2\Delta t^2}$ - it will increase with time. One improvement that can be made is to use a central difference algorithm to expand the following:

$$\psi_{l,m}^{n+1} - \psi_{l,m}^{n-1} = (e^{-i\Delta t\hat{H}} - e^{i\Delta t\hat{H}})\psi_{l,m}^{n} \cong -2i\Delta t\hat{H}\psi_{l,m}^{n}. \tag{2.25}$$

In conjunction with Eq. 2.23 this gives:

$$\begin{aligned}\psi_{l,m}^{n+1} \cong \psi_{l,m}^{n-1} - 2i\big[\big\{\big(\frac{1}{m_1}+\frac{1}{m_2}\big)\lambda + \Delta t V_{l,m}\big\}\psi_{l,m}^{n} - \lambda\big\{\frac{1}{m_1}(\psi_{l+1,m}^{n}+\psi_{l-1,m}^{n}) \\ + \frac{1}{m_2}(\psi_{l,m+1}^{n}+\psi_{l,m-1}^{n})\big\}\big],\end{aligned} \tag{2.26}$$

where $\Delta x_1 = \Delta x_2$ is assumed for simplicity (this will usually be the case, and always is in this work), and $\lambda = \frac{\Delta t}{\Delta x^2}$. This equation gives an explicit solution, where only two past time values must be stored simultaneously. While Eq. 2.26 produces a stable and accurate solution, it is found that it does not conserve probability very well. A final improvement is made by separating the wave function into real and imaginary parts to make use of that extra degree of freedom:

$$\psi_{l,m}^{n} = u_{l,m}^{n} + iv_{l,m}^{n} \tag{2.27}$$

The algorithm in Eq. 2.26 then separates into a pair of coupled equations:

$$\begin{aligned}u_{l,m}^{n+1} = u_{l,m}^{n-1} + 2\big[\big\{\big(\frac{1}{m_1}+\frac{1}{m_2}\big)\lambda + \Delta t V_{l,m}\big\}v_{l,m}^{n} + \lambda\big\{\frac{1}{m_1}(u_{l+1,m}^{n}-u_{l-1,m}^{n}) \\ + \frac{1}{m_2}(u_{l,m+1}^{n}+u_{l,m-1}^{n})\big\}\big].\end{aligned} \tag{2.28}$$

$$\begin{aligned}v_{l,m}^{n+1} = v_{l,m}^{n-1} - 2\big[\big\{\big(\frac{1}{m_1}+\frac{1}{m_2}\big)\lambda + \Delta t V_{l,m}\big\}u_{l,m}^{n} - \lambda\big\{\frac{1}{m_1}(v_{l+1,m}^{n}-v_{l-1,m}^{n}) \\ + \frac{1}{m_2}(v_{l,m+1}^{n}+v_{l,m-1}^{n})\big\}\big].\end{aligned} \tag{2.29}$$

The eponymous staggering comes from evaluating the real and imaginary parts at staggered times (off by half the time step):

$$[u_{l,m}^{n}, v_{l,m}^{n}] = [Re\psi(x,t), Im\psi(x,t+\frac{1}{2}\Delta t)]. \tag{2.30}$$

The definition of probability density is then also different for integer and half-integer time steps:

$$\rho(x,t) = \mid Re\psi(x,t) \mid^2 + Im\psi(x,t+\frac{1}{2}\Delta t)Im\psi(x,t-\frac{1}{2}\Delta t), \tag{2.31}$$

$$\rho(x,t+\frac{1}{2}\Delta t) = Re\psi(x,t+\frac{1}{2}\Delta t)Re\psi(x,t-\frac{1}{2}\Delta t) + \mid Im\psi(x,t) \mid^2. \tag{2.32}$$

These definitions reduce to the standard one for infinitesimal Δt, and provide an algebraic cancellation of errors so that probability is conserved - the algorithm is self-correcting and (small) errors will oscillate between a positive and negative value rather than build up. Overall, the algorithm is very stable if sufficiently small values of Δt and Δx are used.

Staggered Leapfrog Algorithm in Momentum Space

Similarly to the method of eigensolving the Hamiltonian described above, we can attempt to solve the time-dependent evolution in momentum space. This can lead to similar benefits - the size of the problem is reduced, which is expected to lead to lower memory and computation time requirements. The prerequisites for this method to be appropriate are also similar to the time-independent case - the wave function should always be well-described by a relatively small number of momentum eigenfunctions at any time, so that the computational size can be reduced while preserving accuracy. Again, it is found that a wide range of low-energy wave functions of realistic potentials meet these requirements. Additionally, if the potential is time-dependent (which is usually the case, as otherwise an iterative solution would usually not be necessary), we would require it to change slowly enough in time such that the evolution is adiabatic, and the wave function does not become too steep or squeezed, as that will make it difficult to approximate it with a limited number of momentum eigenfunctions.

To time-evolve the momentum-space wave function $\phi(\boldsymbol{k})$, Equations 2.28, 2.29 need to be Fourier-transformed. Let us first split the momentum wave function $\phi(\boldsymbol{k})$, which we will discretise in time and momentum as $\phi^n_{o,j}$, into the real and imaginary parts:

$$\phi^n_{o,j} = f^n_{o,j} + ig^n_{o,j}. \tag{2.33}$$

Then, we transform Equations 2.28 and2.29 into momentum space:

$$f^{n+1}_{o,j} = f^{n-1}_{o,j} + \Delta t\{2\mathscr{I}\left(\mathscr{F}(V(\boldsymbol{r})\psi(\boldsymbol{r}))_{o,j}\right) + \left(\frac{k_o^2}{m_1} + \frac{k_j^2}{m_2}\right) g^n_{o,j}\}, \tag{2.34}$$

$$g_{o,j}^{n+1} = g_{o,j}^{n-1} - \Delta t\{2\mathscr{R}\big(\mathscr{F}(V(\boldsymbol{r})\psi(\boldsymbol{r}))_{o,j}\big) + \left(\frac{k_o^2}{m_1} + \frac{k_j^2}{m_2}\right) f_{o,j}^n\}, \tag{2.35}$$

where k_o, k_j are momentum values corresponding to discretised points o, j, and $\mathscr{R}/\mathscr{I}$ denote the real/imaginary part of their argument. The momentum is discretised as per Eq. 2.15.

A significant issue arises with this method - at each time step, we need to calculate $\mathscr{F}(V(\boldsymbol{r})\psi(\boldsymbol{r}))_{o,j}$, a Fourier transform element of a point-wise product of real-space potential and wave function. Both of these in general change in time, so this value needs to be recalculated. This forces us to keep the real wave function and potential in memory, and to spend significant overhead on performing the Fourier transforms. Alternatively, the convolution theorem can be used to simplify this calculation, however calculating the convolution at every time step is a significant task as well. The momentum-space method turns out to be significantly advantageous only in some cases. Specifically, the number of momentum modes must be small, meaning that the wave function remains close to the first few energy eigenstates. Secondly, we must have explicitly closed boundary conditions- for example, implementing an infinite square well with this method would require us to actually enforce very high potential on the boundaries, unlike the real-space case, where this is enforced by default. In a very general case, the overhead of performing the discrete Fourier transform every time step tends to outweigh the benefit of working in a smaller computational space. However, if the problem is scaled up to a large enough size, the momentum method will win eventually in terms of performance. We find that overall, the real space method is more general, while the momentum space method has specific cases where it is superior.

I.IV Indexing of the wave function in computer memory

To store the wave function in memory, it is simplest and most efficient to use a one-dimensional array, irrespective of the actual arrangement of dimensions of the physical wave function. We choose to use a row-major order indexing, which means the last dimension is contiguous in memory. We choose to arrange the "dimensions" in the following order, assuming a general case of N_p particles with N_s spin states and N_d spatial positions in d spatial dimensions ($d = x, y, z$) : N_p, N_s, N_x, N_y, N_z, with some of these possibly being equal to 1. This means that z- position is contiguous in memory, while number of particles has the largest stride. The total size of the wave function is $2N_pN_sN_xN_yN_z$ *floats* (the factor of two is

due to the real and imaginary parts of a complex number). The wave function is indexed by a tuple (p,s,i,j,k) of zero-based indices. Given this tuple, we can calculate the corresponding position in memory n:

$$n = k + jN_z + iN_zN_y + sN_zN_yN_x + pN_zN_yN_xN_s. \tag{2.36}$$

Sometimes we wish to do the inverse - calculate the tuple of indices given n. This can be achieved by the following algorithm:

$$\begin{aligned} p &= \lfloor n/(N_zN_yN_xN_s) \rfloor \\ n_1 &= n - p(N_zN_yN_xN_s) \\ s &= \lfloor n_1/(N_zN_yN_x) \rfloor \\ n_2 &= n_1 - s(N_zN_yN_x) \\ i &= \lfloor n_2/(N_zN_y) \rfloor \\ n_3 &= n_2 - i(N_zN_y) \\ j &= \lfloor n_3/N_z \rfloor \\ k &= n_3 - jN_z, \end{aligned} \tag{2.37}$$

where $\lfloor a/b \rfloor$ denotes integer division of a by b (rounding the result down to an integer).

II Inclusion of spin - Second Quantisation

Second quantisation is a formalism used to analyse many-particle systems. It is sometimes, perhaps more descriptively, called the occupation number representation. In this approach, quantum many-body states are represented in the Fock state basis. A Fock (number) state is a state of a well defined number of particles. So-called creation and annihilation operators are introduced to construct the Fock states [7]. This approach is useful, as it accounts for Fermi exclusion principle - Fock space representation explicitly gives the number of particles in a given state. Fermi exclusion, which says that no two fermions can occupy the same quantum state, is applicable in this work, as we are dealing with electrons, which are fermions (spin $\frac{1}{2}$ particles). According to Fermi-Dirac statistics, the fermionic many-body wave function must be anti-symmetric under exchange of any two particles:

$$\psi(\boldsymbol{r_1},...,\boldsymbol{r_i},...\boldsymbol{r_j},...) = -\psi(\boldsymbol{r_1},...,\boldsymbol{r_j},...\boldsymbol{r_i},...) \tag{2.38}$$

We can see how this requirement enforces Fermi exclusion: if particles i and j are identical, then if they occupy the same state, i.e. $\boldsymbol{r_i} = \boldsymbol{r_j}$, then this implies that the whole wave function must be equal to zero - such a state is prohibited. Usually, this symmetrisation is handled by representing the many-particle state as a linear combination of determinants of single-particle states (for fermions). However, in second quantisation, symmetry is automatically enforced by the creation ($c^\dagger$) and annihilation (c) operators. For a single mode state, these operators act the following way:

$$\begin{aligned} c^\dagger |0\rangle &= |1\rangle \\ c|0\rangle &= 0 \\ c^\dagger |1\rangle &= 0 \\ c|1\rangle &= |0\rangle \end{aligned} \tag{2.39}$$

where $|0\rangle$ is a vacuum state, and $|1\rangle$ is a state with 1 particle in it. Note how particle statistics are enforced by operator action - trying to add a particle to already occupied state quenches it (gives a 0). Note that this is not equal to the vacuum state $|0\rangle$ -rather, it is a scalar 0 denoting lack of a quantum state. Similarly, trying to remove a particle from a vacuum state also quenches it. For a more meaningful example, consider a situation where there are N possible spatial positions, denoted i and j for particles 1 and 2 respectively, for a spin-$\frac{1}{2}$ particle that can be spin up ($\uparrow$) or down ($\downarrow$) in some basis. Then, the following many particle basis states are possible:

$$\begin{aligned} |\uparrow\uparrow\rangle_{ij} &= c^\dagger_{i\uparrow} c^\dagger_{j\uparrow} |0\rangle, i \neq j \\ |\downarrow\downarrow\rangle_{ij} &= c^\dagger_{i\downarrow} c^\dagger_{j\downarrow} |0\rangle, i \neq j \\ |\uparrow\downarrow\rangle_{ij} &= c^\dagger_{i\uparrow} c^\dagger_{j\downarrow} |0\rangle, \\ |\downarrow\uparrow\rangle_{ij} &= c^\dagger_{i\downarrow} c^\dagger_{j\uparrow} |0\rangle, \end{aligned} \tag{2.40}$$

where $c^\dagger_{is_1}$ is the creation operator for a given position and spin. These fermionic creation operators obey the anti-commutation relation $\{c^\dagger_{is_1}, c^\dagger_{js_2}\} = 0$. Therefore, the basis states are also related by $|\downarrow\uparrow\rangle_{ij} = -|\uparrow\downarrow\rangle_{ji}$. States with double spatial occupancy and same spin, like $|\uparrow\uparrow, 0\rangle$ are not allowed. By extending the above to N spatial sites, corresponding to a chosen spatial discretisation, we can construct valid basis states that allow us to include spin in the numerical simulation. For 2 particles and N spatial sites, without considering spin, there would be N^2 two-particle basis states, as there are N possibilities of placing the first particle, and again N for the second. With second quantisation, there will be $\frac{4N(N-1)}{2}$

single-occupancy states (states where particles are at different positions), as there are N ways of placing the first particle, $N-1$ of placing the second to avoid double occupancy, there are $2^2 = 4$ spin combinations, and we have to divide by 2 because particles are indistinguishable. In addition, there will be N double occupancy $|\uparrow\downarrow\rangle$ states, one for each site. In total, there are therefore $2N^2 - N$ second quantisation states, approximately double the amount without considering spin for large N. By indexing the wave function array with respect to these states, spin can be included at the cost of only a constant factor of 2 in the size of wave function (for 2 particles). This can be used both in the eigensolver and the time-dependent solver, allowing to find two-particle eigenstates with spin, and include explicit spin dependence in the Hamiltonian, extending the possible scope of scenarios to be simulated.

II.I Generalising to many particles

The work presented in this thesis will be followed up by simulations of more than two spin-$\frac{1}{2}$ particles - this is computationally possible owing to continuous advances in GPU hardware. As the cost of including spin will extend the computation time be a factor of approximately 4, it will still be doable in around a day or two. In this general case, a wave function can be constructed and manipulated in the following way.

The wave function is separated into spin $|s\rangle$ and spatial $|x\rangle$ components. The spin component for a single particle is a vector in the z-basis:

$$\begin{aligned} |\uparrow\rangle &= \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \\ |\downarrow\rangle &= \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \end{aligned} \tag{2.41}$$

The complete one-particle wave function is obtained by performing the outer product of the spin and space components:

$$|\psi_{P1}\rangle = |s\rangle \otimes |x\rangle . \tag{2.42}$$

The many-particle wave function is then obtained by simply performing the outer product of the single-particle wave functions, for N particles.

$$|\psi\rangle = |\psi_{P1}\rangle \otimes |\psi_{P2}\rangle \ldots \otimes |\psi_{PN}\rangle . \tag{2.43}$$

To then recover the spatial wave function for a specific particle of a given spin, for example to plot it or calculate single-particle observables, we need to trace out all the other degrees of freedom.

III Using the density matrix formalism

To include mechanisms such as noise, relaxation, dephasing, which cause loss of coherence by randomly affecting the wave function in a non-unitary (non-reversible) way, a statistical description using density matrices is required. This is also needed to describe thermalisation processes.

A density matrix is needed to include the "classical' (as opposed to quantum) probability/randomness. A quantum state $|\psi\rangle$ on Hilbert space can be in a superposition of some eigenstates of the measurement operator we're applying to it. It's a familiar result that there is then some probability of measuring the state in each of the eigenstates, with the process being random, as far as the observer is concerned. However, this quantum randomness stemming from $|\psi\rangle$ being in a superposition is very different from classical randomness, such as tossing a coin. Here, we also observe a random outcome, but a coin cannot be used for quantum computation. So-called "pure" states, described by vectors $|\psi\rangle$, cannot be used to express statistical (incoherent, as opposed to a coherent superposition) mixtures. These mixtures, which could result from loss of coherence of an initial pure state, or otherwise, always occur in a realistic experiment with sources of noise. We need the density matrix formalism to incorporate not only the quantum coherence (like in a superposition), but also the classical lack of knowledge about a state (like a coin that was tossed, but the result is unknown to us).

A density matrix ρ for a pure state $|\psi\rangle$ is given by:

$$\rho = |\psi\rangle\langle\psi|. \tag{2.44}$$

It has the following properties:

$$\begin{aligned} \rho^2 &= \rho(\text{projection}), \\ \rho^\dagger &= \rho(\text{hermiticity}), \\ \rho &\geq 0(\text{positivity}), \\ Tr\rho &= 1(\text{normalisation}). \end{aligned} \tag{2.45}$$

For clarity, the trace of an operator O is (where ρ is technically a projection operator, so the following applies):

$$TrO = \sum_n \langle n| O |n\rangle , \tag{2.46}$$

where n are the elements of some complete basis. For example, for some projection operator $O = |\psi\rangle \langle\phi|$:

$$TrO = \sum_n \langle n|\psi\rangle \langle\phi|n\rangle = \sum_n \langle\phi|n\rangle \langle n|\psi\rangle = \langle\phi|\psi\rangle , \tag{2.47}$$

using the resolution of identity $1 = \sum_n |n\rangle \langle n|$.

The positivity property in Eq. 2.45 means that the eigenvalues of ρ are greater or equal to zero:

$$\langle\phi|\rho|\phi\rangle = \langle\phi|\psi\rangle \langle\psi|\phi\rangle = |\langle\phi|\psi\rangle|^2 \geq 0. \tag{2.48}$$

This essentially means that probabilities should not be negative. We can calculate expectation values of observables w.r.t. a quantum state described by a density matrix:

$$< O >_\rho = Tr(O\rho) = Tr(\rho O), \tag{2.49}$$

where we note that the trace has a cyclic property.

Density matrices are used for statistical description of ensembles of quantum states. Imagine we are working with an ensemble of N states, each of which is described by a familiar vector on a Hilbert space $|\psi\rangle$. The density matrix $\rho = |\psi\rangle \langle\psi|$ will then accurately describe what we observe when drawing states from that ensemble and making measurements on them. An important property of, and a test for *pure* states, is:

$$Tr\rho^2 = 1. \tag{2.50}$$

This so far has not given us much more than using Hilbert space vectors. However, density matrix formalism is truly useful when we are working with *mixed* states. Mixed states describe a situation when we have an ensemble of quantum states, and they are *not* all described by the same Hilbert vector. If we have a total of N states in the ensemble, suppose N_i are in some quantum state $|\psi_i\rangle$, and $\sum_i N_i = N$. When drawing from the ensemble randomly, the chance to find a state i that is described by a vector $|\psi_i\rangle$ is given by $p_i = \frac{N_i}{N}$. This mixed state can then be described as a density matrix that is the convex sum of the pure density matrices $\rho_i = |\psi_i\rangle \langle\psi_i|$:

$$\rho_{mix} = \sum_i p_i \rho_i. \tag{2.51}$$

The expectation values of observables are still given by Eq. 2.49:

$$< O >_{\rho_{mix}} = Tr(O\rho_{mix}). \tag{2.52}$$

These mixed states have the following property:

$$Tr\rho^2 < 1. \tag{2.53}$$

$Tr(\rho^2)$ measures the "purity" of a state, since it is equal to 1 for pure states, and is strictly less for mixed states. For a maximally mixed state, which is an equal mixture of all possible states in the Hilbert space:

$$Tr\rho^2 = \frac{1}{d} > 0, \tag{2.54}$$

where d is the dimension of the Hilbert space. Such a maximally mixed state expresses complete lack of knowledge about what states are actually in our ensemble.

Just as a wave function evolves in time according to the Schrödinger equation, we can derive a similar equation for the density matrix. For $|\psi\rangle$:

$$i\hbar\frac{\partial}{\partial t}|\psi\rangle = \hat{H}|\psi\rangle. \tag{2.55}$$

A similar time derivative for ρ leads to:

$$i\hbar\frac{\partial}{\partial t}\rho = i\hbar\sum_i p_i(\frac{\partial}{\partial t}|\psi\rangle\langle\psi||\psi\rangle\frac{\partial}{\partial t}\langle\psi|) = \sum_i p_i(\hat{H}\rho_i^{pure} - \rho_i^{pure}\hat{H}) = [\hat{H},\rho], \tag{2.56}$$

where ρ_i^{pure} are the constituent pure matrices that ρ can be decomposed into.

The solution of the above von Neumann equation yields a unitary time shift operator $U(t,t_0)$ that is able to evolve the quantum state in time, similar to unitary time evolution for a wave function $|\psi\rangle$:

$$U(t,t_0) = e^{\frac{-i}{\hbar}\hat{H}(t-t_0)}, \tag{2.57}$$

where t_0 is some initial starting time. This operator is then applied to ρ the following way:

$$\rho(t) = U(t,t_0)\rho(t_0)U^{\dagger}(t,t_0). \tag{2.58}$$

Note that as ρ is a matrix, the Hermitian conjugate of the time shift operator needs to also be applied on the R.H.S. This is the case for all unitary operators acting on ρ.

As an example illustrating the difference between state vectors and density matrices, let us look at a single spin-$\frac{1}{2}$ qubit that is represented in the z-basis. Suppose we have the following maximally mixed state:

$$\rho_{mix} = \frac{1}{2}(|\uparrow\rangle\langle\uparrow| + |\downarrow\rangle\langle\downarrow|). \tag{2.59}$$

This state has equal probabilities ($\frac{1}{2}$) of being measured in the $|\uparrow\rangle$ and $|\downarrow\rangle$ states, but is not in a superposition state like $|\psi\rangle = \frac{1}{\sqrt{2}}(|\downarrow\rangle + |\uparrow\rangle)$- it is instead in the state $|\uparrow\rangle$ half of the time it is measured, and in the state $|\downarrow\rangle$ the remaining half of the time. This represents our lack of knowledge about which state we sample from the ensemble, rather than a quantum measurement randomness due to superposition. Note that this maximally mixed matrix is just diagonal.
Instead, a superposition state density matrix would look like this:

$$\rho_{sup} = \frac{1}{2}(|\downarrow\rangle + |\uparrow\rangle)(\langle\downarrow| + \langle\uparrow|) = \frac{1}{2}(|\uparrow\rangle\langle\uparrow| + |\downarrow\rangle\langle\downarrow| + |\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|). \tag{2.60}$$

The off-diagonal matrix elements $|\downarrow\rangle\langle\uparrow|, |\uparrow\rangle\langle\downarrow|$ represent the quantum correlation due to superposition.

So far the use of density matrices has allowed us to describe statistical ensembles of quantum states, but we were still restricted to unitary operations. However, the formalism truly shines because it allows us to use non-unitary operators that can describe non-unitary, irreversible processes such as decoherence. A Lindblad master equation is an extension of the von Neumann equation 2.56 that allows for non-unitary time evolution:

$$\frac{\partial}{\partial t}\rho = -\frac{i}{\hbar}[\hat{H}, \rho] + \sum_{i=1}^{N^2-1} \gamma_i (L_i \rho L_i^\dagger - \frac{1}{2}\{L_i^\dagger L_i, \rho\}). \tag{2.61}$$

The new non-unitary operators L_i are called Lindblad or jump operators that represent some non-unitary processes, while the γ_i are the rates at which these processes occur. For example, in a case where one allows the system to undergo relaxation and excitation (emission/absorption) due to an interaction with a thermal bath, jump operators $L_1 = a$, $L_2 = a^\dagger$ would be used, where $a/a^\dagger$ is the annihilation/creation operator (in the energy basis).

Now we can appreciate why the density matrix formalism is required to describe a realistic quantum experiment. Imagine that we begin some quantum computation in a known fiducial pure state. The state will experience noise, dephasing, relaxation of energy etc.

during the experiment. All these processes are random and irreversible, and thus cannot be described as unitary operators on a Hilbert space. They can however be described by L operators acting on density matrices. Our state will gradually become more mixed as a result of these effects. The resulting mixed density matrix will express our lack of knowledge about what exactly happened to the quantum state during the experiment, as even if we know all we can about the noise processes, they are still random in nature, and make our result more random. By using correct noise models, we can calculate the time evolution of the density matrix, and compare our theory or simulations with the experimental results, taking the noise into account.

IV GPU Acceleration

The time-dependent solver handling the time evolution of the initial wave function is written as a CUDA kernel running on a GPU (see Appendix). This GPU acceleration is desirable here, as the process is very calculation intensive, but can be easily parallelised. This is exactly what GPUs are designed to do optimally. A CPU is architecturally composed of a couple of cores, typically 2-8 for modern processors. GPUs on the other hand have hundreds or thousands of cores, allowing them to process thousands of software threads. The GTX 1080 Ti GPU, 7 of which were used in majority of calculations in this work, have 3584 cores each. The difference in ability to process parallel threads is clear and staggering. The staggered leapfrog algorithm can be parallelised as follows. Looking at Eqs. 2.28, 2.29, each time step must be calculated sequentially, i.e. one entire time step must be finished before proceeding to calculate the next one. Thus this process cannot be parallelised. However, each spatial node for a given time step can be calculated in parallel, as they are independent. This is where the GPU acceleration becomes invaluable.

Another factor to consider for numerical simulations is data precision. Real numbers are stored as floating point numbers in the computer. Floating point numbers use formulaic representation of real numbers as an approximation to support a trade-off between range and precision. A number is represented to some fixed number of significant digits. Most used data types are single precision ($float$ in C++) and double precision ($double$ in C++) floating point numbers. A $float$ uses 4 bytes of data for 7 significant figures, and a $double$ uses 8 bytes for 16 significant figure. Using a $double$ has an increased memory and computation time cost compared to $float$. Additionally, Nvidia GPUs used in this work are better optimised for using $floats$. Since the estimated error in the initial wave function, and the oscillating normalisation error in the time-dependent solver are both greater than $float$ precision, and for the reasons mentioned above, we have decided to use $floats$ throughout the simulations.

Chapter 3

Single qubit control

I Solid-state quantum dots as qubits

The aim of this work is to find an optimal definition of a semiconductor DQD charge qubit, and find an experimentally realistic way of controlling it by applying a voltage bias across the DQD. We wish for this method to be simple to understand and implement, as well as to be able to correct for rise time τ of the experimental setup. While the semiconductor DQD charge qubit is not the most promising qubit implementation, observations on how experimentalists tend to control this system has served as an inspiration for this work. However, the results and conclusions are applicable to a wide variety of qubit implementations.

In this work, we demonstrate that in a realistic semiconductor DQD structure, controlled by voltage bias gate pulses, a reliable qubit can be defined and operated with high fidelity. We first model a generic effective potential for a solid-state DQD system to define the qubit basis states as bonding and anti-bonding states. We show how to initialise a single electron into one of the logical qubit basis states and how to perform a set of mutually orthogonal rotations on the Bloch sphere, thus an arbitrary rotation, using shaped pulses that correct for pulse rise time. The impact of charge noise is considered, and the proposed control scheme's resilience is compared to a square wave pulse. We finally discuss our results and give conclusions on the practicability of the scheme.

II Two-site localised state model

Let us begin with a review of the two-site localised state model that is commonly used to describe the charge qubit. Within the two-state model, one has to solve the time dependent Schrödinger equation with the effective Hamiltonian $\hat{H}_{\text{eff}}$ defined as

$$\hat{H}_{\text{eff}}(t) = -\frac{1}{2}\,\varepsilon(t)\,\sigma_x + \frac{1}{2}\,\Delta\,\sigma_z + \frac{1}{2}(E_{\text{B}} + E_{\text{AB}}). \tag{3.1}$$

Here E_{B} and E_{AB} are the energies of the bonding and antibonding states of the DQD system, i.e. the two lowest energy states, at $\varepsilon = 0$ whereas Δ is the 'hybridisation energy' between the two localised states. At zero detuning, the bonding state $\psi^{\text{B}}(x)$ is symmetric, while the antibonding state $\psi^{\text{AB}}(x)$ is antisymmetric. Therefore, their equal superpositions produce maximally localised left/right states:

$$\psi^{L}(x) = \frac{1}{\sqrt{2}}(\psi^{\text{B}}(x) + \psi^{\text{AB}}(x)), \tag{3.2}$$

$$\psi^{R}(x) = \frac{1}{\sqrt{2}}(\psi^{\text{B}}(x) - \psi^{\text{AB}}(x)). \tag{3.3}$$

The linear detuning breaks the left/right symmetry, however as it is expressed by the Pauli σ_x matrix in the Hamiltonian, it doesn't make the system leave the $\psi^{\text{B}}(x)/\psi^{\text{AB}}(x)$ two-state basis (if done adiabatically), resulting simply in a coordinate rotation of the Bloch sphere. Therefore, we can still think in terms of the left/right localised wave functions even at non-zero detuning, and varying ε is a viable way of performing single qubit rotations. The two-site localised state model describes a DQD well. The effective potential in an experimental DQD system can be found using density functional theory [166, 213] and will be a complex function of all three spatial coordinates x, y, z. By careful design, the dynamics in two of the directions y and z can be confined to the lowest energy subbands so that only the potential in the x direction, $V_{\text{DQD}}(x,t)$ needs be considered. For example in a GaAs/AlGaAs heterostructure, the z direction is the growth direction and modulation doping can be used to create a triangular quantum well in that direction with subband energies two orders of magnitude larger than either ε or Δ. In the y direction, parabolic confinement with energies an order of magnitude larger than ε or Δ can be produced either by etching [73], fabricating a thin gate wrapping the conducting channel [106, 233] or using split-gates [230]. In order to create a DQD potential in the x direction, gates [228, 84, 142, 157] or etching [235, 73] can also be used.

The aim is to create a potential $V_{\text{DQD}}(x,t)$ that has two minima separated by a tunnel barrier. A convenient potential that has this property and is defined by three parameters A, B and σ is given by

$$V_{\text{DQD}}(x) = Ax^2 + B\exp\left(\frac{-x^2}{2\sigma^2}\right) \tag{3.4}$$

This form for V_{DQD} allows us to control both the depth of the dots and the barrier between them directly, by varying the harmonic confinement A, barrier height B, and barrier width σ. This potential will obey the two-site localised state model. For a specific set of parameters, this static potential will define a value for Δ which is the energy difference between the bonding ground state E_{B} and the antibonding first excited state E_{AB}. Detuning is introduced by adding a linear Stark shift of the form

$$V_{\mathrm{linear}}(x) = V_{\mathrm{bias}}\frac{x}{2w}. \tag{3.5}$$

Here, w is half the width of the DQD. By comparing the dependences of E_{B} and E_{AB} on V_{bias} with the expected dependences from two-site Hamiltonian we can define the detuning parameter for V_{DQD} through a linear relation $\varepsilon = e\lambda V_{\mathrm{bias}}$ with λ being constant. We find this linear relationship holds with an accuracy of one part in 10^6 across the range of required values of ε for single-qubit operations. The total potential is $V_{\mathrm{tot}}(x) = V_{\mathrm{DQD}}(x) + V_{\mathrm{linear}}(x)$ and Fig. 3.13a shows this potential at three different detunings.

The DQD dynamics under time-dependent detuning will be given by the TDSE

$$\hat{H}(x,t)\psi(x,t) = i\hbar\frac{\partial}{\partial t}\psi(x,t) \tag{3.6}$$

with

$$\hat{H}(x,t) = -\frac{\hbar^2}{2m^*}\frac{\partial^2}{\partial x^2} + V_{\mathrm{DQD}}(x) + V_{\mathrm{bias}}(t)\frac{x}{2w}. \tag{3.7}$$

Time dependence is included in Eq. 3.6 by varying the potential slope with time: $V_{\mathrm{bias}}(t)$. An example plot of the energies of the two lowest instantaneous solutions (the bonding and antibonding states) as function of V_{bias} is shown in Fig. 3.13b.

Analytic solutions to the TDSE in Eq. 3.7 can only be found in special cases. In this paper we solve Eq. 3.6 numerically using a GPU-accelerated version of the staggered-leapfrog method.

Throughout the paper we avoid using specific numerical values to keep our results general and not tied to a specific experiment. However, here we give the actual values used for reproducability. We've used a total DQD length of 460 nm, with parameter values: $w = 230$ nm, $A = 1.276$ meV nm^{-1}, $B = 4.08$ meV so that $\Delta = 11.7\mu$eV, and the linear coefficient $\lambda = 0.421$. The values given are experimentally realistic [228, 80, 81]. We have also tested various non-symmetric potentials with the two dots having different sizes, but in all the cases the general conclusions were the same as for the symmetric potential of Eq. 3.4.

III Single-Electron Charge Qubit definition

It is generally assumed either for simplicity or ease of experimental manipulations that the logic basis-state wave functions of a DQD, $|0\rangle$ and $|1\rangle$, are fully localised in the left or right side of the DQD [102, 93, 56]. This assumption is convenient for "brute force" initialisation via applying a high bias voltage. The readout is realised by measuring the probability of the electron being in the left or right dot (this is a source of error in itself, as there is some ambiguity in the readout owing to the state overlap - this is discussed separately). However, in a realistic DQD potential, these quantum states necessarily contain contributions from higher energy eigenstates which give rise to additional composite oscillations, typically on timescales faster than the qubit oscillation itself [119]. They ultimately induce a loss of fidelity in gate operations. This issue is critical for practical implementations of quantum computation and schemes like bang-bang pulse sequences (that switch abruptly) have been proposed in order to mitigate this effect [232, 24]. Such sequences involve additional gate operations that could be detrimental to the overall operation time. Consequently, optimizing the qubit basis states is a necessary preliminary requirement before any other attempts at extending coherence or improving the gate fidelity.

If a linear combination of the two lowest eigenstates of the DQD system is used instead of assuming a fully localized state, a true two-level system is formed. A qubit control framework that doesn't involve energy states outside of the computational space would greatly improve the fidelity compared to the method above.

It is optimal to define the qubit states as equal combinations of the ground and first excited states at zero bias, because it produces well-localized qubits that can be measured by detecting charge localisation, while also preserving symmetry between the two logical states. This is demonstrated within the two-site localised state model [162]. A zero-bias potential also makes the qubit first-order insensitive to electrical noise, improving fidelity [122]. Moreover, as described in section IV, having zero detuning as a default achieves a high fidelity $R_{\vec{x}}$ rotation without any pulsing. The coefficients of the energy eigenstates must be equal in order to have symmetry between the qubit states. Therefore, for a given DQD potential $V_{\mathrm{DQD}}(x)$, we define the logical states, which correspond to the $psi^L(x)$ and $psi^R(x)$ localised states described before, as:

$$\begin{aligned} |0\rangle &= \frac{\psi^{\mathrm{B}}(x) + \psi^{\mathrm{AB}}(x)}{\sqrt{2}}, \\ |1\rangle &= \frac{\psi^{\mathrm{B}}(x) - \psi^{\mathrm{AB}}(x)}{\sqrt{2}}, \end{aligned} \tag{3.8}$$

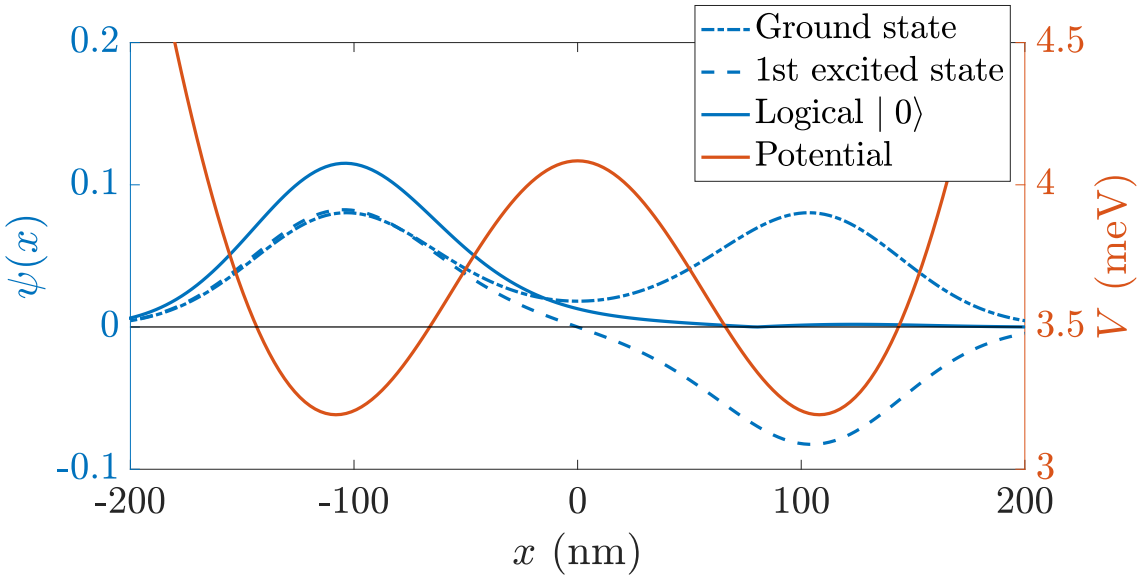

Fig. 3.1 Wave function of the two first excited states. The logical $|0\rangle$ and $|1\rangle$ qubits are formed using Eq. 3.8. The values of the DQD spacing and the electrostatic potential amplitude were chosen for illustrative purposes and the scheme presented here works for a wide range of configurations.

where $\psi^{(A)B}(x)$ is the (anti)bonding state wave function.

While these states are not completely localized on a single dot, as their probability density is tailing to the other side of the dot (Fig. 3.1), they maximize the average probability of successful readout [162]. Further localization of the states would introduce higher-energy states that would consequently not obey the ideal two-site Hamiltonian we aim to model (Eq. 3.10). This can be understood by imagining that if there was no overlap between the states, the inter-dot barrier would have to be so high that the left and right sides would be effectively separated, thus disabling the interaction needed for controlling the qubit. Some overlap, which will result is readout errors, is therefore unavoidable. However, by tuning the barrier height, we can find a balance between the speed of operations, proportional to the energy gap Δ, and readout errors. We show in Sec. VII that for our choice of parameters, these errors are not significant. Since there is no reference to the underlying effective potential of the DQD in our definition, this qubit is well defined for potentials that are not symmetric and more generally, for any dot shape.

IV Single Qubit Control

In the energy eigenbasis, the Hamiltonian of the qubit system reads:

$$\hat{H}(t) = -\frac{1}{2}\varepsilon(t)\,\sigma_x + \frac{1}{2}\Delta\,\sigma_z + \frac{1}{2}(E_{\mathrm{B}} + E_{\mathrm{AB}}). \tag{3.9}$$

Here E_{B} and E_{AB} are the energies of the bonding and antibonding states, i.e. the two lowest energy states, at a linear detuning $\varepsilon = 0$, Δ is the "hybridisation energy" between the two localised states, and $\sigma_{x/z}$ are the Pauli x/z matrices.

Using the basis defined in Eq. 3.8, where $|0\rangle$ and $|1\rangle$ are on the poles of a Bloch sphere, the Hamiltonian in Eq. 3.9 is written as:

$$\hat{H}_{\mathrm{eff}}(t) = -\frac{1}{2}\varepsilon(t)\,\sigma_z + \frac{1}{2}\Delta\,\sigma_x. \tag{3.10}$$

We have neglected the constant factor here. The time-dependent wave function can then be written in terms of the standard θ and ϕ, polar and azimuthal angles respectively, on the Bloch sphere:

$$\psi(x,t) = \cos\left(\frac{\theta(t)}{2}\right)|0\rangle + e^{i\phi(t)}\sin\left(\frac{\theta(t)}{2}\right)|1\rangle\,. \tag{3.11}$$

With no bias voltage, $\varepsilon = 0$, the wave function will undergo a constant rotation around the z-axis on the Bloch sphere. When applying a non-zero bias, the axis of rotation is shifted.

For the Hamiltonian in Eq. 3.10, a general rotation on the Bloch sphere by an angle α around a direction $\vec{n}$ is given by the solution to the time-dependent Schrödinger equation (TDSE):

$$R_{\vec{n}}(\alpha(t)) = \mathscr{T}\exp\left(\frac{1}{i\hbar}\int_0^t \hat{H}_{\mathrm{eff}}(t')\mathrm{d}t'\right) \tag{3.12}$$

where $\mathscr{T}$ is the time-ordering operator.

Rotations are performed by sending a bias voltage pulse of amplitude $V_{\mathrm{bias}} = \frac{\varepsilon}{e\lambda}$ and duration t_p to the double dot where ε and λ are respectively the detuning and voltage amplitude proportionality constant for a given potential. An instantaneous switch between the $V_{\mathrm{bias}} = 0$ and $V_{\mathrm{bias}} = \frac{\varepsilon}{e\lambda}$ bias states is generally preferred as this simplifies the dynamics and avoids spurious qubit rotations [130]. In this case, the detuning $\varepsilon(t)$ is described as a set of step-functions and $R_{\vec{n}}(\alpha)$ is expressed analytically as a rotation of the qubit state around the axis on the Bloch sphere which passes through the eigenstates of $H(t')$ at a rate proportional to the difference in energy of these two eigenstates. Such a pulse requires a linear potential along the axis of the double dot, as in Eq. 3.9, which is achieved by applying

voltages to a set of metallic surface gates. During a square-wave pulse of detuning ε, the system will evolve according to the Hamiltonian in Eq. 3.10, that will be constant during the on-time of the pulse, giving an unitary time evolution (rotation):

$$U(t) = R_{\vec{n}}(\alpha(t)) = \exp\left(-i\frac{\vec{n}\cdot\vec{\sigma}}{2\hbar}t\right), \tag{3.13}$$

where $\vec{n} = (\Delta, 0, \varepsilon)$ is the axis of rotation, with rotation frequency given by its magnitude.

Implementing such a square pulse isn't technically possible owing to practical limitations. Current and most commonly used pulse pattern generators have a built-in rise time τ of about 40ps to 500ps depending on the brand and characteristics. The Keysight 81134A Pulse Pattern Generator has a τ =60ps between 20% and 80% of target amplitude. The Agilent 81130A and the Anritsu MP1763C have $\tau = 500$ps and ~ 40ps respectively, both between 10% and 90% of target amplitude. (Fig. 3.2a). For a pulse to be considered square, the rise time has to be negligible compared to the total pulse time, which would create very slow operations with the rise time numbers given here.

In this case, the step-function decomposition is not possible and, in general, Eq. 3.12 must be solved numerically. If the detuning can be described in terms of linear ramp functions, then Eq. 3.12 can be written analytically as a Landau-Zener-Stuckelberg transition [135, 244, 215] but the resulting expression becomes a function of parabolic cylinder functions which makes understanding the rotation $R_{\vec{n}}(\alpha)$ more complex [199, 94].

In order to investigate the consequences that follow from this technical limitation, we have solved Eq. 3.12 numerically for a pulse with finite τ using a GPU-accelerated version of the staggered-leapfrog method [15, 165, 140, 14, 217], as described in Ch. 2.

For such a pulse, the path of an individual qubit state on the Bloch sphere during the time evolution in Eq. 3.12 differs from the one induced by a square pulse [75] (Fig. 3.3). In order to implement a high-fidelity rotation on the Bloch sphere, an effective $R_{\vec{n}}(\alpha)$ is found by accounting for the aforementioned equipment limitations, such that the path traced on the Bloch sphere is different, but the resulting rotation remains the same as one induced by a perfect square pulse. We find that this can always be done by tuning the pulse duration and amplitude, depending on τ and desired angle of rotation. The details of this correction are outlined in IV.IV. One can question whether such an adjusted operation including transient rotations is a proper rotation, i.e. independent of the initial state. The answer is yes, because while the precise path on the Bloch sphere may be difficult to describe analytically, the instantaneous Hamiltonian is still always expressed in terms of σ_x and σ_z matrices, therefore the effective operation is composed of rotations and is itself an actual rotation. We show that our pulses have the desired effect on any input state in Sec. VIII. Additionally, it is worth

noting that having a finite τ can have a desirable effect on the qubit, as it makes the pulsing operation more adiabatic compared to using square pulses.

IV.I General rotation scheme

To perform an arbitrary qubit rotation, we propose a scheme of concatenating square pulses of alternating amplitudes. We set the bias voltage to produce a detuning $\varepsilon = \pm\Delta$, which gives the axes of rotation during pulsing to be in directions $(\frac{1}{\sqrt{2}}, 0, \pm\frac{1}{\sqrt{2}})$ on the Bloch sphere. We will call these axes $\vec{z}'$ $(\frac{1}{\sqrt{2}}, 0, \frac{1}{\sqrt{2}})$ and $\vec{x}'$ $(\frac{1}{\sqrt{2}}, 0, -\frac{1}{\sqrt{2}})$ respectively, as they are both rotated by $\frac{\pi}{4}$ around $\vec{y}$ w.r.t. the usual $\vec{z}, \vec{x}$ axis of the Bloch sphere. An arbitrary rotation can be performed by combining up to five rotations around any two perpendicular axes, simply by aligning $\vec{x}'$ with the desired axis of rotation $\vec{n}$, performing the rotation, and then reversing the first step. An arbitrary rotation by angle α can thus be performed around axis $\vec{n}$ in the following way:

$$\begin{aligned} R_{\vec{n}}(\alpha) =& R_{\vec{x}'}\left(\frac{\pi}{2} - \phi\right) R_{\vec{z}'}(\theta) R_{\vec{x}'}(\alpha) \\ & \cdot R_{\vec{z}'}(-\theta) R_{\vec{x}'}\left[-\left(\frac{\pi}{2} - \phi\right)\right], \end{aligned} \tag{3.14}$$

where θ, ϕ are the angles of $R_{\vec{y}}(\frac{\pi}{4})\vec{n}$ on the Bloch sphere. The argument angles of the composite rotations correspond to durations of the composite pulses, with 2π corresponding to $T_{\text{rot}} = \frac{2\pi\hbar}{\sqrt{\Delta^2+\varepsilon^2}}$, the period of a full rotation around $\vec{x}'$ or $\vec{z}'$ while bias voltage is on. Since the rotation around $\vec{x}'$ or $\vec{z}'$ is always in the positive direction, any negative angles have to be replaced by a positive complement of 2π.

We present some examples of how such rotations are performed in Fig. 3.4. Here we show R_z, R_x and R_y rotations, all by $\frac{\pi}{2}$, performed with no rise time with great fidelity (>99.99 %). The 5-pulse control sequence as per Eq. 3.14 was used.

This method is simple to implement, but not optimal in operation time - as can be seen in Fig. 3.4, the paths taken on Bloch sphere are far from optimal, wasting a lot of precious operation time. It is known [121] that three rotations are sufficient, resulting in a faster operation:

$$R_{\vec{n}}(\alpha) = e^{i\beta} R_{\vec{x}'}(\Theta_1) R_{\vec{z}'}(\Theta_2) R_{\vec{x}'}(\Theta_3). \tag{3.15}$$

Here, Θ_1, Θ_2 and Θ_3 each depend on the angle and axis of the rotation. The above scheme was suggested by Echo Kexin Zhang, and proven and implemented by myself.

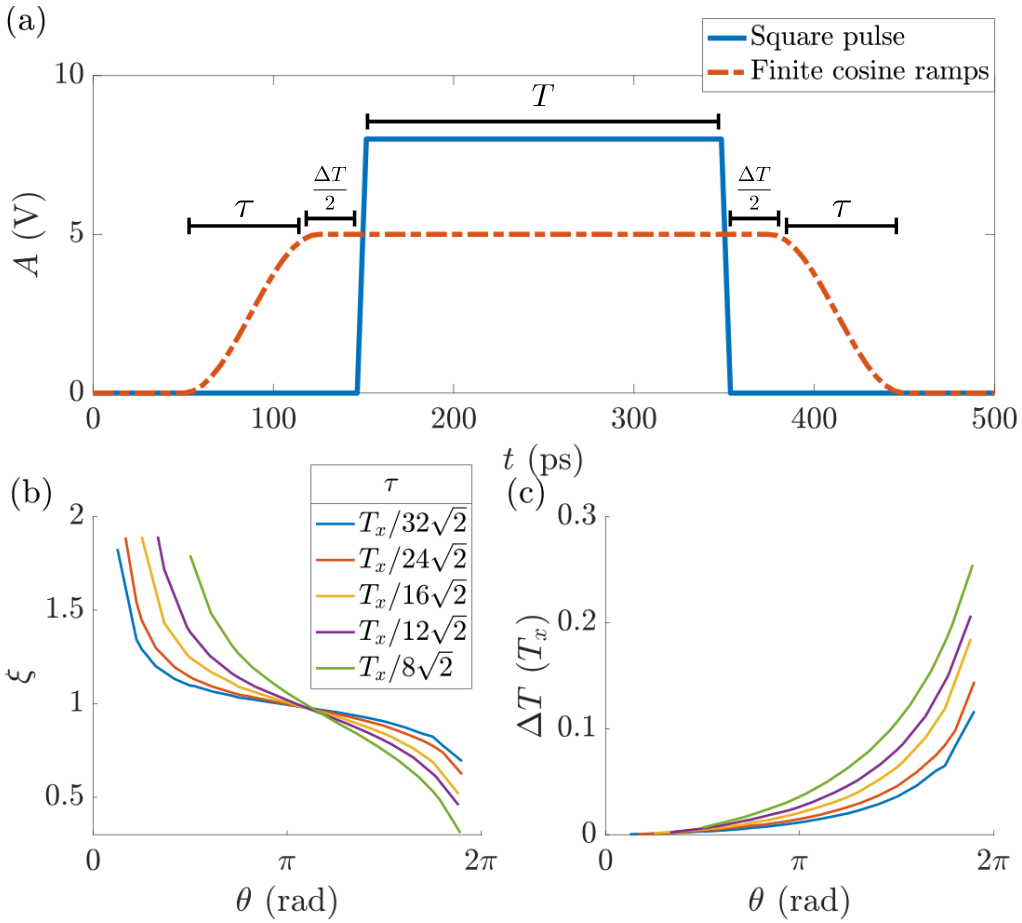

Fig. 3.2 **(a)** Amplitude profile of the ideal square pulse (solid blue line) to apply the linear bias given by Eq. 3.13. The pulse amplitude and duration are adjusted when τ is finite. **(b)** Multiplicative amplitude adjustment factor ξ given a target rotation angle θ. Each coloured line corresponds to a different τ (see legend). **(c)** Additive pulse duration adjustment ΔT with respect to the original square pulse time (see panel **(a)**). The rise times are not included in the additional pulse duration. Each coloured line corresponds to a different τ (see legend)

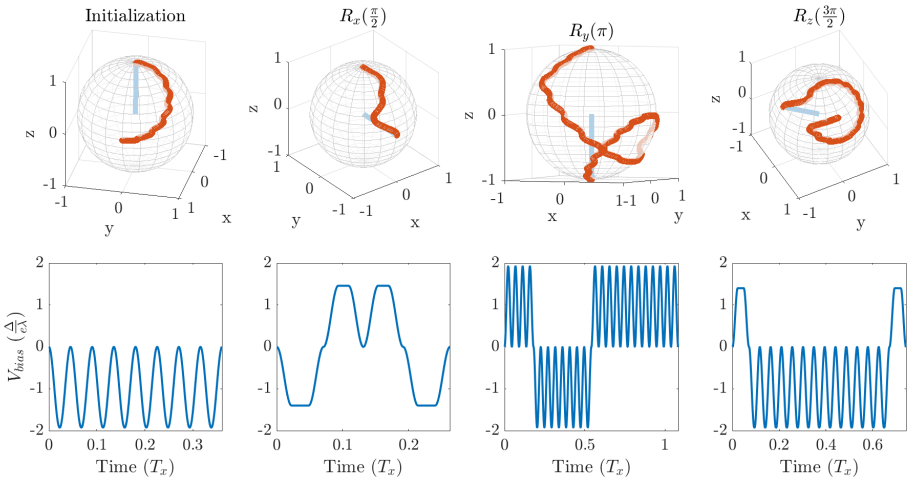

Fig. 3.3 Example pulse sequences and associated qubit rotations. Top: rotation path on the Bloch sphere. Bottom: Optimized pulse sequence where $T_x = \frac{2\pi\hbar}{\Delta}$. To remain general, values of time and voltage are quoted as fractions of T_x and $\frac{\Delta}{e\lambda}$ respectively. Exact experimental values will vary from one setup to the other. See the discussion for more details. All pulse sequences lead to a final state with a fidelity of >99.99%. Furthermore, the same pulse sequence can be used for any initial state on the Bloch sphere without significant loss in fidelity.

