

SEMICONDUCTOR-BASED QUANTUM DEVICES:

MATERIALS, DESIGN AND OPERATION

*by*

Adam Frees

A dissertation presented in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Physics)

*at the*

UNIVERSITY OF WISCONSIN – MADISON

2019

Date of final oral examination: 23 April 2019

The dissertation is approved by the following members of the Final Oral Committee:

Susan N. Coppersmith · Professor of Physics

Mark G. Friesen · Senior Scientist of Physics

Mark A. Eriksson · Professor of Physics

Jason Kawasaki · Assistant Professor of Materials Science and Engineering

# Abstract

This thesis describes work concerning the performance of semiconductor-based devices intended for quantum computation. We begin by focusing on the materials used to fabricate semiconductor-based qubit devices. Specifically, we study the strain induced on the semiconductor domain of a simulated device due to the cooling of Al electrodes in the vicinity of the studied domain. We introduce the theory behind stress and strain, as well as how these quantities can affect the electronic properties of a semiconductor. We then present the results of a simulation modeling the expected strain in a simulated device, and find the corresponding shift in the conduction band of the region intended to contain the quantum dots.

We next consider the design of semiconductor-based qubit devices. We introduce a protocol, which we call the Compressed Optimization of Device Architectures (CODA). This protocol both efficiently identifies sparse sets of voltage changes that control quantum systems, and introduces a metric which can be used to compare device designs. As an example of the former, we apply this method to simulated devices with up to 100 quantum dots and show that CODA automatically tunes devices more efficiently than other common nonlinear optimizers. To demonstrate the latter, we determine the optimal lateral scale for a triple quantum dot, yielding a simulated device that can be tuned with small voltage changes on a limited number of electrodes.

Finally, we study the operation of semiconductor-based qubit devices. We examine this subject in two ways. First, we consider a system comprised of two capacitively-coupled

quantum dot hybrid qubits, and develop a simple electronic pulse sequence that achieves a high-fidelity two-qubit gate under realistic levels of quasistatic charge noise. Then, we consider the performance of a quantum algorithm for PageRank which takes as input a graph of the World Wide Web. We show that the run-time of this algorithm depends on features of the graphs other than the degree distribution, and can be altered sufficiently to rule out a general exponential speedup.

I would like to acknowledge support from the Vannevar Bush Faculty Fellowship program sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering and funded by the Office of Naval Research through Grant No. N00014-15-1-0029.

## Acknowledgements

First and foremost, I would like to thank my advisors, Prof. Susan Coppersmith and Dr. Mark Friesen, without whom none of this work would be possible. Throughout my graduate career, both Sue and Mark have been endless sources of guidance and wisdom. Their advice has shaped me in ways too numerous to fit in this section. In particular, I would like to acknowledge their dedication to teaching me how to be a more effective communicator, a skill which I value deeply and will continue to use in all of my endeavors. I would also like to thank them for teaching me valuable problem-solving skills: Sue for demonstrating how to gain intuition for the problem at hand, and Mark for instilling in me the importance of rigorously testing any possible solution thereof. It is only with their advice that I have been able to grow as a student, as a communicator, and as a scientist.

I would like to thank Prof. Mark Eriksson, whose experimental insights have played a major role in the direction of multiple projects throughout my graduate career. Mark's positivity and knack for explaining complex topics in simple terms are a constant inspiration to me.

I would also like to extend my sincere gratitude to Dr. John Gamble. When I first started working for Sue and Mark, I was an eager, but naive undergrad. John, who at the time was a graduate student, took me under his wing; in many ways, I never left. John continued to be an invaluable resource long after he left Wisconsin, always providing clever insights into challenging problems. John also acted as a mentor through internships at Sandia National Laboratories and at Microsoft, both of which were amazing experiences for me, professionally

and personally.

I thank all of my collaborators for all of their contributions to this work and to my education: Dr. Kenneth Rudinger, Prof. Eric Bach, Prof. Robert Joynt, Dr. Daniel R. Ward, Dr. Robin Blume-Kohout, and Dr. Sebastian Mehl. I would also like to thank the members of the Wisconsin “silicon qubits” group for their insights and their camaraderie: Dr. Don Savage, Dr. Ryan Foote, Dr. Trevor Knapp, Dr. Cameron King, Yuan-Chi Yang, Zhenyi Qi, Ekmel Ercan, Leah Tom, Vickram Premakumar, Brandur Thorgrimsson, Sam Neyens, Nathan Holman, J.P. Dodson, Tom McJunkin, and Joelle Baer.

I thank my thesis committee, Prof. Susan Coppersmith, Dr. Mark Friesen, Prof. Mark Eriksson, and Prof. Jason Kawasaki for the time that they have dedicated to reading and reviewing this dissertation.

And finally, I would like to thank my family. To my mom and dad, thank you both for always supporting me and instilling in me the importance of education. I cannot describe how lucky I feel to have such amazing parents. To Grandma Jeanne, thank you for your continued interest in quantum computing – I hope that this dissertation clears up any lingering questions you may have! To my brother Nathan, thank you for reminding the importance of owning property. And to my wife Claire, whose love and support I cherish every day. Thank you for all that you have done to help make this possible.