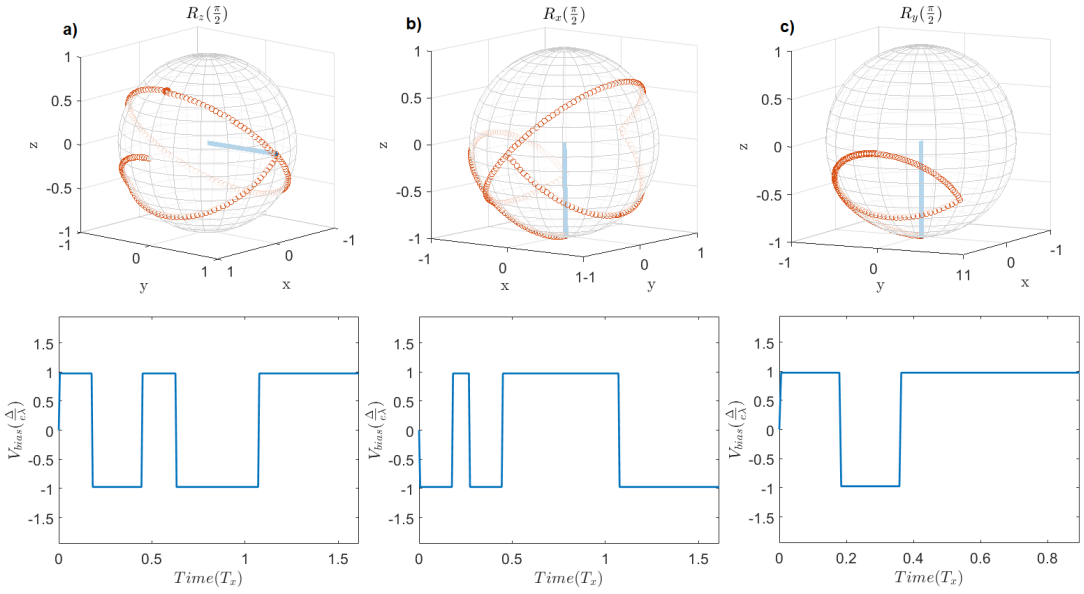

Fig. 3.4 Examples of 5-pulse sequences Top: rotation path on the Bloch sphere. Bottom: Optimized pulse sequence where $T_x = \frac{2\pi\hbar}{\Delta}$. To remain general, values of time and voltage are quoted as fractions of T_x and $\frac{\Delta}{e\lambda}$ respectively.

IV.II State preparation

Before any quantum computation is performed, each qubit has to be initialized to a fiducial state, usually $|0\rangle$ or $|1\rangle$. For a generic operation involving a charge qubit, before any initialisation pulsing, we would expect the initial state of the electron to be the ground state of the DQD (see Fig. 3.1). Such a state is not part of the qubit's logical basis and an initial rotation is needed. In order to rotate the wave function from the ground energy eigenstate to the qubit $|0\rangle$ state, we can take advantage of knowing the initial state to simplify the operation. A $R_{\vec{z}'}(\pi)$ rotation will initialise to the $|0\rangle$ state, while a $R_{\vec{x}'}(\pi)$ will do so to the $|1\rangle$ state. Both are achieved with a single pulse, thus simplifying the initial state preparation.

IV.III Single axis rotations

Any single-qubit operation can be expressed in terms of rotations around two perpendicular axes. Here we provide the control sequence for rotations around the usual $\vec{x},\vec{y},\vec{z}$ Bloch sphere axes from an arbitrary point on the Bloch sphere.

The $R_{\vec{y}}$ rotation consists of only 3 pulses owing to angle cancellation in Eq. 3.14 (as $\phi = \frac{\pi}{2}$):

$$R_{\vec{y}}(\alpha) = R_{\vec{z}'}\left(\frac{\pi}{2}\right) R_{\vec{x}'}(\alpha) R_{\vec{z}'}\left(\frac{3\pi}{2}\right). \tag{3.16}$$

To rotate in the opposite direction, one simply has to invert this pulse (swap $\vec{x}'$ and $\vec{z}'$) to get:

$$R_{\vec{y}}(-\alpha) = R_{\vec{x}'}\left(\frac{\pi}{2}\right) R_{\vec{z}'}(\alpha) R_{\vec{x}'}\left(\frac{3\pi}{2}\right). \tag{3.17}$$

R_z and R_x rotations would require five pulses if done as per Eq. 3.14. Instead, we solve Eq. 3.15 for the angles to also perform them with just three pulses. Owing to symmetry, the first rotation is the same as the third one.

We find a fast and simple general rotation scheme based on creating two perpendicular axes $\vec{x}'$ and $\vec{z}'$, by setting the detuning $\varepsilon = \pm\Delta$. Then we observe that one should be able to perform a rotation around an axis at $\frac{\pi}{4}$ w.r.t the two axes above, which would be $\vec{x}$ and $\vec{z}$. This is achieved by rotating by some angle Θ_1 around the first axis, then by Θ_2 around the second one, and finally by Θ_1 around the first one again.

We will find the relationship between Θ_1, Θ_2, and the net angle rotated around $\vec{x}$ or $\vec{z}$ named α, by analytically comparing the rotation matrix elements with the straightforward $R_{\vec{x}}$ and $R_{\vec{z}}$ rotations.

Looking at $R_{\vec{x}}$ first:

$$R_{\vec{x}}(\alpha) = \begin{pmatrix} \cos\frac{\alpha}{2} & -i\sin\frac{\alpha}{2} \\ -i\sin\frac{\alpha}{2} & \cos\frac{\alpha}{2} \end{pmatrix}. \tag{3.18}$$

In our scheme,

$$R_{\vec{x'}}(\alpha) = R_{\vec{y}}\left(\frac{-\pi}{4}\right) R_{\vec{x}}(\alpha) R_{\vec{y}}\left(\frac{\pi}{4}\right), \tag{3.19}$$

$$R_{\vec{z'}}(\alpha) = R_{\vec{y}}\left(\frac{-\pi}{4}\right) R_{\vec{z}}(\alpha) R_{\vec{y}}\left(\frac{\pi}{4}\right), \tag{3.20}$$

and we need the following to always hold:

$$R_{\vec{x}}(\alpha) = R_{\vec{x'}}(\Theta_1) R_{\vec{z'}}(\Theta_2) R_{\vec{x'}}(\Theta_1). \tag{3.21}$$

Comparing the (1,1) matrix elements:

$$\begin{aligned} \cos\alpha = \cos\frac{\Theta_2}{2}(2\cos^2\frac{\Theta_1}{2} - 1) \\ -j\sqrt{2}(\frac{1}{2}\sin\frac{\Theta_2}{2} + \cos\frac{\Theta_1}{2}\cos\frac{\Theta_2}{2}\sin\frac{\Theta_1}{2}). \end{aligned} \tag{3.22}$$

Since the imaginary part on the LHS is zero, we have:

$$(\frac{1}{2}\sin\frac{\Theta_2}{2} + \cos\frac{\Theta_1}{2}\cos\frac{\Theta_2}{2}\sin\frac{\Theta_1}{2}) = 0. \tag{3.23}$$

Solving the above allows us to find Θ_2 in terms of Θ_1:

$$\Theta_2 = 2\arctan(\sin\Theta_1). \tag{3.24}$$

Now coming back to the real part of Eq. 3.22 and substituting for Θ_2, we have:

$$\cos(\arctan(\sin\Theta_1))\cos\Theta_1 = \cos\frac{\alpha}{2}, \tag{3.25}$$

which gives

$$\Theta_1 = \arccos\left(\frac{\sqrt{2}\cos\frac{\alpha}{2}}{\sqrt{\cos\frac{\alpha}{2}^2 + 1}}\right). \tag{3.26}$$

The above satisfies Eq. 3.21 for all matrix elements, and is therefore equivalent. It allows us to find a three pulse train that performs the $R_{\vec{x}}(\alpha)$ rotation by an arbitrary angle α. We

repeat the above procedure for $R_{\vec{z}}$ to find the following:

$$\Theta_2 = -2\arctan(\sin\Theta_1) + 2\pi, \tag{3.27}$$

$$\Theta_1 = \arccos\left(\frac{\sqrt{2}\cos\frac{\alpha}{2}}{\sqrt{\cos\frac{\alpha}{2}^2 + 1}}\right). \tag{3.28}$$

Therefore, we can perform arbitrary rotations around $\vec{z}$ and $\vec{x}$ this way. However, this scheme is unable to perform the $R_{\vec{y}}$ rotation, which is achieved as shown in Eq. 3.16.

To summarise, we optimally perform the x- and z-axis rotations as follows:

$$R_{\vec{x}/\vec{z}}(\alpha) = R_{\vec{x}'}(\Theta_1) R_{\vec{z}'}(\Theta_2) R_{\vec{x}'}(\Theta_1), \tag{3.29}$$

where

$$\Theta_1 = \arccos\left(\frac{\sqrt{2}\cos\frac{\alpha}{2}}{\sqrt{\cos(\frac{\alpha}{2})^2 + 1}}\right), \tag{3.30}$$

and

$$\Theta_2 = 2\arctan(\sin\Theta_1) \tag{3.31}$$

for $R_{\vec{x}}$, and

$$\Theta_2 = 2(\pi - \arctan(\sin\Theta_1)) \tag{3.32}$$

for $R_{\vec{z}}$. Additionally, we note that $-R_{\vec{z}}(\alpha) = R_{\vec{z}}(-\alpha)$, which allows us to shorten operation time for rotations with $\alpha \geq \frac{\pi}{2}$ by inverting the pulse profile to perform the complementary rotation instead.

We note that one of the effects of defining the qubit as in Eq. 3.8 is that $R_{\vec{x}}$ rotation will occur automatically due to the Hamiltonian, with the rotation period $T_x = \frac{2\pi\hbar}{\Delta}$. In the many qubit case, all the qubits rotate at their respective frequencies, and one would usually work in the rotating basis, therefore an $R_{\vec{x}}$ rotation still needs to be performed as per Eq. 3.29.

Instead of using the usual $\vec{x},\vec{y},\vec{z}$ basis, we can instead use the $\vec{x}',\vec{y},\vec{z}'$ (note the primes) basis which is more natural for the detuned system, and can be used to define logic gates with fewer pulses. A single $R_{\vec{y}}(\frac{\pi}{4})$ rotation is required to move into this basis. $R_{\vec{x}'}, R_{\vec{z}'}$ are then achieved with a single pulse, while $R_{\vec{y}}$ requires three, as in Eq. 3.16. This way, any computation can be performed in the rotated basis, where operations are quicker. At the

end, one would need to rotate back to $\vec{x},\vec{y},\vec{z}$ using a $R_{\vec{y}}(-\frac{\pi}{4})$ rotation, for optimal readout of localised states.

Some logic gate examples are:

$$X = R_{\vec{z}'}(\pi), \tag{3.33}$$

$$Y = R_{\vec{y}}(\pi), \tag{3.34}$$

$$Z = R_{\vec{x}'}(\pi), \tag{3.35}$$

$$H = R_{\vec{y}}(\frac{\pi}{2})R_{\vec{x}'}(\pi), \tag{3.36}$$

$$R_\phi = R_{\vec{x}'}(\phi). \tag{3.37}$$

IV.IV Correcting for rise time

To account for the actual experimentally realisable pulses not being square due to rise time and limited bandwidth, the bias voltage and pulse duration have to be adjusted. This adjustment depends on the target rotation angle and τ, but not on the input state (if we design it correctly). Therefore, it is sufficient to optimize a single pulse for the instrument rise time and range of desired rotations - these single pulses can then be concatenated into three-pulse trains to achieve arbitrary qubit rotations of high fidelity. Here we numerically find the correct adjustments. This allows experimentalists to apply the ideal control sequence by simply changing the amplitude and duration of each square pulse in the train, avoiding complicated pulse shapes while retaining high fidelity. We assume that since applying an electrical bias pulse to the DQD has the effect of a σ_z operator on the Hamiltonian, the resulting operation should be a proper U(2) rotation, and therefore its' effect is not dependent on the initial state (when described as a rotation on the Bloch sphere around some axis by some angle). We assume that by varying only the amplitude and duration of a pulse that is distorted from square by a sinusoidal rise/fall time, we can adjust it to be equivalent to said square pulse, at least within some restrictions on which rotations we are able to perform within this scheme. For example, it is evident that very short rotations are not possible, since just during the rise and fall time the state vector on the Bloch sphere will move by a significant amount. However, such small rotations could still potentially be performed by extending the pulse to last a whole rotation period T_{rot} (which corresponds to a rotation by 2π) to start with, to buffer out the rise time effects.

We use a gradient ascent method to find the required amplitude ξ and pulse duration ΔT adjustments. We start in an initial state $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and then apply the pulse

including rise time using an initial guess for ξ and ΔT. We then calculate its' fidelity w.r.t. the square pulse, and take small steps in the phase space of ξ , ΔT to calculate a derivative of fidelity w.r.t these adjustment parameters. We continue to move in the direction of greatest increase in fidelity, and with a step size that is proportional the magnitude of that derivative and a scaling γ, until a very low error (at least 10^{-6}) is reached (See Fig. 3.5 for an illustration of this method) . There exists a fine balance between number of steps required, and maximum fidelity reachable - decreasing γ increases the number of steps to reach a given fidelity, but also increases the maximum fidelity reachable before diverging. We have tried strategies of using adaptive γ, tuning it based on step number and fidelity reached with some success, but ultimately some hand tuning and multiple runs were needed to reach very high levels of fidelity here. Fig. 3.6 shows an example of this gradient ascent process for one value of θ and τ. We sweep over the requested rotation angle for a few values of set rise times. We find that as long as the required rotation angle is not smaller than a threshold value determined by the rise time, it is always possible to find the adjusted pulse of high fidelity. We later show that the fidelity remains sufficiently high independent of the initial state's location on the Bloch sphere.

We present the numerical results for required amplitude ξ and pulse duration ΔT adjustments, depending on τ and angle of rotation α, all expressed in terms of the physical system parameters. Here, ξ is a multiplicative factor adjusting the amplitude with respect to the square pulse amplitude ($\xi = 1$), and ΔT is the additive time adjustment with respect to the square pulse duration as well, as per Fig. 3.2 (a) - it is always greater than zero. We use generalised rise times expressed in terms of a fraction of generalised time T_x (period of a full rotation without any pulsing), as seen in the legend of Fig. 3.2. We have chosen these values to correspond to minimum possible rotation angles of $\frac{\pi}{8}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{\pi}{2}$, from shortest to longest. These are the minimum possible rotations, because they are given by a pulse that consists only of rising/falling time, with no flat top, and is therefore the shortest pulse of desired amplitude that is possible. Of course, it is still be possible to rotate by an arbitrarily small angle indirectly by adding a 2π rotation. As can be seen in Fig. 3.2 (b,c) the required time adjustment rises exponentially with desired rotation angle. Therefore, it is optimal to compose any pulse of the smallest possible rotations, as this will result in shorter overall rotation time. If the target rotation angle does not subdivide into an integer number of shortest possible rotations, one needs to use somewhat longer sub-pulses appropriately. Assuming a sine-shaped rise ramp, this short pulse is a sine wave, which is straightforward to generate experimentally. Single qubit control can be achieved by sending sine waves, with frequency as high as experimentally possible, and amplitude given by ξ in Fig. 3.2 (b). Note that the only system-specific quantity is the energy gap Δ - the signal frequency is independent of

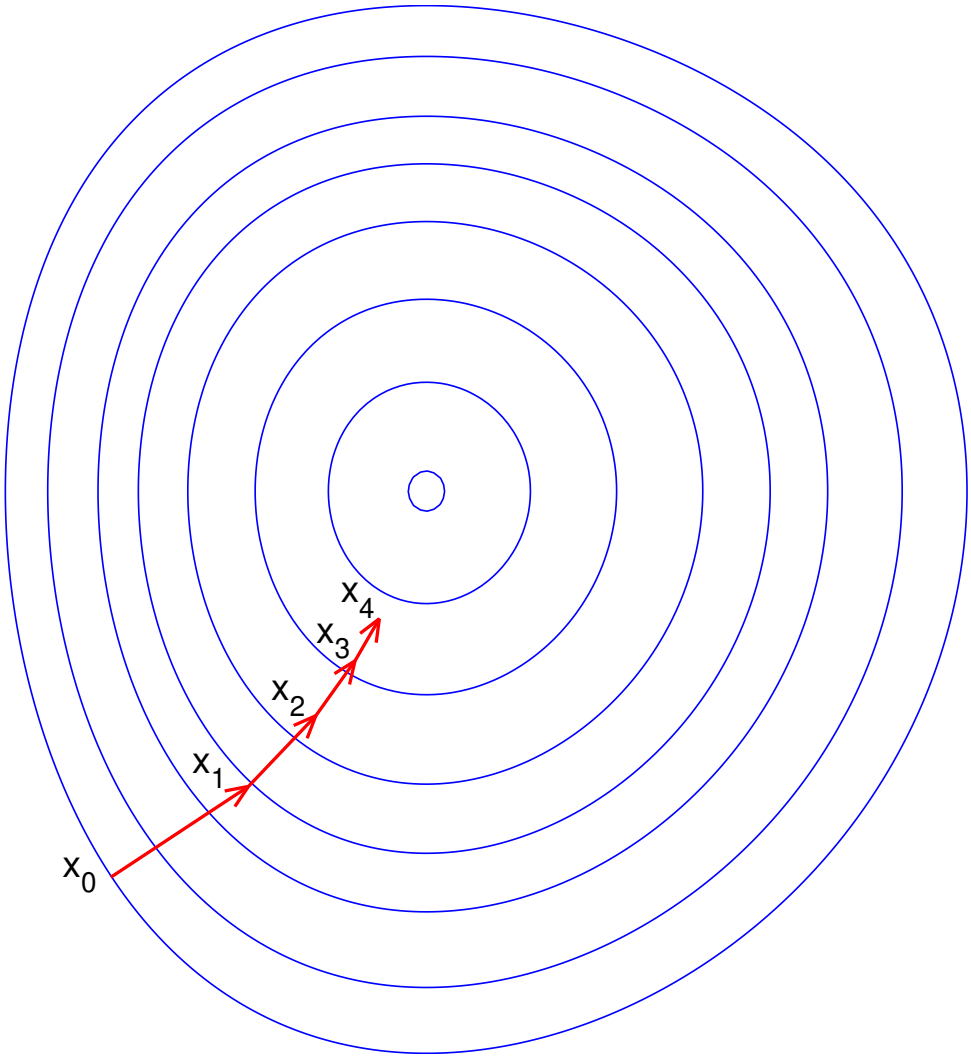

Fig. 3.5 Illustration of the principle of gradient ascent. Starting from some initial guess, derivative of the function (fidelity in this case) is evaluated w.r.t. the parameters, and a step is made in the direction of greatest difference. This ensures that we always move towards at least a local maximum. Image from [1].

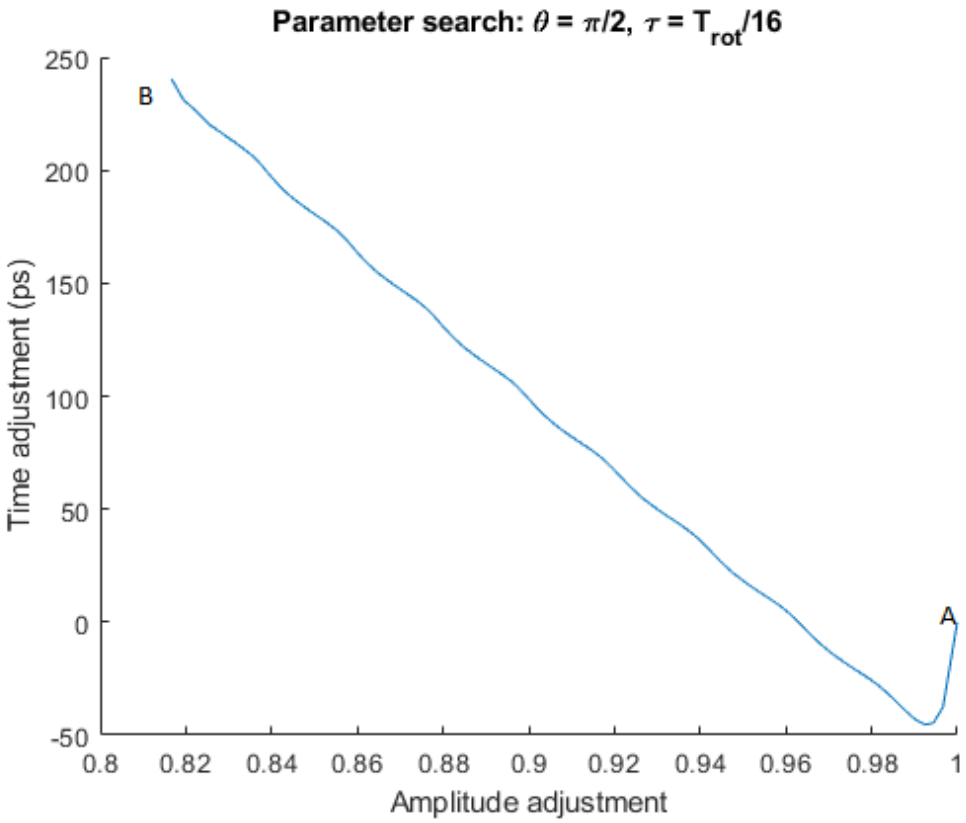

Fig. 3.6 Finding the adjustment parameters for angle of rotation $\frac{\pi}{2}$ with $\tau = \frac{T_{rot}}{16}$. We start at point A with an initial guess of no adjustment, which gives a fidelity error of 0.0019, and end at point B with fidelity error 10^{-5}. Continuing with the same step size scaling γ will eventually cause the search to diverge, therefore it should be reduced and the search iterated again, starting with the best adjustment parameters that the previous search managed to find.

the qubit system and not resonant with the two-level system, and instead purely defined by experimental limitations of the equipment ($\tau > 0$). We present examples of rotations performed with this scheme in Fig. 3.3, which summarises our main results.

As the resulting fidelity varies significantly with even small deviations from the parameters found here, we find that trying to fit analytical expressions to the data is not very useful if high fidelity is required. While ΔT as a function of rotation angle θ seems to be an exponential, while ξ is a rotated S-curve, attempts to fit it results with unacceptably low fidelity for a large θ range. Therefore, we suggest the gradient ascent search procedure described here be performed for the system of interest, taking into account the specificity of the experimental setup. This could be done using numerical simulations like in this work, or directly by taking actual measurements in an experiment. However, the latter might not be practical, as we find that thousands of fidelity evaluations are necessary to find good enough adjustment parameter values. If significant measurement error is present, the required number of experimental runs necessary might not be possible to realise, further highlighting the need for numerical simulations. Pseudocode of the gradient ascent procedure is provided in App. B - it should enable anyone to find the optimal parameters in a general case, for rise time and angles that are required. While this code uses the GPU-accelerated software described in this thesis, it is not strictly necessary, as a single-particle simulation is not heavily taxing on the hardware. Instead, a variety of simpler simulation methods should be successful instead.

IV.V Noise

Noise is always an important source of loss of fidelity in any qubit platform. If unaccounted for, the randomness of noise will lead to gradual loss of quantum information during a computation. While noise mitigation is not the goal of this work, we nonetheless investigate its impact here for completeness. In a quantum-dot-based charge qubit we base our simulations on, charge noise is one of the main sources of noise. It arises from fluctuations of charge states that lead to fluctuations of electric field a qubit experiences [133].
Here, we use a simple model where the charge noise is low-frequency and can be assumed to be constant during a single quantum operation [41]. In practice, this could result from some charge trapped temporarily on one side of the DQD, imparting an electric field gradient, effectively adding an unwanted random bias voltage. Therefore, to calculate the resulting fidelity loss, we average the resulting fidelity from many simulations, each with a random amplitude, but constant during a single simulation. The effective Hamiltonian has an additional noise term:

$$\hat{H}_{\text{noise}}(t) = -\frac{1}{2}\left[\varepsilon(t) + \delta_{\text{noise}}\right]\sigma_x + \frac{1}{2}\Delta\sigma_z, \tag{3.38}$$

where δ_{noise} is the noise amplitude randomly drawn from a normal distribution with mean $\mu = 0$ and standard deviation σ_{noise}, which can be varied to explore different noise strengths.

A large number (order of 100) of simulations are run with this randomised noise for some example operations, and the effects of this noise are compared between a square wave, and adjusted pulses accounting for rise time that are the result of this work. The random number generator seed is the same for both cases, so that they experience exactly the same noise and thus can be compared fairly. The results are presented in Fig.3.7. It can be seen that using the adjusted pulse is effectively the same as a square one for noise resistance. This is to be expected, as the pulse was not designed with noise in mind. At the very least, we confirm that our proposed pulse is not any worse than an idealised square wave, and further error mitigation techniques can be applied to it, as they would be to a square pulse, without it causing any loss of fidelity, while the problems associated with rise time are solved.

Interestingly, we find that there is a subset of cases where an optimisation of our pulse sequence does produce a reduction in noise-related errors. When performing an $R_{\vec{x}}$ rotation, it is possible to sub-divide the pulse into further smaller sub-pulses that add up to the total angle of rotation θ. This is only possible when the rise time constraint allows such a division, as there will exist a minimum angle θ_{min} that you cannot subdivide further. This can be understood by looking at how angles Θ_1, Θ_2 as per Eq. 3.15 depend on required total angle of rotation α - this information is presented in Fig. 3.8. For the $R_{\vec{z}}$ rotation however, this method doesn't work well, as the angle Θ_2 is always relatively large, even for small total rotation angle α. Therefore, attempting to subdivide a larger rotation would result in a very long total operation time, as the total angle that needs to be rotated is no longer (approximately) proportional to α. This data is presented in Fig. 3.9. Additionally, the dependence of total rotated angle (which approximately corresponds to total operation time) on the required rotation angle α is presented in Fig. 3.10 for both $R_{\vec{x}}$ and $R_{\vec{z}}$ rotations. The case for $R_{\vec{y}}$ suffers from similar issues as $R_{\vec{z}}$, therefore one cannot use this optimisation by subdivision to improve resilience against noise there.

An example of noise reduction owing to subdivision into smaller pulses for an $R_{\vec{x}}$ rotation is presented in Fig. 3.11.

This beneficial effect of subdividing the pulse can be understood by investigating the pulse sequence that achieves the rotation. As seen in Fig. 3.12, which shows a pulse shape of

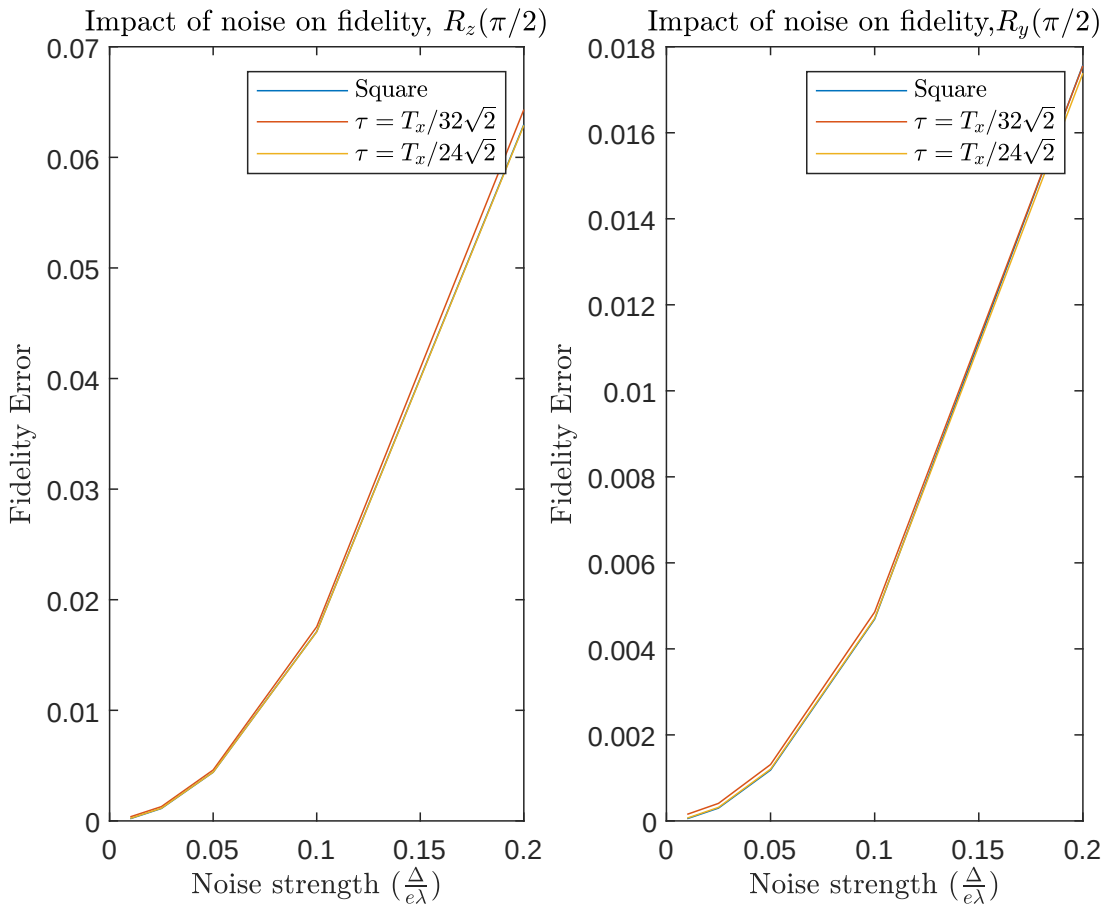

Fig. 3.7 Impact of noise on fidelity is investigated for two example operations, comparing an idealised square wave to adjusted pulses for two examples of rise time τ. Noise strength (standard deviation) is expressed in units of reference pulse strength $\sigma_{\text{noise}} = \frac{\Delta}{e\lambda}$, and varied from $0.01\frac{\Delta}{e\lambda}$ to $0.2\frac{\Delta}{e\lambda}$.

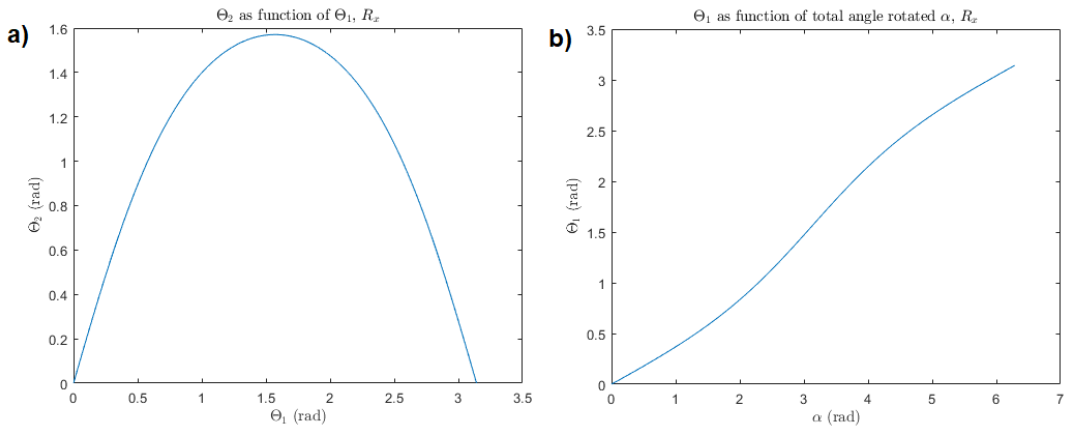

Fig. 3.8 $\Theta_1, \Theta_2, \alpha$ angles co-dependence for $R_{\vec{x}}$. For small α, both Θ_1, Θ_2 are relatively small, enabling one to subdivide a larger rotation into a series of smaller ones, provided rise time τ is small enough that all these angles can be realised.

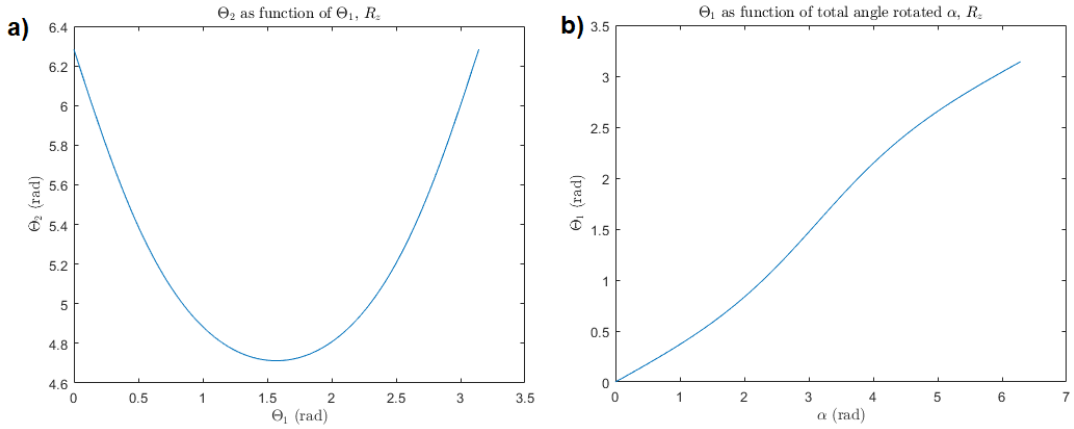

Fig. 3.9 $\Theta_1, \Theta_2, \alpha$ angles co-dependence for $R_{\vec{z}}$. Even for a small α, Θ_2 is large, which prevents one from subdividing a larger rotation into a series of smaller ones, as it would drastically increase the total rotation time, and thus time over which noise affects the operation.

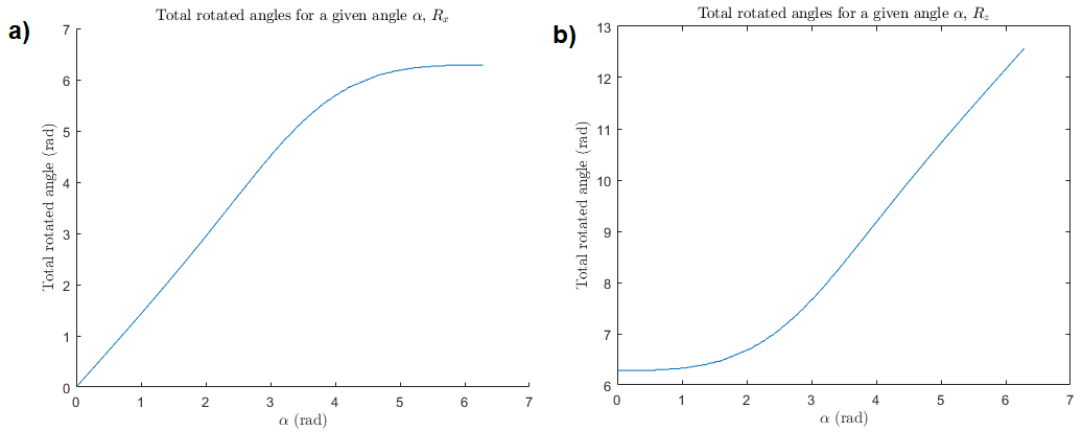

Fig. 3.10 Total rotation angle dependence on α for $R_{\vec{z}}$ and $R_{\vec{z}}$. While the total rotation angle is approximately proportional to α for $R_{\vec{x}}$, is already starts at 2π for $R_{\vec{z}}$, making the subdivision into smaller pulses not practical.

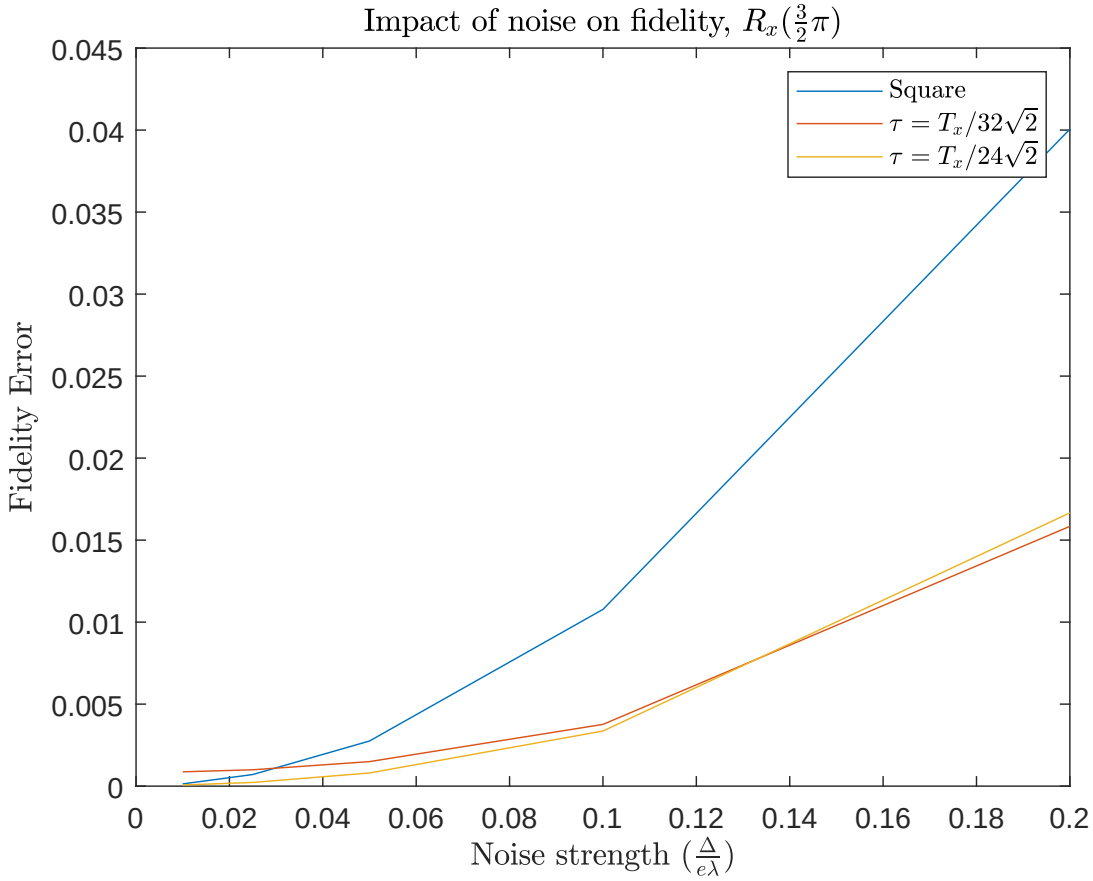

Fig. 3.11 Up to 250 % gain in fidelity is observed when performing an R_x rotation by $\theta = \frac{3}{2}\pi$. Noise strength is expressed in units of reference pulse strength $\sigma_{\text{noise}} = \frac{\Delta}{e\lambda}$, and varied from $0.01\frac{\Delta}{e\lambda}$ to $0.2\frac{\Delta}{e\lambda}$.

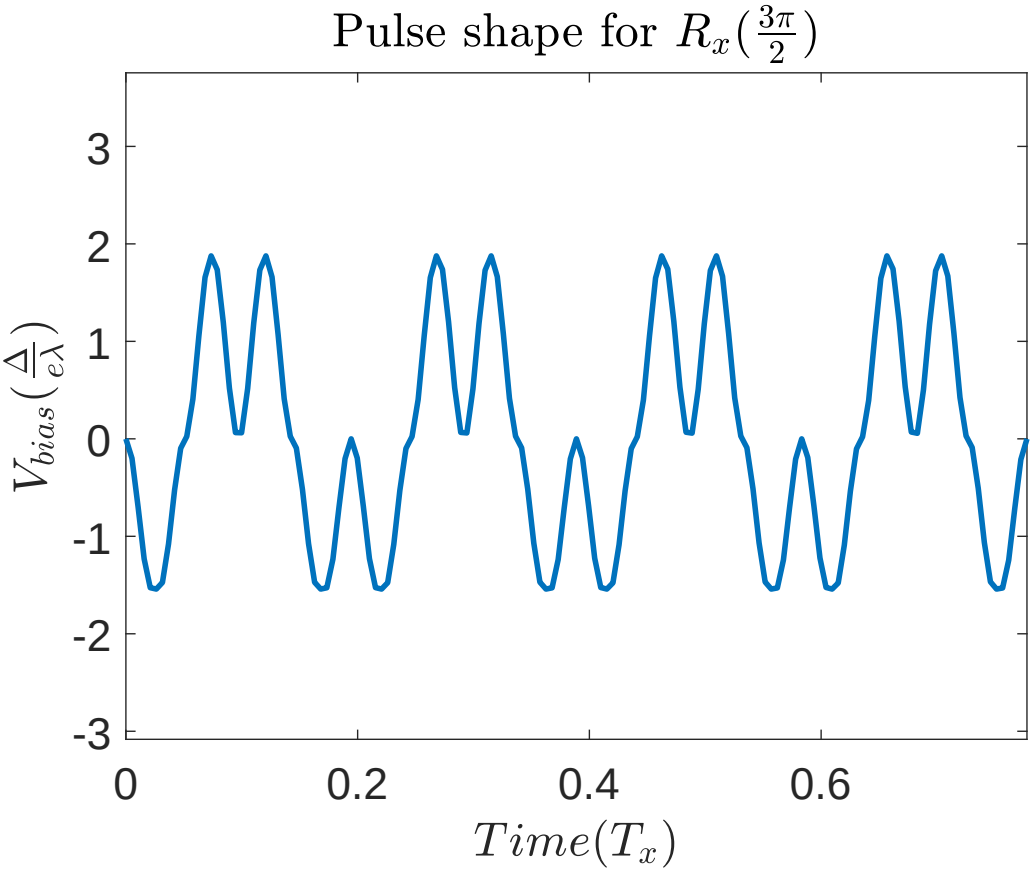

Fig. 3.12 A gain in fidelity is observed when performing an R_x rotation by $\theta = \frac{3}{2}\pi$. The resulting pulse shape oscillates from negative to positive multiple times, effectively canceling out some of the noise. Noise strength is expressed in units of reference pulse strength $\sigma_{\text{noise}} = \frac{\Delta}{e\lambda}$, and varied from $0.01\frac{\Delta}{e\lambda}$ to $0.2\frac{\Delta}{e\lambda}$.

a noise-reducing sequence, the oscillating nature of the pulse takes it from being negative to positive frequently. This will average out the influence of noise to a significant degree, while keeping the total operation time close to the one for an ideal square wave.

Overall, we conclude that the control techniques presented here are at least as good in resisting noise as using a square wave, and can improve upon it under certain conditions. Therefore, they are suitable to replace the square wave, and to have further noise-reducing methods applied upon them, while they offset any errors due to rise time. The optimised $R_{\vec{x}}$ rotation is able to mitigate charge noise up to almost threefold in the fidelity error (this gain increases with noise strength), given that rise time τ enables one to perform multiple smaller rotations that add up to a required total angle.

V Discussion

When experimentally optimising qubit rotations, voltage pulses are usually considered as square while rise and fall times from instrument limitations and other filtering effects due to the finite bandwidth of coaxial cables are neglected. While the voltage is gradually rising to some intended amplitude, the qubit will undergo transient rotations, and will not reach the expected position on the Bloch sphere. These errors accumulate over long operations, leading to poor fidelity. Moreover, applying very sharp pulses of high amplitude, with the intent of performing an R_z rotation, can lead to unwanted energy excitations due to non-adiabacity, causing further fidelity loss [135, 244, 215]. The control scheme presented here overcomes both problems by explicitly adjusting the pulses for rise time, and by using relatively low pulse amplitudes, making the operations adiabatic. By using a specific amplitude giving us two perpendicular rotation axes, we achieve single-qubit control without the need for strong non-adiabatic pulses, or the requirement for perfectly square ones. The disadvantage of this scheme is operation time. As the pulse amplitude is tied to the energy gap between the first two eigenstates of the DQD, there is little control of the rotation speed, at least in the case of a semiconductor DQD system. However, careful engineering of the DQD allows for the operation time to be tailored or optimized [140]. As long as the system energies can be tuned so that the operation time is much less than qubit coherence time, the benefits of increased operation fidelity will outweigh the cost of increased duration.

In this work, we simulate a semiconductor GaAs-based DQD using finite difference methods. The parameters for our simulations were chosen to be experimentally realistic in terms of energy, time scales and pulse generation. We kept these values general since specific rise/fall times and inter-dot energies will depend on each experimental implementation. Current systems are capable of generating pulses with $\tau = 40$ps-500ps. Experimental work by Fujisawa *et al.* [102, 80] contain gate pulses with $\tau \sim 100$ps, with total pulse time of 600ps and $V_{\text{bias}} = 40\mu$eV. More recent work achieves at least 40ps pulse resolution with advanced techniques [152].

The groups cited above as well as other semiconductor-based quantum dot research [220, 173] could practically eliminate errors due to rise time and pulse-induced excitations outside of the computational space by using our proposed pulse sequences.

While the semiconductor charge qubit system was used in simulations in this work, our results are easily generalisable to other types of qubits, as long as the Hamiltonian is of a similar form to Eq. 3.9. For example, the same scheme can be used to control a spin qubit by varying the magnetic field B instead of a voltage bias. In this particular case, it is easier to adjust the energy splitting $\Delta = \frac{1}{2}\gamma B$ by applying a strong reference magnetic field. Increasing Δ will result in faster operation. However, in the charge qubit case, it is achieved by lowering

the DQD barrier. This will increase the overlap between eigenstates, decreasing localization and thus readout fidelity. No such issue arises for the spin qubit, overcoming the slower operation time of our framework. Our results can then be directly translated to the spin qubit case, by applying a magnetic field B' in some perpendicular direction to B.

VI Conclusions

We have described quantum control of the optimal charge qubits for a double-quantum dot system. We presented pulse sequences for state preparation and arbitrary qubit rotations, and show how to account for the experimental control suffering from finite rise/fall times. Owing to hybridization of the eigenstates in a double-dot system, the spatial wave function of the two lowest energy eigenstates cannot be confined exclusively to the left and the right dot. The optimal qubit was found to be defined in terms of the two lowest energy eigenstates of a zero-bias system. This allowed us to reduce our model to a two-state system.

We show that it is possible to prepare the qubit in such a state when it is initially in the ground state of a DQD. Combining theory and numerical techniques yields an optimal pulse sequence that accomplishes arbitrary single-qubit rotation even with non-zero rise time τ. We demonstrate how our framework results in high fidelity despite $\tau > 0$, while avoiding unwanted excitation to higher energy states. Indeed, we show that square pulses are not only unnecessary, but also undesirable, as the sharp rise can induce unwanted oscillations, while being simple to account for. Since our proposed pulse sequence reduces to sine waves to minimize total pulse duration, it is straightforward to implement experimentally. As our numerical fitting parameters depend only on the energy splitting Δ, the results are easily scalable to any particular system. Our scheme is easily generalizable to other qubit systems with similar Hamiltonians, such as spin qubits. Applying our results will lead to increased operation fidelity in many systems, making them viable for practical quantum computing applications.

We also investigate the impact of charge noise on the control proposed protocol, and conclude that using at the very least keeps noise-related error not larger than they would be for a square pulse. In the case of $R_{\vec{x}}$ rotation for certain angles, we find that the noise is actually significantly mitigated by the oscillating waveform, which leads to improvements in fidelity. Overall, we conclude that our protocol is a candidate for further error correction that could be applied to a square wave, but can provide some enhancement on its own.