*If I have seen further than others, it is by standing upon the shoulders of giants.*  
— *Sir Isaac Newton*

*WOOF. WOOF. WOOF.*  
— *Newton, my dog*

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Why quantum computing? . . . . .	2
1.1.1	Qubits . . . . .	3
1.1.2	Circuit-based quantum computing . . . . .	5
1.1.3	Adiabatic quantum computing . . . . .	7
1.2	Why semiconductors? . . . . .	8
1.2.1	Conduction-band electrons in Si . . . . .	9
1.2.2	Creating and controlling quantum dots . . . . .	11
1.3	Thesis outline . . . . .	12
1.4	Publications . . . . .	14
<b>2</b>	<b>Strain in semiconductor quantum devices</b>	<b>17</b>
2.1	Stress . . . . .	18
2.2	Strain . . . . .	18
2.2.1	Mechanical strain . . . . .	19

2.2.2	Thermal strain . . . . .	20
2.3	Effect of strain on band structure . . . . .	21
2.3.1	$\Gamma_1$ strain . . . . .	23
2.3.2	$\Gamma_3$ strain . . . . .	24
2.3.3	$\Gamma_4$ strain . . . . .	25
2.4	Finite-element modeling of stress and strain . . . . .	26
2.5	Acknowledgments . . . . .	30
<b>3</b>	<b>Compressed Optimization of Device Architectures for Semiconductor Quantum Devices</b>	<b>31</b>
3.1	Introduction . . . . .	32
3.2	Automatic tuning of simulated devices . . . . .	34
3.3	Extensibility . . . . .	39
3.4	Device design optimization . . . . .	44
3.5	Conclusion . . . . .	47
3.6	Funding acknowledgments . . . . .	48
<b>4</b>	<b>Adiabatic two-qubit gates in capacitively coupled quantum dot hybrid qubits</b>	<b>50</b>
4.1	Introduction . . . . .	51
4.2	Results . . . . .	53
4.3	Discussion . . . . .	64



4.4	Methods . . . . .	65
4.5	Funding acknowledgements . . . . .	67
<b>5</b>	<b>Power law scaling for the adiabatic algorithm for search engine ranking</b>	<b>69</b>
5.1	Introduction . . . . .	70
5.2	Network growth models . . . . .	72
5.3	Algorithm description . . . . .	75
5.4	Numerical results . . . . .	79
5.5	Discussion . . . . .	80
5.6	Funding acknowledgements . . . . .	81
<b>6</b>	<b>Conclusion</b>	<b>83</b>
<b>A</b>	<b>Supplemental information for Ch. 3</b>	<b>85</b>
A.1	CODA protocol . . . . .	85
A.2	Simulation details . . . . .	87
A.3	Tunnel rates and transmission coefficients . . . . .	89
<b>B</b>	<b>Supplemental information for Ch. 4</b>	<b>90</b>
B.1	Derivation of an effective Hamiltonian for two coupled QDHQS . . . . .	90
B.2	Optimizing single qubit dispersions . . . . .	94
B.3	Pulse sequences . . . . .	96
B.4	Optimization of detuning pulse parameters . . . . .	98

B.5	Process fidelity . . . . .	100
B.6	Dependence of the fidelity on $g$ and $\Delta E_{ST}$ . . . . .	101
B.7	Approximate formula for the infidelity . . . . .	105
B.8	Dynamical sweet spot analysis for the tunnel coupling pulse sequence . . . . .	106
B.9	Relaxation time estimate . . . . .	107
<b>C</b>	<b>Supplemental information for Ch. 5</b>	<b>112</b>
C.1	Parameters of web graph models . . . . .	112
C.1.1	GZL preferential attachment . . . . .	112
C.1.2	GZL copying model . . . . .	115
C.1.3	$\alpha$ -preferential attachment . . . . .	116
C.2	Initial conditions . . . . .	117
C.3	Adaptive binning . . . . .	117

# List of Figures

1.1	Diagram of Bloch sphere . . . . .	4
1.2	Electronic band structure in Si . . . . .	9
2.1	Conduction band shift due to Al electrodes . . . . .	27
3.1	Example implementations of the CODA protocol . . . . .	35
3.2	Demonstrating the extensibility of the CODA protocol using a simple model involving up to 100 quantum dots . . . . .	41
3.3	Protocol used to compare the “tunability” of device designs . . . . .	44
3.4	Using CODA to optimize triple quantum dot designs . . . . .	45
4.1	Implementing a CZ gate between capacitively coupled QDHQs . . . . .	54
4.2	Optimized infidelities of sub-50 ns adiabatic CZ gates . . . . .	55
4.3	Dynamical sweet spot (DSS) and a tunnel-coupling pulse sequence . . . . .	62
5.1	Illustrations of the three network generation models used . . . . .	73
5.2	Comparison of the scaling of the inverse energy gap $\delta^{-1}$ . . . . .	77

5.3	Inverse energy gap scaling of WWW-like networks . . . . .	78
B.1	Energy spectrum of two capacitively-coupled QDHQs as a function of qubit detunings . . . . .	93
B.2	Adjusting tunnel couplings to optimize single qubit dispersions . . . . .	95
B.3	Key steps in the detuning-pulse optimization procedure . . . . .	100
B.4	Optimized process infidelities of adiabatic CZ gates as a function of $\Delta E_{ST}$ and $g$	101
B.5	Adiabatic ramp time diverges for small $\Delta E_{ST}$ . . . . .	103
B.6	Pulse sequences involving both the tunnel coupling and detuning parameters are able to improve the fidelity of a CZ gate . . . . .	104
B.7	Estimated error due to relaxation is smaller than error due to dephasing . .	109
C.1	Degree distributions for the GZL preferential attachment model . . . . .	114

# List of Tables

2.1	Materials parameters used in strain simulations. . . . .	28
B.1	Tunnel coupling pulse sequence parameters . . . . .	97

# Chapter 1

## Introduction

Forget cryptocurrency. Quantum energy is the future. It's the next gold rush.

— Sonny Burch (Walton Goggins), *Ant-Man and the Wasp*

The field of quantum mechanics has captured the imagination of the public for nearly a century. Since its inception in the early 1900s, philosophers, artists, and even politicians have pondered the implications of quantum superposition and entanglement. Quantum mechanics has found its place in popular culture as well. In comedies, a character may be hit in the head and suddenly understand the intricacies of quantum entanglement; in science fiction stories, quantum technologies are often used to explain how the hero can save the day at the last minute.