Our method of accounting for rise/fall times bears resemblance to the GRAPE (Gradient Ascent Pulse Engineering) algorithm [190], however there are important differences. Our method specifically works to cancel the rise/fall times of assumed profile (sinusoidal in

this work, but the method can be used for any shape), resulting in a simple lookup of two parameters ξ and ΔT depending on required angle of rotation and τ itself. GRAPE instead is a more general "black box" technique that tries to optimise a pulse sequence by constructing it from slices of piecewise constant amplitudes, by tuning these amplitudes via gradient ascent methods. We find that the method used here is simpler to implement for experimentalists, and outputs a waveform composed of sinusoids, which can be described analytically, and is by design not limited by the device rise/fall time.

VII Readout

For completeness, we discuss a potential procedure for the readout process. In experimental setups, it is the probability of finding the electron in one of the dots which is measured rather than the qubit superposition weighting coefficients. We can express both qubits defined in Sec. III in terms of their right and left dots parts:

$$\psi_0(x) = \langle x|0\rangle = f_{0L}(x) + f_{0R}(x) \tag{3.39}$$

$$\psi_1(x) = \langle x|1\rangle = f_{1L}(x) + f_{1R}(x) \tag{3.40}$$

Because the qubits $|0\rangle$ and $|1\rangle$ are orthogonal, we have:

$$\begin{aligned} 0 = \int \psi_0^*(x)\psi_1(x)dx = \int f_{0L}^*(x)f_{1L}(x)dx + \\ \int f_{0L}^*(x)f_{1R}(x)dx + \int f_{0R}^*(x)f_{1L}(x)dx + \\ \int f_{0R}^*(x)f_{1R}(x)dx = \int f_{0L}^*(x)f_{1L}(x)dx + \\ \int f_{0R}^*(x)f_{1R}(x)dx. \end{aligned} \tag{3.41}$$

The qubits are mirror images of each other, such that $\langle x|0\rangle$ has the same spatial distribution in the left (right) dot as $\langle x|1\rangle$ has in the right (left) one. We also know that there is some non-zero overlap, unless the DQD barrier is completely separating the dots. Therefore Eq. 3.41 implies that :

$$\int f_{0R}^*(x)f_{1R}(x)dx = -\int f_{0L}^*(x)f_{1L}(x)dx = \eta. \tag{3.42}$$

Any arbitrary state can be written as a linear combination of the two qubits right and left dot components

$$\begin{aligned}\psi(x) = \alpha\psi_0(x) + \beta\psi_1(x) = \\ \alpha\Big(f_{0L}(x) + f_{0R}(x)\Big) + \beta\Big(f_{1L}(x) + f_{1R}(x)\Big),\end{aligned} \tag{3.43}$$

The probability P_R of finding the particle in the right dot is then:

$$\begin{aligned}P_R = \int_0^\infty \psi^*(x)\psi(x)dx = \int_0^\infty \Big(\alpha^* f_{0R}^*(x) + \\ \beta^* f_{1R}^*(x)\Big)\Big(\alpha f_{0R}(x) + \beta f_{1R}(x)\Big)dx.\end{aligned} \tag{3.44}$$

Using Eq. 3.42, this reduces to:

$$\begin{aligned}P_R = |\alpha|^2 \int_0^\infty f_{0R}^*(x) f_{0R}(x)dx + \\ |\beta|^2 \int_0^\infty f_{1R}^*(x) f_{1R}(x)dx + \eta(\alpha^*\beta + \alpha\beta^*) = \\ |\alpha|^2 P_{0R} + |\beta|^2 P_{1R} + 2\eta\mathscr{R}(\alpha^*\beta),\end{aligned} \tag{3.45}$$

where the integrals P_{0R} and P_{1R} can be obtained initialising the qubit in the $\psi_0(x)$ or $\psi_1(x)$ state, respectively, and measuring the probability of finding it in the right dot. Combining Eq. 3.45 with the normalisation condition for $\psi(x)$, we obtain an equation relating $|\beta|$ to the probability P_R of finding the particle in the right dot, up to an error term proportional to η, which quantifies the uncertainty of determining whether the qubit is in the left or right side of the DQD:

$$|\beta|^2 = \frac{P_R - P_{0R}}{P_{1R} - P_{0R}} + \delta. \tag{3.46}$$

A similar expression exists for $|\alpha|^2$, with P_L being the probability of finding the particle in the left dot :

$$|\alpha|^2 = \frac{P_L - P_{0R}}{P_{1R} - P_{0R}} - \delta, \tag{3.47}$$

where $\delta = 2\eta\frac{\mathscr{R}(\alpha^*\beta)}{P_{0R}-P_{1R}}$ is the effective error. Since $P_{1R} \approx 1$, $P_{0R} \approx 0$, we can estimate the maximum readout error, which would occur for a maximally entangled state:

$$| \delta | \lessapprox \eta. \tag{3.48}$$

For the parameters used in this paper, $| \delta | \leq 8 \cdot 10^{-4}$. This magnitude of readout error is not very significant compared to other sources of errors in a quantum computation [141, 111], such as two-qubit gates, relaxation, or dephasing, especially since it's only applied once as the final step. Additionally, it was shown [140] that in a similar situation, adiabatically increasing the inter-dot barrier of the DQD preserves coherence, while greatly reducing this type of' "overlap" error -this technique should be used when possible if the readout error is noticable. Alternatively, as this error is a result of lack of knowledge of $\mathscr{R}(\alpha^*\beta)$, a full state tomography could be performed to eliminate it completely (assuming that errors of operations associated with the tomography do not outweigh the readout error). Therefore, we conclude that measurement of the charge distribution is a viable way of reading out the qubit in our scheme, as it is not a significant source of error compared to others, provided that the experimental fidelity is not a limiting factor.

VIII Fidelity as a function of initial state

Error is found by calculating the fidelity - overlap between the target state and the iterated state. Although some variation in fidelity is dependent on the initial state of the electron, any errors are below 10^{-4}, and as low as 10^{-8} for some initial positions. This error could be reduced further if necessary by fine-tuning the adjustment parameters $\xi, \Delta T$. Figures 3.14, 3.15, and 3.16 show a fidelity map for the R_x, R_y and R_z rotations respectively, as a function of Bloch sphere angles θ, ϕ. A rotation angle of π was chosen in each case, but the results are similar for all angles. Each plot corresponds to 500 simulations of the rotation starting from different initial states equally distributed over the Bloch sphere (with equal sphere surface area for each point) .

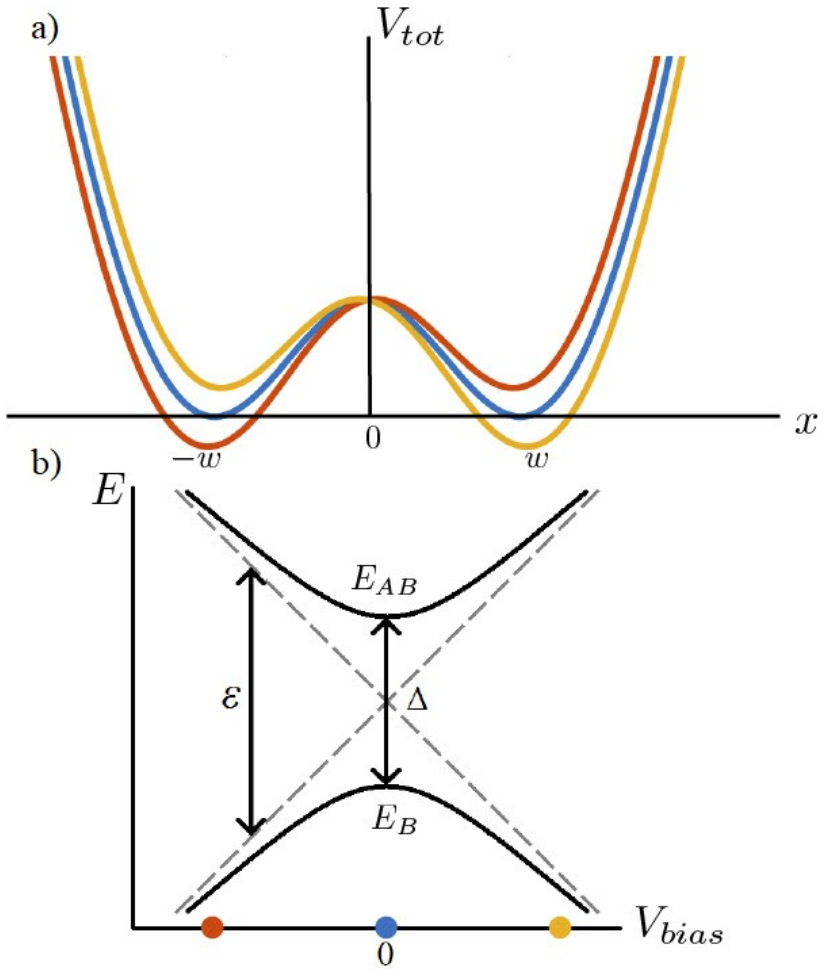

Fig. 3.13 (a) The DQD potential V_{tot} at zero (blue), lowest (red) and highest (orange) detuning values. (b) Energies E of the bonding (E_{B}) and anti-bonding (E_{AB}) eigenstates. The coloured dots mark potential shapes from part (a).

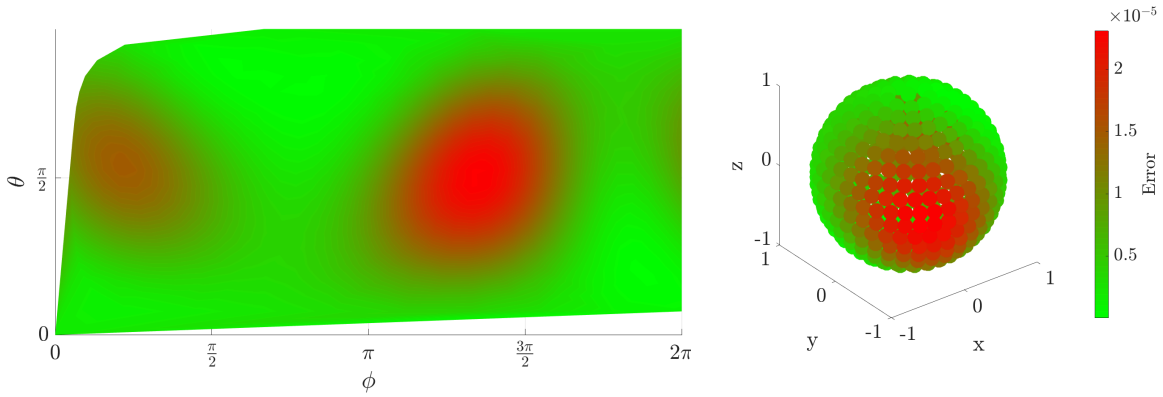

Fig. 3.14 Error in fidelity for an $R_x(\pi)$ rotation, as a function of initial position on the Bloch sphere. The variation in fidelity is relatively small, and the error is expected to approach 0 as parameters are tuned to an ideal value.

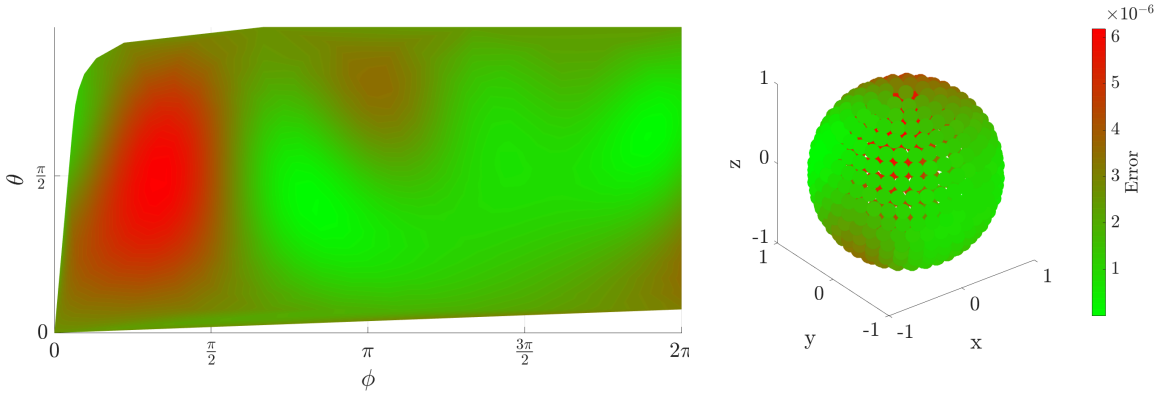

Fig. 3.15 Error in fidelity for an $R_y(\pi)$ rotation, as a function of initial position on the Bloch sphere. The variation in fidelity is relatively small, and the error is expected to approach 0 as parameters are tuned to an ideal value.

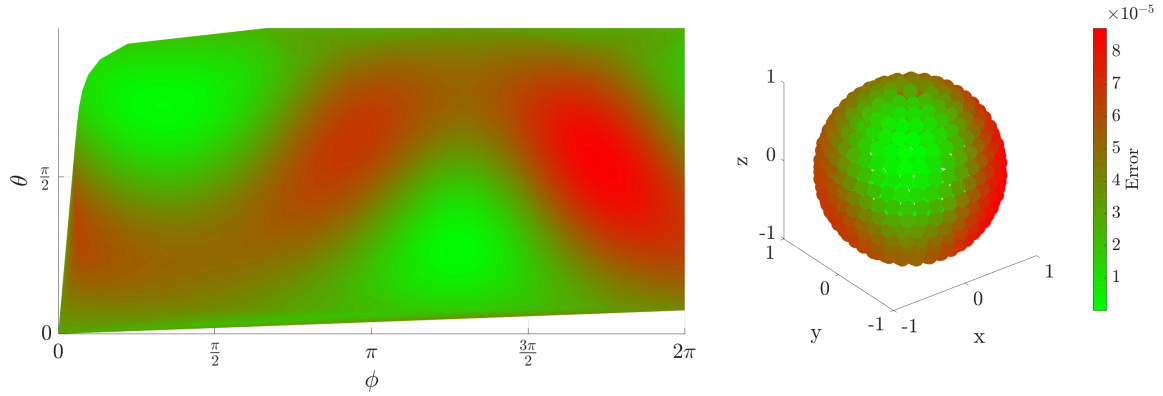

Fig. 3.16 Error in fidelity for an $R_z(\pi)$ rotation, as a function of initial position on the Bloch sphere. The variation in fidelity is relatively small, and the error is expected to approach 0 as parameters are tuned to an ideal value.

Chapter 4

Root-of-SWAP operation with single electron qubits

I Introduction

The development of a universal semiconductor quantum computer hinges on the ability to entangle qubits. One promising method is to use the exchange interaction between electron-spins. This concept was first introduced by Loss *et al.*, for static qubits [146, 32], and Barnes *et al.*, for dynamic, also called *flying*, qubits carried by surface acoustic waves (SAWs) [19]. The use of flying qubits trapped in SAWs is a particularly favorable platform for quantum computation for two reasons. First, the dynamic nature of the qubits enables on-chip operations to be controlled by static electric and magnetic fields from surface gates and magnetic microstructures [83, 159]. Not having to vary surface gate potentials reduces associated errors. Second, the confinement caused by the SAW potential prevents spatial dispersion of the fermionic wave packets [14]. The framework is especially promising for building a universal quantum transducer—a bus that transports entangled qubits between spatially separated parts of a quantum computer that could itself be implemented in a different technology [196, 192].

There are some limitations associated with this architecture - the downside of embedding the computation in the layout of the device means that for arbitrarily long computation, the electrons would have to be reflected back at some point, as they move at the speed of sound through a device, and every gate would take up some finite lengths. Dynamic pulsing is also required to initialise the electrons. However, these potential drawbacks are not fundamental, and not necessary even relevant, depending on the use scenario (for example, a quantum transducer does not have to perform long computations).

The last decade has seen significant developments in the achievement of SAW technologies [53, 155]. Advances include the reliable control of single-electron transport [158, 105, 76, 63] and the increase in electron-qubit coherence times [214, 193, 137, 173]. The SAW framework for manipulating electron-spin qubits has shown promise for realizing optics-like quantum processes with readily interacting particles. Experimentally, single-qubit operations [63], beam-splitters [186, 217] and spin-polarization readout devices [131] have been realized in GaAs heterostructures, and a spin-qubit toolkit for the implementation of generalized measurements has been presented [14].

Previous works on flying electron systems have been restricted to either single particle scenarios [34, 25], or analytical models with a limited number of sites [32, 19] and simplified simulations in 1D for two particles[165, 86].

In the latter case, it was suggested that entanglement generation could be achieved using a single-shot root-of-SWAP operation in which two electrons collide in a harmonic potential. Attempts to simulate realistic devices in layered 3D systems have faced problems owing to the space- and time-domain scaling associated with solutions to the many-particle time-dependent Schrödinger equation (TDSE). However, recent advances in graphics-processing-unit (GPU) performance [8, 191] have made previously demanding problems readily solvable. It is now possible to perform fine-grained two-particle 2D simulations of such systems with realistic device potentials.

In this work, we utilize our state-of-the-art GPU hardware to run a customized *staggered-leapfrog* algorithm [89, 15, 153], as described in Ch. 2. In particular, we study SAW-based flying qubits, interacting via the Coulomb interaction in a 2D double-dot potential. Our results demonstrate the experimental viability of entanglement generation via root-of-SWAP operations. Furthermore, we show that the single-shot method [165] is not experimentally feasible. Not only are our simulations useful to gain insight into quantum logic operations, they also shed new light on simpler analytical models. Specifically, we compare our simulation results to two commonly used two-site models: the Hubbard approach [19], and the Hund-Mulliken method for molecular orbitals [32]. We establish the limits and applicability of these models. We use experimentally realistic parameters for the interaction duration, the device potential, and geometry.

Moreover, the method we present is easily tunable to achieve not only the Root-of-SWAP operation, but a more general Power-of-SWAP. The matrix representation of the more general Power-of-SWAP operation in the two-qubit basis $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ is:

$$SWAP^n = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{2}(1+e^{i\pi n}) & \frac{1}{2}(1-e^{i\pi n}) & 0 \\ 0 & \frac{1}{2}(1-e^{i\pi n}) & \frac{1}{2}(1+e^{i\pi n}) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{4.1}$$

For Root-of-SWAP, $n = \frac{1}{2}$. This power n is essentially a continuous parameter that we can control in the experiment.

To ensure parameter realism, we calculate the potential profile of the heterostructure with voltages applied to the metallic gates. We use a Poisson-Schrödinger self-consisted solver to calculate the range of values that are possible with current semiconductor technologies [217]. Since this work demonstrates a proof-of-concept for the SAW-driven entangling operation, we use analytical equations to reproduce the potentials calculated by our solver. In doing so, we avoid simulating a specific device implementation and ensure that these simulations are reproducible and adaptable to different experimental needs. Our work is a vital step towards constructing the fundamental building blocks of a SAW-based quantum computer. The simulations we present are based on the GaAs/AlGaAs SAW-based heterostructure but our methodology, results, and conclusions are applicable to other semiconductor quantum systems, including static quantum dots.

II Device Description

Let us begin with a description of the device structure and the electron dynamics that allow us to model an entangling operation between two spin-qubits. The physical spin-qubits are electrons, and their spatial dynamics are controlled by SAWs. In each operation, two qubits travel through channels separated by a potential barrier. At the locus of the entangling operation, this barrier is lower, allowing the electron qubits to swap via the exchange interaction.

Fig. 4.1 shows a SAW device designed to carry out a power-of-SWAP entangling operation on two electrons. The device is an adaptation of the one presented in [19]. Sinusoidal SAWs are generated by interdigitated transducers and propagate as transverse plane waves in the positive y-direction. The SAWs modulate the electric potential of a piezoelectric substrate to produce a train of quantum dots propagating along channels defined by metallic gates. The SAWs trap pairs of electrons from a two-dimensional electron gas in the same minimum [185], with one electron in each channel (separated in the x-direction by the tunnel barrier). The SAW then carries the electrons through their respective channels. In the center

of the device, where the barrier is lower, the tunneling rate can be controlled by voltages on TB_L, TB_R.

The quantum dynamics of the system are generated by the travelling SAWs, therefore the voltages on the surface gates can be held constant throughout the entangling operation. This gives the SAW-based system a significant advantage over static qubit systems, which are controlled by generating voltage pulses that introduce charge noise and can induce stray SAWs, causing decoherence.

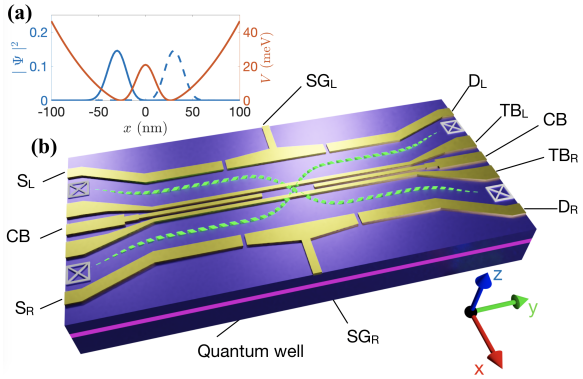

Fig. 4.1 **(a)** Cross section of the potential layout in the region of high barrier alongside a trace of the initial state of the wave function along the x-dimension. **(b)** Schematic of a SAW-based power-of-SWAP device. Electrons are carried by SAWs from bottom to top (positive y-coordinate) along two channels, undergoing a power-of-SWAP operation in the central gate region. Dotted lines show the path electrons can take through the device. Figure from [140], made by H. Lepage.

In what follows we will consider a device consisting of a GaAs/AlGaAs heterostructure containing a single layer of two-dimensional electron gas trapped in a quantum well. On the top surface, a pattern of Schottky gates creates the two channels running in the y-direction, separated by a central barrier. We define the barrier's center as the origin of the x-direction and label the two channels with subscript L (left) and R (right) for negative or positive x, respectively. Negative voltages on the gates labelled by S_L, S_R, SG_L, SG_R, D_L, D_R generate the outer walls of the two channels. Voltages on the gates labeled CB, TB_L, and TB_R control the profile of the central barrier, and ensure that it strongly separates the two channels, except at the middle part of the device in the y-direction, where the barrier is lower. It is in the region of lower central barrier that the entangling operation occurs. TB_L and TB_R are sufficiently close that they only produce a single potential maximum, in the x-direction, between the two channels (Fig. 4.1(a)).

III Analytical Model

To describe quantum dynamics in this device, we use a two-particle Hamiltonian of the form

$$\hat{H} = \sum_{i=1,2} \left(\frac{\hat{p}_i^2}{2m_i} + \hat{V}_{\mathrm{D}}(r_i) + \hat{V}_{\mathrm{SAW}}(t, r_i) \right) + \hat{V}_{\mathrm{C}}(r_1, r_2), \tag{4.2}$$

where $\hat{V}_{\mathrm{C}}(r_1, r_2)$ is the two-particle Coulomb potential, $\hat{V}_{\mathrm{SAW}}(t, r)$ is the SAW potential carrying the electrons along the channels and $\hat{V}_{\mathrm{D}}(r)$ is the device potential. This potential is made up of two parallel harmonic channels running along the y-dimension, coupled in the central gate region by a Gaussian tunnel barrier, forming a double quantum dot with harmonic confinement perpendicular to the channels, along the $x-$dimension. An explicit expression is given in Appendix VII. By boosting our reference frame to match the velocity of the SAW, which is constant, we can treat $\hat{V}_{\mathrm{SAW}}(t, r)$ as a time-independent confining potential along the channel direction. Finding the eigenstates of the boosted time-independent Hamiltonian using a number basis derived from second quantization allows us to obtain the two-particle wave functions when the barrier between both channels is static. Since the potential does not have any explicit spin dependence, because of a weak Lorentz term, single-qubit spin rotations do not occur.

We assume that the electrons in both channels of the device in Fig. 4.1(b) are in a separable spin state initially. At this stage, there is a high potential barrier between the channels and they are too far apart to interact. We also assume they are in eigenstates of the z-axis spin. The spin part of the wave function can thus be labeled $|s_1\rangle|s_2\rangle$, meaning that the first electron is in spin state s_1, and the second one is in spin state s_2. For a double-dot potential, the two-particle ground state is symmetric in spatial coordinates, described by a spatial wave function $|\Psi^{\mathrm{S}}(r_1, r_2)\rangle$, while the first excited state is anti-symmetric, with a spatial wave function $|\Psi^{\mathrm{A}}(r_1, r_2)\rangle$. We call the spin-antisymmetric combination a singlet state $|\mathrm{S}\rangle$, which corresponds to the ground state with energy E_{S}, and the symmetric state a triplet state $|\mathrm{T}\rangle$, which is the first excited state with energy E_{T} (see Fig. 4.2):

$$|\mathrm{S}\rangle = \frac{1}{\sqrt{2}} |\Psi^{\mathrm{S}}(r_1, r_2)\rangle \left(|\uparrow\rangle|\downarrow\rangle - |\downarrow\rangle|\uparrow\rangle \right) \tag{4.3}$$

$$|\mathrm{T}\rangle = \frac{1}{\sqrt{2}} |\Psi^{\mathrm{A}}(r_1, r_2)\rangle \left(|\uparrow\rangle|\downarrow\rangle + |\downarrow\rangle|\uparrow\rangle \right) \tag{4.4}$$

We choose the double-dot potential of the gate region such that an equal linear combination of these states has both particles well localized in different channels. This results in the eigenstates of initial high tunnel barrier and those of the gate region having a high overlap.

The disturbance introduced by adiabatically changing the tunnel barrier in the SAW reference frame is thus minimized. We can write down combined space and spin states as $|s_1 s_2\rangle_{\mathrm{LR}}$, with particle 1 being in the left channel with spin s_1 and particle 2 being in the right channel with spin s_2. They are linear combinations of the triplet and singlet states:

$$\begin{aligned} |{\downarrow\uparrow}\rangle_{\mathrm{LR}} &= \frac{1}{\sqrt{2}}\left(|\mathrm{T}\rangle + |\mathrm{S}\rangle\right) \\ &= \frac{1}{\sqrt{2}}\left(|\Psi^{\mathrm{RL}}(r_1,r_2)\rangle\,|{\uparrow}\rangle|{\downarrow}\rangle - |\Psi^{\mathrm{LR}}(r_1,r_2)\rangle\,|{\downarrow}\rangle|{\uparrow}\rangle\right), \end{aligned} \tag{4.5}$$

$$\begin{aligned} |{\uparrow\downarrow}\rangle_{\mathrm{LR}} &= \frac{1}{\sqrt{2}}\left(|\mathrm{T}\rangle - |\mathrm{S}\rangle\right) \\ &= \frac{1}{\sqrt{2}}\left(|\Psi^{\mathrm{LR}}(r_1,r_2)\rangle\,|{\uparrow}\rangle|{\downarrow}\rangle - |\Psi^{\mathrm{RL}}(r_1,r_2)\rangle\,|{\downarrow}\rangle|{\uparrow}\rangle\right), \end{aligned} \tag{4.6}$$

where $|\Psi^{\mathrm{LR}}(r_1,r_2)\rangle$ denotes a spatial state with particle 1 in the left channel (negative x) and particle 2 in the right channel (positive x). These take the form

$$|\Psi^{\mathrm{RL}}(r_1,r_2)\rangle = \frac{1}{\sqrt{2}}\left(|\Psi^{\mathrm{S}}(r_1,r_2)\rangle + |\Psi^{\mathrm{A}}(r_1,r_2)\rangle\right), \tag{4.7}$$

$$|\Psi^{\mathrm{LR}}(r_1,r_2)\rangle = \frac{1}{\sqrt{2}}\left(|\Psi^{\mathrm{S}}(r_1,r_2)\rangle - |\Psi^{\mathrm{A}}(r_1,r_2)\rangle\right). \tag{4.8}$$

A system placed in such a linear superposition oscillates coherently with the period, $2\pi\hbar/(E_{\mathrm{T}} - E_{\mathrm{S}})$, determined by the energy difference between the ground state and first excited state. A full SWAP operation takes half of this period whilst the root-of-SWAP operation takes a quarter of it, i.e. half the duration of a SWAP. In the limit where the on-site Coulomb energy is much greater than the hopping energy, the doubly-occupied states have vanishingly small probability amplitudes and can be ignored [19]. The state during the time evolution is

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}}\left\{|\mathrm{T}\rangle + \exp\left(\frac{-it}{\hbar}\Delta E\right)|\mathrm{S}\rangle\right\}, \tag{4.9}$$

where $\Delta E = E_{\mathrm{T}} - E_{\mathrm{S}}$. This description of the power-of-SWAP operation allows us to calculate the probabilities of observing spin-up (spin-down) particles in the left (right) channels after the operation. The probability of measuring a swapped state, assuming an initial state $|{\uparrow\downarrow}\rangle_{\mathrm{LR}}$

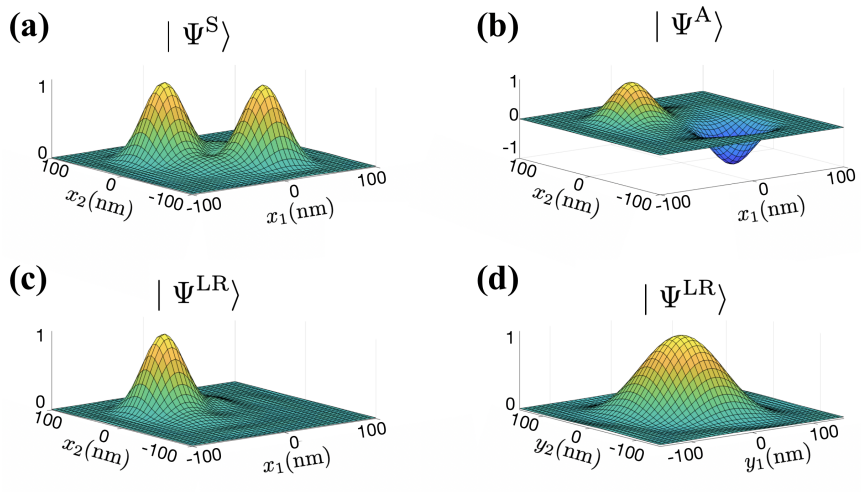

Fig. 4.2 **Two-particle spatial wave functions.** **(a)** Ground state $|\Psi^{S}\rangle$ ($y_1 = 0$ and $y_2 = 0$). **(b)** First excited state $|\Psi^{A}\rangle$ ($y_1 = 0$ and $y_2 = 0$). **(c)** Combination of the ground state and first excited state $|\Psi^{LR}\rangle$. The first particle is localized in the left channel and the second particle is localized in the right channel ($y_1 = 0$ and $y_2 = 0$). **(d)** Gaussian spread of both particles in the y-dimension ($x_1 = 0$ and $x_2 = 0$). All four panels show the wave function divided by its extremum, with the z-axis in arbitrary units.

and a fixed time of interaction τ, is given by

$$\begin{aligned} P_{\mathrm{SWAP}}(J) &= |\,\langle \psi(t=\tau)|\downarrow\uparrow\rangle_{\mathrm{LR}}\,|^2 \\ &= \sin^2\left(\frac{1}{2}J\cdot\tau\right), \end{aligned} \tag{4.10}$$

where $J = \Delta E/(2\pi\hbar)$. This probability, given an input state, depends only on the energy difference between the triplet and singlet states, defined by exchange energy, which in turn is a function of device potential. While any double-dot type potential will result in eigenstate structure to be qualitatively similar to what is described above, some fine-tuning is needed to get the overlap and energy difference ΔE just right for our application. Having the inter-dot energy barrier too high will result in a fully separable degenerate state with $\Delta E = 0$, which will prevent any interesting interaction from happening. This is desirable when we do want the channels/qubits to be separate, but not during the quantum operation. On the other hand, having the barrier too low will result in a high overlap between the left and right parts of the wavefunction, as well as in a high ΔE, causing their interaction to be strong and fast. This would make fine control of the operation challenging. Additionally, the overlap between the qubits in the middle of the dot would cause a loss of fidelity when the barrier is raised again to end the SWAP operation. The potential values we've found here are realistic and present a

good balance between speed of operation, fidelity, and ability to control the interaction as desired (by being able to stop at a required power-of-SWAP).

IV Simulations and results

Numerical simulations of the two-particle dynamics over two dimensions are computationally expensive. The simulations ran here on the hardware described in Ch. 2 take around 10 hours to complete, and we need to sweep over many parameter values to find the right ones, requiring us to run many simulations in parallel. However, the ability to model a complete set of energy eigenstates reveals a more complicated behavior, in contrast with analytical two-site models or one-dimensional simulations - we find ourGPU-accelerated numerical simulations indispensable. In this section, we present the numerical results of entanglement generation via two different root-of-SWAP implementations [32, 19, 165]. In both cases, we find that the realistic dynamics deviates from the simpler models.

IV.I Numerical methods

The eigenstates of the double-dot system are obtained by numerically solving the Hamiltonian built using the allowed two-particle basis states (Ch. 2 Sec. II). To reduce the size of our matrix representation, we can find the initial state of the two-particle system efficiently by using a momentum-space (rather than a position-space) eigensolver (Ch. 2 Sec. I.II). Since the system of interest is very close to the ground state, only a small number of momentum basis states are needed. The ground and first excited spatial states found using this method are presented in Fig. 4.2.

These time-independent methods are sufficient to find the initial state of the system and to describe its time evolution in a constant potential. However, when the potential varies as the electron travels across the device, time-dependent simulations need to be used. We evolve the TDSE iteratively using the staggered-leapfrog algorithm [89, 15, 153]. For the time-dependent evolution, we also find that using the momentum representation is optimal. While the decrease in simulation time due to reduced number of points is modest, as the overhead of transforming to momentum space is still significant at this problem size, we obtain increased fidelity this way. Similarly to the time-independent solution, this can be explained by the system staying very close to its ground and excited states, which in turn are well-represented by low-lying momentum modes. To time-evolve in momentum space, the staggered-leapfrog algorithm is transformed into momentum-space as described in detail in Ch. 2.

IV.II Coulomb tunneling entanglement generation

Building on a proposal from Ref. [19], we explore an exchange-interaction based method for the generation of entanglement between two electrons in a SAW system. As described in Section III, the two electrons occupy adjacent channels separated by a tunneling barrier, suppressing any wave function overlap. When the electrons enter the low-barrier gate region, they can tunnel through to the other channel at a rate that is determined in part by the barrier height (which controls the wave function overlap) and in part by the Coulomb force (which is indirectly controlled by barrier width), thus allowing for the control of the power-of-SWAP gate by tuning the appropriate Schottky gates. Fig. 4.4 shows snapshots of the wave function during an entangling operation with realistic experimental parameters. When the potential barrier is low, the two-particle state undergoes coherent oscillations between the initial state and the fully swapped state. The duration of the two-particle operation is determined by the length of the tunnel-coupled region. Since the speed of a SAW is constant in a given material, the operation is identical for all incoming electron pairs.

Starting with Eq. 4.10, and assuming that J is exponentially dependent on the tunnel barrier height A_{TB}, and time of interaction τ is fixed, the probability of the final state being swapped with respect to the initial state has the following dependence on the tunnel barrier:

$$P_{\mathrm{SWAP}}\left(A_{\mathrm{TB}},\tau\right)=\sin^2\left(\frac{1}{2}J_0\cdot e^{-b\cdot A_{\mathrm{TB}}}\cdot\tau\right), \tag{4.11}$$

where J_0 and b are numerically determined parameters. Fig. 4.3 shows a fit of our time-dependent numerical simulation data (See Appendix VII for parameter values used) with the analytical prediction from Eq. 4.11. It is important to note that although Eq. 4.11 can describe the behaviour of a power-of-SWAP under ideal conditions, a numerical approach is required to account for more realistic scenarios. These can include the presence of impurities in the quantum channels, other sources of noise/decoherence, as well as a finite transition length between the low and high tunnel barrier heights. The inset in Fig. 4.3 shows the probability amplitude of each computational basis state as the electrons travel through a root-of-SWAP gate. Interactions between the electrons are initially prohibited by the high potential barrier separating them. As they are carried through the interaction region, the electrons become entangled. Upon leaving the region of low potential barrier, the particles can no longer interact and the probability amplitudes become constant. We find that the SWAP probability around $P_{\mathrm{SWAP}}=0.5$ varies with the tunnel barrier height at a rate of $8.07\times10^{-4}\ (\mu\mathrm{eV})^{-1}$. This allows for an experimentally viable tunability of the quantum gate via the control of the tunnel barrier height. Assuming a device temperature of 300 mK,

tunnel barrier variations due to thermal fluctuations (of magnitude k_BT) will decrease the root-of-SWAP fidelity by $< 0.1\%$, a value that would not be a main source of error, and in line with other qubit implementations. This error could be reduced by increasing the height of the tunnel barrier, at the cost of extending the operation time.

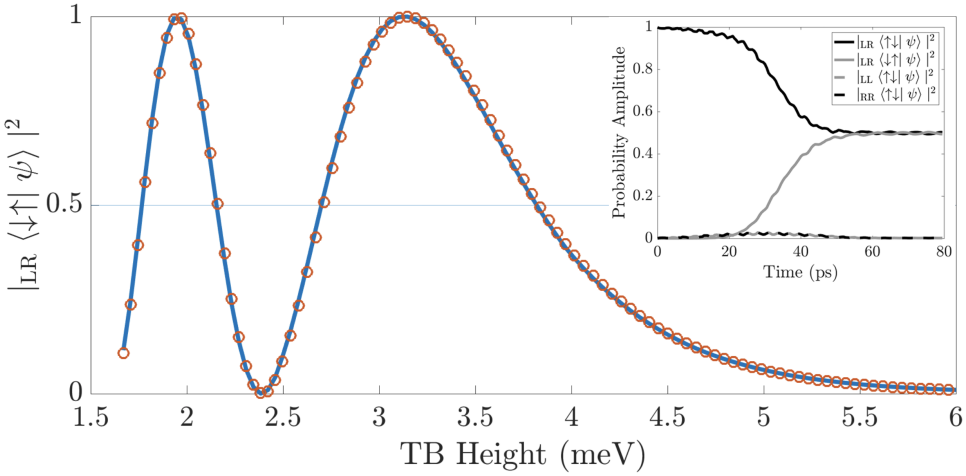

Fig. 4.3 **Probability of SWAP as a function of tunnel barrier height for fixed interaction duration.** Time evolution simulation results (circles) are fit using Equation 4.10 (solid line). The parameters $J_0 = 2.888\text{ps}^{-1}$ and $b = 0.933\text{meV}^{-1}$ were found numerically. The inset figure illustrates the occupation of the computational basis states as well as the double occupancy states. In this example, the input state $|\uparrow\downarrow\rangle_{\text{LR}}$ undergoes a root-of-SWAP operation with finite tunnel barrier potential ramps.

IV.III Comparison to Models

To solve the dynamics of the power-of-SWAP operation in our heterostructure SAW-based device, including the 2D spatial extent of the wave function and a time-dependent potential, numerical simulations must be used. However, to avoid lengthy and complicated computations, ΔE can be estimated using simplified two-site models, thus getting an approximation for the power-of-SWAP extent via Eq. 4.10.

Assuming a tight-binding-like model, where electrons can tunnel between the quantum dots, we can estimate the full 2D time evolution by applying the Hund-Mulliken model for molecular orbitals [32].This model builds a two-particle basis from right- and left-localized single-particle states $|\phi_\pm\rangle$. These states are orthonormalised to $|\Phi_\pm\rangle = (|\phi_\pm\rangle - g\,|\phi_\mp\rangle)/(\sqrt{1-2Sg+g^2})$, where $S = \langle\phi_\pm \mid \phi_\mp\rangle$ is the wave function overlap and $g = (1-\sqrt{1-S^2})/S$. The singly- and doubly-occupied two-particle basis is constructed with direct

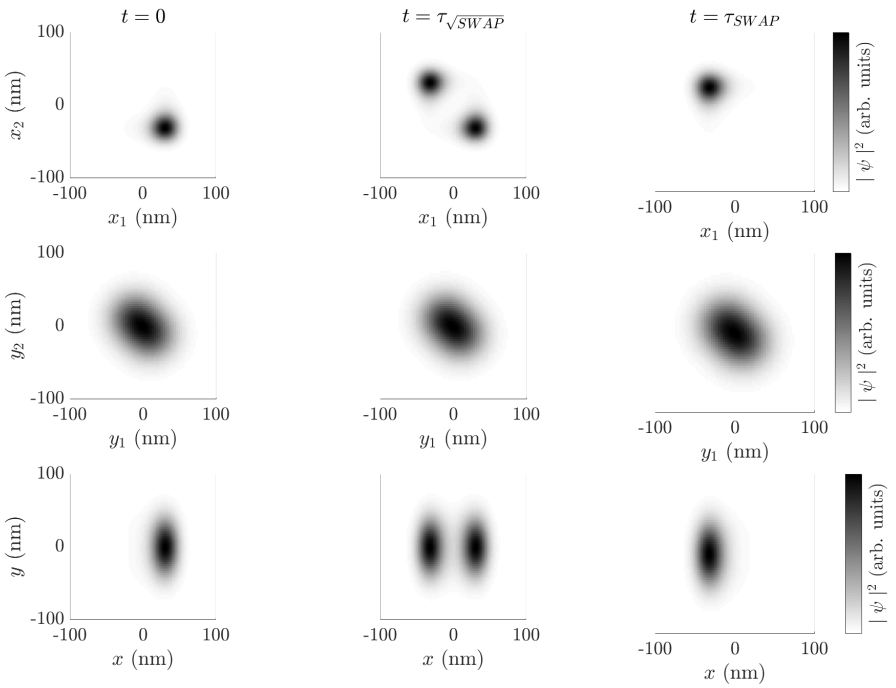

Fig. 4.4 **Entanglement generation using the Coulomb tunneling method.** Top and middle row: trace over the x-dimension and y-dimension respectively for the initial state (left), root-of-SWAP state (centre), and SWAP state (right) of the wave function. Bottom row: trace over the second particle for the initial state (left), root-of-SWAP state (centre), and SWAP state (right) of the wave function. Coordinates are chosen to be in the SAW frame of reference with $y = 0$ corresponding to a SAW minimum and $x = 0$ the peak of the tunnel barrier.

products:

$$\begin{aligned} |\Psi^s_{\mp}\rangle &= \frac{1}{\sqrt{2}}\left(|\Phi_+\rangle|\Phi_-\rangle \mp |\Phi_-\rangle|\Phi_+\rangle\right), \\ |\Psi^d_{\mp}\rangle &= |\Phi_{\mp}\rangle|\Phi_{\mp}\rangle. \end{aligned} \tag{4.12}$$

The Hamiltonian in this basis has the form:

$$\hat{H} = \begin{pmatrix} V_- & 0 & -\sqrt{2}t_\mathrm{h} & 0 \\ 0 & V_+ & -\sqrt{2}t_\mathrm{h} & 0 \\ 0 & -\sqrt{2}t_\mathrm{h} & U & X \\ 0 & -\sqrt{2}t_\mathrm{h} & X & U \end{pmatrix}, \tag{4.13}$$

where each entry is defined as:

$$U = \frac{e^2}{4\pi\varepsilon}\langle\Psi^d_{\pm}|\frac{1}{r}|\Psi^d_{\pm}\rangle, \tag{4.14}$$

$$X = \frac{e^2}{4\pi\varepsilon}\langle\Psi^d_{\pm}|\frac{1}{r}|\Psi^d_{\mp}\rangle, \tag{4.15}$$

$$V_+ = \frac{e^2}{4\pi\varepsilon}\langle\Psi^s_{+}|\frac{1}{r}|\Psi^s_{+}\rangle, \tag{4.16}$$

$$V_- = \frac{e^2}{4\pi\varepsilon}\langle\Psi^s_{-}|\frac{1}{r}|\Psi^s_{-}\rangle, \tag{4.17}$$

and t_h is the hopping term

$$t_\mathrm{h} = \langle\Phi_{\pm}|\frac{\hat{p}^2}{2m}|\Phi_{\mp}\rangle - \frac{e^2}{4\sqrt{2}\pi\varepsilon}\langle\Psi^s_{+}|\frac{1}{r}|\Psi^d_{\pm}\rangle. \tag{4.18}$$

Solving the Hund-Mulliken Hamiltonian in Eq. 4.13, we find the eigenenergies associated with the singlet and triplet states and define the SWAP frequency in terms of U and t_h:

$$J = \frac{1}{2\pi\hbar}\left[V_- - V_+ + \frac{1}{2}\left(\sqrt{U_\mathrm{h}^2 + 16t_\mathrm{h}^2} - U_\mathrm{h}\right)\right] = \frac{E_\mathrm{T} - E_\mathrm{S}}{2\pi\hbar}, \tag{4.19}$$

where $U_\mathrm{h} = U - V_+ + X$.

Alternatively, the evolution of the two-particle state can be modelled with the Hubbard approach for short range Coulomb interaction without magnetic fields [19]. The simplified

2-site Hamiltonian in the second quantization basis from Eq. 4.12 then has the form

$$\hat{H} = \begin{pmatrix} V & 0 & 0 & 0 & 0 & 0 \\ 0 & V & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -t_{\mathrm{LR}} & -t_{\mathrm{LR}} \\ 0 & 0 & 0 & 0 & t_{\mathrm{LR}} & t_{\mathrm{LR}} \\ 0 & 0 & -t_{\mathrm{LR}} & t_{\mathrm{LR}} & U & 0 \\ 0 & 0 & -t_{\mathrm{LR}} & t_{\mathrm{LR}} & 0 & U \end{pmatrix}, \tag{4.20}$$

where t_{LR} is the hopping term

$$t_{\mathrm{LR}} = {}_{\mathrm{L}}\langle\uparrow| \frac{\hat{p}^2}{2m} |\uparrow\rangle_{\mathrm{R}} = {}_{\mathrm{L}}\langle\downarrow| \frac{\hat{p}^2}{2m} |\downarrow\rangle_{\mathrm{R}}, \tag{4.21}$$

U is the on-site energy

$$U = \frac{e^2}{4\pi\varepsilon} {}_{\mathrm{LL}}\langle\uparrow\downarrow| \frac{1}{r} |\uparrow\downarrow\rangle_{\mathrm{LL}} = {}_{\mathrm{RR}}\langle\uparrow\downarrow| \frac{1}{r} |\uparrow\downarrow\rangle_{\mathrm{RR}}, \tag{4.22}$$

and

$$V = \frac{e^2}{4\pi\varepsilon} {}_{\mathrm{LR}}\langle\uparrow\uparrow| \frac{1}{r} |\uparrow\uparrow\rangle_{\mathrm{LR}} = {}_{\mathrm{LR}}\langle\downarrow\downarrow| \frac{1}{r} |\downarrow\downarrow\rangle_{\mathrm{LR}}. \tag{4.23}$$

Solving this Hamiltonian, we find the eigenenergies associated with the singlet and triplet states and define the SWAP frequency in terms of U and t_{LR}:

$$J = \frac{1}{4\pi\hbar}\left(-U + \sqrt{U^2 + 16t_{\mathrm{LR}}^2}\right) = \frac{E_{\mathrm{T}} - E_{\mathrm{S}}}{2\pi\hbar}. \tag{4.24}$$

To summarise, by solving the Hund-Mulliken Hamiltonian we find the eigenenergies associated with the singlet and triplet states and define the SWAP frequency in terms of the on-site energy U and the hopping term t_{h}:

$$J = \frac{1}{2\pi\hbar}\left[V_- - V_+ + \frac{1}{2}\left(\sqrt{U_{\mathrm{h}}^2 + 16t_{\mathrm{h}}^2} - U_{\mathrm{h}}\right)\right] = \frac{\Delta E}{2\pi\hbar}, \tag{4.25}$$

where $U_{\mathrm{h}} = U - V_+ + X$.

For the simplified Hubbard Hamiltonian, this expression reduces to

$$J = \frac{1}{4\pi\hbar}\left(-U + \sqrt{U^2 + 16t_{\mathrm{LR}}^2}\right) = \frac{\Delta E}{2\pi\hbar}. \tag{4.26}$$

For realistic Hamiltonians, it is impossible to obtain U analytically. Instead, we numerically calculate this parameter. To avoid unphysical results introduced by the $1/r$ factor, a

softened (avoiding division by 0) Coulomb potential is used [165] both in the models and the numerical simulations throughout this work. We implement this softening by assuming that the wave function has a Gaussian spread in the third dimension, with a standard deviation of Δ_z.

We compare both the Hubbard model and the Hund-Mulliken method described above to our simulation results for a range of Δ_z. We find a close match between the frequency of the SWAP operation as calculated by our time-dependent numerical solver and these obtained by solving the eigenvalue problem directly. Both models (Eq. 4.25 and Eq. 4.26) show significant discrepancy for most values of Δ_z. Moreover, the Hund-Mulliken model predicts negative frequencies for $\Delta_z < 1$ nm (which are unphysical in this context). We conclude that although both models provide a reasonable qualitative prediction of the two-particle dynamics for Gaussian spread of $\Delta_z \sim 10 - 100$ nm (on the same order as 2DEG width), a more sophisticated numerical approach, such as the one used in this work, is required to obtain precise quantitative dynamics. A quantitative comparison of both analytical models with our simulations can be seen in Fig. 4.5.

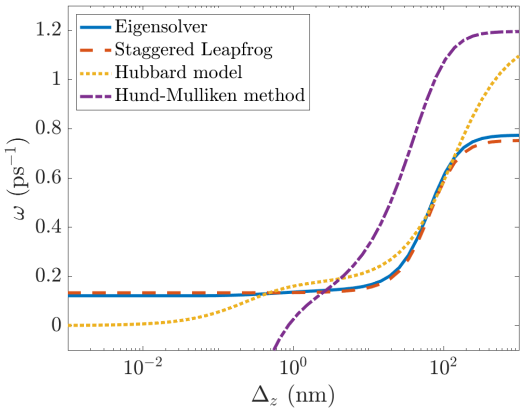

Fig. 4.5 **Comparison to analytical models.** Power-of-SWAP frequency as a function of effective wave function spread in the z-dimension. Coulomb softening accounts for the finite z-dimension and plays an important role in determining the rate of the exchange interaction.

IV.IV Entanglement generation via electron collisions

In a previously suggested root-of-SWAP scheme [165], two electrons are travelling in individual channels separated by a high potential barrier, such that there is no wave-function overlap. The potential barrier *abruptly* (or diabatically) changes in the SAW reference frame such that the two channels are joined to create a global potential minimum between them. Without the presence of the barrier, both electrons fall towards one another in a harmonic

oscillator potential and interact via the Coulomb force. Once the operation is completed, the central barrier is reintroduced, causing the reappearance of separate decoupled channels. As the quantum states of particles in layered semiconductor technologies are confined in the dimension perpendicular to the quantum wells, which has a constant potential throughout the device, the third dimension does not significantly affect the operation. However, we find that the previous reduction to 1D is an oversimplification, as the possible spatial dynamics in the second dimension strongly affects the electron-electron interactions.

Here, we simulate this *single-shot* (i.e. in a single collision) entanglement generation, and we find that under current experimentally realistic parameters, it is impossible to generate a root-of-SWAP, or any significant entanglement over the x-dimension.

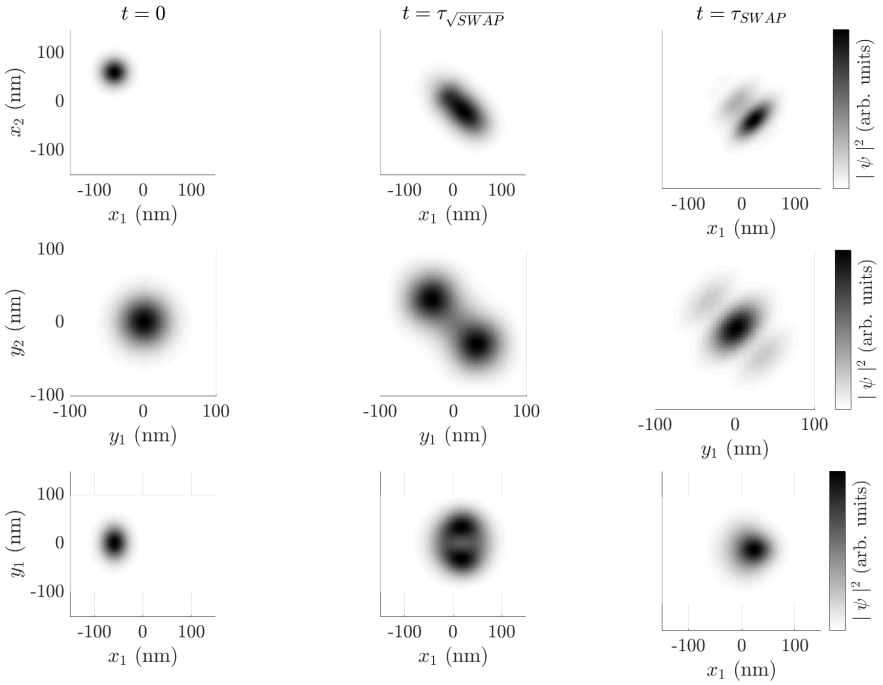

Fig. 4.6 **Entanglement generation via the collision of two electrons.** Top and middle row: trace over the x-dimension and y-dimension respectively for the initial state (left), root-of-SWAP state (centre), and SWAP state (right) of the wave function. Bottom row: trace over the second particle for the initial state (left), root-of-SWAP state (centre), and SWAP state (right) of the wave function. Coordinates are chosen to be in the SAW frame of reference with $y = 0$ corresponding to a SAW minimum and $x = 0$ the middle of the harmonic channel.

Fig. 4.6 shows snapshots of the two-electron wave function undergoing a single collision in two dimensions. The wave function remains fully separable along the x-dimension. However, in the y-dimension, it transitions from a Gaussian-like low-energy state to a more spread-out entangled state. This is conflicting with the desired outcome of generating a

maximally entangled state in the x-dimension. The third column of the figure, which shows the wave function at the time of a full SWAP, should look like the first column, only reflected aroung the $x = y$ axis, but this is not the case. The cause of this is the splitting of the wave function seen in the lower middle panel - the particles split in the y-dimension and avoid each other, which modifies the interaction compared to one dimension. The operation is effectively a SWAP instead of a root-of-SWAP, with the additional downside of exciting higher-energy states in the y-dimension. These unwanted spatial excitations of the wave function lead to lowering the spatial fidelity of the operation and thus it is not possible to concatenate multiple operations for useful quantum information processing. This also prevents the restoration of the wave function to its original state by applying the SWAP twice, a fundamental property of this operation. We find that increasing the y confinement does not prevent this behaviour until the SAW amplitude is increased by a factor of order 10^3, where the problem effectively reduces to 1D. However, this would require SAW amplitudes on the order of 10^4 meV, which is experimentally unrealistic [218]. Varying the x confinement over a wide range also does not solve the issue. Therefore we conclude that the collision method is unable to produce the root-of-SWAP operation in a realistic 2D scenario.

V Experimental Sensitivity

Here, we investigate the power-of-SWAP operation's sensitivity to disturbances in A_{TB} and τ, which the output probabilities depend on. Methods used here are similar to what is used in Ch. 8, borrowing from information theory. From an information-theoretical perspective, an experiment's sensitivity to an unknown physical parameter, θ, is quantified by the Fisher information:

$$F(\theta) = \sum_i P(\mathscr{M}_i|\theta)\left[\frac{\partial}{\partial\theta}\ln P(\mathscr{M}_i|\theta)\right]^2, \tag{4.27}$$

where $\mathscr{M}_i$ denotes the i^{th} measurement's outcome [104]. Given N experimental runs, the precision of an estimate $\hat{\theta}$ of θ is bounded by the Cramér-Rao inequality: $\mathrm{Var}(\hat{\theta}) \geq [NF(\theta)]^{-1}$, such that the greater the Fisher information, the smaller the estimator's variance can be [78]. Using the output probabilities from Sec. III, we find that $F(A_{\mathrm{TB}}) = b^2\tau^2 J_0^2 e^{-2bA_{\mathrm{TB}}}$: the ability to estimate A_{TB} decreases exponentially with A_{TB} itself. However, for parameters that yield the root-of-SWAP operation ($A_{\mathrm{TB}} \approx 3.86$ meV) we find that $F(A_{\mathrm{TB}}) \approx 2.15\ (\mathrm{meV})^{-2}$. This value of $F(A_{\mathrm{TB}})$ lower bounds the standard deviation of A_{TB}: $\sigma_{A_{\mathrm{TB}}} \gtrapprox 0.012$ meV in an experiment with $N = 3000$ trials. Despite the exponentially decreasing sensitivity of tunnel-barrier heights, the relevant values for a root-of-SWAP operation are within the experimentally viable regime [217] specified in Appendix B. The Fisher information about τ is

$F(\tau) = J_0^2 e^{-2bA_{\mathrm{TB}}}$ is constant with respect to τ itself. For the ideal root-of-SWAP parameters, we find that $F(\tau) \approx 6.17 \times 10^{-3}$ ps^{-2}, which gives a lower bound on the standard deviation of τ: $\sigma_\tau \gtrapprox 0.23$ ps, when $N = 3000$.

VI Discussion And Conclusion

The two-qubit entangling operation is an essential building block of a quantum information processor. We find that realising the power-of-SWAP operation via electron collision [165] suffers from significant problems, whilst an implementation based on tunnelling [19, 32] is shown to be realizable with high fidelity even when experimental control of the tunneling barrier is imperfect. We find that this entangling operation governed by the exchange interaction and coherent tunneling of electrons offers a more stable approach and makes high gate fidelities possible. We have shown that surface-gate-controlled flying electron-spin qubits are able to generate entanglement through the power-of-SWAP operation in a reliable and stable fashion using the exchange interaction (tunneling) method. We show that the problem of wave function dispersion can be solved through the use of SAWs, which generate the potential confinement needed to preserve the wave function's profile. We present accurate numerical solutions to the time-dependent Schrödinger equation using a staggered-leapfrog method and we investigate previously proposed schemes for generating entanglement between electron-spin qubits.

Our two-particle simulations use experimentally realistic parameters and potential layouts and show that such devices are readily realizable using current semiconductor fabrication techniques. While the behaviour of an ideal system can be predicted exactly by solving the Hamiltonian and assuming that the electrons are initialized to and remain in a combination of triplet and singlet states, the advantage of our numerical methods is to simulate realistic entangling operations. Although these simulations were focused on the experimental parameters of GaAs-based devices, the same behaviour is expected in other SAW-based semiconductor devices. Moreover, our findings can be generalized to systems that do not include SAWs. Static quantum dots, confined in every dimension and separated by a tunnel barrier interact in the same way. Such a tunnel barrier can be modulated using fast microwave pulses [30] to reproduce the two-particle dynamics presented in this paper. A static root-of-SWAP gate was recently realised with high fidelity using phosphorus donors in silicon [103], proving that such systems are achievable experimentally. Coherent spin state SWAP operations between electron-spin qubits in a quadruple array of semiconductor quantum dots were also lately achieved [117].

Our results provide new evidence that an entangling root-of-SWAP gate based on the exchange interaction is experimentally viable in SAW-based semiconductor heterostructures.

VII Parameter values

Here we present exact parameter values used in the simulations in this chapter, as well as the details of functional potentials and Coulomb softening used. This should enable anyone to reproduce this work, as well as compare the values to some real experiment.

Parameter	Value range
Distance between channels	80 nm
Tunnel coupled region start	y_d = 36 nm
Tunnel coupled region end	y_u = 144 nm
Interaction time	$\tau = 36$ ps
SAW amplitude	A_{SAW} = 25 meV
SAW wavelength	$\lambda = 1\mu$m
SAW velocity	v = 3 nm / ps
Harmonic channel confinement	$\omega_x^2 = 0.002\ \frac{\mathrm{meV}}{\mathrm{nm}^2 m_e}$
Electron effective mass	0.067 m_e
Relative permitivitty (GaAs)	13.1
Gaussian tunneling barrier amplitude	A_1 = 15.3 meV
Gaussian tunneling barrier width	σ_1 = 30 - 40 nm
Gaussian barrier amplitude	A_2 = 510 meV
Gaussian barrier width	σ_2 = 0.8 nm
Transition between barrier heights	σ_y = 10 nm
Coulomb softening	Δ_z = 10 - 100 nm

Table 4.1 Ranges of parameter values used in simulations.

Explicit form of potentials used in Eq. 4.2, in terms of the parameters above, in reference frame of the SAW:

$$V_{\mathrm{D}}(x,y) = \frac{m_e}{2}\omega_x^2 x^2 + A_1 \exp\left(\frac{-x^2}{2\sigma_1^2}\right) + \frac{A_2}{2}\exp\left(\frac{-x^2}{2\sigma_2^2}\right) \times \left(2 - \tanh\left(\frac{y-y_d}{\sigma_y}\right) - \tanh\left(-\frac{y-y_u}{\sigma_y}\right)\right). \tag{4.28}$$

$$V_{\mathrm{SAW}}(x,t) = \frac{A_{\mathrm{SAW}}}{2}\left(1-\cos\left(\frac{x-tv}{\lambda}\right)\right) \tag{4.29}$$

$$V_C(r) = \frac{e^2}{4\sqrt{2}\pi\varepsilon\Delta_z U(-\frac{1}{2},0,\frac{r^2}{2\Delta_z^2})} \tag{4.30}$$

where U is the confluent hypergeometric function of the second kind, which encapsulates a Gaussian spread with standard deviation Δ_z in the z-dimension:

$$U(a,b,z) = \frac{1}{\Gamma(a)}\int_0^\infty e^{-zt}t^{a-1}(1+t)^{b-a-1}dt. \tag{4.31}$$

If this kind of "Coulomb softening" is not performed, the potential will suffer from divide-by-zero errors when the two particles overlap. This is a result of restricting the calculation to 2D. Doing the full 3D calculation would be very computationally expensive. To implement an effective spread in the z-dimension in a physical way, we assume that the third dimension of some thickness Δ_z exists, and then integrate over the z-dimension for an effective Coulomb strength, assuming that the charge distribution is Gaussian in this dimension. It turns out that this integral can be expressed in terms of the confluent hypergeometric function of the second kind, as per Eq. 4.31. This mathematical model has the desired properties - when the inter-particle distance r is small, the effective distance for the Coulomb force calculation is around Δ_z, while at a large distance, the effective distance is very close to r itself. Therefore, this approximation is physical and realistic, and the softening parameter Δ_z is easy to interpret - it's simply the spread of charge in the "extra" z-dimension.

Chapter 5

Mapping the charge states of semi-isolated DQDs using a single donor device

I Introduction

Recent progress in lithography techniques like focused helium ion beam milling [229] or extreme ultraviolet lithography [67] as well as advances in single ion implantation [113] have allowed for shrinking the transistor feature size to the atomic level without compromising on the quality of the surfaces and interfaces, thus opening the way to new applications in high speed computing and high density information storage [246, 161]. Due to the simplicity of their energy level structure, such ultimate devices are generally easier to control, easier to model and ultimately better candidate for ultralow energy-consuming electronics compared with standard Metal-Oxide-Semiconductor (MOS) structures. Currently, such a downscaling can be achieved by implanting deterministically single ions [113] or low concentration dopants (non-deterministic approach) into a nanoscale architectures. However, the first requires high technical expertise and world-class equipment whereas the second relies on randomness and consequently lacks scalability. Such a device can also be fabricated by depositing single atoms with nanometre precision using the tip of a scanning transmission microscope [195]. In this case, requirements for the mechanical stability of the tip implies either a long processing time or the purchase of expensive corrective and stabilization softwares [247, 182]. However, the need for both a cost-effective, fast and reliable solution is still required for carrying out experimental investigations to devise atom-based applications.

We first demonstrate how to electrostatically isolate a single donor from the large ensemble of dopants before investigating the system characteristics and dynamics under gate voltages. We finally show this device can be used as a single atom detector by tuning the detection level to the D_- state and sensing the charge occupancy of a nearby capacitively coupled double quantum dot [73].

II Device fabrication and isolation of single donors

Most of nanoscale silicon transistors currently used in industries make use of a metal-oxide-semiconductor structure which allows for controlling the number of electrons down to the single electron regime with great accuracy. Fin-field effect transistors (FinFETs) are one of many successful examples of such a nano-engineering. Further down-scaling to sub-10 nm structures have been demonstrated [201, 21] and reliable Coulomb oscillations have been obtained up to room temperature. However, the repeatability of the process has not yet been demonstrated at the industrial scale. Still, some of these structures have been proposed as a basic element in quantum information architectures [23].

In contrast, doped devices are often overlooked, despite offering a reduction in the number of processing steps by avoiding the use of top- or back-gates. This lack of interest is partly due to the intrinsic localization effects as well as the randomness in dopant positions and ionization energy that contribute to the $1/f$ noise (beyond undoped device levels) and decrease the detection efficiency. In this type of architecture down-scaling is challenging owing to the high dopant density and the difficulty in realising a dopant free-tunnel barrier. However, in much larger devices, one can take advantage of the resulting glass properties of such a material, in particular the charge rearrangement and the long relaxation time T_1, to electrostatically isolate a single donor from a large number of dopants. This can be achieved by modifying the electrostatic potential of the quantum dot at room temperature and then altering charge relaxation mechanisms while the device is cooled down to the lowest temperatures. To this end, a 60 nm diameter single electron transistor was patterned from a highly phosphorous doped silicon-on-insulator (SOI) material using standardized fabrication process that is described elsewhere [219]. The dopant concentration is about 3×10^{19} cm^{-3}, giving an average donor separation of about 2 nm. The various elements of the devices are defined by etching the SOI down to the underlying silicon oxide in selected areas leaving side gates capacitively coupled to the quantum dot and tunnel barriers forming at constriction points (Fig. 5.1 **a)**). An additional gate connected to a double quantum dot is patterned in the same way.

The experimental part of this work and all measurements were conducted by Thierry Ferrus. A positive source-drain bias V_{SD0} larger than the donor ionization energy ($eV_{SD} \gg 45\,\text{meV}$) is first applied to the device at room temperature, creating a steep potential profile across it without affecting its operability or performance. The device is then slowly cooled down to the base temperature.

By decreasing the temperature, the states located at the edge of the structures start localising first, creating electrostatic tunnel barriers preferentially at the constriction locations. This is mainly owing to the presence of high density non-(100) surfaces, dielectric screening [224, 2], and the formation of P_2O_5 compound within 10 nm of the edges. These states always exist in nanoscale structures, even if highly doped [73].

Under this condition, a charge imbalance appears in the quantum dot, with donors at the center and at the source-side of the island ionizing and creating P positive centers, while a small number $N \ll P$ of electrons accumulate at the edge of the dot near the drain (left side) barrier (Fig. 5.1 **b)** and **c)**). Raising V_{SD0} allows for increasing the electron tunneling rate but also facilitates for the higher energy confined electrons to escape to the drain contact due to the increase energy windows for conduction, thus decreasing the number of accumulated electrons. Due to the inherent insulating behaviour of the devices [1] (Fig. 5.2), the shape of the tunnel barriers and consequently the tunneling rates are affected when the temperature is lowered. Consequently, most donors remain ionized during the cooling process, providing that the flow of electrons through the dot, and consequently the tunneling rate through the device remain steady. This is achieved by monitoring the current I_{SD} and maintaining it at the room temperature value by continuously adjusting the source-drain bias during the cooldown process accordingly.

Once at the base temperature, the bias is slowly decreased to zero. Electrons are then allowed to recombine with ionized donors on the source side by flowing both from the source lead at a rate Γ_S, and from the high electron density region near the drain barrier at $\Gamma_A \ll \Gamma_S$. The initial situation in the quantum dot ($N \ll P$) is then the one of a low electron regime. The variation of conductivity in temperature clearly shows an Efros-Shklovskii variable range hopping is taking place in a large range of temperatures. The small deviation from the usual $\sim T^{-1/2}$ dependence [64] has already been observed in silicon MOS device and can easily be explained by partial screening in the structure due to weak disorder, e.g. the Coulomb gap is getting partially filled [72]. This behaviour demonstrates the role of electron-electron interaction during the cooldown. At high temperatures, charge rearrangement is then preferentially done with next-neighbour donors with high energy level difference. On the

[1] Doping concentration is about 3×10^{19} cm^{-3} which provides metallic behavior in the large areas of the device (source and drain leads). However, localization at the edge of the structure, near the Si-SiO_2 interface.

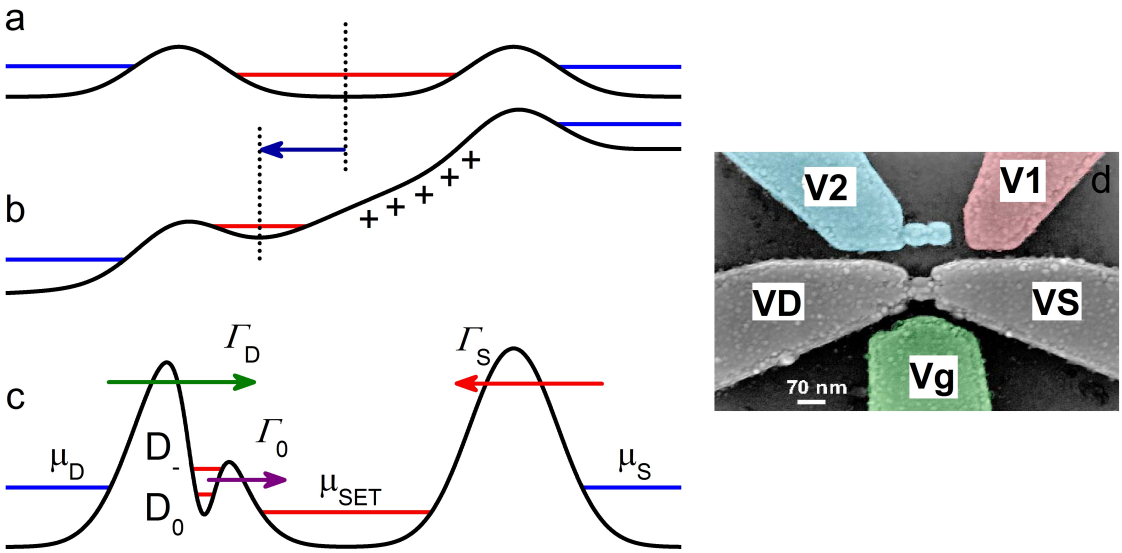

Fig. 5.1 Electronic process involved in the isolation of the donor : **a)** device prior to biasing, **b)**, during bias and **c)**, after the creation of the electrostatic tunnel barrier isolating the donor. **d)** Device under study with the quantum dot (bottom) and the semi-connected double quantum dot (top).

contrary, at low temperatures, hopping occurs between donor pairs well separated spatially but with a small difference in energy.

Consequently, electron dynamics are governed by electron glass behaviour [55] (Fig. 5.2), with shallow traps filled first while occupancy of deeper traps remains unsettled for much longer times. Localization and trap density are defined by the device structure and intrinsic randomness due to the fabrication process, including doping. Because of the spatial location of the traps in the structure, donors at the center of the quantum dot are the first to be neutralised, then followed by the edges states. Traps located in the transport pathway near or at the tunnel barrier are less susceptible to electron screening due to the presence of the tunnel barrier, and so are filled the latest. In the absence of sufficient thermal energy, traps are filled at timescales much longer than potential fluctuations and long-term electrostatic deformation is present in the system. This implies that already ionised donors will remain ionized for a significant time, allowing for the creation of an additional electrostatic tunnel barrier between the dot and the drain tunnel barrier with height greater than $k_B T$ (Fig. 5.1 **c)**). The width and height of the newly created barrier will depend on the competition between Γ_S and Γ_A but also from the ratio N/P, the latter defining the average settling time $\propto C_p^n \propto 1/\Gamma_A$.

The resulting device structure is then similar to a donor-dot system in series [92]. Such an hybrid system has now been widely studied in particular for its ability to act as a spin readout device [209] or memory in quantum information applications [204]. However its'

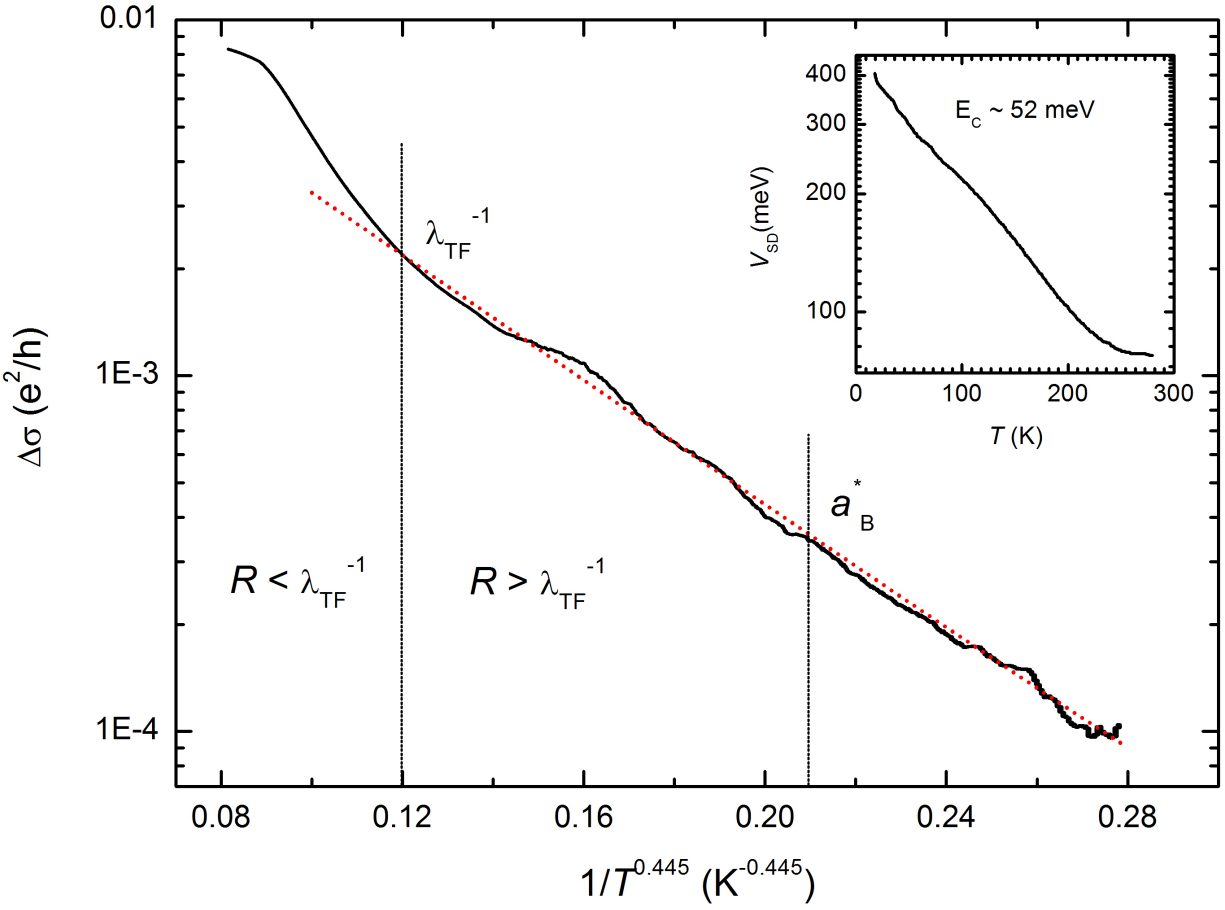

Fig. 5.2 Variation of the conductivity with temperature, showing a screened correlated hopping conduction mechanism. Inset shows the variation of the applied source-drain bias during cooldown with $V_{SD} \gg E_C$, with E_C being the quantum dot charging energy. The $T^{-1/2}$ law demonstrates the presence of electron-electron interaction in an insulating system compatible with $R > \lambda_{TF}^{-1}$, with TF the Thomas-Fermi screening length and R the hopping length.

realization has mostly been uncontrolled with accidental trapping occurring at the tunnel barriers, unlike the proposed method.

III Device characterization

The shape of the Coulomb diamonds (Figs 5.3 **a)**, **b)**) is quite distinct from the one obtained on the same device but following normal cooling conditions, e. g. all connections grounded (5.3 **c)**).The most noticeable difference is the appearance of large diamonds indicating the presence of impurities in the conduction path (Fig. 5.3 **a)**). The standard $D_+ - D_0$ state transition is also clearly observable at $V_g \sim 3.5$ V. The three main diamonds have different sizes, as well as shifts, indicated by red arrows in Fig. 5.3 **a)**. The manifestation of these feature differences is reproducible.

III.I Main quantum dot

The first unusual feature is the presence of low intensity, periodic and reproducible conduction lines at the edge of the diamonds, as shown on high resolution scans. The periodicity in gate voltage $\Delta V_g \sim 75$ mV and the gate capacitance $C_g \sim 2.2$ aF are similar to the ones obtained under normal conditions ($\Delta V_g \sim 78$ mV and $C_g \sim 2.1$ aF), which points towards the main quantum dot being the origin of these features. In the absence of an active donor in the barrier, the corresponding lever arm was found to be $\alpha_g = \frac{C_g}{C_\Sigma} \sim 0.052$, giving a charging energy of 4.1 meV. The corresponding overall dot diameter was estimated to be $D = e/(2\pi\varepsilon\varepsilon_0\alpha_g\Delta V_g) \sim$ 72 nm, with a 58 nm of doped silicon surrounded by a 7 nm silicon oxide. Estimations were based on an effective permittivity $\varepsilon = 9.9$ for the dot structure (see supporting information). These values were in excellent agreement with the dot diameter observed after the electron beam lithography, but before oxidation of ~ 60 nm as well as a target protective oxide of 10 nm.

However, after the cooling procedure, the lever arm is found to be significantly larger ($\alpha_g \sim 0.113$), indicating a significant decrease in the capacitance of the drain. This can only occur if the main transport mechanism is governed by a donor in series with the quantum dot with a much weaker coupling capacitance compared to the drain lead. In that specific case, the estimation of the dot diameter from the total capacitance is no longer valid.

III.II Single donor and ionised trap

The second effect relates to the appearance of shifts in the large Coulomb diamonds both for the D_+, D_0 and D_- states (marked by arrows in Fig. 5.3 **a)**). Understanding the origin

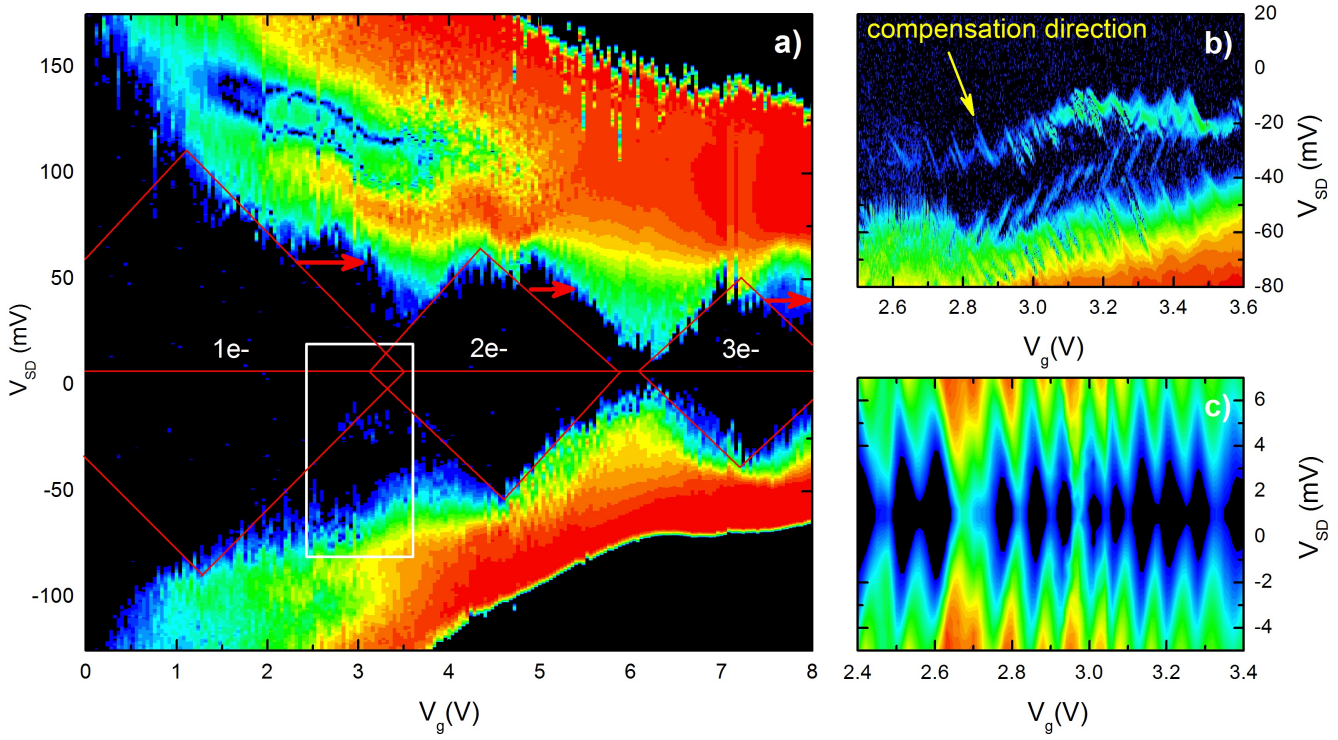

Fig. 5.3 **a)** Coulomb diamonds as observed at 300 mK, **b)** Detailed region around the $D_+ - D_0$ state transition, with fine features resulting from the main quantum dot, **c)** Coulomb diamonds at 300 mK before thermal cycle and cooldown under the strong bias condition, showing the SET operating as a single quantum dot.

of these shifts is important in order to precisely determine the donor energies, as well as to adequately perform the charge detection carried in the next section. Such shifts have already been observed in FinFets where two donors have been diffusing from the contacts into the channel region [177]. In that case, transport was found to be mainly dominated by a single donor (D1) whereas a second donor (D2) had its occupancy modified under certain voltage conditions. A similar situation happens here.

Charging energies E_0 and E_- for both the D_0 and D_- states can be estimated when the donor D2 is ionised by either measuring the size of the diamonds along the V_{SD} axis or along the V_g axis, by converting voltages into energies using the measured lever arm $\alpha_{gD} \sim 0.022$. The corresponding D1 level energies are found to be $E_0 \sim 67$ meV and $E_- \sim 57$ meV, consistently by both methods. These values are far greater than the one expected for isolated donors in bulk silicon [96], but have already been experimentally observed in silicon quantum dots [92]. However, these differences can easily be explained by the device structure and doping. In nanostructures, the presence of interfaces reduces the extension of the wavefunction and in the case of the device studied, the bound state wavefunction is elongated in the current direction, e.g., along the source-drain axis, while being reduced in the transverse direction [178]. This significantly increases the separation between the ground and excited levels, and so between the D_0 and D_- states. This is particularly true as the isolated dopants are located at the constriction, e.g., the narrower part of the device. Also, energy levels are sensitive to the electrostatic environment, in particular the electric field at boundaries, which leads the effective surrounding permittivity to be renormalised and the localisation to be enhanced. We did actually estimate the effective permittivity to be $\varepsilon \sim 9.9$ in the previous sections [224, 2, 211]. This leads naturally the charging energy of the D_0 states to be raised to 63 meV without any other adjustment.

When D2 is occupied by an electron, the local electric field at D1 is modified and its energy levels shifts by the screened Coulomb interaction between the tunneling electron and the electron localised on D2. This allows us to estimate the average distance between the two donors in the tunnel barrier to be $d \sim 1.4$ nm, a value close to the effective Bohr radius in silicon (see Sec. VI.II). Such a cluster state has already been identified in highly-doped devices by Kelvin probe force microscopy (KPFM) [225]. The shifts observed on the D_0 and D_- states are a direct consequence of this interaction and are related to the shift of the D_+ state. These features are well described quantitatively and qualitatively by simulations including a double donor (two donors in parallel with one another) near the drain barrier and in series with the main quantum dot (see Sec. IV).

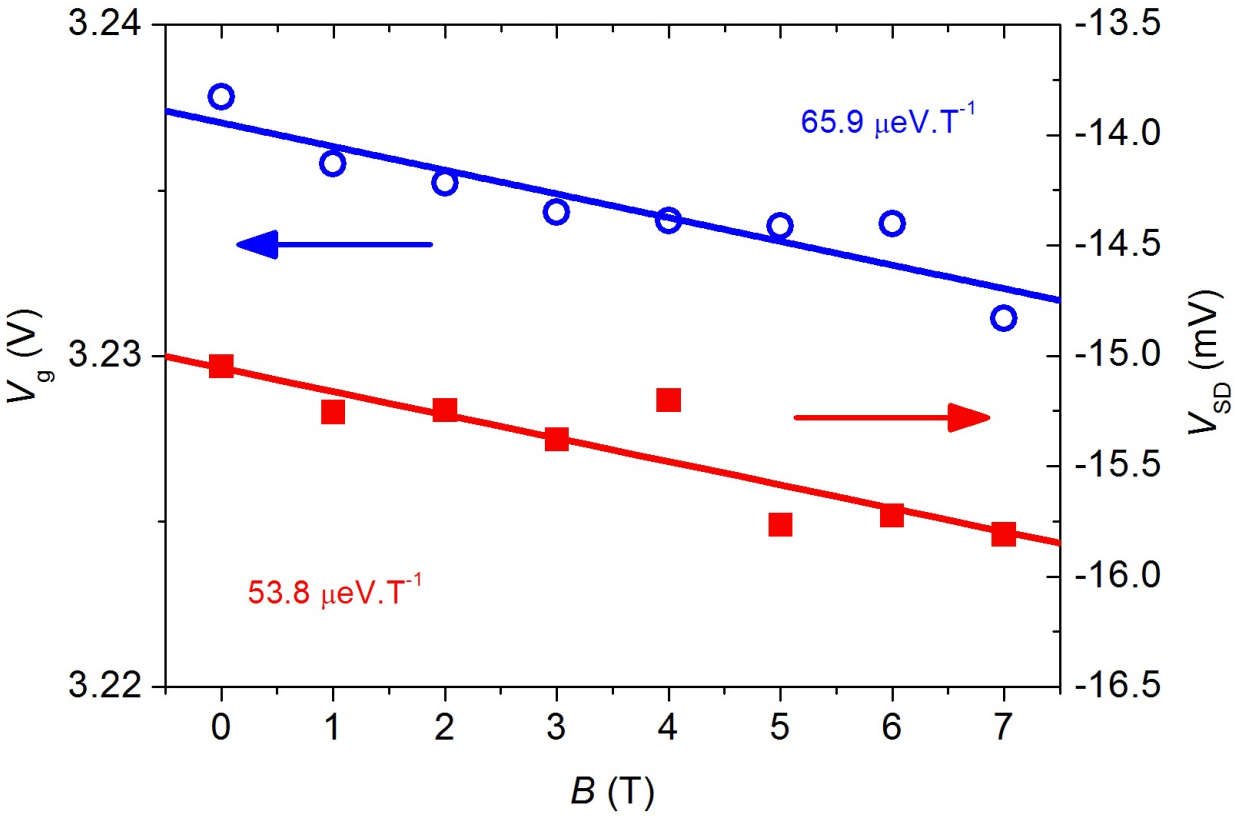

Fig. 5.4 Displacement in V_{SD} and in V_g of the chosen detection point at the D_0 state, as a function of the perpendicular magnetic applied.

III.III Additional features

Additional features are present in the data. In particular, at large $\mid V_{SD} \mid$ and $1V < V_g < 4V$, one can observe the presence of negative differential conductivities (NDC) with isoenergy lines following the edges of Coulomb diamonds. Such a behaviour has already been observed and previously discussed in details in similar doped silicon devices [188]. In the latter case, the NDC resulted from the modification of the shape of the tunnel barrier and so, of the tunneling rates due to the presence of ionised donors near or at the barriers. This shows the NDC is an intrinsic property of doped devices, rather than an effect of the cooldown process. However, the distribution of dopants near the barrier, and so, the shape of the tunnel barrier favour the formation of isolated but transport-active donors.

We also observed the absence of conductivity at negative source-drain bias and $2V < V_g < 4V$, in the region where D2 is singly occupied and a second electron is tunneling through the D_0 of D1. Such an effect is commonly associated with spin-related phenomena. Although, full study of spin states and spin interaction in that device has not been carried out, this observation signifies the role of spin in that region of gate and source-drain biases.

Only a small region at the transition point between the D_+ and D_0 state shows active transport. These features can be displaced linearly by the application of a perpendicular magnetic field both along the V_{SD} and V_g axis (Fig. 5.4), suggesting that the Zeeman effect is active. Results show an average energy level displacement of $\sim 60\mu$eV T^{-1}, consistent with the expected value of $58\mu eVT^{-1}$ for a 1s state electron with an effective Landé g-factor of 2.

Consequently, the transition point at $V_g \sim 2.8$ V (Fig. 5.3) corresponds well to the state where an electron is tunneling via the lowest available level of the D_0 state of the first donor D1 whereas the second donor D2 remains ionised in the D_+ state.

III.IV Using the D_0 state as a charge detector

The previous results allow for the possibility of using the donor as a detector rather than the main quantum dot, as this would normally be the case in most experiments. Such a scheme would present a significant advantage in terms of detection efficiency. In an hydrogen-like donor, only two main states are allowed : the D_0 and D_- states. If the detection point is chosen to be at the D_+-D_0 transition, then tunneling events are restricted to the D_0 state only, as, at low temperatures, the D_- is not accessible without significantly modifying V_g and V_{SD}. The difference in measured currents between the blocked and allowed tunnel events then makes the detection of capacitively coupled structures easily observable. On the contrary, in a quantum dot, multi-tunneling processes are allowed, including inelastic ones like cotunneling. Consequently, the detection efficiency is reduced to a fraction of the one of

a single donor. Additionally, the donor levels are known to be very sensitive to the electric field. Any modification of the electrostatic potential would then shift the levels of the donor and the detection point significantly, improving detection.

In the present device, a semi-connected double quantum dot (SDQD) is capacitively coupled to the SET and the isolated donor. In the SDQD one dot (QD2) is connected to a side gate that act both as an electron reservoir and a gate, whereas the second (QD1) is capacitively coupled to QD2 and the gate V_1. In order to properly use the donor as a detector, a double gate compensation has to be performed. Such a technique consists in modifying simultaneously a pair of gate voltages in order to eliminate the influence of the main dot (SET) on the detection. The charge diagram of the double dot is then obtained by sweeping the gate V_2 and stepping the gate V_1, while keeping the detection point at the D_0 level stable. This implies modifying the SET gate voltage,

$$V_g = V_{g0} + \alpha V_1 + \beta V_2, \tag{5.1}$$

where α and β are the compensating coefficients and V_{g0} the initial value of the gate voltage. As the detection point is at the edge of the D_+-D_0 transition, the source-drain bias V_{SD} also has to be modified in the same way by following the compensation line in the Coulomb peak (Inset, Fig. 5.3 **b)**). Optimal detection is obtained with $\alpha \sim 0.284$ and $\beta \sim 0.373$.

The resulting stability diagram between V_1 and V_2 (Fig. 5.5), despite noisy data, shows characteristics of a weakly coupled DQD structure, in particular the double periodicity ($\Delta V_1 \sim 219$ mV and $\Delta V_2 \sim 599$ mV), and transition lines nearly parallel to the gate axis. Periodic current peaks along V_1 are consistent with electrons tunneling to the dot closer to the V_1 gate. The absence of sharp edges in the transitions is explained by the small difference in capacitances between the individual dots of the double quantum dot and the donor, i.e. $C_{m1} \sim C_{m2}$(Fig. 5.5 **b)**). On the contrary, V_2 both acts as a gate and an electron reservoir to the double dot structure. This gives the possibility for the electrons to tunnel out of the double dot, consequently causing a larger change in the conductivity at the donor site, and so, abrupt jumps in V_2 (Fig. 5.5 **d)**). This allows for estimating the gate capacitances for V_1 and V_2, e.g, $C_{11} \sim e/\Delta V_1 \sim 0.73$ aF and $C_{22} \sim e/\Delta V_2 \sim 0.27$ aF respectively.

We also notice a shift in the pattern along V_1 when V_2 is increased, e.g. when QD2 acquires an additional electron. In this case, adding an electron to DQ1 from the reservoir now requires overcoming the interaction energy between the two dots, which leads to a shift in $V_1 \sim C_m/C_{\Sigma 1}C_{\Sigma 2}$.

Bistable regions are present at around $V_2 = 1.25$ V, 2 V and 2.75 V (Fig. 5.5 **a)**). One has to remember that if the value of the conductivity depends on the tunneling events at

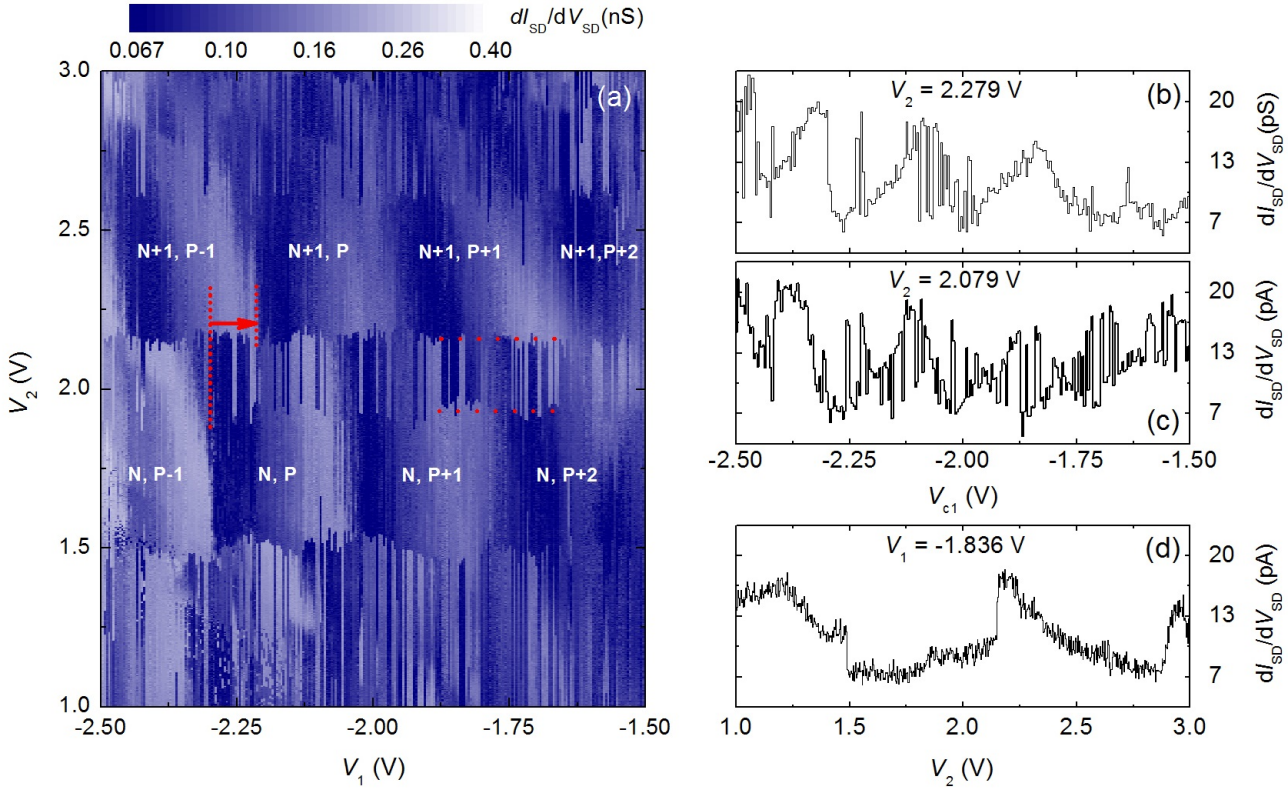

Fig. 5.5 **a)** Compensated $V_1 - V_2$ map of the double dot at 300 mK, **b)** Current profile along V_1 showing individual electron tunneling events inside the double dot, **c)** Current profile at the bistable region involving tunneling events from individual dots to the gate V_1, **d)** Current profile along V_2 indicating electrons tunneling out of the double dot into the gate reservoir.

the detector, the location of the electronic transitions in the compensated diagram result from electron tunneling in the DQD structure. When V_2 increases, three tunneling events are possible, i) an electron can enter QD2 directly, (ii) an electron can enter QD1 by overcoming the interdot coupling energy, (iii) an electron enter QD2 while another electron tunnels between QD2 and QD1. Given the randomness of these processes and the measurement being time-averaged ($t_{\text{measure}} \gg t_{\text{tunnel}}$), this phenomenon is reflected through the observed bistability. By direct measurements of the periodicities in V_1 and V_2, the shift in V_1 and the width of the bistable region in V_2, it is possible to extract capacitance parameters, in particular $C_2^* \sim 0.12$ aF, $C_1^* \sim 0.27$ as well as the coupling capacitance between QD1 and QD2 $C_{\text{m}}^{'} \sim 0.14$ aF. Such a low value for $C_{\text{m}}^{'}$ explains the approximate square shape of the charge stability diagram in Fig. 5.5.

IV Simulations

To confirm the interpretation of the experimental results presented in this work, numerical simulations were performed. Here, I present the results and compare them to experiments, as well as describe the methods and theory used.

IV.I Simulation methods and theory

I have used a constant interaction model that solves a master equation to find the equilibrium states of current and charge occupation at given external voltages. I model the device as a set of nodes that are able to hold quantised amounts of charge, with external electrodes that have voltages imposed on them. Additionally, source and drain electrodes serve as reservoirs of charge that are able to inject and remove electron charges from the system. To solve for the steady state at given external voltages, one first has to know the potential energy of the system, depending on charging state of every element [35, 45].
Firstly, I calculate the electrostatic potential energy of the system by summing over all connected capacitative elements:

$$U_E(N) = \sum_i U_i = \sum \frac{C_i V_i^2}{2}. \tag{5.2}$$

To give a concrete example, imagine the simplest system with an SET serving as a node, connected to (S)ource, (D)rain and (G)ate electrodes with respective capacitances C_S, C_D, C_G . In this case, the electrostatic potential energy is:

$$U_E(N) = \frac{1}{2C_\Sigma}(-eN + C_S V_S + C_D V_D + C_G V_G)^2, \tag{5.3}$$

where e is the electron charge, $C_\Sigma = \sum_i C_i$ is the total capacitance, and V_i are voltages on respective elements.

If more than one node is present, their interactions have to be accounted for. This situation can still be described by the capacitive model. Here I present a general method of obtaining the total electrostatic potential energy of a system of M nodes. We start with a charge vector $\boldsymbol{Q}(N_1, N_2 ... N_M)$:

$$\boldsymbol{Q}(N_1, N_2 ... N_M) = \begin{pmatrix} Q_1 \\ Q_2 \\ \dots \\ Q_M \end{pmatrix}, \tag{5.4}$$

where each $Q_j = -eN_j + \sum_i C_i V_i$ is the electrostatic charge of a single j-th node, where N_j is the number of charges on it, and the sum goes over all connected electrodes (but not other charge nodes). Then, we write down a cross-capacitance coupling matrix $\boldsymbol{C}$ that holds the information of how the nodes are connected:

$$\boldsymbol{C} = \begin{pmatrix} C_{11} & -C_{12} & \dots & -C_{1M} \\ -C_{21} & C_{22} & \dots & -C_{2M} \\ \dots & \dots & \dots & \dots \\ -C_{M1} & \dots & C_{MM-1} & C_{MM} \end{pmatrix}, \tag{5.5}$$

where the off-diagonal elements $C_{ij} = C_{ji}$ are the capacitance between nodes i and j if $i \neq j$, and the on-diagonal elements C_{ii} are the total capacitance of node i, which is the sum over all connected elements. The total electrostatic potential energy is then just the familiar formula $U = \frac{1}{2}Q^2C$ expressed in matrix form:

$$U_E(N_1, N_2 ... N_M) = \frac{1}{2}\boldsymbol{Q}^T\boldsymbol{C}\boldsymbol{Q}. \tag{5.6}$$

However, electrostatic energy is not the only possible contribution in a general case. For the simulations presented in this work, I also include Coulomb interaction energy, as well as the energy stemming from quantised energy levels that the nodes can have. The details of Coulomb energy calculation are presented later in Eq. 5.15. These additional contributions, that in general depend on charges in the system, are added to the electrostatic potential energy to get the total potential energy of the system.