It is therefore perhaps unsurprising that the prospect of a “quantum computer” should engender excitement in the public. A machine that leverages the perhaps mysterious effects of quantum mechanics to solve hard problems sounds like a plot contrivance from a bad sci-

fi film. In reality, a quantum computer cannot transport its user across time and space; it cannot create a bridge into a parallel universe, nor is it a source of unlimited energy. While the true advantages of quantum computing may not be cinematic, they are nonetheless incredibly exciting.

## 1.1 Why quantum computing?

It is believed that certain computational problems can be solved more efficiently with a quantum computer than a classical computer. One such problem is determining the prime factors  $p$  and  $q$  given only their product  $N = p \cdot q$ . While this problem may sound esoteric, it is the foundation of the widely-used RSA public-key cryptoscheme [1]. In Ref. [2], Peter Shor proposes a quantum algorithm whose runtime scales polynomially with the number of digits of  $N$ : a super-polynomial speed-up over the best-known classical algorithm.

While any problem solved by a quantum computer can in principle also be solved using a classical computer, the time and energy required to do so may be prohibitive. It is estimated that it would take a modern classical computer centuries to find the prime factors of a 1000-bit number  $N$  [3]. In contrast to this, a quantum computer with a clock rate of more than 1 MHz is expected to be able to solve this problem in less than a day. Even if advances in classical computing technology were to yield a computer that could factor a 1000-bit number in a day, it would take this machine tens of thousands of years to factor a 2000-bit number; the comparable quantum computer could factor a 2000-bit number in less than a month [3].

The prospect of a quantum speedup motivates the interest in quantum computers. In this section, I introduce the fundamental building blocks of quantum computers: qubits. Then I introduce two potential quantum computing architectures, circuit-based and adiabatic quantum computation, and compare the two approaches. Good introductions of this material with further details are given in Refs. [4, 5, 6].

### 1.1.1 Qubits

In the same way that a classical computer is comprised of bits, variables that can either take on the value 0 or 1, a quantum computer is comprised of quantum bits, or *qubits*. A qubit is a two-level quantum system which can exist in the  $|0\rangle$  or  $|1\rangle$  state. In contrast to a classical bit, a qubit can also exist in a quantum superposition between the  $|0\rangle$  and  $|1\rangle$  states. In general, the state of a single qubit  $|\psi\rangle$  can be represented as

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \tag{1.1}$$

where  $\alpha$  and  $\beta$  are complex numbers such that  $|\alpha|^2 + |\beta|^2 = 1$ .  $|\psi\rangle$  is typically represented using the Bloch sphere, shown in Fig. 1.1.

While a qubit can exist anywhere on the surface of the Bloch sphere, its exact state cannot be measured directly. A measurement of a qubit is performed along an axis of the Bloch sphere, and can yield one of two possible outcomes. For instance, a qubit measured along the  $z$ -axis has the possible measurement outcomes “0” and “1.” In this case, a qubit






Figure 1.1: Diagram of Bloch sphere. A coherent qubit state  $|\psi\rangle$  is represented as a point on the surface of the Bloch sphere; the  $|0\rangle$  and  $|1\rangle$  states are represented as the “north” and “south poles” of the sphere respectively. An arbitrary coherent state  $|\psi\rangle$  can be fully parameterized via the polar angle  $\theta$  and azimuthal angle  $\phi$ :  $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\phi}\sin(\theta/2)|1\rangle$ .

with the state  $\psi$  from Eq. (1.1) has a probability  $|\alpha|^2$  of yielding a “0” measurement and a probability  $|\beta|^2$  of yielding “1.” After this measurement, the state of the qubit *collapses*: if a “0” is measured, the qubit’s state becomes  $|\psi\rangle = |0\rangle$ ; if a “1” is measured, the qubit’s state becomes  $|\psi\rangle = |1\rangle$ . The concept of measurements affecting the qubit state is an important distinction between classical and quantum computing.

Another crucial difference between qubits and classical bits is that qubits can be entangled with one another. For instance, a two-qubit system can be placed into the state  $|\psi\rangle = 1/\sqrt{2}(|00\rangle + |11\rangle)$ . In this state, a measurement performed on one of the qubits has an equal probability of yielding a “0” or a “1.” However, once one of the qubits has been measured, the state of the second qubit collapses; the result of measuring the second qubit will always be the same as the first. Quantum entanglement is at the core of many quantum

algorithms and is a crucial part of quantum computation [7, 8].

### 1.1.2 Circuit-based quantum computing

While every proposed quantum computer is comprised of qubits, there are multiple different proposals for the architecture of the computer. The most widely-studied proposed architecture is *circuit-based quantum computing*. The circuit-based quantum computing architecture draws inspiration from classical computers, in which logical gates act on classical bits. Similarly, algorithms performed on circuit-based quantum computers are represented in terms of logical gates acting on qubits.

While in principle one could conceive of quantum gates that act on any number of qubits, any algorithm in a circuit-based quantum computer can be implemented using a combination of single- and two-qubit gates [9]. A single-qubit gate can be visualized as the rotation of a state around the Bloch sphere. For instance, a Pauli-X gate (sometimes referred to as a *bit-flip* gate) maps  $|0\rangle$  to  $|1\rangle$  and vice-versa; this is a rotation of  $\pi$  radians around the  $x$ -axis of the Bloch sphere [5]. While the Pauli-X gate is a clear analogue to the NOT gate in classical computing, in general a single-qubit gate can be any rotation around the Bloch sphere, many of which do not have an obvious counterpart in classical computing. Two-qubit gates are required to construct entangled states such as the previously-mentioned Bell state  $|\psi\rangle = 1/\sqrt{2}(|00\rangle + |11\rangle)$ . A commonly considered two-qubit gate is the controlled-Z (CZ)

gate, which in the  $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$  basis is represented as the unitary

$$CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad (1.2)$$

The CZ gate can be understood as an operation in which a “target” qubit undergoes a Z-gate if the “control” qubit is in the  $|1\rangle$  state (otherwise the “target” qubit undergoes an identity operation). While there are other common two-qubit gates (SWAP gate, CNOT gate, *etc.*), all of these are equivalent to the CZ gate up to single-qubit operations [5].

Just as with classical computers, gates performed on quantum computers are not always perfect. The success rate of a quantum gate is quantified in terms of the *gate fidelity*, which ranges from 1 (a perfectly successful gate) to 0. The error rate of a quantum computer can be suppressed by encoding several *physical* qubits into one *logical* qubit. There is a veritable zoo [10, 11, 12, 13] of proposed quantum error correction techniques that leverage this idea into making a fault-tolerant quantum computer. All quantum error correction schemes have an associated *error threshold*: a maximum gate infidelity required for the scheme to be able to reduce error [14, 15]. Typical error thresholds for modern correction schemes are around 1% [13], making this an important benchmark in assessing the accuracy of a quantum gate. In Chapter 4, I propose a method to improve the fidelity of a two-qubit gate performed in a particular physical system.

### 1.1.3 Adiabatic quantum computing

A separate method for achieving quantum computation is *adiabatic quantum computing*. In quantum mechanics, an adiabatic process is one which happens slowly enough such that the system remains in its instantaneous eigenstate [16]. For instance, if a system begins in the ground state of an initial Hamiltonian, at the end of an adiabatic process the system will be in the ground state of the final Hamiltonian. This process can be leveraged for quantum computation: one can slowly perturb a system from a simple initial Hamiltonian into a final Hamiltonian for which the ground state is the solution to the problem of interest [17]. The efficiency of this adiabatic quantum algorithm is determined by the speed at which one can adiabatically perturb the system, which is limited by the minimum energy gap between the ground state and the first excited states [18]. In Chapter 5, I analyze the efficiency of an adiabatic quantum algorithm by considering the scaling of this energy gap with the problem size.

While both circuit-based and adiabatic quantum computing are composed of qubits, there are many differences between these approaches. Crucially, while there are several proposals [10, 11, 12, 13] for methods to construct a fault-tolerant circuit-based quantum computer, true fault-tolerance for adiabatic quantum computers remains elusive [19, 20]. For this reason, the majority of experimental quantum computing research is currently focused on circuit-based implementations. While adiabatic quantum computing is not the focus of many experimental efforts, it is worth noting that it has been shown that adiabatic quantum algorithms have been shown [21] to be polynomially equivalent to circuit-based quantum

algorithms. In other words, the existence of an adiabatic quantum algorithm which presents an exponential speedup over the best known classical algorithm does implies the existence of a similarly efficient circuit-based quantum algorithm, and vice-versa.

## 1.2 Why semiconductors?

In the previous section, I introduced the fundamental building blocks of quantum computing, qubits, and gave a brief overview of how one can leverage these qubits into performing quantum computation. In that description, I did not address how qubits are physically implemented. There are multiple parallel research efforts that are working towards building a large-scale quantum computer, such as superconducting circuits [22], trapped ions [23], and Majorana-based approaches [24]. Here, I focus on electrically-gated semiconductor-based qubits [25].

Semiconductor-based qubits are especially attractive due to their potential for scalability. The devices used to realize these qubits (the details of which will be introduced in this section) bear a striking resemblance to transistors used in classical computation. This similarity implies that the knowledge gained over several decades of fabricating large-scale semiconductor devices may be leveraged to help build large-scale quantum computers. While the systems currently achieved by researchers still consist of a small number of qubits, the promise of scalability makes this approach competitive with approaches that have achieved systems with tens of qubits [26, 27].

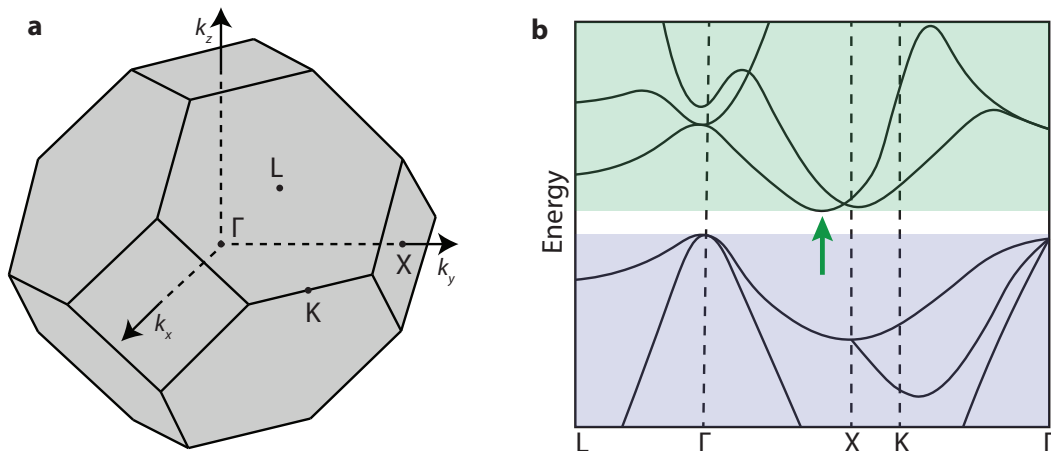


Figure 1.2: Electronic band structure in Si. **a**, The Brillouin zone for Si. The high-symmetry points  $L$ ,  $\Gamma$ ,  $X$ , and  $K$  are labeled according to convention. **b**, Schematic of Si band structure. The two domains of allowed energies shown are referred to as the conduction band (shaded green) and the valence band (shaded blue). The white region in between is referred to as the band gap. The minimum of the conduction band (indicated with green arrow) is located between the  $\Gamma$  point and the  $X$  point.