I then use a master equation approach to determine the probabilities of a given (charge) state of the system being occupied. Upon applying the external voltages on the electrodes, the system should quickly settle in a steady state, where probabilities do not change anymore. I define a probability vector $\boldsymbol{P} = (p_1, p_2, ..., p_K)$, where each probability corresponds to some overall charge state. With M nodes, that can hold a maximum number of charges N_i each, the total number of states is $K = \Pi_{i=1}^{i=M} N_i$, and we map the charge state to its' index in the same way one would map a tensor to a one-dimensional index, for example using the row-major order here. I then define a matrix of transition rates between these states, $\boldsymbol{R}$. The off-diagonal elements R_{ij} are transition rates from state i to j, and the diagonal ones are the negative of the sum of the respective row. This ensures normalisation $\sum_i p_i = 1$, as it means that the increase in probability of moving to a different charge state is accompanied by an equal decrease of the probability of staying in the same one. The following equation defines the rate of change of probabilites:

$$\dot{\boldsymbol{P}} = \boldsymbol{R}\boldsymbol{P}. \tag{5.7}$$

To find the steady state, in which $\dot{\boldsymbol{P}} = 0$, one needs to find an eigenvector of $\boldsymbol{R}$ with eigenvalue 0 - the matrix $\boldsymbol{R}$ needs to be eigensolved. However, one first needs to calculate the elements of $\boldsymbol{R}$. As they are transition rates from one charge state to another, one needs to identify all the processes that can lead to such a change in a given system, and sum over their rates. We assume that electrons have a thermal Fermi distribution, and that a transition only occurs if there is an electron with energy such that the transition will be energetically favourable for the whole system. Therefore, the transition rate from state i to j can be written down as:

$$R_{ij} = \sum_k \Gamma_k f_F(\Delta U), \tag{5.8}$$

where Γ_i is some ' intrinsic' rate associated with the process k, and $f_F(\Delta U) = \frac{1}{e^{\Delta U/k_B T}+1}$ is the Fermi distribution, with T being the temperature, k_B the Boltzmann constant, and ΔU the potential energy change associated with the transition (defined to be negative if we are moving to a state of lower potential energy). ΔU can be calculated using the result of Eq. 5.6. To give a concrete example, consider a quantum dot connected to a source electrode. If the electrode is at a voltage V_S, it will be at chemical potential $\mu_S = -eV_S$. If the dot currently holds charge N, the transition rate to a state with one extra electron (coming from the source) is:

$$R_{N,N+1} = \Gamma_S f_F(\Delta U - \mu_S). \tag{5.9}$$

Meanwhile, the opposing transition element is:

$$R_{N+1,N} = \Gamma_S[1 - f_F(\Delta U - \mu_S)], \tag{5.10}$$

where in both equations I have labeled the charge states with their charge number, and ΔU is the potential energy change of the dot.

In the simulations presented here, I additionally model the tunneling resistance of contacts between the nodes. Assuming a square potential barrier of height U_T and length L is present and has to be tunneled through to transport the electron, the transition probability standard formula is:

$$p_{\mathrm{T}} = 16(\mid \Delta U \mid /U_T)(1 - \mid \Delta U \mid /U_T) exp(-2L\sqrt{U_T - \mid \Delta U \mid}), \tag{5.11}$$

where ΔU is the energy difference between the initial and final states. As a reminder, we have defined ΔU to be negative if the transition is energetically favourable. If it is positive, I instead set $p_{\mathrm{T}} = 0$, and if the energy difference is sufficient to completely jump over the barrier ($\Delta U > U_T$), I set $p_{\mathrm{T}} = 1$. Here, I incorporate this effect into the transition rate Γ.

Eigensolving the matrix $\boldsymbol{R}$ yields the probability vector of the steady state. One can then calculate the current flowing through the system by summing over appropriate transitions through the source (or drain) electrode:

$$I_S = \sum_i -e\Gamma_i[f_F(\Delta U - \mu_S)]p_i^{\mathrm{init}}, \tag{5.12}$$

where the sum i is over all transitions that leave the source electrode, and p_i^{init} is the probability that the initial state one transitions from is occupied. The more useful differential current $\frac{dI}{dV}$, which is usually measured in experiments, can then be calculated by simply taking the derivative with respect to the source-drain bias voltage.

IV.II Simulation results

To explain the features observed in Fig. 5.3 **a)**, we model the device as two donors that are in parallel with each other, and in series with a large quantum dot. Therefore, the tunneling through one of the donors must necessarily occur for the charge to flow through the device. Owing to computational restrictions, we do not actually include the large quantum dot in series with the donors. This is justified, as the effects of such a dot in series are well understood to produce fine oscillatory features modulating the Coulomb diamonds observed in Fig. 5.3 **b)**. These features are small enough to not be significant or even visible on the scale of subplot **a)**. Additionally, we confirm that the effects of this large dot are as expected by simulating a simplified system, with results presented in Fig. 5.9.

Fig. 5.7 shows the simulation results of tunneling through the two parallel donors D1 and D2. By directly comparing with Fig. 5.3 **a)**, it can be seen that all the main features of the Coulomb diamonds are well reproduced. There is some mismatch of the finer features for the negative source-drain voltage, however we attribute them to higher order effects, the experimental data having a lot of noise in that region, and possible spin interactions, as discussed in Sec. III.III. The experimental values of capacitances and voltages matches the values used in the simulation very well. The donors are modeled as having discrete energy levels on top of their capacitive potential energy. These discrete energy levels reflect the D_0 and D_- donor charge states. A repulsive screened Coulomb interaction between the charges that the donors hold is needed to reproduce the observed features, which is expected. Overall,

Fig. 5.7 shows that the interpretation of experimental data that there exist two isolated donors that must be tunneled through is correct, as this configuration is necessary to reproduce the data observed.

Additionally, I simulate a small quantum dot in series with a larger one, to confirm that the effects of filtering through this large dot produce only fine oscillations modulating the smaller dot features as expected, and therefore to justify not including this filtering in the main simulation. The results are presented in Fig. 5.8. The small dot has a total capacitance of $C_\Sigma = 1.8aF$, and the larger one $C_\Sigma = 12.8aF$. The large scale features match that of tunneling through the small dot only, while the superimposed fine oscillations are due to filtering through the larger dot.

I also simulate the system before cooling, such that only a single SET is present. The results are presented in Fig. 5.9, and are to be compared with Fig. 5.3 **c)**. There is a very good agreement, with a simple Coulomb diamond pattern present. The scale of the diamonds matches the experimental data well, with the total capacitance used for simulations $C_\Sigma = 33.6aF$ also matching the experimental measurement of capacitance.

Finally, I simulate how this SET can be used as a detector of charge for a nearby DQD. By placing the SET nearby the DQD, the charge state of the DQD with capacitively influence the SET. If tuned correctly and compensated, the charge state of the DQD will be detectable in the current flowing through the SET. This is especially useful of the DQD is only semi-connected, so that there is no current flowing through it to infer the charge from. The setup is presented in Fig. 5.6 - we are measuring differential current through the SET, trying to observe signatures of charge state changes in the nearby DQD. For this simulation, it was not feasible to include the isolated donors for computational reasons, therefore it serves as an illustration of principle for the detection. The differential current going through the SET is measured as a function of the DQD gates voltage (V_1, V_2 on Fig. 5.6) . The effect of varying voltage on the DQD gates is compensated by varying the SET gate voltage V_G accordingly, while keeping a constant source-drain bias that allows the SET current to flow, as per Eq. 5.1. The results are presented in Fig. 5.10. A checkerboard pattern that can be compared with Fig. 5.5 is observed, where each square matches a combined charge state of the DQD. This can clearly be seen by comparing the differential current profile with the calculated charge states of the DQD, presented in Fig. 5.11. Additionally, it can be seen on Fig. 5.12 that the charge on the SET stays largely constant over the voltage range explored, which means that the compensation procedure was successful. The charge is not completely constant over this

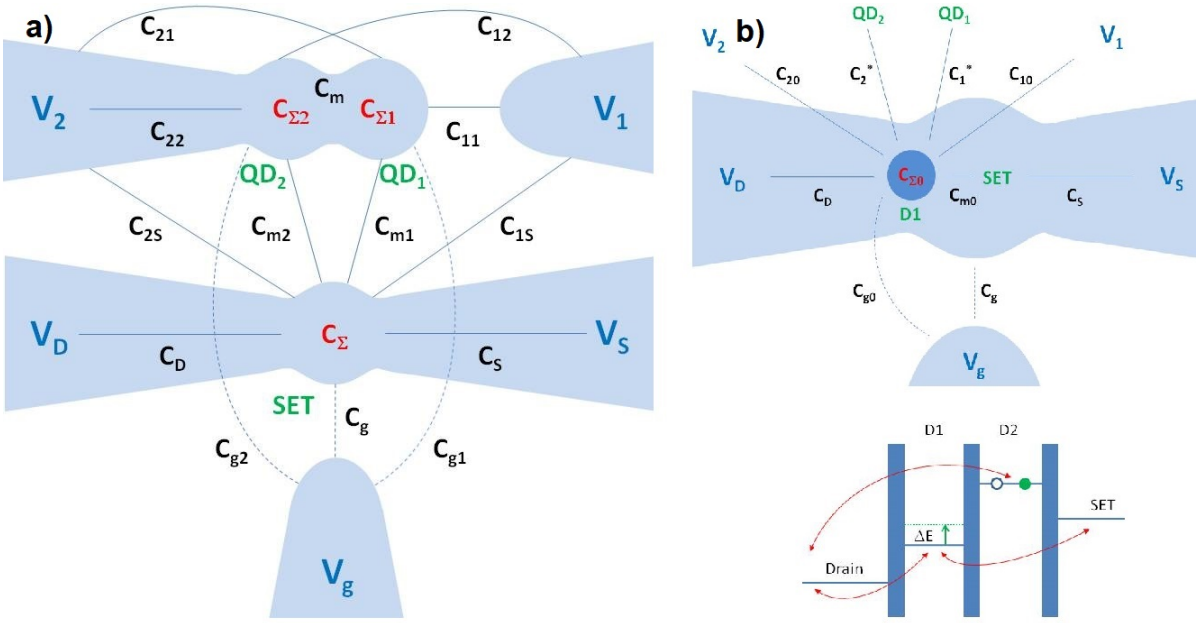

Fig. 5.6 **a)** Capacitance model for a semi-connected double quantum dot and its detector, **b)** Schematics of the energy levels and tunneling processes involving the two donors. Tunneling from and to the drain contact and the SET are allowed via D1. However, due to the energy difference, such a process is not allowed in D2 that can only be charged and discharged via the drain.

voltage range owing to the influence of charges on the DQD, which introduce second-order effects, first-order being the direct capacitance from V_1, V_2 to the SET.

Overall, the interpretations and conclusions drawn in this work are supported by numerical simulations.

V Conclusions

We have shown that a single donor could be electrostatically isolated from a large ensemble by making use of the intrinsic glassy behaviour of doped semiconductor device and reshaping the quantum dot potential. The creation of an extra potential barrier allows for realising a donor-quantum dot hybrid system. By adjusting the gate and source-drain voltages, the electrostatic influence of the quantum dot can be canceled out and one can use the D_0 state to turn the device into a single atom detector and map the charge states of a nearby but capacitively coupled double quantum dot. Such a method is largely applicable to all doped semiconductors quantum dots and provides a reliable, low-cost and fast way of

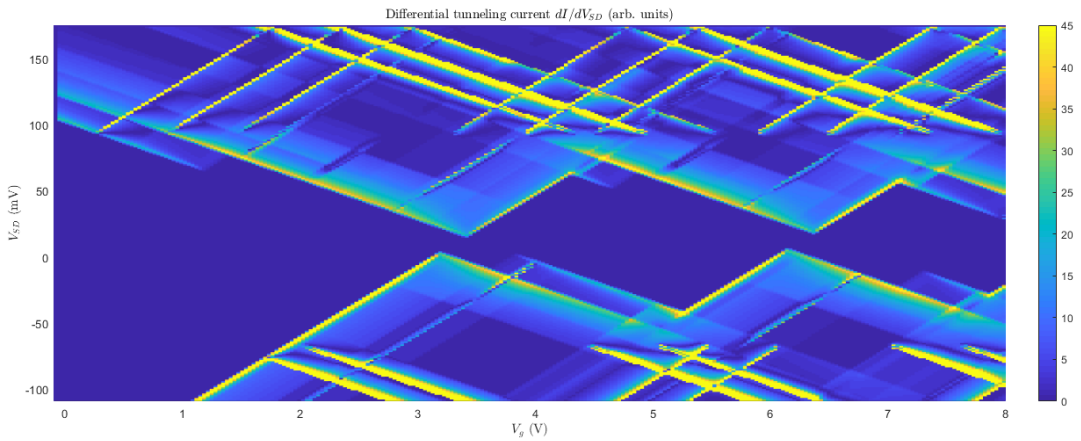

Fig. 5.7 Differential current through parallel donors D1 and D2, in series with another quantum dot. This simulation accurately reproduces subplot **a)** in 5.3, confirming the interpretation of the data that two donors with discrete energy levels participate in charge transport.

realising single donor structures that could be utilised for quantum computing applications and quantum states detection.

VI Supporting information

VI.I Experimental setup

Measurements have been carried in an Oxford Instrument Heliox^{TM} cryostat with a base temperature of 300 mK and a superconducting magnet of 7 T. The device temperature was controlled by an Oxford Instrument ITC 503 and measured by either a calibrated RuO_2 thermometer in the range 0.3 K to 1.8 K or a Cernox^{TM} thermometer from 1.8 K up to 290 K. For the DC measurements, voltages applied to the devices (source and gates) were provided by Hewlett Packard HP 3245A voltage sources with the current being read out by a HP 3458. The differential conductivity was measured by a Stanford Research SR 830 lock-in amplifier after adding a small AC signal of about 150μV to the source contact. In all cases, Stanford Research SR570 current amplifiers were used. Low-pass 10 MHz BNC filters were fitted on all lines, with a 50 Hz and 100 Hz rejection on AC lines.

VI.II Effective permittivity and self calculation of the dot layer thicknesses

In the spherical approximation, the dot diameter D is given by :

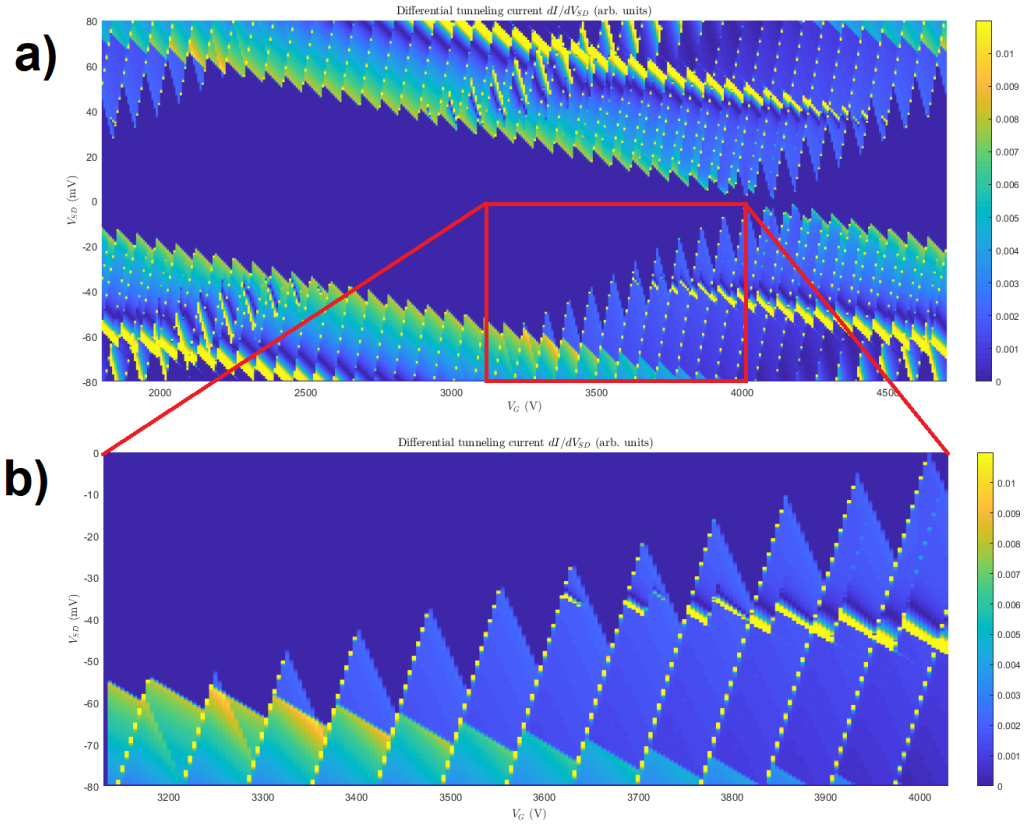

Fig. 5.8 Effects of filtering through a large quantum dot on an SET. The small dot has a total capacitance of $C_\Sigma = 1.8$ aF, and the larger one $C_\Sigma = 12.8$ aF. **a)** The large scale features match that of tunneling through the small dot only. **b)** The inset shows the detail of superimposed fine oscillations, due to filtering through the larger dot. The scale of these fine oscillations is much smaller than the large Coulomb diamonds visible in **a)**, because the large dot capacitance is much greater than the small dot capacitance. Some numerical artifacts arise in this figure, which should be ignored.

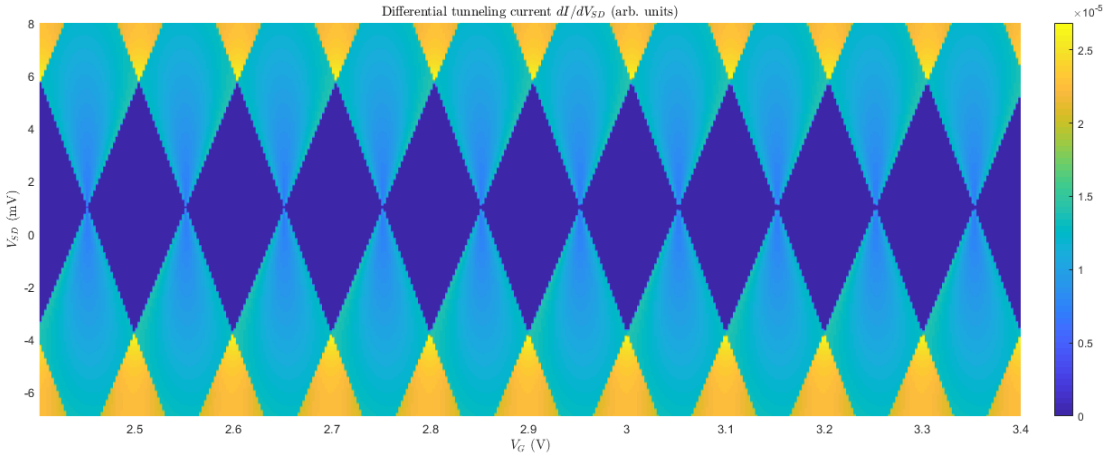

Fig. 5.9 Single quantum dot transport simulation. Standard Coulomb diamond pattern is observed as expected. The capacitance values used match with the ones measured experimentally, and the result matches well with differential current observed experimentally in Fig. 5.6 **c)**.

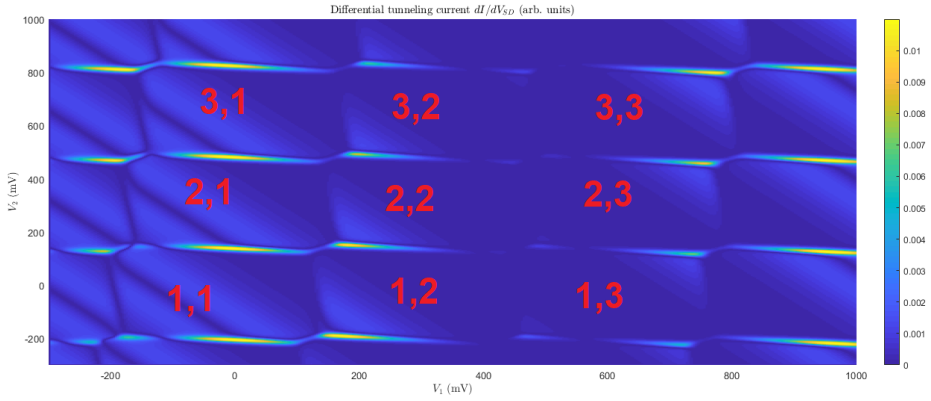

Fig. 5.10 Signatures of DQD charge detectable through SET differential current. The checkerboard pattern observable matches the charge stability diagram of the DQD exactly. The corresponding charge states of the DQD are labeled in red - charge in the right dot, charge in the left dot.

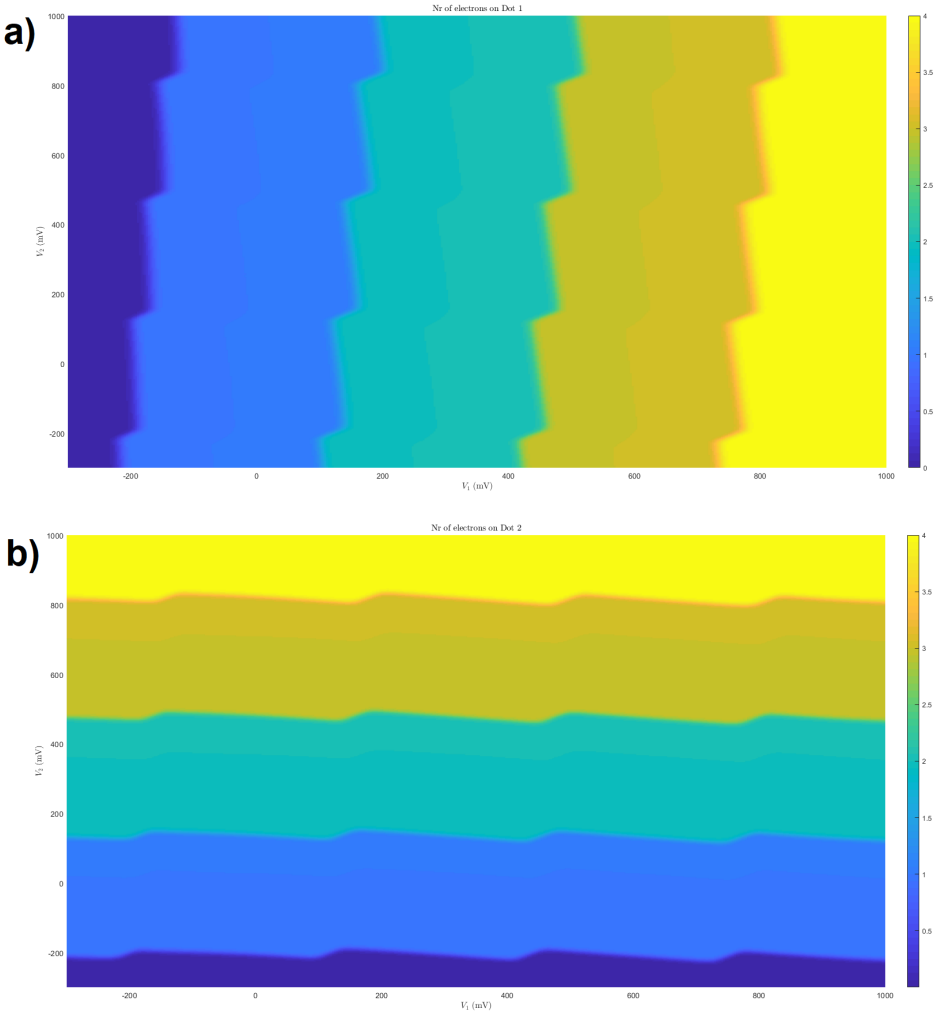

Fig. 5.11 Charge states of the DQD, **a)** for Dot one (left) , **b)** for Dot two (right), as a function of the DQD voltages V_1, V_2 as per the schematic in Fig. 5.6.

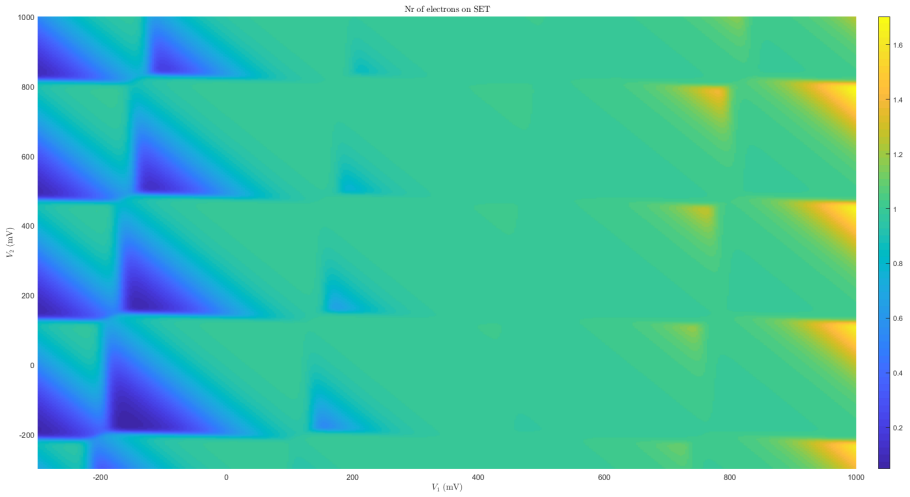

Fig. 5.12 Charge states of the SET, as a function of the DQD voltages V_1, V_2 as per the schematic in Fig. 5.6. The charge stays almost constant over the voltage range explored, unlike the charges on the DQD presented in Fig. 5.11. The variation in charge is caused by charge transitions in the DQD, which is capacitively coupled.

$$D = 2(R+L) = \frac{\mathrm{e}}{4\pi\varepsilon_{\mathrm{eff}}\varepsilon_0\alpha_{\mathrm{g}}\Delta V_{\mathrm{g}}}, \tag{5.13}$$

where R and L are the silicon and oxide thicknesses respectively, while the effective permittivity is given by :

$$\frac{1}{\varepsilon_{\mathrm{eff}}} = \frac{R}{R+L}\frac{1}{\varepsilon_{\mathrm{Si}}} + \frac{L}{R+L}\frac{1}{\varepsilon_{\mathrm{SiO_2}}}, \tag{5.14}$$

with $\varepsilon_{\mathrm{Si}}$ and $\varepsilon_{\mathrm{SiO_2}}$ the permittivity of silicon and silicon oxide respectively.

In these devices, the high doping concentration and dielectric screening due to the proximity of a SiO_2 layer modify significantly the effective permittivity $\varepsilon_{\mathrm{eff}}$ of the medium. In the high doping region, the permittivity of silicon $\varepsilon_{\mathrm{Si}}$ increases with concentration, and reaches the value $\varepsilon_{\mathrm{Si}} \sim 16$ at $N_{\mathrm{d}} \sim 3\times10^{19}\mathrm{cm}^{-3}$.

For a given Coulomb oscillation period and lever arm, we can obtain a set of solution for R and L. From the measurement of the total dot diameter by SEM, one can determine $R \sim 30$ nm and $L \sim 7$ nm, which leads to $\varepsilon_{\mathrm{eff}} \sim 9.9$.

Additional analysis has been carried out by considering the flat disk approximation, but obtained values for both R and L were incompatible with the observed final dimensions of the device. Finally, the ellipsoid approximation gave very similar values to the spherical one.

VI.III Estimation of the donor separation

When an electron tunnels via D1 while the second donor D_0 state is occupied, the resulting Coulomb interaction shifts the position of the energy levels by :

$$\Delta E_+ = \frac{e^2}{4\pi\varepsilon\varepsilon_0 d}\mathrm{e}^{-\lambda_{\mathrm{TF}}d}. \tag{5.15}$$

The exponential term accounts for the screening due to the other surrounding donors and electrons while d is the average donor separation, ε the effective permittivity of the medium calculated previously, and λ_{TF} the Thomas-Fermi screening length.

At low temperature and high doping concentration N_d, λ_{TF} is given by :

$$\lambda_{\mathrm{TF}} = 4\left(\frac{3N_{\mathrm{d}}}{\pi}\right)^{1/3}. \tag{5.16}$$

Direct readout of the shift value could be performed on the V_{SD} axis, and we find the $D_+ - D_0$ transition shifts by $\Delta D_+ \sim 19.0$ meV. A similar analysis can be performed on the V_{g} axis by utilising the lever arm $\alpha_{\mathrm{gD}} \sim 0.022$ found previously from the Coulomb diamonds. In this case we obtain $\Delta D_+ \sim 20.2$ meV.

Using an average donor concentration $N_{\mathrm{d}} \sim 3\times10^{19}\mathrm{cm}^{-3}$, we obtain $\lambda_{\mathrm{TF}} \sim 0.8$ nm^{-1} and $d \sim 1.4$ nm. This value is of the same order as the λ_{TF} and effective Bohr radius in silicon $a^*_{\mathrm{B}} \sim 2.3$ nm. This clearly indicates that the strong unscreened correlation does exist between the two donors but without forming a cluster as a helium molecule.

When the donor D1 is singly occupied, and so in the D_0 state, and a second electron tunnels via the D_-, energy levels are expected to shift by ΔE_0 when the second donor D2 is occupied. Shifts can also be related via a capacitive model via

$$\Delta E_0 \sim \frac{C_0}{C_-} \tag{5.17}$$

where C_0 and C_- are the total capacitance for the D_0 and D_- levels.

From the size of the Coulomb diamonds we deduce $\Delta E_0 \sim 16.5$ meV in perfect agreement with the direct measured value of 16.9 meV.

Chapter 6

Conclusion

The main goal of my work presented here was to show how numerical simulations are an invaluable tool for analysing complex quantum-mechanical systems. The research described, while it required theoretical analysis and deep knowledge of quantum mechanics, would not be feasible to complete without these simulations due to its' complex and unintuitive nature. The physicist can come up with clever and novel ways of solving a given problem, but the fine-tuning of parameters in realistic systems is often difficult, if not impossible to do without the aid of computers. While quantum systems, owing to their very nature, scale up the computational requirements for simulating them exponentially fast, I show here that with the latest hardware, exploiting GPU paralellisation and smart algorithms, it is still possible to accurately model the time dynamics of complicated many-particle systems described on a fine-grained lattice. The state-of-the-art reached here is two particles in two spatial dimensions, which is already able to cover many experimental devices. With further advancements in GPU technology, even three particles should be achievable soon. It needs to be stressed that these are not merely two-level systems, but spatially extended quantum particles possessing many degrees of freedom and experiencing physically realistic potentials.

To summarise, in Chapter 3 I have used my simulation software to design control pulses that are able to optimally control a single charge qubit, taking pulse rise/fall time into account. The qubit is fully modelled on a spatial lattice, therefore all non-adiabatic effects that could reduce fidelity, such as fast pulses driving Landau-Zener transitions to eigenstates outside the computational basis, are included. The adjustment parameters that allow these pulses to rotate the qubit with significant rise/fall time are found using gradient ascent methods. This requires running many iterations of the numerical simulation, for which my software is critical. This iterative search allows one to find pulse parameters that enable reaching fidelities arbitrarily close to 1. I therefore show that this novel control scheme is feasible,

and is a significant upgrade in operation time and fidelity compared to simply assuming the experimentalists can use square pulses. I also show how in certain cases, it can significantly decrease the effects of charge noise a qubit experiences. This result can be easily generalised to many other types of qubits, as long as they share a similar form of the Hamiltonian.

In Chapter 4, I have shown how to design a SAW-based two-qubit gate to perform the Root-of-SWAP operation. In a model device, two spin-qubit electrons travel along electrostatically defined channels carried by surface-acoustic waves. The channels, originally completely separated, meet for some length, separated only by a modest tunnel barrier. This in principle allows them to swap via the exchange interaction. However, given effects such as Coulomb repulsion, the exact shape of the potential, electron velocity, and the steepness of the transition from separated to interacting channels, it is very difficult to tune the device layout to achieve a good operation fidelity. I show that previously considered collision method of performing this task, described theoretically in one dimension, breaks down in two dimensions. My GPU-accelerated software is essential to run the many simulations necessary to tune the device potential so that the operation is achieved with very high fidelity. Moreover, I show this is possible with parameters currently realistic in experiments, therefore demonstrating the SAW-based device as potentially competitive quantum computing technology.

The computational size of this problem, owing to the presence of two quantum particles in two dimensions, stresses the capabilities of my software to the limit, and shows its merits. Despite this, the software is able to find the initial state and time-evolve it in a reasonable amount of time. In particular, the momentum-space eigensolver is critical for finding the initial state (constructed from the ground and first-excited states). Fourier-transforming into momentum space allows me to drastically reduce the size of the problem, and find the initial state efficiently. It is then time-evolved by the staggered-leapfrog algorithm in a matter of hours. Here again, I find that it is beneficial to work in momentum space to reduce the problem size. However, for time evolution, much of the benefits of reduced size are offset by the overhead of moving between momentum- and real-space. The end results is still more accurate using momentum space for a comparable runtime.

In Chapter 5, I have shown that a single donor could be electrostatically isolated from a large ensemble. My simulations confirm that the presence of such donors explains the experimentally observed differential current characteristics of the device. I show how a parallel arrangement of donors that must be tunneled through explains the shifts observed in the Coulomb diamond pattern recorded in the experiment. Secondly, I explain how such a device can be used to detect charge states of a nearby capacitatively coupled double quantum dot, while keeping the charge on the detector itself constant by using a voltage compensation

procedure. The device described here could serve as a reliable and affordable component for quantum computing applications and quantum state detection.

The above shows how simulations, in tandem with theory, can open new avenues of research, as well as help improve the design of real-life experimental devices. This work is only the beginning, as with new hardware advances and further development, there is a huge potential to use the techniques described here for further research. For example, the next step for improved qubit control would be the inclusion of realistic noise by using the density matrix formalism with Lindblad operators- the numerical software will be invaluable for such tasks. For the SAW-based Root-of-SWAP gate, the next step would also include realistic noise, for which numerical simulations are practically necessary. Then the simulation outcome would be directly comparable to experimental results, as I could model both a specific device and noise channels present in it. Other more ambitious goals and possibilities are described in the “Further Work” section in Chapter 7.

Chapter 7

Further work and preliminary results

The following further work that aims to build on and expand upon the research completed in this thesis is currently planned, or in progress already. I describe the planned development of our GPU-accelerated simulation software, as well as planned research projects where it would be applicable. I aim to complete the goals laid out here over the course of the next year.

I Further simulation software development

We continue to improve our GPU-accelerated simulation software, as well as qualitatively extend its capabilities. While most of the simulations performed for the work presented here were done in the C++ version of the software, we have recently rewritten it in Python, as it makes it easier to manage. Development of new features is faster in Python, and it allows us to avoid issues with explicit memory management and garbage collection. While Python is generally slower and less efficient than C++, we've concluded that because the most computationally intensive part of the program, the CUDA kernel, is still written in C++, the slowdown on the main program is a price worth paying. It indeed turns out that while the initial setup steps of the simulation take slightly longer now, this does not noticeably extend the total simulation time if it is complex enough. We continue to fine-tune the software, and have plans to add many features that will open up possibilities for new research.

I.I Spin solver

A full spin solver that not only accounts for the spin of each particle, but also implements general spin-spin and spin-field interaction via the Pauli or Dirac equations will be implemented. With the ability to simulate an arbitrary number of particles possessing spin (up to

computational restrictions) will expand our options dramatically. While quantum solvers which treat each spin as a two-level system are readily available (for example QuTiP [181]), they do not simulate the spatial dimension, and do not take advantage of GPU acceleration. Here, we want to distinguish our work by focusing specifically on projects where the spatial dynamics is critical. Some potential applications to research are described below in Sec. II. As most immediate use cases would be non-relativistic, for example spins in an ion trap or embedded in a semiconductor, implementing the Pauli equation is the obvious choice:

$$\left[\frac{1}{2m}\left(\vec{\sigma}\cdot\left(\vec{p}-q\vec{A}\right)\right)^2+q\phi\right]|\psi\rangle = i\hbar\frac{\partial}{\partial t}|\psi\rangle\,. \tag{7.1}$$

Here $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ is a vector of Pauli matrices, $\vec{p}$ is the momentum operator, $\vec{A}$ is the magnetic vector potential, q is the particle electric charge and ϕ is the scalar electric potential.

I.II Density solver with decoherence

To study open quantum systems that interact with their environment and undergo decoherence, the simulation of the full density matrix is needed, as described in Ch. 2 Sec. III. While extending from simulating a quantum state vector $|\psi\rangle$ to a full matrix ρ would require more memory and be more computationally expensive, potentially up to a factor N^2, where N is the size of $|\psi\rangle$, simulating at least two particles in a reasonable number of spatial points would still be possible with more advanced hardware, which we will have access to. The von Neumann equation will be simulated, allowing us to use various Lindblad operators to express the noise or decoherence a qubit experiences. This will make our simulations even more experimentally realistic, making our results directly comparable to experiments, given that noise is described correctly. We plan to use this advancement to simulate SAW-assisted single electron transport in collaboration with Christopher Bäuerle's group [217]. This would help them diagnose the source of their errors, identifying what to focus on to achieve a better device or experiment design. These experimental advances could elevate SAW-driven quantum computing to the next level, allowing for experimental simulations of two-qubit operations, such as the Root-of-SWAP described in Ch. 4, and open up new avenues of research and opportunities for collaboration.

We have already implemented a prototype of the above, which, while not optimal, is able to simulate a simple one-particle quantum system and show how it reacts to noise channels. Here, I present some results of proof-of-concept simulations of a single particle trapped in a

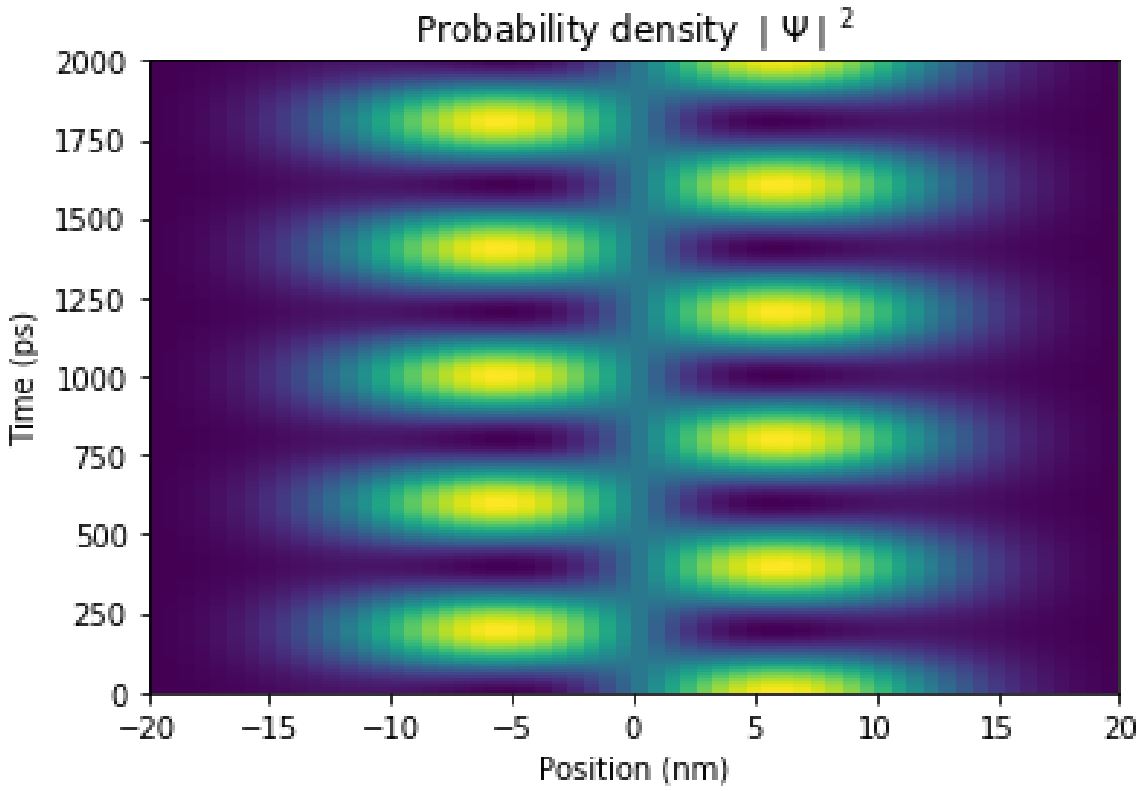

Fig. 7.1 Evolution of a superposition state in a DQD without decoherence. The state can be seen localised on the right side of the well at $t = 0$, and then proceeds to coherently oscillate as predicted by unitary evolution. As there are no non-unitary processes present, this oscillation repeats forever.

double quantum well.

Let us initialise the quantum state in a superposition of the ground and first excited eigenstates:

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle). \tag{7.2}$$

This way, we will be able to observe evolution in time, as well as quantum coherence. Evolving this initial state for a number of periods yields results presented in Fig. 7.1. It can be seen that the superposition state oscillates from left to right as expected from the relative phase accumulation between the different energy eigenstates. This result should be used as reference to compare with figures that include various noise channels.

Firstly, let us introduce a relaxation mechanism - the particle can emit a quantum of energy to the environment and relax to a lower energy state. The dissipation operator associated

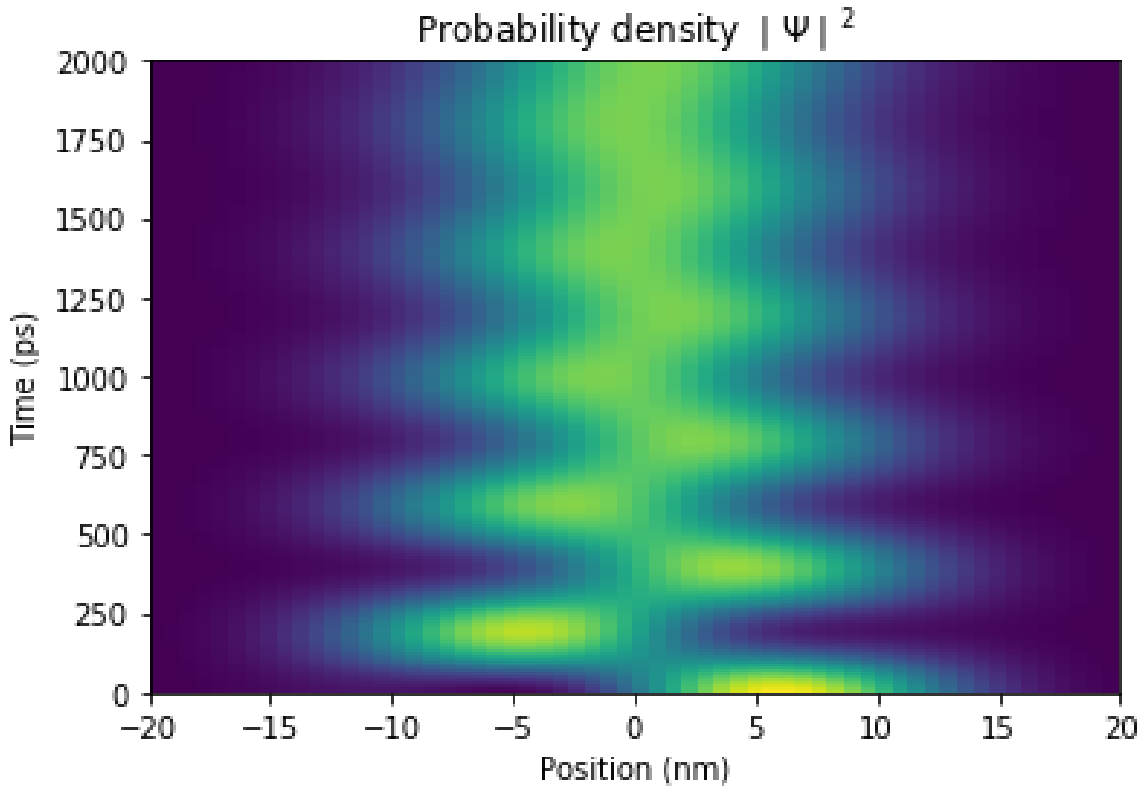

Fig. 7.2 Evolution of a superposition state with dissipation present. The initial state gradually dissipates energy, evolving towards the ground state in time.

with this process is simply proportional to the annihilation operator a, as per Eq. 7.3. The characteristic timescale for this process is called T_1.

$$L_d = \frac{1}{\sqrt{T_1}} a \tag{7.3}$$

However, L_d in Eq. 7.3 is expressed in an energy basis, and my simulations are conducted in position basis on a discrete lattice. Therefore, the operator has to be transformed into position basis - a matrix V of energy eigenvectors in position basis serves as a unitary transformation between the bases:

$$L_d^X = V L_d^E V^\dagger, \tag{7.4}$$

where L_d^X is the position representation of the operator and L_d^E is the energy representation. The effects of including the dissipation operator are presented in Fig. 7.2.

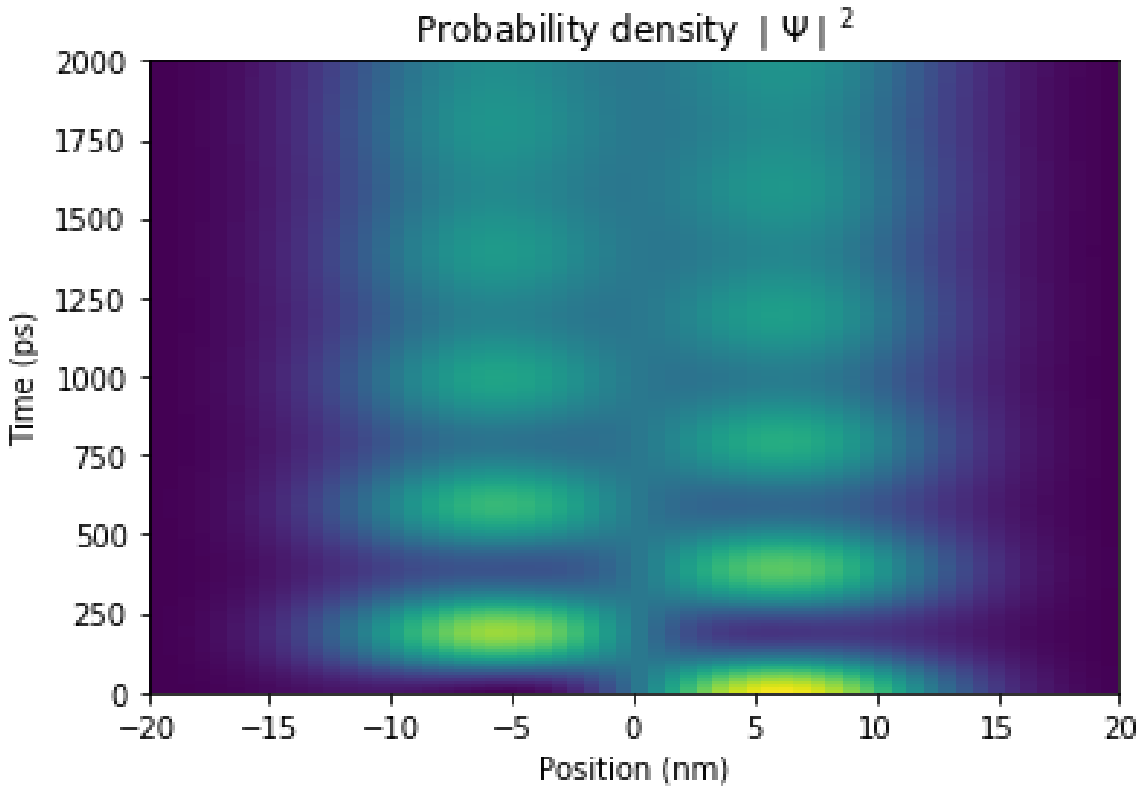

Fig. 7.3 Evolution of a superposition state with dephasing present. Unlike in the dissipation case, no energy is lost, but owing to loss of coherence between the two energy eigenstates, the state is gradually washed out, and the coherent oscillation is lost. While initially we are able to predict whether the particle is on the left or right side, this information is gradually lost, and the position measurement results become completely random.