All semiconductor-based qubits rely on confining and manipulating a small number of charge carriers in a semiconductor; here, I focus on qubits constructed from conduction-band electrons in Si. In this section, I first review the properties of conduction-band electrons in Si, then I explain how these electrons are confined and manipulated to form qubits.

### 1.2.1 Conduction-band electrons in Si

The properties of conduction-band electrons in semiconductor lattices are best understood in the context of the periodic structure of the lattice. In an ideal semiconductor lattice, an electron’s wavefunction must be periodic in momentum-space as well as real-space. The primitive cell in momentum-space is called the *Brillouin zone*. The Brillouin zone for “diamond cubic” lattices (the lattice structure for Si) is shown in Fig. 1.2a. For each value of  $k$

in the first Brillouin zone, there are multiple allowed electron energies, as shown in Fig. 1.2b. A corollary to this is that an electron in a lattice has ranges of allowed energies; these energy ranges are referred to as *bands*. In semiconductors, the Fermi level is located in between two bands, in what is called a *band gap*. The energy band directly below the Fermi level (*i.e.*, the highest energy band populated at zero temperature) is called the *valence band*, and the band directly above this is called the *conduction band*. In Fig. 1.2b, the conduction band is shaded green, the valence band is shaded blue, and the band gap is white. The bands of a material can be shifted by applying an external electrical potential, a process referred to as *band bending*. In semiconductors, the band gap between the valence and conduction bands is on the order of an eV; this means that it is practical to apply a potential which shifts the bottom of the conduction band below the Fermi level. In this situation, the semiconductor becomes conductive, with the charge carriers located in momentum-space at the conduction band minimum, indicated with green arrow in Fig. 1.2b.

Many electronic properties of semiconductors are determined by the conduction band minimum. For instance, the *effective mass* of conduction-band electrons is determined by the curvature of the conduction band around its minimum. The location of the minimum in momentum space can also impact the properties of the charge carriers in a semiconductor. For instance, in Si the conduction band minimum occurs 85% of the way in between the  $\Gamma$  point and the  $X$  point, as shown in Fig. 1.2b. Due to the symmetry of the Brillouin zone, this means that there are six degenerate conduction band minima (denoted as  $\pm x$ ,  $\pm y$ , and  $\pm z$ ), referred to as *valleys*. This degeneracy can be broken by strain, as discussed in

Chapter 2. Both the effective mass and the valley degeneracy must be taken into account when performing device simulations of semiconductor-based quantum devices, as in Chapter 3.

### 1.2.2 Creating and controlling quantum dots

In order to manipulate a small number of conduction-band electrons, these electrons must be confined to a small volume within the device. While some semiconductor-based qubits use electrons confined to donor atoms within the lattice [28], here I focus on electrons confined electrostatically in *quantum dots* (QDs).

While electrons in QDs are confined in all three dimensions, the confinement mechanism in these directions are different: the confinement in the  $z$  direction is provided by the heterostructure, while the confinement in the  $x$  and  $y$  directions is provided via voltages on metallic electrodes. There are several heterostructures used for semiconductor-based qubit devices – a common heterostructure, and the focus of this work, is a Si/SiGe heterostructure. This arrangement consists of a thin ( $\sim 10$  nm) layer of strained Si in between two layers of relaxed SiGe, where “SiGe” is an alloy consisting of  $\sim 70\%$  Si and  $\sim 30\%$  Ge. The conduction band minima of the strained Si are lower than the conduction band minima of the relaxed SiGe, which leads to the confinement of low-energy electrons along the  $z$  direction, in the thin Si layer. Nanoscale metallic electrodes are fabricated above the electrons. Voltages above (below) the conduction band minima leads to the accumulation (depletion) of electrons underneath the electrode via band bending. Band bending also contributes to



the confinement of the electrons in the  $z$  direction, particularly in MOS devices [29]. By manipulating the voltages on the electrodes, individual electrons can be confined in the  $x$  and  $y$  directions, leading to a QD.

The voltages on the metallic electrodes in the device can be manipulated to confine an electron to a QD; they can also be used to control the encoded qubit. There are many types of qubits that can be formed by using a small number of electrons in a small number of QDs [25, 30, 31, 32, 33, 34, 35]. The state of these qubits depends on the physical properties of the constituent QDs, such as the inter-dot energy difference (referred to as the *detuning*) and tunneling rates, as well as the intra-dot orbital energy splittings. In Chapter 3, I go into detail as to how the voltages on the metallic electrodes can be used to manipulate these quantities. In Chapter 4, I introduce the quantum dot hybrid qubit (QDHQ), and show how pulses on the detunings and tunnel couplings can be used to improve the performance of a two-QDHQ system.

### 1.3 Thesis outline

As indicated by this introduction, leveraging quantum mechanics to answer computational questions efficiently is a complex problem. Researchers in this field often make use of abstraction, allowing one to focus on a particular problem without worrying about the details of the system being studied. This thesis is organized by the level of abstraction of the problem being considered. Chapter 2 focuses on the physics of the materials used in a device.

Next, Chapter 3 considers how to optimize the design of a device comprised of these materials. Then, Chapter 4 addresses how such a device could be operated to produce qubits and high-fidelity gates. Finally, Chapter 5 studies the performance of an algorithm running on many such qubits.

In Chapter 2, I focus on the materials used to fabricate semiconductor-based qubit devices. Specifically, I study the strain induced on the semiconductor domain of a simulated device by aluminum electrodes that are cooled in the vicinity of the studied domain. I introduce the theory behind stress and strain, as well as how these quantities can affect the electronic properties of a semiconductor. I then model the expected strain in a simulated device, and find the corresponding shift of the conduction band in the region intended to contain quantum dots.

Moving from materials properties to device design optimization, in Chapter 3 I introduce a protocol, which I call the Compressed Optimization of Device Architectures (CODA). This protocol both efficiently identifies sparse sets of voltage changes that control quantum systems, and introduces a metric which can be used to compare device designs. As an example of the former, I apply this method to simulated devices with up to 100 quantum dots and show that CODA automatically tunes devices more efficiently than other common nonlinear optimizers. To demonstrate the latter, I determine the optimal lateral scale for a triple quantum dot, yielding a simulated device that can be tuned with small voltage changes on a limited number of electrodes.

Abstracting away from device-specific physics, in Chapter 4 I propose a method to entan-

gle two capacitively-coupled quantum dot hybrid qubits using pulses on the qubits' detunings and tunnel couplings. I show that by pulsing only the qubits' detunings, CZ gates can be optimized to have fidelities of  $\sim 99\%$  for a typical level of quasistatic charge noise. I then identify a pulse sequence involving both the detunings and tunnel couplings that achieves a CZ gate with a fidelity of  $\sim 99.8\%$  under the same level of charge noise.

In Chapter 5, I transition from studying a system with a small number of qubits to considering the operation of a large-scale adiabatic quantum computer. I consider the performance of a quantum algorithm for PageRank first proposed in Ref. [36]. This algorithm takes as input a graph of the World Wide Web. I show that the run-time of this algorithm depends on features of the graphs other than the degree distribution, and can be altered sufficiently to rule out a general exponential speedup. According to these simulations, for a sample of graphs with degree distributions that are scale-free, with parameters thought to closely resemble the Web, the proposed algorithm for eigenvector preparation does not appear to run exponentially faster than the classical case.

## 1.4 Publications

Several of the chapters in this thesis are based on either published or submitted work; each is the product of the work of many individuals. Here, I document my contribution to each work, as well as the contributions of others.

Chapter 2 is based on preliminary, unpublished work. The simulation was performed by

myself, with assistance from Dr. Evan MacQuarrie, Prof. Mark Eriksson, Dr. Mark Friesen, and Prof. Susan Coppersmith. The electrode layout was designed by Mr. Sam Neyens. Mr. Neyens additionally provided the details of the heterostructure used in this simulation.

Chapter 3 is based on Ref. [37], titled *Compressed Optimization of Device Architectures for Semiconductor Quantum Devices*, which I completed with Dr. John King Gamble, Dr. Daniel R. Ward, Dr. Robin Blume-Kohout, Prof. Mark Eriksson, Dr. Mark Friesen, and Prof. Susan Coppersmith. I principally carried out the numerical calculations, assisted by Dr. Gamble, Prof. Eriksson, Dr. Friesen and Prof. Coppersmith. All authors provided input to the final manuscript.

Chapter 4 is based on Ref. [38], titled *Adiabatic two-qubit gates in capacitively coupled quantum dot hybrid qubits*, which I completed with Dr. Sebastian Mehl, Dr. John King Gamble, Dr. Mark Friesen, and Prof. Susan Coppersmith. All authors contributed to the idea of using adiabatic gate sequences to entangle capacitively-coupled quantum dot hybrid qubits. I performed the analytic work. Dr. Gamble and I performed the numerical simulation, and analyzed the data with Dr. Friesen and Prof. Coppersmith. Dr. Gamble, Dr. Friesen, Prof. Coppersmith and I wrote the manuscript and prepared the figures, with input from all the authors.

Chapter 5 is based on Ref. [39], titled *Power-law scaling for the adiabatic algorithm for search-engine ranking*, which I completed with Dr. John King Gamble, Dr. Kenneth Rudinger, Prof. Eric Bach, Dr. Mark Friesen, Prof. Robert Joynt, and Prof. Susan Coppersmith. I performed the numerical calculations under the guidance and direction of Dr. Gam-

ble and Dr. Rudinger, with assistance from all authors.

## Chapter 2

# Strain in semiconductor quantum devices

It is well-known that different materials can contract different amounts as their temperature decreases; in particular, aluminum contracts far more than many semiconductors. Thus, semiconductor-based devices containing aluminum components will become strained when cooled, which can affect the electronic band structure of the semiconductor. As some of the quantum dot devices used for qubit experiments contain aluminum and are cooled to cryogenic temperatures, the energy of electrons in the semiconductor domains of these devices will be affected by strain. Here, we introduce the concepts of stress and strain, and describe how these quantities can affect the energies of conduction band electrons in semiconductors. We then describe a calculation of the expected energy shift for conduction-band electrons in a semiconductor device used for quantum computation.

## 2.1 Stress

The force per unit area that a material undergoes is defined by the *stress*. This quantity can be represented in multiple ways, the most common of which is in terms of the Cauchy stress tensor  $\sigma$ . Considering the stresses applied to a cube, the nine components of this tensor  $\sigma_{ij}$  can be interpreted as the force per unit area in the  $i$  direction on a face whose normal is in the  $j$  direction. It is assumed that the cube undergoes no net torque, meaning that  $\sigma$  is symmetric,  $\sigma_{ij} = \sigma_{ji}$ . The objects we consider here are in equilibrium, and therefore the net force acting on them must be zero. From this, it follows that

$$\nabla \cdot \sigma + \mathbf{f} = 0 \tag{2.1}$$

where  $\mathbf{f}$  is the external force per unit volume that is applied to the object.

## 2.2 Strain

The amount by which an object deforms is defined by the *strain*. As an intuitive example, if a bar of length  $L_0$  is elongated by  $\Delta L$ , then the strain of the object is  $\Delta L/L_0$ . As with stress, strain can be represented in multiple ways. Here, we will introduce what is often referred to as the “infinitesimal strain” tensor  $\varepsilon$ . This tensor is defined as

$$\varepsilon = \frac{1}{2} \left( (\nabla \mathbf{u})^T + \nabla \mathbf{u} \right), \tag{2.2}$$

where  $\mathbf{u}$  is the *displacement* vector of the material.  $\varepsilon$  is symmetric by construction – the corresponding antisymmetric tensor  $\frac{1}{2} \left( (\nabla \mathbf{u})^T - \nabla \mathbf{u} \right)$  describes the rotation of the object [40]. In the situation considered here, the strain can be divided into two categories: *mechanical strain* and *thermal strain*.