Secondly, I will show the effects of a dephasing operator, which represents the mechanism of preserving the energy of the system, but losing phase coherence between different energy eigenstates. This operator, L_{ph}, is proportional to the Hamiltonian:

$$L_{ph} = \frac{1}{\sqrt{T_2 \hbar \omega}} H, \tag{7.5}$$

where T_2 is the characteristic dephasing time, and ω is the harmonic oscillator frequency of our potential. The results are presented in Fig. 7.3

Finally, I investigate the effects of a thermal excitation operator L_{th}, which represents a mechanism of the particle randomly absorbing an energy quantum from the environment and being promoted to a higher energy state. This operator is proportional to the creation operator $a^\dagger$:

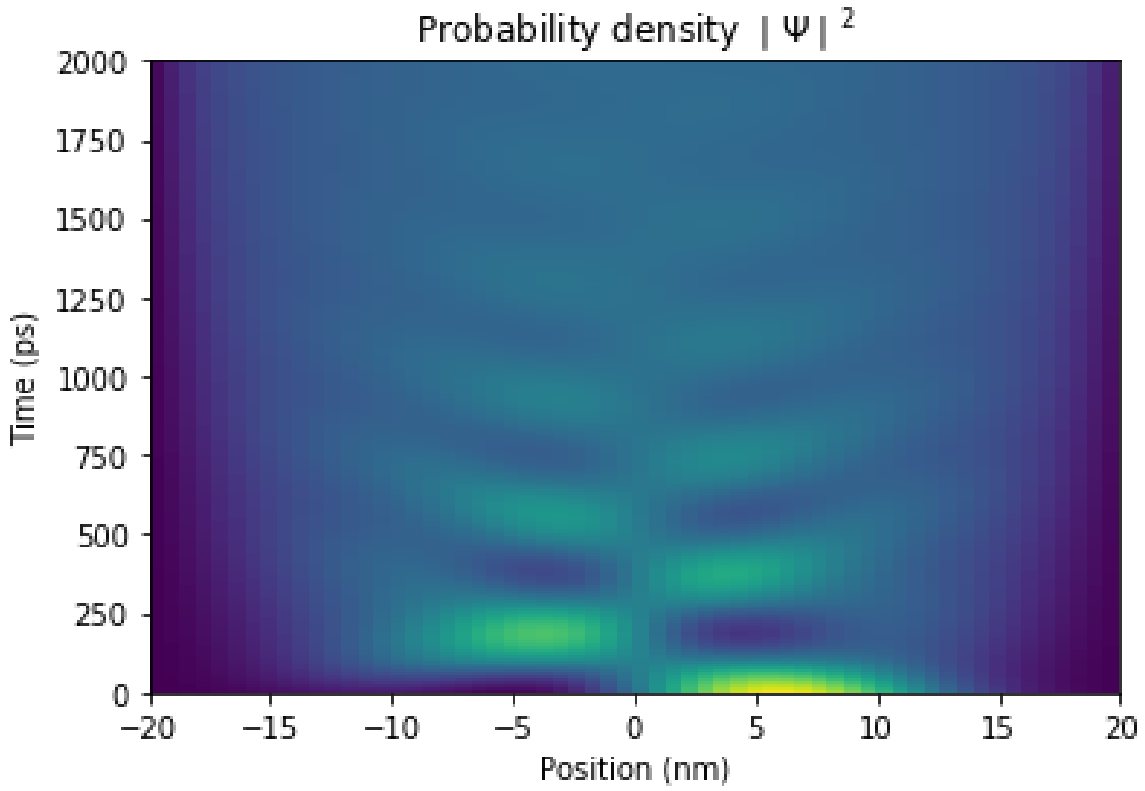

Fig. 7.4 Evolution of a superposition state with dissipation and excitation present. The competing processes of emission and absorption move the state towards a thermal equilibrium.

$$L_{th} = \frac{1}{\sqrt{T_1}} a^{\dagger}. \tag{7.6}$$

Generally, the form of this operator would be more complicated and depend on the temperature difference between the qubit and the environment, however this simple model is sufficient for the proof of principle. Together with the dissipation operator defined in Eq. 7.3, the quantum state will move towards an equilibrium thermal state, and any quantum coherence will be destroyed in the process. The example results are presented in Fig. 7.4.

This initial model, while simple, is still able to illustrate how noise channels lead to loss of quantum information - decoherence. While each of the noise mechanisms investigated here has its own signature, it can be seen that they all cause unwanted, irreversible disturbance to the quantum state. Noise like this is a major source of errors in quantum computing, and learning how to reduce or negate it is key to improving quantum computing performance in the future.

Further development of this feature of the software should be aimed at simulating noise in a two-particle system.GPU-acceleration can make these calculations feasible. A major hurdle making this calculation difficult is associated with the need to constantly recalculate the annihilation and creation operators $a, a^\dagger$ that define the dissipation and excitation noise channels as modelled in Eqs. 7.3, 7.6. As per Eq. 7.4, the position representation of these operators requires the knowledge of energy eigenvectors, which will constantly change in a time-dependent potential, which is the case I am ultimately most interested in. Fully eigensolving the Hamiltonian at every time step is not feasible here. A potential solution that will be explored is iteration - as the potential is expected to not change drastically between time steps, an algorithm that approximates $a, a^\dagger$ based on previous values and changes in the Hamiltonian seems to be promising.

II Planned research

II.I Quantum Newton's cradle

Being able to simulate more than two particles with spin allows for interesting investigation of how the Pauli exclusion principle manifests itself as a repulsive force. This work would look at a collision of three or four spin-$\frac{1}{2}$ particles arranged in a line, where the left- or rightmost one has some initial velocity. Within one spatial dimension, this setup is similar to the classical Newton's cradle toy, where a ball bearing collides and transfers momentum through a chain of identical ball bearings. The dynamics of the quantum version will be compared with the classical one. The dependence of the outcome on the initial spin state will be investigated - it is assumed that the behaviour would be similar to the Newton's cradle if all the particles are initialised in the same spin eigenstate owing to Pauli exclusion principle acting like effective contact repulsion. However, when particles are not all in the same spin state, and are perhaps even in superpositions of spin eigenstates, rich and interesting dynamics not illustrated before are expected to emerge. The ability to simulate the spatial dimension of this problem in at least dozens of points is predicted to yield interesting behaviour as the particles are brought close to "collide". At the very least, this research would be an enlightening illustration on how the exclusion principle influences dynamics when spatial position is effectively continuous.

II.II Simulating qubits controlled by geometric pulse design

The theory and techniques for controlling a qubit with a pulse sequence that accounts explicitly for rise time described in Chapter 3 are compared with somewhat similar work

done by Edwin Barnes' (not to be confused with Crispin Barnes) group [36]. Their control pulses are designed based on geometric curves in the phase space of the pulse. They claim to be able to eliminate noise to second order. While this method doesn't take the equipment rise time into account explicitly, it claims to suffer less from it as it's not attempting to generate square pulses. We think that there are advantages and disadvantages to both methods, and instead of trying to directly compare them, we want to use this opportunity to collaborate. We plan to use our software to simulate the geometrically designed pulses [62, 245, 20] performing single- and two-qubit operations with realistic noise implemented, and ascertain their advantage over simpler control methods. As we will be able to include both the spatial dynamics of the qubits, as well as general decoherence mechanisms using the density matrix formalism, we are in a position to show how advantageous this pulse design methodology really is.

II.III Quantum thermodynamics

Quantum thermodynamics is a rapidly developing area of research that's studying the relationship between classical thermodynamics with quantum properties such as superposition, entanglement and non-commutation. It deals with issues such as heat exchange, interactions with heat baths, equilibrium and non-equilibrium in a quantum setting. The study of open quantum systems, such as with qubit decoherence, falls in the scope of quantum thermodynamics. This project looks at non-commutation of conserved quantities, and how is impacts thermalisation specifically.

In quantum thermodynamics, a small system can exchange some general conserved quantities, such as particles, heat, spin, etc, with a large heat bath, sometimes called the environment in quantum terms. The system of interest is predicted to thermalise (reach equilibrium with the environment) to a canonical (C) or grand canonical (GC) ensemble, depending on the conserved quantities [136]:

$$\rho_C = e^{-\beta H^S}/Z_C^S, \tag{7.7}$$

where C stands for canonical, $\beta = 1/T$ is the inverse temperature of the bath, H^S is the Hamiltonian of the system of interest and Z_C^S is the system's partition function that normalizes the quantum state ($Z_C^S = Tr(e^{-\beta H^S})$). For simplicity, we set Boltzmann's constant to one. If the system exchanges not only heat, but also particles with the bath, it will be best described by the grand canonical ensemble, which is achieved simply by replacing H^S with $H^S - \mu N$ to account for the energy of exchanging the particles, where μ is a chemical potential and N the number of particles.

The operators that represent these conserved quantities (e.g. N) are generally assumed to commute in the derivations of these states, which is not the case when they are quantum. For example, if the conserved "charges" in question are components of spin σ_α , none of them commute with each other. A so called Non-Abelian thermal state (NATS) [241] was proposed to take the above into account. It is defined as such:

$$\rho_{NATS} = e^{-\beta(H^S - \Sigma_\alpha \mu_\alpha Q^S_\alpha)}/Z^S_{NATS}, \tag{7.8}$$

where α's index the generalised charges Q_α and their chemical potentials μ_α. The above would reduce to the GC ensemble if we set $Q_x, Q_y = 0$ and $Q_z \neq 0$.

I am able to use our GPU-accelerated software to efficiently simulate many-particle quantum systems expected to be described by the NATS. I have already performed some initial work that is using simulations to illustrate how the NATS is distinguished from other thermodynamical predictions. In this work, a spin chain with a nearest-neighbour Heisenberg interaction is used. This is because this system is well-studied and obviously possesses non-commuting charges (σ_α's). We divide the spin chain into a smaller "system" with size n ($n = 2$ throughout this work), and a larger "bath" with N spins. The total Hamiltonian of the system then reads:

$$H^{tot} = J\left(\sum_{j=1}^{2N-1} \vec{\sigma}^j \cdot \vec{\sigma}^{j+1} + \sum_{j=1}^{2N-2} \vec{\sigma}^j \cdot \vec{\sigma}^{j+2}\right), \tag{7.9}$$

where J is the coupling strength, $\vec{\sigma}^j$ is the Pauli vector for site j, and I have set $n = 2$. Note that closed boundary conditions are used - this is mainly due to collaboration with experimentalists that use trapped ions [79, 116], so that these results can be directly comparable.

This Hamiltonian is also known to be non-integrable, which is thought to aid thermalisation - it will have an easier time evolving into the thermal state $\rho_{C/NATS}$ from an initial state, given certain requirements are satisfied [241]. Namely, the system should be initialised in an approximate microcanonical subspace (AMC) [242].

As the total spin charges σ_α do not commute with the Hamiltonian, the system cannot be in an eigenstate of energy E and spin components σ_α simultaneously. However, a subspace can exist where the expectation values $\langle H\rangle = E, \langle\sigma_\alpha\rangle = S_\alpha$ are at least approximately well-defined, that is, their standard deviation grows at most as $\sqrt{N}$ with the total system size N. That means that measuring a total spin σ_α will have a large probability of yielding an outcome close to S_α Secondly, the initial state should not be an energy eigenstate (for example, all spins aligned), as it would not evolve in time. In particular, product states of

single particle spin eigenstates in different directions are in such an AMC.

The goal of this research is to find initial states that best show the difference between NATS and GC predictions. The expectation is that in the presence of many non-commuting charges, the NATS prediction would be better than GC.

To find the required initial state, it is necessary to be able to calculate β and μ_α as in Eq. 7.8. These are properties of a given state under the Hamiltonian, and so each state will have different values. While it is possible to calculate them analytically for low μ_α and β, these results are not the directly useful, since the regime on the opposite end is the most interesting - high μ_α and β. This is because low μ_α means that the non-commutative aspect of the charges is weak, whereas we expect the NATS to be a better description when non-commutation is significant. Likewise, low β means that the temperature is high, which will quickly move the initial state towards the maximally mixed state, and mask the differences between the NATS and GC predictions. However, the analytical results hint that states with a low $\langle E \rangle$ will have a high β - intuitively we might associate low energy with low temperature. Similarly, states with high total spin expected value $\langle \sigma_\alpha \rangle$ should have a high S_α (in the same direction) - a large expectation value of spin means that the "cost" of exchanging it is also high.

A general way of calculating β and μ_α numerically for any initial state is found based on [6, 3]. We need to minimise the following function F by varying β and μ_α's :

$$F = logTr\left[exp\left(-\beta\left(H^{\text{tot}} - E\mathbb{I} + \sum_\alpha(\sigma_\alpha - S_\alpha\mathbb{I})\right)\right)\right]. \tag{7.10}$$

It is worth noting that β and μ_α are in principle different quantites for the NATS, GC and C predictions. I will therefore subscript them, e.g. β_{NATS}, for clarity moving forwards. As the only state-related inputs to the above Eq. 7.10 are the energy and spin expectation values, it is evident that these are the important properties of the initial state to focus on.
Using this method, I've numerically investigated various initial states with 4, 6, 8, 10 and 12 qubits in total. To summarise, I find that low energy states indeed result in high β, while high E states can easily reach negative β, which signifies negative temperatures - states that are "hotter" than infinitely high positive temperature ones. This is undesirable as mentioned before - we need to look at lower energy states (however, energy eigenstates are not suitable candidates, as they will not thermalise). Meanwhile, μ_α's are directly proportional to their respective S_α's.

I then attempt to construct a good initial state based on these observations :

- As I want to maximise μ_α's, high S_α's are desirable.
- As I care about the non-commutation between all three charges, I want the μ_α's to be of similar magnitude to preserve symmetry.
- As I want the inverse temperature β to be high, I want the expected energy E of my state to be as low as possible.

I therefore build the product state from single qubit spin eigenstates $|\alpha\pm\rangle$, meaning a spin eigenstate in $\alpha = x, y, z$ direction with value ± 1. I want all the spins for a given direction α to point in the same direction to maximise S_α, arbitrarily picking the positive eigenvalue ones (the results would be the same if we flip it to the negative one). To keep the μ_α's to be of similar magnitude, I want the same number of every direction spin state, if possible for a given number of qubits (this is only realisable exactly for N divisible by 3). Finally, I want to pick a permutation of the above that minimises the total expected energy. By numerically checking different permutation, the following state emerges:

$$|\psi\rangle = |x+\rangle\,|y+\rangle\,|z+\rangle\,...\,|x+\rangle\,|y+\rangle\,|z+\rangle\,. \tag{7.11}$$

We can check that it satisfies the AMC conditions by checking the scaling of spin standard deviation. The results are presented in Fig. 7.5. It can be seen that the standard deviations of all charges increase slower than $\sqrt{N}$ with the total system size N, satisfying the condition for being in an AMC.

I then compare the NATS, GC and C predictions for the thermal state using the following strategy. We evolve the initial state unitarily for a time sufficient for thermalisation, which scales with total system dimension 2^N. We then divide the total system into a two-particle "system of interest" and a "bath". We calculate the relative entropy between this state and the theoretical predictions. The quantum relative entropy between two quantum states ρ_1 and ρ_2 is defined as follows:

$$D(\rho_1 \,||\, \rho_2) = Tr\rho_1(log\rho_1 - log\rho_2). \tag{7.12}$$

This is averaged over choosing all possible neighboring pairs as the "system of interest', as well as averaged over a window of time $\pm 5\%$ around the value decided on above. The results are presented in Fig. 7.6.

Indeed, it can be seen that the NATS prediction is the best for all numbers of qubits, with GC being second-best, and C the worst. This holds true for all numbers of qubits explored

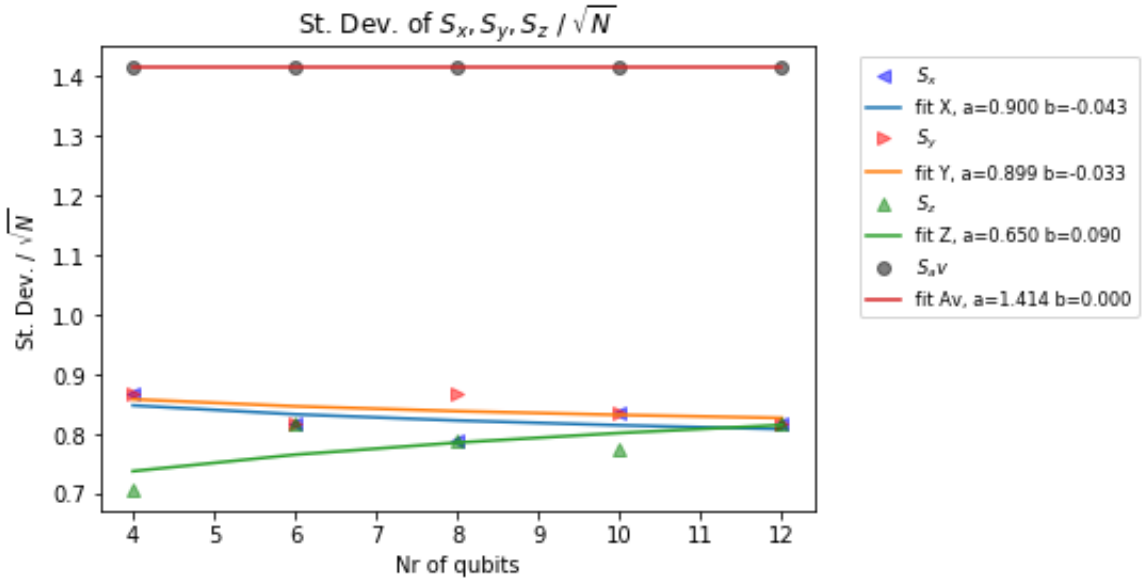

Fig. 7.5 Spin charges standard deviation scaling with total system size N. The total spin, as well as each direction separately, is presented. Data is fitted with an equation of the form $y = ax^b$, with least-squares fitted parameters provided. The total spin scales exactly as $\sqrt{N}$, as expected from a product state of this form. Each spin charge also scales at most as fast as $\sqrt{N}$, satisfying the AMC condition. While the standard deviation of S_z seems to scale faster on this plot, we conclude that it is an artefact of small number of qubits explored here, as the chosen state is symmetric w.r.t the $x/y/z$ axes, given that the state is not truncated, i.e. N is divisible by 3.

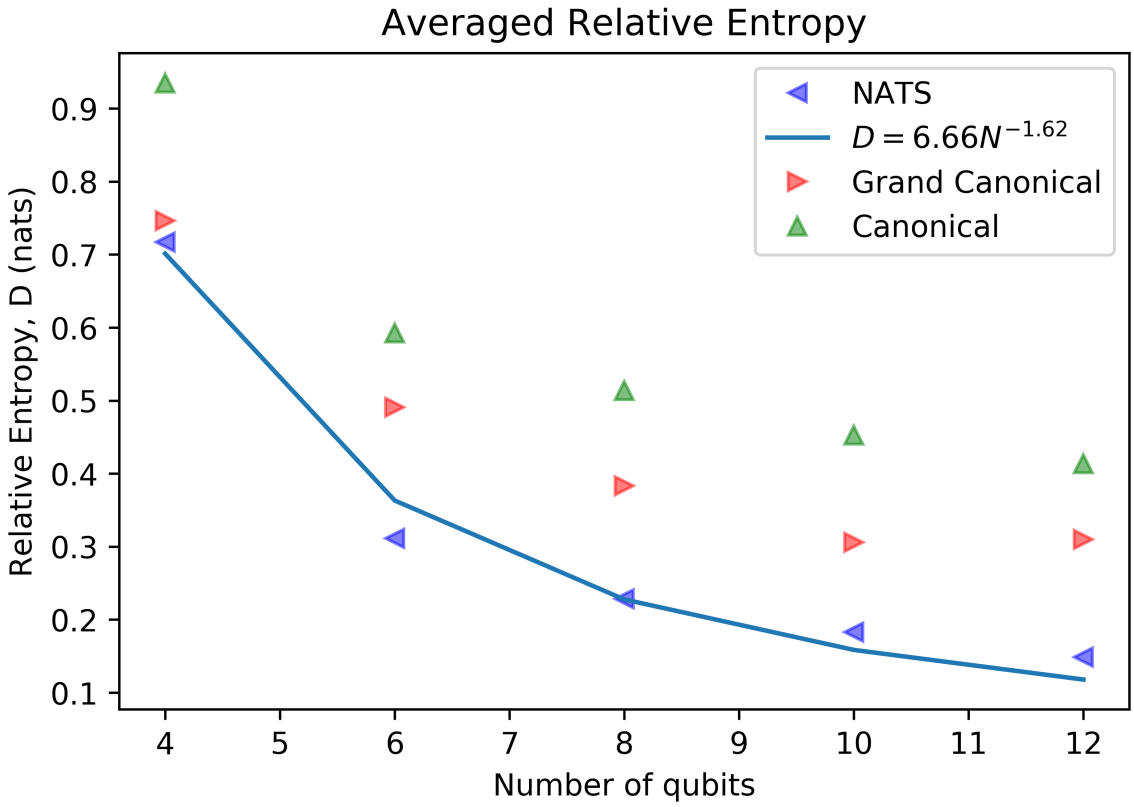

Fig. 7.6 Relative entropy comparison of the three thermodynamical predictions. For all data points, the NATS prediction is the best one, followed by GC and C. Within the number of qubits explored here, the gap between the predictions does not seem to be closing, although more data would be needed to draw definitive conclusions.

here, and within the regime explored, the gap between the predictions does not seem to be closing. This shows that the state found in Eq. 7.11 is a good initial state to show how the NATS is the right choice when multiple non-commuting charges are present. Within the regime explored, this repeating 3-particle product state remains the best one for any number of qubits, even when the total system size is not a multiple of three (in which case the state is truncated).

I have received measurement results from a 12-qubit trapped ion chain experiment, that implement an effective Heisenberg Hamiltonian and measure the spins repeatedly to generate data as per Fig. 7.6. The data seems to provide experimental confirmation of the NATS theory laid out here. It should be noted that the experiment uses a different interaction model - all-to-all coupling between qubits is present, with strength proportional to inverse distance between them. Additionally, only the x-dimension interaction is present. However, the latter can be overcome by using Trotterisation steps with basis rotations in-between, so that all directions are coupled equally on average. With small enough time step for this Trotterisation, my simulations shows that the experimental setup is qualitatively equivalent to the next-nearest-neighbour Heisenberg Hamiltonian described here. The same initial state found in this work is also the best to use in this experiment. The comparison of relative entropy distance between the exact theory and experimental measurement, NATS, GC and C predictions is presented in Table 7.1. Here, the relative entropy was calculated for all 11 neighbour pairs that each serve as the (S)ystem to the remaining 10 ions forming a (B)ath, and then averaged over all pairs (there were no significant outliers among the pairs). The data shows that the NATS is the best prediction out of the three by far, reaching a relative entropy w.r.t. the exact theory of 0.151 nats. The experiment's relative entropy w.r.t. the exact theory, and the Trotterised version (as the experiment was Trotterised) show a value that is very close to the NATS one. The small increase is mostly attributed to noise and measurement errors. The experimental initial state was also immediately measured and compared to theory for benchmarking, yielding a low entropy difference of 0.097 nats. The experiment is also closest to the NATS compared to GC and C predictions. Overall, I conclude that the experiment was performed well, with only small errors, and that the NATS gives a prediction that is close to reality, and gives a very notable improvement over the GC and C predictions.

Further work will be done on the NATS, possibly looking for different classes of good initial states. I am also planning to explore the regime of couplings other than XXX (where spin-spin interactions have equal strength on all axes). Changing these coupling strengths to be unequal will break the symmetry of the Hamiltonian under U(2) rotation, possibly

Table 7.1 Relative entropy D comparison

Comparison	**Average Relative entropy D (nats)**
Exact theory/NATS	0.151
Exact theory/GC	0.337
Exact theory/C	0.451
Experiment final state/exact theory	0.184
Experiment final state/Trotterized theory	0.148
Experiment final state/NATS	0.251
Experiment final state/GC	0.385
Experiment final state/C	0.403

furthering the difference between the NATS and GC predictions. This is because under the XXX coupling, a quantum state could be rotated to have the total spin S pointing in the z-direction - doing this would make the NATS and GC equal, as only μ_z for the NATS would be non-zero. However, under XYZ coupling, this is no longer the case, and all three μ's are required to describe the quantum state. There is also evidence that couplings other than XXX would change non-commutation properties of the charges [241]. I will explore all these exciting avenues in further research.

Part II

Other work

The second auxiliary part of this thesis shows other work done during my PhD. This part is presented as separate as to not detract from the focus of the main work on GPU-accelerated simulations, however, it is still closely related to quantum information processing and realistic quantum devices. The work on quantum metrology in Ch. 8 below proves how the negativity of Kirkwood-Dirac distribution describing a quantum system is a sign of non-classical phenomena and can be used to improve experimental measurements. Our ability to deliver new quantum-mechanical improvements to technologies relies on a better understanding of the foundation of quantum theory: When is a phenomenon truly nonclassical? We take noncommutation as our notion of nonclassicality and we quantify this nonclassicality with negativity: Quantum states can be represented by quasiprobability distributions, extensions of classical probability distributions. Whereas probabilities are real and nonnegative, quasiprobabilities can assume negative and nonreal values. Quasiprobabilities' negativity stems from the impossibility of representing quantum states with joint probability distributions [149, 206, 71]. The distribution we use, an extension of the Kirkwood-Dirac distribution [123, 50, 243], signals nonclassical noncommutation through the presence of negative or nonreal quasiprobabilities.

Chapter 8

Postselected metrology using Quantum Fisher information

I Introduction

One field advanced by quantum mechanics is metrology, which concerns the statistical estimation of unknown physical parameters. Quantum metrology relies on quantum phenomena to improve estimations beyond classical bounds [88]. A famous example exploits entanglement [87, 132, 48]. Consider using N separable and distinguishable probe states to evaluate identical systems in parallel. The best estimator's error will scale as $N^{-1/2}$. If the probes are entangled, the error scaling improves to N^{-1} [151]. As Bell's theorem rules out classical (local realist) explanations of entanglement, the improvement is genuinely quantum.

A central quantity in parameter estimation is the Fisher information, $\mathcal{I}(\theta)$. The Fisher information quantifies the average information learned about an unknown parameter θ from an experiment [**?** 27, 42]. $\mathcal{I}(\theta)$ lower-bounds the variance of an unbiased estimator θ_e via the Cramér-Rao inequality: $\mathrm{Var}(\theta_e) \geq 1/\mathcal{I}(\theta)$ [43, 183]. A common metrological task concerns optimally estimating a parameter that characterizes a physical process. The experimental input and the final measurement are optimized to maximize the Fisher information and to minimize the estimator's error.

Classical parameter estimation can benefit from postselecting the output data before postprocessing. Postselection can raise the Fisher information per final measurement or postprocessing event (Fig. 8.1). Postselection can also raise the rate of information per final measurement in a quantum setting. But classical postselection is intuitive, whereas an intense discussion surrounds postselected quantum experiments [4, 139, 223, 5, 58, 227, 69, 179, 180, 10, 12, 194, 11, 39]. The ontological nature of postselected quantum states, and the extent to

which they exhibit nonclassical behavior, is subject to an ongoing debate. Particular interest has been aimed at pre- and postselected averages of observables. These *weak values* can lie outside an observable's eigenspectrum when measured via a weak coupling to a pointer particle [4, 61]. Such values offer metrological advantages in estimations of weak-coupling strengths [223, 59, 168, 179, 115, 101, 134, 239].

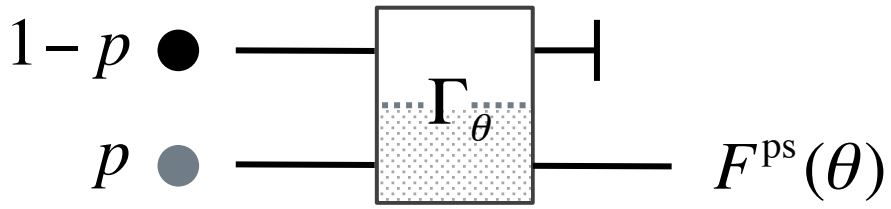

Fig. 8.1 **Classical experiment with postselection.** A nonoptimal input device initializes a particle in one of two states, with probabilities p and $1-p$, respectively. The particle undergoes a transformation Γ_θ set by an unknown parameter θ. Only the part of the transformation that acts on particles in the lower path depends on θ. If the final measurement is expensive, the particles in the upper path should be discarded: they possess no information about θ.

In this article, we go beyond this restrictive setting and ask, can postselection provide a nonclassical advantage in *general* quantum parameter-estimation experiments? We conclude that it can. We study metrology experiments for estimating an unknown transformation parameter whose final measurement or postprocessing incurs an experimental cost [144, 143]. Postselection allows the experiment to incur that cost only when the postselected measurement's result reveals that the final measurement's Fisher information will be sufficiently large. We express the Fisher information in terms of a quasiprobability distribution. Quantum negativity in this distribution enables postselection to increase the Fisher information above the values available from standard input-and-measurement-optimized experiments. Such an anomalous Fisher information can improve the rate of information gain to experimental cost, offering a genuine quantum advantage in metrology. We show that, within a commuting theory, a theory in which observables commute classically, postselection can improve information-cost rates no more than a strategy that uses an optimal input and final measurement can. We thus conclude that experiments that generate anomalous Fisher-information values require noncommutativity.

II Results

II.I Postselected quantum Fisher information

As aforementioned, postselection can raise the Fisher information per final measurement. Figure 8.1 outlines a classical experiment with such an information enhancement. Below, we show how postselection affects the Fisher information in a quantum setting.

Consider an experiment with outcomes i and associated probabilities $p_i(\theta)$, which depend on some unknown parameter θ. The Fisher information about θ is [42]

$$\mathscr{I}(\theta) = \sum_i p_i(\theta)[\partial_\theta \ln(p_i(\theta))]^2 = \sum_i \frac{1}{p_i(\theta)}[\partial_\theta p_i(\theta)]^2. \tag{8.1}$$

Repeating the experiment $N \gg 1$ times provides, on average, an amount $N\mathscr{I}(\theta)$ of information about θ. The estimator's variance is bounded by $\mathrm{Var}(\theta_e) \geq 1/[N\mathscr{I}(\theta)]$.

Below, we define and compare two types of metrological procedures. In both scenarios, we wish to estimate an unknown parameter θ that governs a physical transformation.

Optimized prepare-measure experiment: An input system undergoes the partially unknown transformation, after which the system is measured. Both the input system and the measurement are chosen to provide the largest possible Fisher information.

Postselected prepare-measure experiment: An input system undergoes, first, the partially unknown transformation and, second, a postselection measurement. Conditioned on the postselection's yielding the desired outcome, the system undergoes an information-optimized final measurement.

In quantum parameter estimation, a quantum state is measured to reveal information about an unknown parameter encoded in the state. We now compare, in this quantum setting, the Fisher-information values generated from the two metrological procedures described above. Consider a quantum experiment that outputs a state $\hat{\rho}_\theta = \hat{U}(\theta)\hat{\rho}_0\hat{U}^\dagger(\theta)$, where $\hat{\rho}_0$ is the input state and $\hat{U}(\theta)$ represents a unitary evolution set by θ. The quantum Fisher information is defined as the Fisher information maximized over all possible generalized measurements [27, 82, 174, 175, 88]:

$$\mathscr{I}_Q(\theta|\hat{\rho}_\theta) = \mathrm{Tr}\left[\hat{\rho}_\theta \hat{\Lambda}^2_{\hat{\rho}_\theta}\right]. \tag{8.2}$$

$\hat{\Lambda}_{\hat{\rho}_\theta}$ is the symmetric logarithmic derivative, implicitly defined by $\partial_\theta\hat{\rho}_\theta = \frac{1}{2}(\hat{\Lambda}_{\hat{\rho}_\theta}\hat{\rho}_\theta + \hat{\rho}_\theta\hat{\Lambda}_{\hat{\rho}_\theta})$ [**?**].

If $\hat{\rho}_\theta$ is pure, such that $\hat{\rho}_\theta = |\Psi_\theta\rangle\langle\Psi_\theta|$, the quantum Fisher information can be written as [168, 167]

$$\mathscr{I}_Q(\theta|\hat{\rho}_\theta) = 4\langle\dot{\Psi}_\theta|\dot{\Psi}_\theta\rangle - 4|\langle\dot{\Psi}_\theta|\Psi_\theta\rangle|^2, \tag{8.3}$$

where $|\dot{\Psi}_\theta\rangle \equiv \partial_\theta|\Psi_\theta\rangle$.

We assume that the evolution can be represented in accordance with Stone's theorem [212], by $\hat{U}(\theta) \equiv e^{-i\hat{A}\theta}$, where $\hat{A}$ is a Hermitian operator. We assume that $\hat{A}$ is not totally degenerate: If all the $\hat{A}$ eigenvalues were identical, $\hat{U}(\theta)$ would not imprint θ onto the state in a relative phase. For a pure state, the quantum Fisher information equals $\mathscr{I}_Q(\theta|\hat{\rho}_\theta) = 4\mathrm{Var}(\hat{A})_{\hat{\rho}_0}$ [88]. Maximizing Eq. 8.1 over all measurements gives $\mathscr{I}_Q(\theta|\hat{\rho}_\theta)$. Similarly, $\mathscr{I}_Q(\theta|\hat{\rho}_\theta)$ can be maximized over all input states. For a given unitary $\hat{U}(\theta) = e^{-i\hat{A}\theta}$, the maximum quantum Fisher information is

$$\max_{\hat{\rho}_0}\{\mathscr{I}_Q(\theta|\hat{\rho}_\theta)\} = 4\max_{\hat{\rho}_0}\{\mathrm{Var}(\hat{A})_{\hat{\rho}_0}\} = (\Delta a)^2, \tag{8.4}$$

where Δa is the difference between the maximum and minimum eigenvalues of $\hat{A}$ [88].[1] To summarize, in an optimized quantum prepare-measure experiment, the quantum Fisher information is $(\Delta a)^2$.

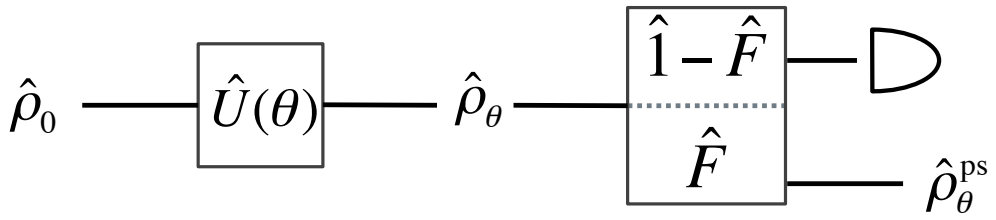

Fig. 8.2 **Preparation of postselected quantum state.** First, an input quantum state $\hat{\rho}_0$ undergoes a unitary transformation $\hat{U}(\theta) = e^{-i\theta\hat{A}}$: $\hat{\rho}_0 \to \hat{\rho}_\theta$. Second, the quantum state is subject to a projective postselective measurement $\{\hat{F}, \hat{1} - \hat{F}\}$. The postselection is such that if the outcome related to the operator $\hat{F}$ happens, then the quantum state is not destroyed. The experiment outputs renormalized states $\hat{\rho}_\theta^{\mathrm{ps}} = \hat{F}\hat{\rho}_\theta\hat{F}/\mathrm{Tr}(\hat{F}\hat{\rho}_\theta)$.

We now find an expression for the quantum Fisher information in a postselected prepare-measure experiment. A projective postselection occurs after $\hat{U}(\theta)$ but before the final measurement. Figure 8.2 shows such a quantum circuit. The renormalized quantum state that passes the postselection is $|\Psi_\theta^{\mathrm{ps}}\rangle \equiv |\psi_\theta^{\mathrm{ps}}\rangle / \sqrt{p_\theta^{\mathrm{ps}}}$, where we have defined an unnormalized

[1] The information-optimal input state is a pure state in an equal superposition of one eigenvector associated with the smallest eigenvalue and one associated with the largest.

state $|\psi_\theta^{\rm ps}\rangle \equiv \hat{F}\,|\Psi_\theta\rangle$ and the postselection probability $p_\theta^{\rm ps} \equiv {\rm Tr}(\hat{F}\hat{\rho}_\theta)$. As before, $\hat{\rho}_\theta = \hat{U}(\theta)\hat{\rho}_0\hat{U}^\dagger(\theta)$. $\hat{F} = \sum_{f\in\mathscr{F}^{\rm ps}} |f\rangle\langle f|$ is the postselecting projection operator, and $\mathscr{F}^{\rm ps}$ is a set of orthonormal basis states allowed by the postselection. Finally, the postselected state undergoes an information-optimal measurement.

When $|\Psi_\theta^{\rm ps}\rangle \equiv |\psi_\theta^{\rm ps}\rangle / \sqrt{p_\theta^{\rm ps}}$ is substituted into Eq. 8.3, the derivatives of $p_\theta^{\rm ps}$ cancel, such that

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) = 4\,\langle\dot{\psi}_\theta^{\rm ps}|\dot{\psi}_\theta^{\rm ps}\rangle\,\frac{1}{p_\theta^{\rm ps}} - 4|\,\langle\dot{\psi}_\theta^{\rm ps}|\psi_\theta^{\rm ps}\rangle\,|^2\frac{1}{(p_\theta^{\rm ps})^2}. \tag{8.5}$$

Equation 8.13 gives the quantum Fisher information available from a quantum state after its postselection. Unsurprisingly, $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$ can exceed $\mathscr{I}_Q(\theta|\hat{\rho}_\theta)$, since $p_\theta^{\rm ps} \leq 1$. Also classical systems can achieve such postselected information amplification (see Fig. 8.1). Unlike in the classical case, however, $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$ can also exceed the Fisher information of an optimized prepare-measure experiment, $(\Delta a)^2$. We show how below.

II.II Quasiprobability representation

In classical mechanics, our knowledge of a point particle can be described by a probability distribution for the particle's position, $\vec{x}$, and momentum, $\vec{k}$: $p(\vec{x},\vec{k})$. In quantum mechanics, position and momentum do not commute, and a state cannot generally be represented by a joint probability distribution over observables' eigenvalues. A quantum state can, however, be represented by a quasiprobability distribution. Many classes of quasiprobability distributions exist. The most famous is the Wigner function [236, 238, 38]. Such a distribution satisfies some, but not all, of Kolmogorov's axioms for probability distributions [163]: the entries sum to unity, and marginalizing over the eigenvalues of every observable except one yields a probability distribution over the remaining observable's eigenvalues. A quasiprobability distribution can, however, have negative or nonreal values. Such values signal nonclassical physics in, for example, quantum computing and quantum chaos [206, 68, 129, 110, 57, 46, 47, 240, 243, 99, 90].

A cousin of the Wigner function is the Kirkwood-Dirac quasiprobability distribution [123, 50, 243]. This distribution, which has been referred to by several names across the literature, resembles the Wigner function for continuous systems. Unlike the Wigner functions, however, the Kirkwood-Dirac distribution is well-defined for discrete systems, even qubits. The Kirkwood-Dirac distribution has been used in the study of weak-value amplification [210, 114, 107, 57, 176, 243], information scrambling [243, 99, 90, 160] and direct measurements of quantum wavefunctions [148, 147, 17, 221]. Moreover, negative and nonreal values of the distribution have been linked to nonclassical phenomena [57, 243, 99, 90]. We cast the quantum Fisher information for a postselected prepare-measure experiment

in terms of a *doubly extended*[2] Kirkwood-Dirac quasiprobability distribution [243]. We employ this distribution due to its usefulness as a mathematical tool: This distribution enables the proof that, in the presence of noncommuting observables, postselection can give a metrological protocol a nonclassical advantage.

Our distribution is defined in terms of eigenbases of $\hat{A}$ and $\hat{F}$. Other quasiprobability distributions are defined in terms of bases independent of the experiment. For example, the Wigner function is often defined in the bases of the quadrature of the electric field or the position and momentum bases. However, basis-independent distributions can be problematic in the hunt for nonclassicality [206, 46]. Careful application, here, of the extended Kirkwood-Dirac distribution ties its nonclassical values to the operational specifics of the experiment.

To begin, we define the quasiprobability distribution of an arbitrary quantum state $\hat{\rho}$:

$$q^{\hat{\rho}}_{a,a',f} \equiv \langle f|a\rangle \langle a| \hat{\rho} |a'\rangle \langle a'|f\rangle . \tag{8.6}$$

Here, $\{|f\rangle\}$, $\{|a\rangle\}$ and $\{|a'\rangle\}$ are bases for the Hilbert space on which $\hat{\rho}$ is defined. We can expand $\hat{\rho}$ [148, 147] as[3]

$$\hat{\rho} = \sum_{a,a',f} \frac{|a\rangle \langle f|}{\langle f|a\rangle} q^{\hat{\rho}}_{a,a',f}. \tag{8.7}$$

Let $\{|a\rangle\} = \{|a'\rangle\}$ denote an eigenbasis of $\hat{A}$, and let $\{|f\rangle\}$ denote an eigenbasis of $\hat{F}$. The reason for introducing a doubly extended distribution, instead of the standard Kirkwood-Dirac distribution $q^{\hat{\rho}}_{a,f} \equiv \langle f|a\rangle \langle a| \hat{\rho} |f\rangle$, is that $\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta)$ can be expressed most concisely, naturally, and physically meaningfully in terms of $q^{\hat{\rho}_\theta}_{a,a',f}$. Later, we shall see how the nonclassical entries in $q^{\hat{\rho}}_{a,a',f}$ and $q^{\hat{\rho}}_{a,f}$ are related. We now express the postselected quantum Fisher information (Eq. 8.13) in terms of the quasiprobability values $q^{\hat{\rho}_\theta}_{a,a',f}$ (Supp. Inf. 1).

$$\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta) = 4 \sum_{\substack{a,a',\\ f\in\mathscr{F}^{\text{ps}}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\text{ps}}_\theta} aa' - 4 \Big| \sum_{\substack{a,a',\\ f\in\mathscr{F}^{\text{ps}}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\text{ps}}_\theta} a \Big|^2, \tag{8.8}$$

[2]The modifier "doubly extended" comes from the experiment in which one would measure the distribution: One would prepare $\hat{\rho}$, sequentially measure two observables weakly, and measure one observable strongly. The number of weak measurements equals the degree of the extension [243].

[3]If any $\langle f|a\rangle = 0$, we perturb one of the bases infinitesimally while preserving its orthonormality.

where a and a' denote the eigenvalues associated with $|a\rangle$ and $|a'\rangle$, respectively.[4] Equation 8.14 contains a conditional quasiprobability distribution, $q^{\hat{\rho}_\theta}_{a,a',f}/p^{\text{ps}}_\theta$. If $\hat{A}$ commutes with $\hat{F}$, as they do classically, then they share an eigenbasis for which $q^{\hat{\rho}_\theta}_{a,a',f}/p^{\text{ps}}_\theta \in [0, 1]$, and the postselected quantum Fisher information is bounded as $\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta) \leq (\Delta a)^2$:

Theorem 1 *In a classically commuting theory, no postselected prepare-measure experiment can generate more Fisher information than the optimized prepare-measure experiment.*

Proof of Theorem 1.—We upper-bound the right-hand side of Eq. 8.14. First, if $\{|a\rangle\} = \{|a'\rangle\} = \{|f\rangle\}$ is an eigenbasis shared by $\hat{A}$ and $\hat{F}$, Eq. 8.6 simplifies to a probability distribution:

$$q^{\hat{\rho}_\theta}_{a,a',f} = \langle a|\,\hat{\rho}_\theta\,|a'\rangle\,[|f\rangle = |a\rangle][|a'\rangle = |f\rangle] \in [0, 1], \tag{8.9}$$

where $[X]$ is the Iverson bracket, which equals 1 if X is true and equals 0 otherwise. Second, summing $q^{\hat{\rho}_\theta}_{a,a',f}/p^{\text{ps}}_\theta$ over $f \in \mathscr{F}^{\text{ps}}$, we find

$$\sum_{f\in\mathscr{F}^{\text{ps}}} q^{\hat{\rho}_\theta}_{a,a',f}/p^{\text{ps}}_\theta = \langle a|\,\hat{\rho}_\theta\,|a'\rangle\,\langle a'|\,\hat{F}\,|a\rangle\,/p^{\text{ps}}_\theta. \tag{8.10}$$

By the eigenbasis shared by $\hat{A}$ and $\hat{F}$, the sum simplifies to $\langle a|\,\hat{\rho}_\theta\hat{F}\,|a'\rangle\,[|a'\rangle = |a\rangle]/p^{\text{ps}}_\theta$. We can thus rewrite Eq. 8.14:

$$\begin{aligned}\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta) =& 4\sum_{a,a'} \frac{\langle a|\,\hat{\rho}_\theta\hat{F}\,|a'\rangle\,[|a'\rangle = |a\rangle]}{p^{\text{ps}}_\theta} aa' \\ &- 4\Big|\sum_{a,a'} \frac{\langle a|\,\hat{\rho}_\theta\hat{F}\,|a'\rangle\,[|a'\rangle = |a\rangle]}{p^{\text{ps}}_\theta} a\Big|^2 \\ =& 4\sum_a q_a a^2 - 4\Big(\sum_a q_a a\Big)^2,\end{aligned} \tag{8.11}$$

where we have defined the probabilities $q_a \equiv \langle a|\,\hat{\rho}_\theta\hat{F}\,|a\rangle\,/p^{\text{ps}}_\theta = \sum_{f\in\mathscr{F}^{\text{ps}}} \langle a|\,\hat{\rho}_\theta\,|a\rangle\,[|f\rangle = |a\rangle]/p^{\text{ps}}_\theta$.

Apart from the multiplicative factor of 4, Eq. 8.11 is in the form of a variance with respect to the observable's eigenvalues a. Thus, Eq. 8.11 is maximized when $q_{a_{\min}} = q_{a_{\max}} = \frac{1}{2}$:

$$\max_{\{q_a\}}\{\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta)\} = (\Delta a)^2. \tag{8.12}$$

[4] We have suppressed degeneracy parameters γ in our notation for the states, e.g., $|a,\gamma\rangle \equiv |a\rangle$.

This Fisher-information bound must be independent of our choice of eigenbases of $\hat{A}$ and $\hat{F}$. In summary, if $\hat{A}$ commutes with $\hat{F}$, then all $q^{\hat{\rho}_\theta}_{a,a',f}/p^{\text{ps}}_\theta$ can be expressed as real and nonnegative, and $\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta) \leq (\Delta a)^2$. □

In contrast, if the quasiprobability distribution contains negative values, the postselected quantum Fisher information can violate the bound: $\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta) > (\Delta a)^2$. In Supp. Inf. 2, we prove a second theorem:[5]

Theorem 2 *An anomalous postselected Fisher information implies that the quantum Fisher information cannot decompose in terms of a nonnegative doubly extended Kirkwood-Dirac quasiprobability distribution.*

Proof: see Supplementary Note 2 for a proof.

This inability to decompose implies that $\hat{A}$ fails to commute with $\hat{F}$. However, pairwise noncommutation of $\hat{\rho}_\theta, \hat{A}$ and $\hat{F}$ is insufficient to enable anomalous values of $\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta)$. For example, noncommutation could lead to a nonreal Kirkwood-Dirac distribution without any negative real components. Such a distribution cannot improve $\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta)$ beyond classical values. Furthermore, the presence or absence of commutation is a binary measure. In contrast, how much postselection improves $\mathscr{I}_Q(\theta|\Psi^{\text{ps}}_\theta)$ depends on how much negativity $q^{\hat{\rho}_\theta}_{a,a',f}/p^{\text{ps}}_\theta$ has. We build on this observation, and propose two experiments that yield anomalous Fisher-information values, in Supp. Infs. 3 and 4.[6]

As promised, we now address the relation between nonclassical entries in $q^{\hat{\rho}}_{a,a',f}$ and nonclassical entries in $q^{\hat{\rho}}_{a,f}$. For pure states $\hat{\rho} = |\Psi\rangle\langle\Psi|$, the doubly extended quasiprobability distribution can be expressed time symmetrically in terms of the standard Kirkwood-Dirac distribution [123, 50, 240, 243, 99, 90]: $q^{\hat{\rho}}_{a,a',f} = \frac{1}{p_f} q^{\hat{\rho}}_{a,f} \left(q^{\hat{\rho}}_{a',f}\right)^*$, where $q^{\hat{\rho}}_{a,f} = \langle f|a\rangle\langle a|\hat{\rho}|f\rangle$ and $p_f \equiv |\langle f|\Psi\rangle|^2$.[7] Therefore, a negative $q^{\hat{\rho}}_{a,a',f}$ implies negative or nonreal values of $q^{\hat{\rho}}_{a,f}$. Similarly, a negative $q^{\hat{\rho}}_{a,a',f}$ implies a negative or nonreal weak value $\langle f|a\rangle\langle a|\Psi\rangle / \langle f|\Psi\rangle$ [4], which possesses interesting ontological features (see below). Thus, an anomalous Fisher information is closely related to a negative or nonreal weak value. Had we weakly measured the observable $|a\rangle\langle a|$ of $\hat{\rho}_\theta$ with a qubit or Gaussian pointer particle before the postselection, and had we used a fine-grained postselection $\{\hat{1} - \hat{F}, |f\rangle\langle f| : f \in \mathscr{F}^{\text{ps}}\}$, the weak measurement would have yielded a weak value outside the eigenspectrum of $|a\rangle\langle a|$. It has been shown that such an anomalous weak value proves that quantum mechanics, unlike classical mechanics, is contextual: quantum outcome probabilities can depend on more than a unique set of underlying physical states [205, 179, 134]. If $\hat{\rho}_\theta$ had undergone the

[5] The theorem's converse is not generally true.

[6] It remains an open question to investigate the relationship between Kirkwood-Dirac negativity in other metrology protocols with noncommuting operators, e.g., [216].

[7] See [5, 138] for discussions about time-symmetric interpretations of quantum mechanics.

aforementioned weak measurement, instead of the postselected prepare-measure experiment, the weak measurement's result would have signaled quantum contextuality. Consequently, a counterfactual connects an anomalous Fisher information and quantum contextuality. While counterfactuals create no problems in classical physics, they can lead to logical paradoxes in quantum mechanics [128, 100, 169, 205]. Hence our counterfactual's implication for the ontological relation between an anomalous Fisher information and contextuality offers an opportunity for future investigation.