### 2.2.1 Mechanical strain

The mechanical strain of an object depends on the stress applied. There are many stress-strain relationships observed in nature; here, we model the materials in our devices as *linear elastic materials*, in which the mechanical strain depends linearly on the stress. For these materials, one can relate the stress and mechanical strain via the *compliance tensor*  $S$ :

$$\varepsilon_{ij}^m = \sum S_{ijkl} \sigma_{kl}, \quad (2.3)$$

where  $\varepsilon^m$  is the mechanical strain. The mapping from mechanical strain to stress is given by the inverse, the *stiffness tensor*  $C$ .

It is more convenient to represent  $\varepsilon^m$  and  $\sigma$  as vectors, and  $S$  as a rank-two tensor. Assuming the material is orthotropic (*i.e.*, it has three orthogonal planes of symmetry),



then

$$\begin{pmatrix} \varepsilon_{xx}^m \\ \varepsilon_{yy}^m \\ \varepsilon_{zz}^m \\ 2\varepsilon_{yz}^m \\ 2\varepsilon_{zx}^m \\ 2\varepsilon_{xy}^m \end{pmatrix} = \begin{pmatrix} \frac{1}{E_x} & -\frac{\nu_{yx}}{E_y} & -\frac{\nu_{zx}}{E_z} & 0 & 0 & 0 \\ -\frac{\nu_{xy}}{E_x} & \frac{1}{E_y} & -\frac{\nu_{zy}}{E_z} & 0 & 0 & 0 \\ -\frac{\nu_{xz}}{E_x} & -\frac{\nu_{yz}}{E_y} & \frac{1}{E_z} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{G_{yz}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{G_{zx}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{G_{xy}} \end{pmatrix} \begin{pmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{xy} \end{pmatrix}, \quad (2.4)$$

where  $E_i$  is the Young's modulus along axis  $i$ ,  $\nu_{ij}$  is the Poisson's ratio corresponding with a contraction in direction  $j$  and an extension in direction  $i$ , and  $G_{ij}$  is the shear modulus in direction  $j$  on the plane which is normal to direction  $i$ . The compliance tensor is symmetric even for anisotropic materials - this means that  $\nu_{xy}$  need not equal  $\nu_{yx}$ , but  $\nu_{xy}/E_x = \nu_{yx}/E_y$ . In the case of materials with cubic symmetry (such as Si and SiGe), there is a single value for  $E$ , a single value for  $\nu$ , and a single value for  $G$ . In fully isotropic materials, the shear modulus is given by  $G = \frac{E}{2(1+\nu)}$ .

### 2.2.2 Thermal strain

When an object is heated or cooled, the change in the kinetic energy of the molecules in the material cause the object to expand or contract, respectively. The resulting strain is given by

$$\varepsilon_{ii}^{th} = \alpha(T_{final} - T_{initial}), \quad (2.5)$$

where  $\varepsilon^{th}$  is the thermal strain,  $i = x, y, z$ ;  $\alpha$  is the *coefficient of thermal expansion* (CTE), and  $T_{final}$  and  $T_{initial}$  are the final and initial temperatures, respectively.

While this source of strain may at first look like a trivial volume dilation, the resulting strain can in fact be quite complex due to CTE mismatch: the CTE is a material-based quantity, meaning that different components of a device will undergo different amounts of volume dilation. In particular, semiconductor devices with aluminum electrodes are thought to be affected by CTE mismatch [41], as the coefficient associated with aluminum is roughly an order of magnitude larger than that of most semiconductors.

## 2.3 Effect of strain on band structure

The electronic band structure of a semiconductor material depends on the details of the lattice. It is therefore unsurprising that deforming the lattice will affect the band structure. In particular, there are three ways in which strain can affect the band structure of semiconductors. In this section, we first explain the origin of these three different types of strain, and then detail each effect separately.

A key feature of semiconductors is the symmetry of the lattice. In particular, the lattices of diamond-type semiconductors (such as Si) have a tetrahedral symmetry. This means that the lattice remain invariant under a group of transformations  $T_d$ , including for instance reflections about the (110) plane. In contrast to this, the strain tensor  $\varepsilon$  does not necessarily have tetrahedral symmetry; the only symmetry imposed is  $\varepsilon_{ij} = \varepsilon_{ji}$ . Therefore, it is instruc-

tive to take linear combinations of the different components of the strain tensor, forming groups which behave the same under the symmetry transformations. In other words, the strain tensor  $\varepsilon$  can be decomposed into the irreducible representations of the tetrahedral symmetry group  $T_d$ .

The result is that the strain tensor can be rewritten as the sum of three matrices,  $\varepsilon = \varepsilon(\Gamma_1) + \varepsilon(\Gamma_3) + \varepsilon(\Gamma_4)$ , where

$$\varepsilon(\Gamma_1) = \frac{1}{3} \begin{pmatrix} \varepsilon_{xx} + \varepsilon_{yy} + \varepsilon_{zz} & 0 & 0 \\ 0 & \varepsilon_{xx} + \varepsilon_{yy} + \varepsilon_{zz} & 0 \\ 0 & 0 & \varepsilon_{xx} + \varepsilon_{yy} + \varepsilon_{zz} \end{pmatrix}, \quad (2.6)$$

$$\varepsilon(\Gamma_3) = \frac{1}{3} \begin{pmatrix} 2\varepsilon_{xx} - (\varepsilon_{yy} + \varepsilon_{zz}) & 0 & 0 \\ 0 & 2\varepsilon_{yy} - (\varepsilon_{xx} + \varepsilon_{zz}) & 0 \\ 0 & 0 & 2\varepsilon_{zz} - (\varepsilon_{xx} + \varepsilon_{yy}) \end{pmatrix}, \quad (2.7)$$

and

$$\varepsilon(\Gamma_4) = \begin{pmatrix} 0 & \varepsilon_{xy} & \varepsilon_{xz} \\ \varepsilon_{xy} & 0 & \varepsilon_{yz} \\ \varepsilon_{xz} & \varepsilon_{yz} & 0 \end{pmatrix}. \quad (2.8)$$