II.III Improved metrology via postselection

In every real experiment, the preparation and final measurement have costs, which we denote $\mathscr{C}_P$ and $\mathscr{C}_M$, respectively. For example, a particle-number detector's *dead time*, the time needed to reset after a detection, associates a temporal cost with measurements [95]. Reference [144] concerns a two-level atom in a noisy environment. Liuzzo *et al.* detail the tradeoff between frequency estimation's time and energy costs. Standard quantum-metrology techniques, they show, do not necessarily improve metrology, if the experiment's energy is capped. Also, the cost of postprocessing can be incorporated into $\mathscr{C}_M$.[8] We define the information-cost rate as $R(\theta) = N\mathscr{I}(\theta)/(N\mathscr{C}_P + N\mathscr{C}_M) = \mathscr{I}(\theta)/(\mathscr{C}_P + \mathscr{C}_M)$. If our experiment conditions the execution of the final measurement on successful postselection of a fraction p_θ^{ps} of the states, we include a cost of postselection, $\mathscr{C}_{\text{ps}}$. We define the postselected experiment's information-cost rate as $R^{\text{ps}}(\theta) = Np_\theta^{\text{ps}}\mathscr{I}^{\text{ps}}(\theta)/(N\mathscr{C}_P + N\mathscr{C}_{\text{ps}} + Np_\theta^{\text{ps}}\mathscr{C}_M) = p_\theta^{\text{ps}}\mathscr{I}^{\text{ps}}(\theta)/(\mathscr{C}_P + \mathscr{C}_{\text{ps}} + p_\theta^{\text{ps}}\mathscr{C}_M)$, where $\mathscr{I}^{\text{ps}}(\theta)$ is the Fisher information conditioned on successful postselection. Generalizing the following arguments to preparation and measurement costs that differ between the postselected and nonpostselected experiments is straightforward.

In classical experiments, postselection can improve the information-cost rate. See Fig. 8.1 for an example. But can postselection improve the information-cost rate in a classical experiment with information-optimized inputs? Theorem 1 answered this question in the negative. $\mathscr{I}^{\text{ps}}(\theta) \leq \max\{\mathscr{I}(\theta)\}$ in every classical experiment. The maximization is over all physically accessible inputs and final measurements. A direct implication is that $R^{\text{ps}}(\theta) \leq \max\{R(\theta)\}$.

In quantum mechanics, $\mathscr{I}_Q(\theta|\Psi_\theta^{\text{ps}})$ can exceed $\max_{\hat{\rho}_0}\{\mathscr{I}_Q(\theta|\hat{\rho}_\theta)\} = (\Delta a)^2$. This result would be impossible classically. Anomalous Fisher-information values require quantum negativity in the doubly extended Kirkwood-Dirac distribution. Consequently, even compared to quantum experiments with optimized input states, postselection can raise information-cost

[8] In an experiment, these costs can be multivariate functions that reflect the resources and constraints. Such a function could combine a detector's dead time with the monetary cost of liquid helium and a graduate student's salary. However, presenting the costs in a general form benefits this platform-independent work.

rates beyond classically possible rates: $R^{\rm ps}(\theta) > \max\{R(\theta)\}$. This result generalizes the metrological advantages observed in the measurements of weak couplings, which also require noncommuting operators. References [109, 54, 207, 29, 208, 65, 108, 154, 150, 156] concern metrology that involves weak measurements of the following form. The primary system S and the pointer P begin in a pure product state $|\Psi_{\rm S}\rangle \otimes |\Psi_{\rm P}\rangle$; the coupling Hamiltonian is a product $\hat{H} = \hat{A}_{\rm S} \otimes \hat{A}_{\rm P}$; the unknown coupling strength θ is small; and just the system is postselected. Our results govern arbitrary input states, arbitrary Hamiltonians (that satisfy Stone's theorem), arbitrarily large coupling strengths θ, and arbitrary projective postselections. Our result shows that postselection can improve quantum parameter estimation in experiments where the final measurement's cost outweighs the combined costs of state preparation and postselection: $\mathscr{C}_M \gg \mathscr{C}_P + \mathscr{C}_{\rm ps}$. Earlier works identified that the Fisher information from nonrenormalized trials that succeed in the postselection cannot exceed the Fisher information averaged over all trials, including the trials in which the postselection fails [70, 40].[9] In accordance with practical metrology, not only the Fisher information, but also measurements' experimental costs, underlie our results.

So far, we have shown that $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$ can exceed $(\Delta a)^2$. But how large can $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$ grow? In Supp. Inf. 3, we show that, if the generator $\hat{A}$ has $M \geq 3$ not-all-identical eigenvalues, there is no upper bound on $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$. If $\mathscr{C}_P$ and $\mathscr{C}_{\rm ps}$ are negligible compared to $\mathscr{C}_M$, then there is no theoretical cap on how large $R^{\rm ps}(\theta)$ can grow. In general, when $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) \to \infty$, $p_\theta^{\rm ps} \times \mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) < (\Delta a)^2$, such that information is lost in the events discarded by postselection. But if $\hat{A}$ has doubly degenerate minimum and maximum eigenvalues, $p_\theta^{\rm ps} \times \mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$ can approach $(\Delta a)^2$ while $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$ approaches infinity (see Supp. Inf. 4). In such a scenario, postselection can improve information-cost rates, as long as $\mathscr{C}_{\rm ps} < (1 - p_\theta^{\rm ps})\mathscr{C}_M$—a significantly weaker requirement than $\mathscr{C}_M \gg \mathscr{C}_P + \mathscr{C}_{\rm ps}$.

III Discussion

From a practical perspective, our results highlight an important quantum asset for parameter-estimation experiments with expensive final measurements. In some scenarios, the postselection's costs exceed the final measurement's costs, as an unsuccessful postselection might require fast feedforward to block the final measurement. But in single-particle experiments, the postselection can be virtually free and, indeed, unavoidable: an unsuccessful postselec-

[9] Reference [167] considered squeezed coherent states as metrological probes in specific weak-measurement experiments. It is shown that postselection can improve the signal-to-noise ratio, irrespectively of whether the analysis includes the failed trials. However, this work concerned nonpostselected experiments in which only the probe state was measured. Had it been possible to successfully measure also the target system, the advantage would have disappeared.

tion can destroy the particle, precluding the triggering of the final measurement's detection apparatus [37]. Thus, current single-particle metrology could benefit from postselected improvements of the Fisher information. A photonic experimental test of our results is currently under investigation.

From a fundamental perspective, our results highlight the strangeness of quantum mechanics as a noncommuting theory. Classically, an increase of the Fisher information via postselection can be understood as the *a posteriori* selection of a better input distribution. But it is nonintuitive that quantum mechanical postselection can enable a quantum state to carry more Fisher information than the best possible input state could. The optimized Cramér-Rao bound, obtained from Eq. 8.4, can be written in the form of an uncertainty relation: $\sqrt{\mathrm{Var}(\theta_e)}(\Delta a) \geq 1$ [88]. Our results highlight the probabilistic possibility of violating this bound. More generally, the information-cost rate's ability to violate a classical bound leverages negativity, a nonclassical resource in quantum foundations, for metrological advantage.

IV Supplementary Information

IV.I Supplementary Note 1 – Expressing the postselected quantum Fisher information in terms of the KD distribution

As shown in the Results section, the postselected quantum Fisher information is given by

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) = 4\,\langle \dot{\psi}_\theta^{\mathrm{ps}}|\dot{\psi}_\theta^{\mathrm{ps}}\rangle \frac{1}{p_\theta^{\mathrm{ps}}} - 4|\,\langle \dot{\psi}_\theta^{\mathrm{ps}}|\psi_\theta^{\mathrm{ps}}\rangle\,|^2 \frac{1}{(p_\theta^{\mathrm{ps}})^2}, \tag{8.13}$$

where nonrenormalized postselected quantum state is $|\psi_\theta^{\mathrm{ps}}\rangle = \hat{F}\hat{U}(\theta)\,|\Psi_0\rangle$, where $|\Psi_0\rangle\,\langle\Psi_0| \equiv \hat{\rho}_0$. $p_\theta^{\mathrm{ps}} = \mathrm{Tr}(\hat{F}\hat{\rho}_\theta)$ is the probability of postselection.

In this supplementary note, we show that Eq. 8.13 can be expressed in terms of the double-extended KD distribution:

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) = 4 \sum_{\substack{a,a',\\ f\in\mathscr{F}^{\mathrm{ps}}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p_\theta^{\mathrm{ps}}} aa' - 4\Big| \sum_{\substack{a,a',\\ f\in\mathscr{F}^{\mathrm{ps}}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p_\theta^{\mathrm{ps}}} a\Big|^2, \tag{8.14}$$

The first term of the quantum Fisher information (Eq. 8.13) is

$$\frac{4}{p_\theta^{\mathrm{ps}}}\langle\dot{\psi}_\theta^{\mathrm{ps}}|\dot{\psi}_\theta^{\mathrm{ps}}\rangle = \frac{4}{p_\theta^{\mathrm{ps}}}\mathrm{Tr}\big(\hat{F}\dot{\hat{U}}(\theta)\hat{\rho}_0\dot{\hat{U}}^\dagger(\theta)\hat{F}^\dagger\big) = \frac{4}{p_\theta^{\mathrm{ps}}}\mathrm{Tr}\Big(\hat{F}\hat{A}\hat{\rho}_\theta\hat{A}\Big) \tag{8.15}$$

$$= \frac{4}{p_\theta^{\mathrm{ps}}}\mathrm{Tr}\Big(\sum_a |a\rangle\langle a|a\hat{\rho}_\theta\sum_{a'}|a'\rangle\langle a'|a'\sum_{f\in\mathscr{F}_{\mathrm{ps}}}|f\rangle\langle f|\Big), \tag{8.16}$$

where, in Eq. 8.16, we have expressed $\hat{A}$ and $\hat{F}$ in their corresponding eigendecompositions. This expression can be rewritten in terms of the doubly extended Kirkwood-Dirac quasiprobability distribution ($q^{\hat{\rho}}_{a,a',f} = \langle f|a\rangle\langle a|\hat{\rho}|a'\rangle\langle a'|f\rangle$):

$$\frac{4}{p_\theta^{\mathrm{ps}}}\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\mathrm{ps}}}}\mathrm{Tr}\left(aa' q^{\hat{\rho}_\theta}_{a,a',f}\frac{|a\rangle\langle f|}{\langle f|a\rangle}\right) = \frac{4}{p_\theta^{\mathrm{ps}}}\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\mathrm{ps}}}} q^{\hat{\rho}_\theta}_{a,a',f}aa'. \tag{8.17}$$

Similarly, the second term of Eq. 8.13 is

$$\frac{4}{(p_\theta^{\mathrm{ps}})^2}\big|\langle\psi_\theta^{\mathrm{ps}}|\dot{\psi}_\theta^{\mathrm{ps}}\rangle\big|^2 = \frac{4}{(p_\theta^{\mathrm{ps}})^2}\big|\mathrm{Tr}\big(\hat{F}\hat{\rho}_\theta\hat{A}\big)\big|^2 = \frac{4}{(p_\theta^{\mathrm{ps}})^2}\Big|\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\mathrm{ps}}}} q^{\hat{\rho}_\theta}_{a,a',f}a\Big|^2. \tag{8.18}$$

Combining the expressions above gives Eq. 8.14:

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) = 4\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\mathrm{ps}}}}\frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p_\theta^{\mathrm{ps}}}aa' - 4\Big|\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\mathrm{ps}}}}\frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p_\theta^{\mathrm{ps}}}a\Big|^2. \tag{8.19}$$

IV.II Supplementary note 2 – Proof of Theorem 2

Here, we prove Theorem 2. First, we upper-bound the right-hand side of Eq. 8.14, assuming that all $q^{\hat{\rho}_\theta}_{a,a',f}/p_\theta^{\mathrm{ps}} \in [0,1]$. We label the M eigenvalues of $\hat{A}$ and arrange them in increasing order: $a_1, a_2, ..., a_M$, such that $a_1 \equiv a_{\min}$ and $a_M \equiv a_{\max}$. Initially, we assume that the 0-point of the eigenvalue axis is set such that $a_1 = 0$ and $a_M = \Delta a$. In this scenario, all the components of the first term of Eq. 8.14 are nonnegative. We temporarily ignore the form of $q^{\hat{\rho}_\theta}_{a,a',f}/p_\theta^{\mathrm{ps}}$, and treat this ratio as a general quasiprobability distribution. Then, $\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})$ maximizes when $q^{\hat{\rho}_\theta}_{a,a',f}/p_\theta^{\mathrm{ps}}$ vanishes at all a' values except $a' = a_{\max}$. We define $q_a \equiv \sum_{a',f\in\mathscr{F}^{\mathrm{ps}}} q^{\hat{\rho}_\theta}_{a,a',f}/p_\theta^{\mathrm{ps}}$, such that all $q_a \in [0,1]$ and $\sum_a q_a = 1$. If $q^{\hat{\rho}_\theta}_{a,a',f}/p_\theta^{\mathrm{ps}}$ is nonzero

only when $a' = a_{\max}$, Eq. 8.14 becomes

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) = 4a_M \sum_a q_a a - 4\Big(\sum_a q_a a\Big)^2. \tag{8.20}$$

Expanding each sum, we obtain

$$\begin{aligned}\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) &= 4a_M(q_{a_1}a_1 + K + q_{a_M}a_M) - 4(q_{a_1}a_1 + K + q_{a_M}a_M)^2 &(8.21)\\ &= 4a_M(K + q_{a_M}a_M) - 4(K + q_{a_M}a_M)^2, &(8.22)\end{aligned}$$

where we used $q_{a_1}a_1 = 0$ and defined $K \equiv \sum_{a\in\{a_2,\ldots,a_{M-1}\}} q_a a \leq a_M$. As $\hat{A}$ is not totally degenerate, $a_M \neq 0$, and Eq. 8.22 is maximized when $q_{a_M} = (a_M - 2K)/(2a_M)$. This yields

$$\max\{\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})\} = a_M^2 = (\Delta a)^2, \tag{8.23}$$

where we have recalled that $a_M = \Delta a$.

We are left with proving that we can always set $a_1 = 0$ and $a_M = \Delta a$. We continue to assume that $q^{\hat{\rho}_\theta}_{a,a',f}/p^{\rm ps}_\theta \in [0,1]$, and we shift all the eigenvalues by a constant real value δ_a. The effect on $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$ is

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) \to 4\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\rm ps}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\rm ps}_\theta}(a+\delta_a)(a'+\delta_a) - 4\Big[\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\rm ps}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\rm ps}_\theta}(a+\delta_a)\Big]^2 \tag{8.24}$$

$$= 4\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\rm ps}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\rm ps}_\theta} aa' - 4\Big[\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\rm ps}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\rm ps}_\theta} a\Big]^2 + 4\delta_a\Big(\sum_{\substack{a,a',\\ f\in\mathscr{F}^{\rm ps}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\rm ps}_\theta} a - \sum_{\substack{a,a',\\ f\in\mathscr{F}^{\rm ps}}} \frac{q^{\hat{\rho}_\theta}_{a,a',f}}{p^{\rm ps}_\theta} a'\Big) = \mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}). \tag{8.25}$$

The last equality holds because $q^{\hat{\rho}}_{a,a',f} = \big(q^{\hat{\rho}}_{a',a,f}\big)^*$ generally and we are assuming that $q^{\hat{\rho}}_{a,a',f} \in \mathbb{R}$. Consequently, if all $q^{\hat{\rho}_\theta}_{a,a',f}/p^{\rm ps}_\theta \in [0,1]$, then $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) \leq (\Delta a)^2$. The second term of Eq. 8.14 cannot be decreased by imaginary values in $q^{\hat{\rho}_\theta}_{a,a',f}$. Moreover, the first term is necessarily real and nonnegative. Thus imaginary elements $q^{\hat{\rho}_\theta}_{a,a',f}$ cannot increase $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps})$. If $\mathscr{I}_Q(\theta|\Psi_\theta^{\rm ps}) > (\Delta a)^2$, then $q^{\hat{\rho}_\theta}_{a,a',f}$ must have negative entries. □

IV.III Supplementary note 3 – Infinite postselected quantum Fisher information

Here, we show that the postselected quantum Fisher information $\mathscr{I}_Q(\theta|\Psi_\theta^{\text{ps}})$ can approach infinity. The proof is by example; other examples might exist.

We assume that the generator $\hat{A}$ has $M \geq 3$ eigenvalues that are not all identical. We also assume that we possess an estimate θ_0 that lies close to the true value of θ: $\delta_\theta \equiv \theta - \theta_0$, with $|\delta_\theta| \ll 1$. (The derivation of the quantum Fisher information also rests on the assumption that one has access to such an estimate [27].)

By Eqs. 8.13, 8.15 and 8.18,

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\text{ps}}) = \frac{4}{p_\theta^{\text{ps}}}\text{Tr}\Big(\hat{F}\hat{A}\hat{U}(\theta)\hat{\rho}_0\hat{U}(\theta)^\dagger\hat{A}\Big) - \frac{4}{(p_\theta^{\text{ps}})^2}\Big|\text{Tr}\Big(\hat{F}\hat{U}(\theta)\hat{\rho}_0\hat{U}(\theta)^\dagger\hat{A}\Big)\Big|^2. \quad (8.26)$$

We now choose $\hat{F}$ and $\hat{\rho}_0$ such that $\mathscr{I}_Q(\theta|\Psi_\theta^{\text{ps}})$ approaches infinity. Crudely, p_θ^{ps} must approach 0 while $\text{Tr}(\hat{F}\hat{A}\hat{U}(\theta)\hat{\rho}_0\hat{U}(\theta)^\dagger\hat{A})$ either stays constant or approaches 0 more slowly. We label the M eigenvalues of $\hat{A}$ and arrange them in increasing order: $a_1, a_2, ..., a_M$, such that $a_1 \equiv a_{\min}$ and $a_M \equiv a_{\max}$.

First, we choose $\hat{F} = |f_1\rangle\langle f_1| + |f_2\rangle\langle f_2|$, where

$$|f_1\rangle \equiv \frac{|a_{\max}\rangle + |a_{\min}\rangle}{\sqrt{2}}, \quad (8.27)$$

$$|f_2\rangle \equiv \frac{\frac{i}{\sqrt{2}}(|a_{\max}\rangle - |a_{\min}\rangle) + |a_k\rangle}{\sqrt{2}}, \quad (8.28)$$

and $|a_k\rangle \neq |a_{\max}\rangle, |a_{\min}\rangle$. We also choose $\hat{\rho}_0 = |\Psi_0\rangle\langle\Psi_0|$ such that

$$|\Psi_0\rangle \equiv |\Psi_0(\theta_0, \phi)\rangle = \hat{U}^\dagger(\theta_0)\frac{1}{\sqrt{2}}\Big\{[\cos(\phi) - \sin(\phi)]\frac{i}{\sqrt{2}}(|a_{\min}\rangle - |a_{\max}\rangle) + [\cos(\phi) + \sin(\phi)]\,|a_k\rangle\Big\}. \quad (8.29)$$

$\phi \approx 0$ is a parameter that can be tuned to maximize the postselected Fisher information for a given approximation accuracy δ_θ. As ϕ is a parameter of the input state, variations in the Fisher information with ϕ will reflect the effects of disturbances to the input state.

Substituting the expressions for $\hat{F}$ and $\hat{\rho}_0$ into Eq. 8.26, we find

$$\begin{aligned}\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) =& 8\big\{5-2\cos(2\phi)\big(\cos[(a_M-a_k)\delta_\theta]+\cos[(a_k-a_1)\delta_\theta]\big)+\cos[(a_M-a_1)\delta_\theta]\\ &\times[\sin(2\phi)-1]-\sin(2\phi)\big\}^{-2}\\ &\times\Big\{2a_M^2-a_Ma_k+a_k^2+2a_1^2-(3a_M+a_k)a_1+(a_M-a_k)(a_k-a_1)\cos(4\phi)\\ &\times\big(\cos[(a_M-a_1)\delta_\theta]-1\big)\\ &+(a_M-a_k)(a_k-a_1)\cos[(a_M-a_1)\delta_\theta]+2(a_M-a_1)\cos(2\phi)\big((a_1-a_k)\cos[(a_M-a_k)\delta_\theta]\\ &+(a_k-a_M)\cos[(a_k-a_1)\delta_\theta]\big)-2(a_M-a_1)^2\sin(2\phi)\\ &+(a_M-a_1)\big((a_k-a_1)\cos[(a_M-a_k)\delta_\theta]\\ &+(a_M-a_k)\cos[(a_k-a_1)\delta_\theta]\big)\sin(4\phi)\Big\}.\end{aligned}$$

The postselection probability is

$$p_\theta^{\mathrm{ps}}=\frac{1}{8}\Big\{5-2\cos(2\phi)\big(\cos[(a_M-a_k)\delta_\theta]+\cos[(a_k-a_1)\delta_\theta]\big)+\cos[(a_M-a_1)\delta_\theta][\sin(2\phi)-1]-\sin(2\phi)\Big\} \tag{8.30}$$

In the limit as our estimate θ_0 approaches the true value of θ, such that $\delta_\theta\to 0$,

$$\lim_{\delta_\theta\to 0} p_\theta^{\mathrm{ps}}=\sin^2(\phi), \tag{8.31}$$

$$\lim_{\delta_\theta\to 0}\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})=\frac{(\cot(\phi)-1)^2}{2}(\Delta a)^2,\ \text{and} \tag{8.32}$$

$$\lim_{\delta_\theta\to 0} p_\theta^{\mathrm{ps}}\times\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})=\frac{1}{2}[1-\sin(2\phi)](\Delta a)^2. \tag{8.33}$$

In the limit as $\phi\to 0$,

$$\lim_{\phi\to 0}\Big[\lim_{\delta_\theta\to 0} p_\theta^{\mathrm{ps}}\Big]=0, \tag{8.34}$$

$$\lim_{\phi\to 0}\Big[\lim_{\delta_\theta\to 0}\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})\Big]=\infty,\ \text{and} \tag{8.35}$$

$$\lim_{\phi\to 0}\Big[\lim_{\delta_\theta\to 0} p_\theta^{\mathrm{ps}}\times\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})\Big]=\frac{1}{2}(\Delta a)^2. \tag{8.36}$$

According to Eq. 8.35, if first δ_θ and then ϕ approaches 0 in Eq. 8.30, $\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})$ approaches infinity.

There are a few points to note. First, $\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})$ diverges in the two ordered limits. In any real experiment, one could not blindly set $\phi=0$, but would have to choose ϕ based

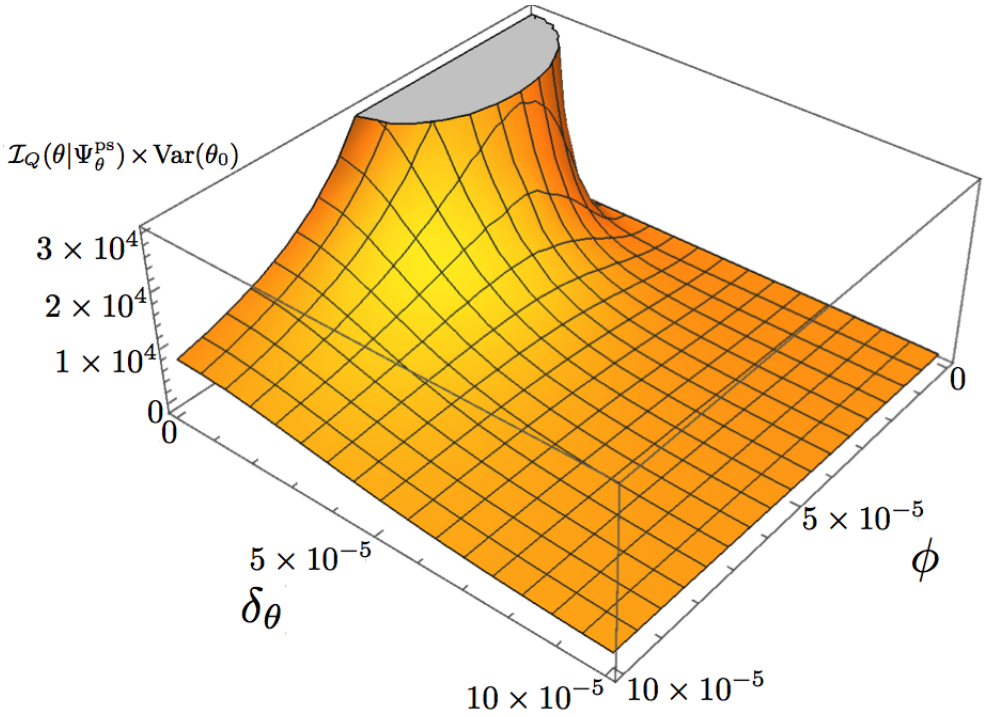

Fig. 8.3 **Scaled postselected quantum Fisher information.** The figure shows the postselected quantum Fisher information (Eq. 8.30) multiplied by the pre-experiment variance ${\rm Var}(\theta_0)$ as a function of ϕ and δ_θ. For small values of δ_θ and ϕ, the value of $\mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps}) \times {\rm Var}(\theta_0)$ diverges. The eigenvalues a_1, a_k and a_M are set to -1, 1 and 3, respectively. ${\rm Var}(\theta_0)$ was set to 1×10^{-6}.

on an estimate of θ. Second, if $\delta_\theta \approx 0$, then $\theta_0 \approx \theta$, and the pre-experiment variance of our initial estimate θ_0, ${\rm Var}(\theta_0)$, must be small. That is, we begin the experiment with much information about θ. Guided by the Cramér-Rao bound, we expect that, in a useful experiment, $\mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps})$ would grow large, while $1/{\rm Var}(\theta_0) < \mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps})$. Figure 8.3 shows $\mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps}) \times {\rm Var}(\theta_0)$ as a function of ϕ and δ_θ for an experiment where $a_1 = -1$, $a_k = 1$, $a_M = 3$ and ${\rm Var}(\theta_0) = 10^{-6}$. If θ_0 is within a few $\sigma_{\theta_0} \equiv \sqrt{{\rm Var}(\theta_0)}$ of θ, then $\mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps}) \times {\rm Var}(\theta_0) \gg 1$. Figure 8.3 shows that large values of $1/\delta_\theta$ can result in even larger values of $\mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps})$. Figure 8.3 also illustrates the effect of input-state disturbances of ϕ on $\mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps}) \times {\rm Var}(\theta_0)$. Third, while the theoretical strategy investigated in this appendix achieves an infinite postselected quantum Fisher information, the postselection also "wastes" information as $\lim_{\phi\to 0}[\lim_{\delta_\theta \to 0} p_\theta^{\rm ps} \times \mathcal{I}_Q(\theta|\Psi_\theta^{\rm ps})] < (\Delta a)^2$. If $\hat{A}$ possesses certain properties, it is possible to avoid wasting information through the postselection; we show how in the following appendix.

IV.IV Supplementary note 4 – Infinite postselected quantum Fisher information without loss of information

If the generator $\hat{A}$ has $M \geq 4$ eigenvalues, and the minimum and maximum eigenvalues are both at least doubly degenerate, then $\mathscr{I}_Q(\theta|\Psi_\theta^{\text{ps}})$ can approach infinity without information's being lost in the events discarded by postselection. We show how below.

First, we assign the orthonormal eigenvectors $|a_{\min_1}\rangle$ and $|a_{\min_2}\rangle$ to the eigenvalues $a_1 = a_{\min}$ and $a_2 = a_{\min}$, respectively. Here, we have reused the eigenvalue notation from Supp. Mat. IV.III. Similarly, we assign the orthonormal eigenvectors $|a_{\max_1}\rangle$ and $|a_{\max_2}\rangle$ to the eigenvalues $a_M = a_{\max}$ and $a_{M-1} = a_{\max}$, respectively. Second, we set $\hat{F} = |f_1\rangle\langle f_1| + |f_2\rangle\langle f_2|$, where

$$|f_1\rangle \equiv \frac{|a_{\max_2}\rangle - |a_{\min_1}\rangle}{\sqrt{2}}, \tag{8.37}$$

$$|f_2\rangle \equiv \frac{|a_{\min_2}\rangle - |a_{\max_1}\rangle}{\sqrt{2}}. \tag{8.38}$$

We also choose $|\Psi_0\rangle$ such that

$$|\Psi_0(\theta_0,\phi)\rangle = \hat{U}^\dagger(\theta_0)\frac{1}{2}\big\{[\cos(\phi)-\sin(\phi)](|a_{\max_2}\rangle + |a_{\min_2}\rangle) + [\sin(\phi)+\cos(\phi)](|a_{\max_1}\rangle + |a_{\min_1}\rangle)\big\}. \tag{8.39}$$

As in App. IV.III, $\phi \approx 0$ is a parameter that can be tuned to maximize $\mathscr{I}_Q(\theta|\Psi_\theta^{\text{ps}})$ for a given approximation accuracy of δ_θ.

Substituting the expressions for $\hat{F}$ and $\hat{\rho}_0$ into Eq. 8.26, we find

$$\mathscr{I}_Q(\theta|\Psi_\theta^{\text{ps}}) = \frac{\sin^2(2\phi)(a_M - a_1)^2}{\big(1-\cos(2\phi)\cos[(a_M-a_1)\delta_\theta]\big)^2}. \tag{8.40}$$

The postselection probability is

$$p_\theta^{\text{ps}} = \frac{1}{2}\Big\{1-\cos(2\phi)\cos[(a_M-a_1)\delta_\theta]\Big\}. \tag{8.41}$$

Again, we investigate the limit as our estimate θ_0 approaches the true value of θ:

$$\lim_{\delta_\theta \to 0} p_\theta^{\mathrm{ps}} = \sin^2(\phi), \tag{8.42}$$

$$\lim_{\delta_\theta \to 0} \mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) = \cot^2(\phi)(\Delta a)^2, \text{ and} \tag{8.43}$$

$$\lim_{\delta_\theta \to 0} p_\theta^{\mathrm{ps}} \times \mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) = \cos^2(\phi)(\Delta a)^2. \tag{8.44}$$

In the limit as $\phi \to 0$,

$$\lim_{\phi \to 0} \left[\lim_{\delta_\theta \to 0} p_\theta^{\mathrm{ps}} \right] = 0, \tag{8.45}$$

$$\lim_{\phi \to 0} \left[\lim_{\delta_\theta \to 0} \mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) \right] = \infty, \text{ and} \tag{8.46}$$

$$\lim_{\phi \to 0} \left[\lim_{\delta_\theta \to 0} p_\theta^{\mathrm{ps}} \times \mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) \right] = (\Delta a)^2. \tag{8.47}$$

In conclusion, the above strategy allows us to obtain an infinite value for $\mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}})$, while $p_\theta^{\mathrm{ps}} \times \mathscr{I}_Q(\theta|\Psi_\theta^{\mathrm{ps}}) = (\Delta a)^2$. No information is lost in the postselection. As in Supp Mat. IV.III, the results hold for the two ordered limits.

References

[1] Wikipedia gradient descent illustration. https://en.wikipedia.org/wiki/Gradient_descent. Accessed: 2021-04-22.

[2] Abernathy, C. R., Gusev, E. P., Schlom, D., and Stemmer, S. (2003). Fundamentals of novel oxide/semiconductor interfaces. volume 786.

[3] Agmon, N., Alhassid, Y., and Levine, R. (1979). An algorithm for finding the distribution of maximal entropy. *Journal of Computational Physics*, 30(2):250–258.

[4] Aharonov, Y., Albert, D. Z., and Vaidman, L. (1988). How the result of a measurement of a component of the spin of a spin- *1/2* particle can turn out to be 100. *Phys. Rev. Lett.*, 60:1351–1354.

[5] Aharonov, Y. and Vaidman, L. (2008). *The two-state vector formalism: an updated review*. Springer.

[6] Alhassid, Y., Agmon, N., and Levine, R. (1978). An upper bound for the entropy and its applications to the maximal entropy problem. *Chemical Physics Letters*, 53(1):22–26.

[7] Altland, A. and Simons, B. (2010). *Condensed Matter Field Theory*. Cambridge University Press, 2^{nd} edition.

[8] Arunkumar, A., Bolotin, E., Cho, B., Milic, U., Ebrahimi, E., Villa, O., Jaleel, A., Wu, C. J., and Nellans, D. (2017). Mcm-gpu: Multi-chip-module gpus for continued performance scalability. *SIGARCH Comput. Archit. News*, 45(2):320–332.

[9] Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G. S. L., Buell, D. A., Burkett, B., Chen, Y., Chen, Z., Chiaro, B., Collins, R., Courtney, W., Dunsworth, A., Farhi, E., Foxen, B., Fowler, A., Gidney, C., Giustina, M., Graff, R., Guerin, K., Habegger, S., Harrigan, M. P., Hartmann, M. J., Ho, A., Hoffmann, M., Huang, T., Humble, T. S., Isakov, S. V., Jeffrey, E., Jiang, Z., Kafri, D., Kechedzhi, K., Kelly, J., Klimov, P. V., Knysh, S., Korotkov, A., Kostritsa, F., Landhuis, D., Lindmark, M., Lucero, E., Lyakh, D., Mandrà, S., McClean, J. R., McEwen, M., Megrant, A., Mi, X., Michielsen, K., Mohseni, M., Mutus, J., Naaman, O., Neeley, M., Neill, C., Niu, M. Y., Ostby, E., Petukhov, A., Platt, J. C., Quintana, C., Rieffel, E. G., Roushan, P., Rubin, N. C., Sank, D., Satzinger, K. J., Smelyanskiy, V., Sung, K. J., Trevithick, M. D., Vainsencher, A., Villalonga, B., White, T., Yao, Z. J., Yeh, P., Zalcman, A., Neven, H., and Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510.

[10] Arvidsson-Shukur, D. R. M. and Barnes, C. H. W. (2016). Quantum counterfactual communication without a weak trace. *Phys. Rev. A*, 94:062303.

[11] Arvidsson-Shukur, D. R. M. and Barnes, C. H. W. (2019). Postselection and counterfactual communication. *Phys. Rev. A*, 99:060102.

[12] Arvidsson-Shukur, D. R. M., Gottfries, A. N. O., and Barnes, C. H. W. (2017a). Evaluation of counterfactuality in counterfactual communication protocols. *Phys. Rev. A*, 96:062316.

[13] Arvidsson-Shukur, D. R. M., Halpern, N. Y., Lepage, H. V., Lasek, A. A., Barnes, C. H. W., and Lloyd, S. (2020). Quantum advantage in postselected metrology. *Nature Comm.*, 11:3775.

[14] Arvidsson-Shukur, D. R. M., Lepage, H. V., Owen, E. T., Ferrus, T., and Barnes, C. H. W. (2017b). Protocol for fermionic positive-operator-valued measures. *Phys. Rev. A*, 96:052305.

[15] Askar, A. and Cakmak, A. S. (1978). Explicit integration method for the time-dependent schrödinger equation for collision problems. *The Journal of Chemical Physics*, 68(6):2794–2798.

[16] Atkinson, K. A. (1989). *An Introduction to Numerical Analysis*. Wiley & Sons, 2^{nd} edition.

[17] Bamber, C. and Lundeen, J. S. (2014). Observing dirac's classical phase space analog to the quantum state. *Phys. Rev. Lett.*, 112:070405.

[18] Barenco, A., Bennett, C. H., Cleve, R., and et al, D. P. D. (1995). Elementary gates for quantum computation. *Physical Review A*, 52(5):3457.

[19] Barnes, C., Shilton, J., and Robinson, A. (2000). Quantum computation using electrons trapped by surface acoustic waves. *Physical Review B*, 62(12):8410.

[20] Barnes, E., Arenz, C., Pitchford, A., and Economou, S. E. (2017). Fast microwave-driven three-qubit gates for cavity-coupled superconducting qubits. *Phys. Rev. B*, 96:024504.

[21] Barraud, S., Coquand, R., Casse, M., Koyama, M., Hartmann, J.-M., Maffini-Alvaro, V., Comboroure, C., Vizioz, C., Aussenac, F., Faynot, O., and Poiroux, T. (2012). Performance of omega-shaped-gate silicon nanowire mosfet with diameter down to 8 nm. *IEEE Electron Device Letters*, 33(11):1526–1528.

[22] Bergli, J., Somoza, A. M., and Ortuño, M. (2011). Effects of many-electron jumps in the relaxation and conductivity of coulomb glasses. *Phys. Rev. B*, 84:174201.

[23] Betz, A. C., Tagliaferri, M. L. V., Vinet, M., Broström, M., Sanquer, M., Ferguson, A. J., and Gonzalez-Zalba, M. F. (2016). Reconfigurable quadruple quantum dots in a silicon nanowire transistor. *Applied Physics Letters*, 108(20):203108.

[24] Bhole, G., Anjusha, V. S., and Mahesh, T. S. (2016). Steering quantum dynamics via bang-bang control: Implementing optimal fixed-point quantum search algorithm. *Phys. Rev. A*, 93:042339.

[25] Bordone, P., Bertoni, A., Rosini, M., Reggiani, S., and Jacoboni, C. (2004). Coherent transport in coupled quantum wires assisted by surface acoustic waves. *Semiconductor Science and Technology*, 19(4):S412–S414.

[26] Bransden, B. H. and Joachain, C. J. (2000). *Quantum Mechanics*. Prentice Hall, 2^{nd} edition.

[27] Braunstein, S. L. and Caves, C. M. (1994). Statistical distance and the geometry of quantum states. *Phys. Rev. Lett.*, 72:3439–3443.

[28] Brock, D. C. and Moore, G. E. (2006). Understanding moore's law: four decades of innovation. *Chemical Heritage Foundation*.

[29] Brunner, N. and Simon, C. (2010). Measuring small longitudinal phase shifts: Weak measurements or standard interferometry? *Phys. Rev. Lett.*, 105:010405.

[30] Brunner, R., Shin, Y.-S., Obata, T., Pioro-Ladrière, M., Kubo, T., Yoshida, K., Taniyama, T., Tokura, Y., and Tarucha, S. (2011). Two-qubit gate of combined single-spin rotation and interdot spin exchange in a double quantum dot. *Phys. Rev. Lett.*, 107:146801.

[31] Bäuerle, C., Glattli, D. C., and Tristan Meunier, e. a. (2018). Coherent control of single electrons: a review of current progress. *Reports on Progress in Physics*, 81(5).

[32] Burkard, G., Loss, D., and DiVincenzo, D. P. (1999). Coupled quantum dots as quantum gates. *Physical Review B*, 59(3):2070.

[33] Burnett, J. J., Bengtsson, A., Scigliuzzo, M., Niepce, D., Kudra, M., Delsing, P., and Bylander, J. (2019). Decoherence benchmarking of superconducting qubits. *npj Quantum Information*, 5(1):54.

[34] Buscemi, F., Bordone, P., and Bertoni, A. (2010). Quantum teleportation of electrons in quantum wires with surface acoustic waves. *Phys. Rev. B*, 81:045312.

[35] Bush, R. A., Ochoa, E. D., and Perron, J. K. (2021). Transport through quantum dots: An introduction via master equation simulations. *American Journal of Physics*, 89(3):300–306.

[36] Buterakos, D., Das Sarma, S., and Barnes, E. (2021). Geometrical formalism for dynamically corrected gates in multiqubit systems. *PRX Quantum*, 2:010341.

[37] Calafell, I. A., Strömberg, T., Arvidsson-Shukur, D., Rozema, L., Saggio, V., Greganti, C., Harris, N., Prabhu, M., Carolan, J., Hochberg, M., et al. (2019). Trace-free counterfactual communication with a nanophotonic processor. *npj Quantum Information*, 5(1):61.

[38] Carmichael, H. J. (2013). *Statistical methods in quantum optics 1 Master equations and Fokker-Planck equations*. Springer Science & Business Media.

[39] Cimini, V., Gianani, I., Piacentini, F., Degiovanni, I. P., and Barbieri, M. (2020). Anomalous values, fisher information, and contextuality, in generalized quantum measurements. 5(2):025007.

[40] Combes, J., Ferrie, C., Jiang, Z., and Caves, C. M. (2014). Quantum limits on postselected, probabilistic quantum metrology. *Phys. Rev. A*, 89:052117.

[41] Connors, E. J., Nelson, J., Qiao, H., Edge, L. F., and Nichol, J. M. (2019). Low-frequency charge noise in si/sige quantum dots. *Phys. Rev. B*, 100:165305.

[42] Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. John Wiley and Sons Inc., Hoboken, New Jersey, USA, 2nd edition.

[43] Cramér, H. (2016). *Mathematical methods of statistics (PMS-9)*, volume 9. Princeton University Press.

[44] Dasgupta, S. (2014). *It Began with Babbage: The Genesis of Computer Science*. Oxford University Press.

[45] Datta, S. (2005). *Quantum Transport: Atom to Transistor*. Cambridge University Press.

[46] Delfosse, N., Allard Guerin, P., Bian, J., and Raussendorf, R. (2015). Wigner function negativity and contextuality in quantum computation on rebits. *Phys. Rev. X*, 5:021003.

[47] Delfosse, N., Okay, C., Bermejo-Vega, J., Browne, D. E., and Raussendorf, R. (2017). Equivalence between contextuality and negativity of the wigner function for qudits. *New Journal of Physics*, 19(12):123024.

[48] Demkowicz-Dobrzański, R. and Maccone, L. (2014). Using entanglement against noise in quantum metrology. *Phys. Rev. Lett.*, 113:250801.

[49] Deutsch, D. and Jozsa, R. (1992). Rapid solution of problems by quantum computation. *Proceedings of The Royal Society A*, 439.

[50] Dirac, P. A. M. (1945). On the analogy between classical and quantum mechanics. *Rev. Mod. Phys.*, 17:195–199.

[51] DiVincenzo, D. P. (1995). Two-bit gates are universal for quantum computation. *Physical Review A*, 51(12):1015.

[52] DiVincenzo, D. P. (2000). The physical implementation of quantum computation. *arXiv:quant-ph/0002077*.

[53] DiVincenzo, D. P., Bacon, D., Kempe, J., Burkard, G., and Whaley, K. B. (2000). Universal quantum computation with the exchange interaction. *Nature*, 408(6810):339.

[54] Dixon, P. B., Starling, D. J., Jordan, A. N., and Howell, J. C. (2009). Ultrasensitive beam deflection measurement via interferometric weak value amplification. *Phys. Rev. Lett.*, 102:173601.

[55] Dobrosavljević, V., Trivedi, N., and Valles, J. (2012). *Conductor Insulator Quantum Phase Transitions*. Oxford Scholarship Online.

[56] Dovzhenko, Y., Stehlik, J., Petersson, K. D., Petta, J. R., Lu, H., and Gossard, A. C. (2011). Nonadiabatic quantum control of a semiconductor charge qubit. *Phys Rev. B*, 84:161302.

[57] Dressel, J. (2015). Weak values as interference phenomena. *Phys. Rev. A*, 91:032116.

[58] Dressel, J., Agarwal, S., and Jordan, A. N. (2010). Contextual values of observables in quantum measurements. *Phys. Rev. Lett.*, 104(24):240401.

[59] Dressel, J., Malik, M., Miatto, F. M., Jordan, A. N., and Boyd, R. W. (2014). Colloquium: Understanding quantum weak values: Basics and applications. *Reviews of Modern Physics*, 86(1):307.

[60] Duan, L.-M., Cirac, J. I., and Zoller, P. (2001). Geometric manipulation of trapped ions for quantum computation. *Science*, 292(5522).

[61] Duck, I. M., Stevenson, P. M., and Sudarshan, E. C. G. (1989). The sense in which a "weak measurement" of a spin-½ particle's spin component yields a value 100. *Phys. Rev. D*, 40:2112–2117.

[62] Economou, S. E. and Barnes, E. (2015). Analytical approach to swift nonleaky entangling gates in superconducting qubits. *Phys. Rev. B*, 91:161405.

[63] Edlbauer, H., Takada, S., Roussely, G., Yamamoto, M., Tarucha, S., Ludwig, A., Wieck, A. D., Meunier, T., and Bäuerle, C. (2017). Non-universal transmission phase behaviour of a large quantum dot. *Nature communications*, 8(1):1710.

[64] Efros, A. L. and Shklovskii, B. I. (1975). Coulomb gap and low temperature conductivity of disordered systems. *Journal of Physics C: Solid State Physics*, 8(4):L49–L51.

[65] Egan, P. and Stone, J. A. (2012). Weak-value thermostat with 0.2 mk precision. *Opt. Lett*, 37(23):4991–4993.

[66] Eigen. Eigen documentation. https://eigen.tuxfamily.org/dox/ Accessed: 2018-07-04.

[67] Elg, D. T., Sporre, J. R., Panici, G. A., Srivastava, S. N., and Ruzic, D. N. (2016). In situ collector cleaning and extreme ultraviolet reflectivity restoration by hydrogen plasma for extreme ultraviolet sources. *Journal of Vacuum Science & Technology A*, 34(2):021305.

[68] Ferrie, C. (2011). Quasi-probability representations of quantum theory with applications to quantum information science. *Rep. Prog. Phys.*, 74(11):116001.

[69] Ferrie, C. and Combes, J. (2014a). How the result of a single coin toss can turn out to be 100 heads. *Phys. Rev. Lett.*, 113:120404.

[70] Ferrie, C. and Combes, J. (2014b). Weak value amplification is suboptimal for estimation and detection. *Phys. Rev. Lett.*, 112:040406.

[71] Ferrie, C. and Emerson, J. (2008). Frame representations of quantum mechanics and the necessity of negativity in quasi-probability representations. 41(35):352001.

[72] Ferrus, T., George, R., Barnes, C. H. W., and Pepper, M. (2010). Disorder and electron interaction control in low-doped silicon metal-oxide-semiconductor field effect transistors. *Applied Physics Letters*, 97(14):142108.

[73] Ferrus, T., Rossi, A., Tanner, M., Podd, G., Chapman, P., and Williams, D. A. (2011). Detection of charge motion in a non-metallic silicon isolated double quantum dot. *New J. Phys.*, 13:103012.

[74] Feynman, R. P. (1986). Quantum mechanical computers. *Foundations of Physics*, 16(6):507.

[75] Foletti, S., Bluhm, H., Mahalu, D., Umansky, V., and Yacoby, A. (2009). Universal quantum control of two-electron spin quantum bits using dynamic nuclear polarization. *Nature Phys.*, 5:903.

[76] Ford, C. J. B. (2017). Transporting and manipulating single electrons in surface-acoustic-wave minima. *Physica Status Solidi (b)*, 254(3).

[77] Freeth, T., Bitsakis, Y., and X. Moussas, e. a. (2006). Decoding the ancient greek astronomical calculator known as the antikythera mechanism. *Nature*, 444.

[78] Frieden, B. R. (2004). *Science from Fisher information: a unification*. Cambridge University Press.

[79] Friis, N., Marty, O., Maier, C., Hempel, C., Holzäpfel, M., Jurcevic, P., Plenio, M. B., Huber, M., Roos, C., Blatt, R., and Lanyon, B. (2018). Observation of entangled states of a fully controlled 20-qubit system. *Phys. Rev. X*, 8:021012.

[80] Fujisawa, T., Hayashi, T., Cheong, H. D., Jeong, Y. H., and Hirayama, Y. (2004). Rotation and phase-shift operations for a charge qubitin a double quantum dot. *Physica E*, 21:1046.

[81] Fujisawa, T., Hayashi, T., and Sasaki, S. (2006). Time-dependent single-electron transport through quantum dots. *Rep. Prog. Phys.*, 69:759.

[82] Fujiwara, A. and Nagaoka, H. (1995). Quantum fisher metric and estimation for pure state models. *Phys. Lett. A*, 201(2):119 – 124.

[83] Furuta, S., Barnes, C. H. W., and Doran, C. J. L. (2004). Single-qubit gates and measurements in the surface acoustic wave quantum computer. *Physical Review B*, 70:205320.

[84] Gardelis, S., Smith, C. G., Cooper, J., Ritchie, D. A., Linfield, E. H., Jin, Y., and Pepper, M. (2003). *Phys .Rev. B*, 67:073302.

[85] Georgescu, I. (2020). 25 years of quantum error correction. *Nature Reviews Physics*, 2(10):519–519.

[86] Giavaras, G., Jefferson, J. H., Fearn, M., and Lambert, C. J. (2007). Generation of einstein-podolsky-rosen pairs and interconversion of static and flying electron spin qubits. *Physical Review B*, 76(24):245328.

[87] Giovannetti, V., Lloyd, S., and Maccone, L. (2006). Quantum metrology. *Physical review letters*, 96(1):010401.

[88] Giovannetti, V., Lloyd, S., and Maccone, L. (2011). Advances in quantum metrology. *Nature photonics*, 5(4):222.

[89] Goldberg, A., Schey, H. M., and Schwartz, J. L. (1967). Computer-generated motion pictures of one-dimensional quantum-mechanical transmission and reflection phenomena. *American Journal of Physics*, 35(3):177–186.

[90] González Alonso, J. R., Yunger Halpern, N., and Dressel, J. (2019). Out-of-time-ordered-correlator quasiprobabilities robustly witness scrambling. *Phys. Rev. Lett.*, 122:040404.

[91] Gonzalez-Zalba, M. F., Barraud, S., Ferguson, A. J., and Betz, A. C. (2015). Probing the limits of gate-based charge sensing. *Nat. Comm.*, 6:6084.

[92] Gonzalez-Zalba, M. F., Heiss, D., Podd, G., and Ferguson, A. J. (2012). Tunable aluminium-gated single electron transistor on a doped silicon-on-insulator etched nanowire. *Applied Physics Letters*, 101(10):103504.

[93] Gorman, J., Hasko, D. G., and Williams, D. A. (2005). Charge-qubit operation of an isolated double quantum dot. *Phys. Rev. Lett.*, 95:090502.

[94] Gradshteyn, I. S. and Ryzhik, I. M. (1994). *Table of Integrals, Series, and Products*. Academic Press.

[95] Greganti, C., Schiansky, P., Calafell, I. A., Procopio, L. M., Rozema, L. A., and Walther, P. (2018). Tuning single-photon sources for telecom multi-photon experiments. *Opt. Express*, 26(3):3286–3302.

[96] Grimmeiss, H. G., Janzén, E., and Larsson, K. (1982). Multivalley spin splitting of 1s states for sulfur, selenium, and tellurium donors in silicon. *Phys. Rev. B*, 25:2627–2632.

[97] Group, C. S. P. Saw wafer. https://www.sp.phy.cam.ac.uk/research/fundamentals-of-low-dimensional-semiconductor-systems/saw Accessed: 2021-07-01.

[98] Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, page 212.

[99] Halpern, N. Y., Bartolotta, A., and Pollack, J. (2019). Entropic uncertainty relations for quantum information scrambling. *Communications Physics*, 2(1):1–12.

[100] Hardy, L. (1992). Quantum mechanics, local realistic theories, and lorentz-invariant realistic theories. *Phys. Rev. Lett.*, 68:2981–2984.

[101] Harris, J., Boyd, R. W., and Lundeen, J. S. (2017). Weak value amplification can outperform conventional measurement in the presence of detector saturation. *Phys. Rev. Lett.*, 118(7):070802.

[102] Hayashi, T., Fujisawa, T., Cheong, H. D., Jeong, Y. H., and Hirayama, Y. (2003). Coherent manipulation of electronic states in a double quantum dot. *Phys. Rev. Lett.*, 91:226804.

[103] He, Y., Gorman, S. K., Keith, D., Kranz, L., Keizer, J. G., and Simmons, M. Y. (2019). A two-qubit gate between phosphorus donor electrons in silicon. *Nature*, 571(7765):371–375.

[104] Helstrom, C. W. (1969). Quantum detection and estimation theory. *Journal of Statistical Physics*, 1(2):231–252.

[105] Hermelin, S., Takada, S., Yamamoto, M., Tarucha, S., Wieck, A. D., Saminadayar, L., Bäuerle, C., and Meunier, T. (2011). Electrons surfing on a sound wave as a platform for quantum optics with flying electrons. *Nature*, 477(7365):435.

[106] Hisamoto, D., Kaga, T., Kawamoto, Y., and Takeda, E. (1990). A fully depleted lean-channel transistor (delta)-a novel vertical ultrathin soi mosfet. *IEEE Electron Device Letters*, 11:36.

[107] Hofmann, H. F. (2012). Complex joint probabilities as expressions of reversible transformations in quantum mechanics. *New Journal of Physics*, 14(4):043031.

[108] Hofmann, H. F., Goggin, M. E., Almeida, M. P., and Barbieri, M. (2012). Estimation of a quantum interaction parameter using weak measurements: Theory and experiment. *Phys. Rev. A*, 86:040102.

[109] Hosten, O. and Kwiat, P. (2008). Observation of the spin hall effect of light via weak measurements. *Science*, 319(5864):787–790.

[110] Howard, M., Wallman, J., Veitch, V., and Emerson, J. (2014). Contextuality supplies the 'magic'for quantum computation. *Nature*, 510(7505):351.

[111] Huang, W., Yang, C. H., Chan, K. W., Tanttu, T., Hensen, B., Leon, R. C. C., Fogarty, M. A., Hwang, J. C. C., Hudson, F. E., Itoh, K. M., Morello, A., Laucht, A., and Dzurak, A. S. (2019). Fidelity benchmarks for two-qubit gates in silicon. *Nature*, 569(7757):532–536.

[112] James D. Siverns, e. a. (2012). Optimization of two-dimensional ion trap arrays for quantum simulation. *New Journal of Physics*, 14.

[113] Jamieson, D. N., Yang, C., Hopf, T., Hearne, S. M., Pakes, C. I., Prawer, S., Mitic, M., Gauja, E., Andresen, S. E., Hudson, F. E., Dzurak, A. S., and Clark, R. G. (2005). Controlled shallow single-ion implantation in silicon using an active substrate for sub-20-kev ions. *Applied Physics Letters*, 86(20):202101.

[114] Johansen, L. M. (2007). Quantum theory of successive projective measurements. *Phys. Rev. A*, 76:012119.

[115] Jordan, A. N., Martínez-Rincón, J., and Howell, J. C. (2014). Technical advantages for weak-value amplification: when less is more. *Physical Review X*, 4(1):011031.

[116] Joshi, M. K., Elben, A., Vermersch, B., Brydges, T., Maier, C., Zoller, P., Blatt, R., and Roos, C. F. (2020). Quantum information scrambling in a trapped-ion quantum simulator with tunable range interactions. *Phys. Rev. Lett.*, 124:240505.

[117] Kandel, Y. P., Qiao, H., Fallahi, S., Gardner, G. C., Manfra, M. J., and Nichol, J. M. (2019). Coherent spin-state transfer via heisenberg exchange. *Nature*, 573(7775):553–557.

[118] Kane, B. E. (1998). A silicon-based nuclear spin quantum computer. *Nature*, 393(6681):133.

[119] Kataoka, M., Astley, M. R., Thorn, A. L., Oi, D. K. L., Barnes, C. H. W., Ford, C. J. B., Anderson, D., Jones, G. A. C., Farrer, I., Ritchie, D. A., and Pepper, M. (2009). Coherent time evolution of a single-electron wave function. *Phys. Rev. Lett.*, 102:156801.

[120] Kataoka, M., Schneble, R., Thorn, A., Barnes, C., Ford, C., Anderson, D., Jones, G., Farrer, I., Ritchie, D., and Pepper, M. (2007). Single-electron population and depopulation of an isolated quantum dot using a surface-acoustic-wave pulse. *Physical Review Letters*, 98:046801.

[121] Kaye, P., Laflamme, R., and Mosca, M. (2007). *An Introduction to Quantum Computing*. Oxford University Press. Theorem 4.2.2.

[122] Kim, D., Ward, D. R., Simmons, C. B., Gamble, J. K., Blume-Kohout, R., Nielsen, E., Savage, D. E., Lagally, M. G., Friesen, M., Coppersmith, S. N., and Eriksson, M. A. (2015). Microwave-driven coherent operation of a semiconductor quantum dot charge qubit. *Nature Nano.*, 10:243.

[123] Kirkwood, J. G. (1933). Quantum statistics of almost classical assemblies. *Phys. Rev.*, 44:31–37.

[124] Kish, L. B. (2002). End of moore's law: thermal (noise) death of integration in micro and nano electronics. *Physical Review A*, 305:144–149.

[125] Kittel, C. (2005). *Introduction to Solid State Physics*. Wiley & Sons, 8^{th} edition.

[126] Knill, E., e. a. (2001). A scheme for efficient quantum computation with linear optics. *Nature*, 409:46–52.

[127] Kobayashi, T., Salfi, J., Chua, C., van der Heijden, J., House, M. G., Culcer, D., Hutchison, W. D., Johnson, B. C., McCallum, J. C., Riemann, H., Abrosimov, N. V., Becker, P., Pohl, H.-J., Simmons, M. Y., and Rogge, S. (2021). Engineering long spin coherence times of spin–orbit qubits in silicon. *Nature Materials*, 20(1):38–42.

[128] Kochen, S. and Specker, E. P. (1975). The problem of hidden variables in quantum mechanics. In *The logico-algebraic approach to quantum mechanics*, pages 293–328. Springer.

[129] Kofman, A. G., Ashhab, S., and Nori, F. (2012). Nonperturbative theory of weak pre- and post-selected measurements. *Physics Reports*, 520(2):43 – 133. Nonperturbative theory of weak pre- and post-selected measurements.

[130] Koppens, F. H. L., Buizert, C., Tielrooij, K. J., Vink, I. T., Nowack, K. C., Meunier, T., Kouwenhoven, L. P., and Vandersypen, L. M. K. (2006). Driven coherent oscillations of a single electron spin in a quantum dot. *Nature*, 442:766.

[131] Kosaka, H., Inagaki, T., Rikitake, Y., Imamura, H., Mitsumori, Y., and Edamatsu, K. (2009). Spin state tomography of optically injected electrons in a semiconductor. *Nature*, 457(7230):702.

[132] Krischek, R., Schwemmer, C., Wieczorek, W., Weinfurter, H., Hyllus, P., Pezzé, L., and Smerzi, A. (2011). Useful multiparticle entanglement and sub-shot-noise sensitivity in experimental phase estimation. *Phys. Rev. Lett.*, 107:080504.

[133] Kuhlmann, A. V., Houel, J., Ludwig, A., Greuter, L., Reuter, D., Wieck, A. D., Poggio, M., and Warburton, R. J. (2013). Charge noise and spin noise in a semiconductor quantum device. *Nature Physics*, 9(9):570–575.

[134] Kunjwal, R., Lostaglio, M., and Pusey, M. F. (2018). Anomalous weak values and contextuality: robustness, tightness, and imaginary parts. *arXiv preprint arXiv:1812.06940*.

[135] Landau, L. (1932). *Physics of the Soviet Union*, 2:46.

[136] Landau, L. D. and Lifshitz, E. M. (1980). *Statistical Physics: Part 1*. Butterworth-Heinemann.

[137] Lazić, S., Violante, A., Cohen, K., Hey, R., Rapaport, R., and Santos, P. V. (2014). Scalable interconnections for remote indirect exciton systems based on acoustic transport. *Physical Review B*, 89(8):085313.

[138] Leifer, M. S. and Pusey, M. F. (2017). Is a time symmetric interpretation of quantum theory possible without retrocausality? *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2202):20160607.

[139] Leifer, M. S. and Spekkens, R. W. (2005). Pre-and post-selection paradoxes and contextuality in quantum mechanics. *Phys. Rev. Lett.*, 95(20):200405.

[140] Lepage, H. V., Lasek, A. A., Arvidsson-Shukur, D. R. M., and Barnes, C. H. W. (2020). Entanglement generation via power-of-swap operations between dynamic electron-spin qubits. *Phys. Rev. A*, 101:022329.

[141] Li, S., Castellano, A. D., Wang, S., Wu, Y., Gong, M., Yan, Z., Rong, H., Deng, H., Zha, C., Guo, C., Sun, L., Peng, C., Zhu, X., and Pan, J.-W. (2019). Realisation of high-fidelity nonadiabatic cz gates with superconducting qubits. *npj Quantum Information*, 5(1):84.

[142] Lim, W. H., Huebl, H., van Beveren, L. H. W., Rubanov, S., Spizzirri, P. G., Angus, S. J., Clark, R. G., and Dzurak, A. S. (2009). Electrostatically defined few-electron double quantum dot in silicon. *Appl. Phys. Lett.*, 94:173502.

[143] Lipka-Bartosik, P. and Demkowicz-Dobrzański, R. (2018). Thermodynamic work cost of quantum estimation protocols. *Journal of Physics A: Mathematical and Theoretical*, 51(47):474001.

[144] Liuzzo-Scorpo, P., Correa, L. A., Pollock, F. A., Górecka, A., Modi, K., and Adesso, G. (2018). Energy-efficient quantum frequency estimation. *New Journal of Physics*, 20(6):063009.

[145] Lloyd, S. (1993). A potentially realizable quantum computer. *Science*, 261:1569.

[146] Loss, D. and DiVincenzo, D. P. (1998). Quantum computation with quantum dots. *Physical Review A*, 57(1):120.

[147] Lundeen, J. S. and Bamber, C. (2012). Procedure for direct measurement of general quantum states using weak measurement. *Phys. Rev. Lett.*, 108:070402.

[148] Lundeen, J. S., Sutherland, B., Patel, A., Stewart, C., and Bamber, C. (2011). Direct measurement of the quantum wavefunction. *Nature*, 474(7350):188.

[149] Lütkenhaus, N. and Barnett, S. M. (1995). Nonclassical effects in phase space. *Phys. Rev. A*, 51:3340–3342.

[150] Lyons, K., Dressel, J., Jordan, A. N., Howell, J. C., and Kwiat, P. G. (2015). Power-recycled weak-value-based metrology. *Phys. Rev. Lett.*, 114:170801.

[151] Maccone, L. (2013). Intuitive reason for the usefulness of entanglement in quantum metrology. *Phys. Rev. A*, 88:042109.

[152] MacQuarrie, E. R., Neyens, S. F., Dodson, J. P., Corrigan, J., Thorgrimsson, B., Holman, N., Palma, M., Edge, L. F., Friesen, M., Coppersmith, S. N., and Eriksson, M. A. (2020). Progress toward a capacitively mediated cnot between two charge qubits in si/sige. *npj Quantum Information*, 6(1):81.

[153] Maestri, J. J. V., Landau, R. H., and Paez, M. J. (2000). Two-particle Schrödinger equation animations of wave packet–wave packet scattering. *American Journal of Physics*, 68(12):1113–1119.

[154] Magaña Loaiza, O. S., Mirhosseini, M., Rodenburg, B., and Boyd, R. W. (2014). Amplification of angular rotations using weak measurements. *Phys. Rev. Lett.*, 112:200401.

[155] Malinowski, F. K., Martins, F., Smith, T. B., Bartlett, S. D., Doherty, A. C., Nissen, P. D., Fallahi, S., Gardner, G. C., Manfra, M. J., Marcus, C. M., and Kuemmeth, F. (2018). Spin of a multielectron quantum dot and its interaction with a neighboring electron. *Phys. Rev. X*, 8:011045.

[156] Martínez-Rincón, J., Mullarkey, C. A., Viza, G. I., Liu, W.-T., and Howell, J. C. (2017). Ultrasensitive inverse weak-value tilt meter. *Opt. Lett.*, 42(13):2479–2482.

[157] Mason, N., Biercuk, M. J., and Marcus, C. M. (2004). Local gate control of a carbon nanotube double quantum dot. *Science*, 303:655.

[158] McNeil, R. P. G., Kataoka, M., Ford, C. J. B., Barnes, C. H. W., Anderson, D., Jones, G. A. C., Farrer, I., and Ritchie, D. A. (2011). On-demand single-electron transfer between distant quantum dots. *Nature*, 477(7365):439.

[159] McNeil, R. P. G., Schneble, R. J., Kataoka, M., Ford, C. J. B., Kasama, T., Dunin-Borkowski, R. E., Feinberg, J. M., Harrison, R. J., Barnes, C. H. W., Tse, D. H. Y., Trypiniotis, T., Bland, J. A. C., Anderson, D., Jones, G. A. C., and Pepper, M. (2010). Localized magnetic fields in arbitrary directions using patterned nanomagnets. *Nano Letters*, 10:1549–1553.

[160] Mohseninia, R., Alonso, J. R. G., and Dressel, J. (2019). Optimizing measurement strengths for qubit quasiprobabilities behind out-of-time-ordered correlators. *Phys. Rev. A*, 100:062336.

[161] Morello, A., Escott, C. C., Huebl, H., Willems van Beveren, L. H., Hollenberg, L. C. L., Jamieson, D. N., Dzurak, A. S., and Clark, R. G. (2009). Architecture for high-sensitivity single-shot readout and control of the electron spin of individual donors in silicon. *Phys. Rev. B*, 80:081307.

[162] Mosakowski, J. (2016). *Quantum Computation in Double Quantum Dots*. PhD thesis, University of Cambridge.

[163] N. Kolmogorov, A. (1951). Foundations of the theory of probability. *The Mathematical Gazette*, 35.

[164] Nvidia. Cuda documentation. https://docs.nvidia.com/cuda/ Accessed: 2018-07-04.

[165] Owen, E., Dean, M., and Barnes, C. (2012). Generation of entanglement between qubits in a one-dimensional harmonic oscillator. *Physical Review A*, 85(2):022319.

[166] Owen, E. T. and Barnes, C. H. W. (2016). Ground-state electronic structure of quasi-one-dimensional wires in semiconductor heterostructures. *Phys. Rev. Appl.*, 6:054007.

[167] Pang, S. and Brun, T. A. (2015). Improving the precision of weak measurements by postselection measurement. *Phys. Rev. Lett.*, 115:120401.

[168] Pang, S., Dressel, J., and Brun, T. A. (2014). Entanglement-assisted weak value amplification. *Phys. Rev. Lett.*, 113:030401.

[169] Penrose, R. (1994). *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford University Press, Inc., New York, NY, USA, 1st edition.

[170] Pepper, M. and Uren, M. J. (1982). The wigner glass and conductance oscillations in silicon inversion layers. *Journal of Physics C: Solid State Physics*, 15(20):L617–L625.

[171] Pepper, M., Uren, M. J., and Oakley, R. E. (1979). Conductance oscillations and source-drain-limited conduction in si MOSFETs. *Journal of Physics C: Solid State Physics*, 12(23):L897–L900.

[172] Petersson, K. D., Petta, J. R., Lu, H., and Gossard, A. C. (2010). Quantum coherence in a one-electron semiconductor charge qubit. *Phys. Rev. Lett.*, 105:246804.

[173] Petta, J. R., Johnson, A. C., Taylor, J. M., Laird, E. A., Yacoby, A., Lukin, M. D., Marcus, C. M., Hanson, M. P., and Gossard, A. C. (2005). Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science*, 309:2180.

[174] Petz, D. (2007). *Quantum information theory and quantum statistics*. Springer Science & Business Media.

[175] Petz, D. and Ghinea, C. (2011). *Introduction to quantum Fisher information*. World Scientific.

[176] Piacentini, F., Avella, A., Levi, M. P., Gramegna, M., Brida, G., Degiovanni, I. P., Cohen, E., Lussana, R., Villa, F., Tosi, A., Zappa, F., and Genovese, M. (2016). Measuring incompatible observables by exploiting sequential weak values. *Phys. Rev. Lett.*, 117:170402.

[177] Pierre, M., Hofheinz, M., Jehl, X., Sanquer, M., Molas, G., Vinet, M., and Deleonibus, S. (2009). Background charges and quantum effects in quantum dots transport spectroscopy. *The European Physical Journal B*, 70(4):475–481.

[178] Pierre, M., Wacquez, R., Jehl, X., Sanquer, M., Vinet, M., and Cueto, O. (2010). Single-donor ionization energies in a nanoscale cmos channel. *Nature Nanotechnology*, 5(2):133–137.

[179] Pusey, M. F. (2014). Anomalous weak values are proofs of contextuality. *Phys. Rev. Lett.*, 113:200401.

[180] Pusey, M. F. and Leifer, M. S. (2015). Logical pre-and post-selection paradoxes are proofs of contextuality. *arXiv preprint arXiv:1506.07850*.

[181] QuTiP. Qutip documentation. https://github.com/qutip/qutip-doc Accessed: 2021-05-11.

[182] Randall, J. N., Lyding, J. W., Schmucker, S., Von Ehr, J. R., Ballard, J., Saini, R., Xu, H., and Ding, Y. (2009). Atomic precision lithography on si. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena*, 27(6):2764–2768.

[183] Rao, C. R. (1992). Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in statistics*, pages 235–247. Springer.

[184] Raussendorf, R. (2012). Key ideas in quantum error correction. *Philosophical Transactions of The Royal Society A*, 370.

[185] Robinson, A. M. and Barnes, C. H. W. (2001). Classical dynamics of electrons in quantized-acoustoelectric-current devices. *Phys. Rev. B*, 63:165418.

[186] Rodriquez, R., Oi, D. K. L., Kataoka, M., Barnes, C. H. W., Ohshima, T., and Ekert, A. K. (2005). Surface-acoustic-wave single-electron interferometry. *Physical Review B*, 72(8):085329.

[187] Rohde, P. (2015). Simple scheme for universal linear-optics quantum computing with constant experimental complexity using fiber loops. *Physical Review A*, 91:012306.

[188] Rossi, A., Ferrus, T., Lin, W., Kodera, T., Williams, D. A., and Oda, S. (2011). Detection of variable tunneling rates in silicon quantum dots. *Applied Physics Letters*, 98(13):133506.

[189] Rossi, A., Ferrus, T., Podd, G. J., and Williams, D. A. (2010). Charge detection in phosphorus-doped silicon double quantum dots. *Appl. Phys. Lett.*, 97:223506.

[190] Rowland, B. and Jones, J. A. (2012). Implementing quantum logic gates with gradient ascent pulse engineering: principles and practicalities. *Phil. Trans. R. Soc. A*, 370:4636–4650.

[191] S., S., Q., W., P., X., and X., C. (2016). Benchmarking state-of-the-art deep learning software tools. In *2016 7th International Conference on Cloud Computing and Big Data (CCBD)*, pages 99–104.

[192] Salari, A. and Darvazehban, A. (2019). A spiral surface acoustic wave transducer for quantum information processing. *Journal of Low Temperature Physics*, 195(1):138–152.

[193] Sanada, H., Kunihashi, Y., Gotoh, H., Onomitsu, K., Kohda, M., Nitta, J., Santos, P. V., and Sogawa, T. (2013). Manipulation of mobile spin coherence using magnetic-field-free electron spin resonance. *Nature Physics*, 9(5):280.

[194] Schmid, D. and Spekkens, R. W. (2018). Contextual advantage for state discrimination. *Physical Review X*, 8(1):011015.

[195] Schofield, S. R., Curson, N. J., Simmons, M. Y., Rueß, F. J., Hallam, T., Oberbeck, L., and Clark, R. G. (2003). Atomically precise placement of single dopants in si. *Phys. Rev. Lett.*, 91:136104.

[196] Schuetz, M. J. A., Kessler, E. M., Giedke, G., Vandersypen, L. M. K., Lukin, M. D., and Cirac, J. I. (2015). Universal quantum transducers based on surface acoustic waves. *Phys. Rev. X*, 5:031031.

[197] Schultz, M. (1999). The end of the road for silicon? *Nature*, 399:729–730.

[198] Sherwin, M., Imamoglu, A., and Montroy, T. (1999). Quantum computation with quantum dots and terahertz cavity quantum electrodynamics. *Physical Review A*, 60.

[199] Shevchenko, S., Ashhab, S., and Nori, F. (2010). Landau–zener–stückelberg interferometry. *Physics Reports*, 492(1):1–30.

[200] Shi, Z., Simmons, C. B., Ward, D. R., Prance, J. R., Mohr, R. T., Koh, T. S., Gamble, J. K., Wu, X., Savage, D. E., Lagally, M. G., Friesen, M., Coppersmith, S. N., and Eriksson, M. A. (2013). Coherent quantum oscillations and echo measurements of a si charge qubit. *Phys Rev. B*, 88:075416.

[201] Shin, S. J., Lee, J. J., Kang, H. J., Choi, J. B., Yang, S.-R. E., Takahashi, Y., and Hasko, D. G. (2011). Room-temperature charge stability modulated by quantum effects in a nanoscale silicon island. *Nano Letters*, 11(4):1591–1597.

[202] Shklovskii, B. and Efros, A. (1984). *Electronic Properties of Doped Semiconductors*. Springer.

[203] Shor, P. W. (1995). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM*, 26(5):1484–1509.

[204] Specht, H. P., Nölleke, C., Reiserer, A., Uphoff, M., Figueroa, E., Ritter, S., and Rempe, G. (2011). A single-atom quantum memory. *Nature*, 473(7346):190–193.

[205] Spekkens, R. W. (2005). Contextuality for preparations, transformations, and unsharp measurements. *Phys. Rev. A*, 71:052108.

[206] Spekkens, R. W. (2008). Negativity and contextuality are equivalent notions of nonclassicality. *Phys. Rev. Lett.*, 101(2):020401.

[207] Starling, D. J., Dixon, P. B., Jordan, A. N., and Howell, J. C. (2009). Optimizing the signal-to-noise ratio of a beam-deflection measurement with interferometric weak values. *Phys. Rev. A*, 80:041803.

[208] Starling, D. J., Dixon, P. B., Jordan, A. N., and Howell, J. C. (2010). Precision frequency measurements with interferometric weak values. *Phys. Rev. A*, 82:063822.

[209] Steger, M., Saeedi, K., Thewalt, M. L. W., Morton, J. J. L., Riemann, H., Abrosimov, N. V., Becker, P., and Pohl, H.-J. (2012). Quantum information storage for over 180 s using donor spins in a 28si "semiconductor vacuum". *Science*, 336(6086):1280–1283.

[210] Steinberg, A. M. (1995). Conditional probabilities in quantum theory and the tunneling-time controversy. *Phys. Rev. A*, 52:32–42.

[211] Stojan, R., Aneta, P., and Zoran, P. (2004). Dependence of static dielectric constant of silicon on resistivity at room temperature. 1(2):237–247.

[212] Stone, M. H. (1932). On one-parameter unitary groups in hilbert space. *Ann. Math*, 33(3):643–648.

[213] Stopa, M. (1996). Quantum dot self-consistent electronic structure and the coulomb blockade. *Phys. Rev. B*, 54:13767.

[214] Stotz, J. A. H., Hey, R., Santos, P. V., and Ploog, K. H. (2005). Coherent spin transport through dynamic quantum dots. *Nature materials*, 4(8):585.

[215] Stueckelberg, E. C. G. (1932). Theory of inelastic collisions between atoms. *Helvetica Physica Acta*, 5:369.

[216] Sun, L., He, X., You, C., Lv, C., Li, B., Lloyd, S., and Wang, X. (2020). Exponentially enhanced quantum metrology using resources of linear complexity. *arXiv preprint arXiv:2004.01216*.

[217] Takada, S., Edlbauer, H., Lepage, H. V., Wang, J., Mortemousque, P.-A., Georgiou, G., Barnes, C. H. W., Ford, C. J. B., Yuan, M., Santos, P. V., Waintal, X., Ludwig, A., Wieck, A. D., Urdampilleta, M., Meunier, T., and Bäuerle, C. (2019). Sound-driven single-electron transfer in a circuit of coupled quantum rails. *Nature Communications*, 10(1):4557.

[218] Talyanskii, V. I., Shilton, J. M., Pepper, M., Smith, C. G., Ford, C. J. B., Linfield, E. H., Ritchie, D. A., and Jones, G. A. C. (1997). Single-electron transport in a one-dimensional channel by high-frequency surface acoustic waves. *Phys. Rev. B*, 56:15180–15184.

[219] Tanner, M. G., Hasko, D. G., and Williams, D. A. (2006). Investigation of silicon isolated double quantum-dot energy levels for quantum computation. *Microelectronic Engineering*, 83(4):1818–1822. Micro- and Nano-Engineering MNE 2005.

[220] Taylor, J. M., Petta, J. R., Johnson, A. C., Yacoby, A., Marcus, C. M., and Lukin, M. D. (2007). Relaxation, dephasing, and quantum control of electron spins in double quantum dots. *Phys. Rev. B*, 76:035315.

[221] Thekkadath, G. S., Giner, L., Chalich, Y., Horton, M. J., Banker, J., and Lundeen, J. S. (2016). Direct measurement of the density matrix of a quantum system. *Phys. Rev. Lett.*, 117:120401.

[222] Thompson, S. E. and Parthasarathy, S. (2006). Moore's law: the future of si microelectronics. *Materials Today*, 9:20–25.

[223] Tollaksen, J. (2007). Pre-and post-selection, weak values and contextuality. *J. Phys. A*, 40(30):9033.

[224] Tsu, R. (2010). *Superlattice to Nanoelectronics, 2nd Edition*. Elsevier.

[225] Tyszka, K., Moraru, D., Samanta, A., Mizuno, T., Jabłoński, R., and Tabe, M. (2015). Effect of selective doping on the spatial dispersion of donor-induced quantum dots in si nanoscale transistors. *Applied Physics Express*, 8(9):094202.

[226] Unanyan, R., Shore, B., and Bergmann, K. (1999). Laser-driven population transfer in four-level atoms: Consequences of non-abelian geometrical adiabatic phase factors. *Physical Review A*, 59.

[227] Vaidman, L. (2013). Past of a quantum particle. *Phys. Rev. A*, 87:052104.

[228] van der Wiel, W. G., Franceschi, S. D., Elzerman, J. M., Fujisawa, T., Tarucha, S., and Kouwenhoven, L. P. (2002). Electron transport through double quantum dots. *Rev. Mod. Phys.*, 75:1.

[229] van Veldhoven, E., Sidorkin, V., Chen, P., Alkemade, P., van der Drift, E., Salemink, H., Zandbergen, H., and Maas, D. (2010). Nanofabrication with a high resolution helium ion beam. *Microscopy and Microanalysis*, 16(S2):202–203.

[230] van Wees, B. J., van Houten, H., Beenakker, C. W. J., Williamson, J. G., Kouwenhoven, L. P., van der Marel, D., and Foxon, C. T. (1988). Quantized conductance of point contacts in a two-dimensional electron gas. *Phys. Rev. Lett.*, 60:848.

[231] Veldhorst, M., Hwang, J. C. C., Yang, C. H., Leenstra, A. W., de Ronde, B., Dehollain, J. P., Muhonen, J. T., Hudson, F. E., Itoh, K. M., Morello, A., and Dzurak, A. S. (2014). An addressable quantum dot qubit with fault-tolerant control-fidelity. *Nature Nanotech.*, 9:981.

[232] Viola, L. and Lloyd, S. (1998). Dynamical suppression of decoherence in two-state quantum systems. *Phys. Rev. A*, 58:2733.

[233] Voisin, B., Nguyen, V.-H., Renard, J., Jehl, X., Barraud, S., Triozon, F., Vinet, M., Duchemin, I., Niquet, Y.-M., de Franceschi, S., and Sanquer, M. (2014). Few-electron edge-state quantum dots in a silicon nanowire field-effect transistor. *Nano Lett.*, 14:2094.

[234] Vrijen, R., Yablonovitch, E., Wang, K., Jiang, H. W., Balandin, A., Roychowdhury, V., Mor, T., and DiVincenzo, D. (2000). Electron-spin-resonance transistors for quantum computing in silicon-germanium heterostructures. *Physical Review A*, 62.

[235] Wei, D., Li, H.-O., Cao, G., Luo, G., Zheng, Z.-X., Tu, T., Xiao, M., Guo, G.-C., Jiang, H.-W., and Guo, G.-P. (2013). Copper nanowires as fully transparent conductive electrodes. *Sci. Rep.*, 3:2323.

[236] Wigner, E. (1932). On the quantum correction for thermodynamic equilibrium. *Phys. Rev.*, 40:749–759.

[237] Wikipedia. Bloch sphere image. https://en.wikipedia.org/wiki/Bloch_sphere Accessed: 2021-07-01.

[238] Wootters, W. K. (1987). A wigner-function formulation of finite-state quantum mechanics. *Ann. Phys.*, 176(1):1–21.

[239] Xu, L., Liu, Z., Datta, A., Knee, G. C., Lundeen, J. S., Lu, Y.-q., and Zhang, L. (2020). Approaching quantum-limited metrology with imperfect detectors by using weak-value amplification. *arXiv preprint arXiv:2005.03629*.

[240] Yunger Halpern, N. (2017). Jarzynski-like equality for the out-of-time-ordered correlator. *Phys. Rev. A*, 95:012120.

[241] Yunger Halpern, N., Beverland, M. E., and Kalev, A. (2020). Noncommuting conserved charges in quantum many-body thermalization. *Phys. Rev. E*, 101:042117.

[242] Yunger Halpern, N., Faist, P., Oppenheim, J., and Winter, A. (2016). Microcanonical and resource-theoretic derivations of the thermal state of a quantum system with noncommuting charges. *Nature Communications*, 7(1):12051.

[243] Yunger Halpern, N., Swingle, B., and Dressel, J. (2018). Quasiprobability behind the out-of-time-ordered correlator. *Phys. Rev. A*, 97:042105.

[244] Zener, C. (1932). Non-adiabatic crossing of energy levels. *Proc. R. Soc. London A*, 137:696.

[245] Zeng, J. and Barnes, E. (2018). Fastest pulses that implement dynamically corrected single-qubit phase gates. *Phys. Rev. A*, 98:012301.

[246] Zhao, R., Dudin, Y. O., Jenkins, S. D., Campbell, C. J., Matsukevich, D. N., Kennedy, T. A. B., and Kuzmich, A. (2009). Long-lived quantum memory. *Nature Physics*, 5(2):100–104.

[247] Zyvex. Zyvex website. http://www.zyvex.com Accessed: 2021-07-01.

Appendix A

GPU-accelerated staggered-leapfrog code

Here I present essential sections of the GPU-accelerated staggered-leapfrog code central to this work. For brevity, only the most interesting parts are shown, and are abridged with pseudocode. While not public at the moment, the full program is available on a reasonable request.

I Staggered-leapfrog CUDA kernel

The CUDA kernel of the C++ program, implementing the Staggered Leapfrog algorithm to time-evolve some initial wave function. The entire program is too large to include here.

```
__global__ void rSpaceSL(int N, float* PsiRe, float* PsiIm, float* V, bool* domain,
                         float t, int iterateRe, int* kernelInds, int Np, int N1, int dim,
                         float ASAW, float kSAW, float SAWoffset, float wSAW, float cSAW,
                         float RS, float DRt0, float URt0, float ATB2, float s2TB2,
                         float fbb, float* faa, int* Sizes, float* Mass, float* dXs, float Xa, float Ya) {

        int index = blockIdx.x * blockDim.x + threadIdx.x; // CUDA index

        if (index < N && domain[index]==1) { //check if index is within the wavefunction ant NOT on a boundary

                float Vt = 0.0; //some potentially time depentent part of V

                for (int p = 0; p < Np; p++) {
                        // Set current coordinates
                        int ind = ((index / (int)powf(N1, Np - p - 1)) % (N1)) * dim;
                        float x = Xa + kernelInds[ind] * dXs[0];
                        float y = dim > 1 ? Ya + kernelInds[ind + 1] * dXs[1] : 0.0;
                        // Add time-dependent SAW potential
                        Vt += ASAW * (1.0 - cosf(kSAW * (x - SAWoffset) - wSAW * t));
      // time-dependent ramp potential
                        float TBRamp = 0.5 * (tanhf(RS * (y + (t - DRt0) * cSAW)) + tanhf(-RS * (y + (t - URt0) * cSAW)));
                        Vt -= ATB2 * Mass[p] * expf(-s2TB2 * powf(x, 2.0) / 2.0) * TBRamp;
                }


                int idxNearestNeighbor=1;
                if (iterateRe) { // Iterate real part of the wave function
```

```
                        PsiRe[index] += (V[index] + Vt) * fbb * PsiIm[index]; // The potential part
                        for (int d = 0; d < dim; d++) {
                                for (int p = 0; p < Np; p++) {
                                        PsiRe[index] += (faa[d] / Mass[p]) * (2.0 * PsiIm[index]); // On-site kinetic term

                                        // Calculate the index offset of the nearest neighbor lattice site
                                        if (d == (dim - 1))
                                                idxNearestNeighbor = 1;
                                        else if (d == (dim - 2))
                                                idxNearestNeighbor = Sizes[dim - 1];
                                        else if (d == (dim - 3))
                                                idxNearestNeighbor = Sizes[dim - 1]* Sizes[dim - 2];

                                        int pp = Np - p -1;
                                        while (pp > 0) {
                                                idxNearestNeighbor *= N1;
                                                pp--;
                                        }

                                        PsiRe[index] -= (faa[d] / Mass[p]) * PsiIm[index - idxNearestNeighbor];
                                        PsiRe[index] -= (faa[d] / Mass[p]) * PsiIm[index + idxNearestNeighbor];

                                }
                        }
                }
                else { // Iterate imaginary part of the wave function
                        PsiIm[index] -= (V[index] + Vt) * fbb * PsiRe[index];
                        for (int d = 0; d < dim; d++) {
                                for (int p = 0; p < Np; p++) {
                                        PsiIm[index] -= (faa[d] / Mass[p]) * (2.0 * PsiRe[index]); // On-site kinetic term

                                        // Calculate the index offset of the nearest neighbor lattice site
                                        if (d == (dim - 1))
                                                idxNearestNeighbor = 1;
                                        else if (d == (dim - 2))
                                                idxNearestNeighbor = Sizes[dim - 1];
                                        else if (d == (dim - 3))
                                                idxNearestNeighbor = Sizes[dim - 1] * Sizes[dim - 2];

                                        int pp = Np - p -1;
                                        while (pp > 0) {
                                                idxNearestNeighbor *= N1;
                                                pp--;
                                        }

                                PsiIm[index] += (faa[d] / Mass[p]) * PsiRe[index - idxNearestNeighbor];
                                        PsiIm[index] += (faa[d] / Mass[p]) * PsiRe[index + idxNearestNeighbor];
                                }
                        }
                }
        }
}
```

II Eigensolver Code

A Python verison of the eigensolver code, using momentum space.

```
#########################################################################################################################################
# Real space eigensolver using the CPU
#########################################################################################################################################

import numpy as np
from src.export.exportEigenvalues import exportEigenvalues
from src.input.potentialKSpace import potentialKSpace
import src.index.index1DtoND as index1DtoND
import src.index.indexNDto1D as indexNDto1D
from src.get_ccbin import get_ccbin
import reikna.fft as rfft
```

```
import pycuda.gpuarray as gpuarray
from src.fft.fftAuxPsi import ToReal
from skcuda import linalg


def kSpaceSolverGPU(params, PsiOut, mod, thr): #Momentum-space eigensolver on a GPU

   #read input parameters from params (containting all user inputs) & declare variables

    KxCO = params["KxECO"] #momentun cuttoffs in each dimension - total number of points is 2*cutoff + 1 for k=0
    KyCO = params["KyECO"]
    KzCO = params["KzECO"]

    Kx = 2 * KxCO + 1
    Ky = 2 * KyCO + 1
    Kz = 2 * KzCO + 1

    dKx = params["dKx"] #step in momentum in each dimension
    dKy = params["dKy"]
    dKz = params["dKz"]

    Np = params["Np"] #Number of pariticles
    dim = params["dim"] #Number of dimensions
    PsiSize = int((Kx * Ky * Kz) ** Np) #WAve function size

    SizesK = []
    SizesKE = []
    KCOs = []
    dKs = []

    for i in range(0,dim):
        if (i == 0):
            SizesK.append(Kx)
            SizesKE.append(4 * KxCO + 1)
            KCOs.append(KxCO)
            dKs.append(dKx)
        elif (i ==1):
            SizesK.append(Ky)
            SizesKE.append(4 * KyCO + 1)
            KCOs.append(KyCO)
            dKs.append(dKy)
        elif (i ==2):
            SizesK.append(Kz)
            SizesKE.append(4 * KzCO + 1)
            KCOs.append(KzCO)
            dKs.append(dKz)

   #generate Hamiltonian matrix for eigenfunction solving

   #Init Hamiltonian
    H = np.zeros(PsiSize*PsiSize, dtype = np.complex64)

    Vk = potentialKSpace(params, mod, thr, True) # Generate the static potential for the problem,
    defined by the user before, and transform it to momentum space

    indsCol = np.zeros(dim*Np, dtype=np.int32)
    indsRow = np.zeros(dim*Np, dtype=np.int32)
    indsDiff = np.zeros(dim*Np, dtype=np.int32)

    for k in range(0, PsiSize**2):# Go over each row

        j = int(k % PsiSize) # column index
        i = int(np.floor(np.float64(k) / np.float64(PsiSize))) #i is row index

        if (j >= i):
            index1DtoND.index(j, indsCol, SizesK, params)
            index1DtoND.index(i, indsRow, SizesK, params)
            #calulate offset from diagonal
            for l in range(0, dim * Np):
                indsDiff[l] = indsCol[l] - indsRow[l] if  (indsCol[l] - indsRow[l]) >= 0 else (4 * KCOs[l % dim] + 1)
                + (indsCol[l] - indsRow[l]) #((indsCol[l] - indsRow[l]) >= 0 ?
                 (indsCol[l] - indsRow[l]) : (4 * KCOs[l % dim] + 1) + (indsCol[l] - indsRow[l]))


                #fill Hamiltonian, making use of hermiticity to only fill half and do conjugate
```

```
            H[j + PsiSize * i] = Vk.pot[indexNDto1D.index(indsDiff, SizesKE, params)]


            if (i == j): #on - diagonal element
                for p in range(0,Np):
                    for d in range(0,dim):
                        H[j + PsiSize * i] +=  (((indsRow[p * dim + d] - KCOs[d]) * dKs[d]) ** 2) / (2 * params["Mass"][p])
            else:
                H[i + PsiSize * j] = H[j + PsiSize * i].conjugate()


                #on-site potential

    H=np.reshape(H.ravel(), [PsiSize,PsiSize], order='F')
    H_gpu = gpuarray.to_gpu(H)
    D=np.zeros(PsiSize, dtype=np.float32)

    linalg.init()
    V_gpu, D_gpu = linalg.eig(H_gpu, 'N', 'V')


    D_gpu.get(D)
    V_gpu.get(H)
    H=np.reshape(H.ravel(), [PsiSize,PsiSize], order='C')


    for i in range(0,5):
        print(D[i])


    ccbin = get_ccbin()
    PsiTemp = np.zeros(PsiOut.size, dtype = np.complex64)
    PsiTemp_gpu = gpuarray.to_gpu(PsiTemp)

    fft_h = rfft.FFT(PsiTemp_gpu)
    fftc=fft_h.compile(thr, compiler_options = ['-ccbin', ccbin])

        # Go over eigenstates to construct the init wave function as specified by Input Params
    for J in range(0,5):
        if ((params["InputStates"][2 * J] != 0.0) or (params["InputStates"][2 * J + 1] != 0.0)):
            #PsiTemp = PsiTemp.ravel();
            indsR = np.zeros(dim*Np)
            indsK = np.zeros(dim*Np)

            for i in range(0, (Kx*Ky*Kz)**Np):
                index1DtoND.index(i, indsK, SizesK, params)
                for j in range (0, dim*Np):
                    indsR[j] = indsK[j] + params["Sizes"][j % dim] - KCOs[j % dim] if indsK[j] < KCOs[j % dim] else
                    indsK[j] - KCOs[j % dim] #indsR[j] = indsK[j] < KCOs[j % params->dim] ?
                    indsK[j] + params->Sizes[j % params->dim] - KCOs[j % params->dim] : indsK[j] - KCOs[j % params->dim];

                indNP = indexNDto1D.index(indsR, params["Sizes"], params)
                PsiTemp.flat[int(indNP)] = H[i,J]

            #ToReal(params, PsiTemp, False)
            #PsiTemp.reshape(PsiOut.size)
            PsiTemp_gpu.set(PsiTemp)
            fftc(PsiTemp_gpu, PsiTemp_gpu, 1)
            thr.synchronize()
            PsiTemp_gpu.get(PsiTemp)
            ToReal(params, PsiTemp, False)


            if (params["TDSESolver"] == "kSpace"):
                PsiOut.wave += (params["InputStates"][2 * J] + 1j * params["InputStates"][2 * J + 1]) * PsiTemp
            else:
                PsiOut.wave += (params["InputStates"][2 * J] + 1j * params["InputStates"][2 * J + 1]) * PsiTemp.ravel()



    if (params["TDSESolver"] == "rSpace"):              # do a final normalisation
        norm = PsiOut.norm()
        PsiOut.wave /= np.sqrt(norm)
        # Write wave and eigenvalues to file as appropriate
        PsiOut.exportFull( params, -1, "wave")
```

```
elif (params["TDSESolver"] == "kSpace"):
    PsiOut.exportFull( params, -1, "wave")
    PsiOut.wave = np.reshape(PsiOut.wave, PsiOut.size)
    PsiOut.htod()
    PsiOut.fft(thr, 0)
    PsiOut.dtoh()
    norm = PsiOut.norm()
    PsiOut.wave /= np.sqrt(norm)
    PsiOut.exportFull( params, -1, "Kwave")


exportEigenvalues(params, D, 5) #Write eigenvalues to file
```

Appendix B

Finding optimal adjustment parameters accounting for rise time τ

Here I present the pseudocode for finding the optimal adjustment parameters accounting for rise time τ for single qubit control, using a gradient ascent method. This supplements the results of Ch. 3, as it gives a general method for finding these parameters for any rise time and required total angle of rotation. As the method is numeric, complete analytical results for any input are not possible to present. However, this code should enable anyone to apply the control methods found in Ch. 3 in a general case.

I Gradient-ascent search pseudocode

The pseudocode for finding the optimal adjustment parameters accounting for rise time τ for single qubit control. The code is written in MATLAB, with the time-dependent evolution relegated to the GPU-accelerated staggered-leapfrog code described in the main work and in App. A.

```
v_si2nu = 1.312342066e-2; % conversion factor from natural units to nm
t_si2nu = 1.157676458e2; %conversion factor from natural units to ps

%% Generate potential & Find initial state
Xa = -230; Xb = 230; Lx = Xb-Xa;  % Define lattice length scales
Ya=-100; Yb=100; Ly=Yb-Ya;

Nx = 201; % Number of lattice points


x = linspace(Xa,Xb,Nx); % 1D Lattice

dx = x(2) - x(1); % Lattice spacing


hbar=1;
m=0.067;%Effective mass in GaAs
```

```
w1 = 130; % (nm)
w2 = 240; % (nm)


%Construct the Doubel DOt potential, with a harmonic and central Gaussian sub-potentials
aHarmX=0.5*m*0.00005;                          %harmonic X potential strength

a1GaussX=(0.8)*m;          %1st gaussian pot in X strength
width1GaussX=0.00015;

Vjm = zeros(Nx,Ny);
for i = 1:length(Vjm)
    Vjm(i)=aHarmX*x(i)^2 + a1GaussX * exp(-width1GaussX*(x(i))^2/2);
end

V = Vjm ;


[eV,D]=eig(H); %Eigensolve the Hamiltonian here
D = diag(D);


psi0; %Get the ground state
psi1; %Get the first excited state


% Solve the TDSE


dE = D(2) - D(1); %Energy difference between the states

epsilonC=0.42066 ; %Epsilon factor of the qubit

Vpulse0 = -*dE/epsilonC; %Reference pulse strength

Trot=2*pi/(sqrt(dE^2+(Vpulse0*epsilonC)^2)); %Time for a full rotation under the reference pulse
Twait=2*pi/dE; %Time for a full rotation without pulsing

gamma2=1; %Set parameters controlling the step size in pulse amplitude and duration
refStep=[0.01 15*Trot];
probeStep=[gamma2*0.01 gamma2*0.01*Trot];
gamma=20;

Vtheta=[pi/6, pi/4, pi/2]; %List of angles for which we want to find the parameters
Ksteps=length(Vtheta);


%List of fidelity, amplitude and duration adjustments
KO=zeros(1,Ksteps);
KA=zeros(1,Ksteps);
Kt=zeros(1,Ksteps);

for K=1:Ksteps

        %initialise variables
        Oprev=0;
        OAp1=0;
        OAm1=0;
        Otp1=0;
        Otm1=0;

        Amod=0;
        Tmod=0;
        theta=Vtheta(K);%Angle being optimised in this loop

        tauFrac=4;%Rise time parameter, as defined in the main text
        tau1=Trot*(1/(tauFrac*2)); %actual rise time value

        N=floor((K/Ksteps)*1000000) + 200000 ;% Number of iterations of the time-dependent code

        Vmod0=[1 0]; %initial guess of adjustment parameters. Here initialised to square pulse values as a first guess
        Vmod=Vmod0;

        steps=200; %Max gradient ascent search steps
```

```
        %Keep track of the progress of adjustment parameters and fidelity
        VmodHistory=zeros(steps+1,2);
        OHistory=zeros(1,steps+1);
        VmodHistory(1,1)=Vmod0(1);
        VmodHistory(1,2)=Vmod0(2);
        OHistory(1)=KO(K);

%Step through the gradient ascent, break if good enough fidelity is reached, increase step size parameters on the way
        for X=1:steps
                if (X==1 && (1-KO(K))<1e-6)
                                break
                end
                        if (X~=1 && (1-O)>0.1)
                        break
                end

                if X<=50
                        gamma=25;
                        gamma2=1;
                elseif X<=100
                        gamma=50;
                        gamma2=0.1;
                else
                        gamma=75;
                        gamma2=0.01;
                end

    %The Gdescent() function usis the GPU-accelerated code to time-evolve the qubit system with given pulse parameters, and returns the final
    %fidelity. This is the quantity we want to maximise. The system is initialised in the [1,1] equal superposition of the ground and
    %first excited states, and then evolved under the adjusted pulse twice - we have found that
    %this setup produces very good results for all cases

                O=Gdescent(tau1,theta,Vmod,dE,V,psi0,psi1,N,steepenPow);% calculate fidelity for current adjustment parameters value


%Save values used and fidelity received
                VmodHistory(X+1,1)=Vmod(1);
                VmodHistory(X+1,2)=Vmod(2);
                OHistory(X+1)=O;


%Calculate the fidelity after taking a step in both directions in pulse amplitude and duration - this calculates the effective gradient
% in fidelity as a function of adjustment parameters
                OAp1=Gdescent(tau1,theta,[Vmod(1)+probeStep(1) Vmod(2)],dE,V,psi0,psi1,N,steepenPow);
                OAm1=Gdescent(tau1,theta,[Vmod(1)-probeStep(1) Vmod(2)],dE,V,psi0,psi1,N,steepenPow);
                Otp1=Gdescent(tau1,theta,[Vmod(1) Vmod(2)+probeStep(2)],dE,V,psi0,psi1,N,steepenPow);
                Otm1=Gdescent(tau1,theta,[Vmod(1) Vmod(2)-probeStep(2)],dE,V,psi0,psi1,N,steepenPow);

%Update the adjustment parameters based on the gradient calculated
                Vmod=Vmod + stepDecay*gamma*[0.5*(OAp1-OAm1)*refStep(1)/probeStep(1) 0.5*(Otp1-Otm1)*refStep(2)/probeStep(2)];

        end

%Save the maximum fidelity obtained - best results
        [Omax,Imax]=max(OHistory);

        KO(K)=Omax;
        KA(K)=VmodHistory(Imax,1);
        Kt(K)=VmodHistory(Imax,2);
        KTTotal(K)=Trot*(theta/(2*pi))+Kt(K);

end

%Save relevant final outcomes to file for further use
save('descSweepBloch.mat', 'KO', 'KA', 'Kt', 'KTTotal','Vtheta');
```