$\Gamma_1$ ,  $\Gamma_3$ , and  $\Gamma_4$  (sometimes referred to as  $A_1$ ,  $E$ , and  $T_1$ , respectively) refer to the irreducible representation with which the matrices are associated. As expected, each of these matrices exhibits different symmetries. For instance,  $\varepsilon(\Gamma_1)$  is invariant under any of the transformations in  $T_d$ , whereas  $\varepsilon(\Gamma_3)$  is only invariant under rotations of  $180^\circ$  about the  $[100]$ ,  $[010]$ , or

[001] axes. Similarly,  $\varepsilon(\Gamma_4)$  is only invariant under rotations of  $90^\circ$  about the [100], [010], or [001] axes.

These three types of strains have different physical interpretations as well.  $\varepsilon(\Gamma_1)$  is the strain associated with volume dilation, whereas the volume does not change under  $\varepsilon(\Gamma_3)$  or  $\varepsilon(\Gamma_4)$ ; these are what are known as *shear strains*. In particular,  $\varepsilon(\Gamma_3)$  and  $\varepsilon(\Gamma_4)$  are shear strains emerging from uniaxial stress in the [100] and [111] direction, respectively.

As the symmetry of these strains differ, it is not surprising that each of these strains lead to different effects on the band structure. In the rest of this section, we describe the effect of each of these strains separately.

### 2.3.1 $\Gamma_1$ strain

The simplest component of strain is the volume dilation, given in Eq. (2.6). The matrix  $\varepsilon(\Gamma_1)$  contains only one degree of freedom:  $\text{tr}(\varepsilon(\Gamma_1)) = \text{tr}(\varepsilon)$ . Due to the symmetry of this component, this strain can affect the energy of any wavefunction in the reciprocal space of the lattice.

Because the strain is assumed to be small, it follows that the energy of an electron in the semiconductor changes linearly with the strain. In general, we define

$$\delta E_{n,\mathbf{k}} = a_{n,\mathbf{k}} \text{tr}(\varepsilon) \quad (2.9)$$

where  $\delta E_{n,\mathbf{k}}$  is the change in energy of an electron in band  $n$  with wavevector  $\mathbf{k}$ , and  $a_{n,\mathbf{k}}$  is the

associated *deformation potential*. Deformation potentials are notoriously difficult to measure [40], with a wide variety of values reported for a given band extremum [42]. However, this is the best-known method for calculating the change in an electron’s energy due to lattice strain, and is used frequently in the community (see Ref. [43] for examples). Typically these values are in the 1-10 eV range [42]. To match the conventions of other works, the deformation potential of an electron at the conduction band minimum of Si is referred to as  $\Xi_d + 1/3\Xi_u$  [44], for reasons that will be explained in the next subsection.

### 2.3.2 $\Gamma_3$ strain

We next consider the “diagonal” component of shear strain, given in Eq. (2.7). As with the  $\varepsilon(\Gamma_1)$  component of the strain, this strain induces a linear shift the band energies. However, unlike  $\varepsilon(\Gamma_1)$ ,  $\varepsilon(\Gamma_3)$  is not invariant under every symmetry transformation of the lattice. For instance, at the  $\Gamma$  point the electronic wavefunction is symmetric in ways that  $\varepsilon(\Gamma_3)$  is not; this means that, to linear order, this component of strain cannot have an effect on the energies at this high symmetry point.

This component of strain can have an effect on energy extrema that are not at the  $\Gamma$  point. Si has a six-fold degenerate conduction band minimum, located at 85% of the way along the  $\Delta$  direction from the  $\Gamma$  point to the  $X$  point (see Fig. 1.2); the symmetry of the wavefunctions at these conduction band minima do not preclude an effect from  $\varepsilon(\Gamma_3)$ . In general, the conduction band energy shifts due to this component of strain can be written

as

$$\delta E_{\mathbf{k}}^{\Gamma_3} = \Xi_u(\hat{\mathbf{k}} \cdot \varepsilon(\Gamma_3) \cdot \hat{\mathbf{k}}), \quad (2.10)$$

where  $\mathbf{k}$  is the wavevector and  $\Xi_u$  is a deformation potential. Considering the case of Si, we can combine this energy shift with the shift due to the  $\varepsilon(\Gamma_1)$  strain, yielding

$$\delta E_i = \Xi_d(\varepsilon_{xx} + \varepsilon_{yy} + \varepsilon_{zz}) + \Xi_u \varepsilon_{ii}, \quad (2.11)$$

where  $i$  denotes the valley.

Eq. (2.11) demonstrates that strain of type  $\Gamma_3$  can break the six-fold valley degeneracy in Si. This phenomena occurs in many situations, including the case of Si wells embedded in SiGe. For a small ( $\sim 10$  nm) well, the Si is lattice-matched to the surrounding SiGe, causing a biaxial stress in the  $xy$  (assuming that  $z$  is the growth direction). This stress leads to a shear strain of type  $\Gamma_3$ , in which the  $\varepsilon_{zz}$  component of the strain tensor is smaller than  $\varepsilon_{xx} = \varepsilon_{yy}$ . This causes the  $\pm x, \pm y$  valleys to increase more in energy, leaving the  $\pm z$  valleys as the doubly-degenerate ground state. In a typical strained Si well, the  $\pm z$  valleys are about 200 meV below the  $\pm x, \pm y$  valleys [45].

### 2.3.3 $\Gamma_4$ strain

Finally, we consider the “off-diagonal” components of the strain tensor, given in Eq. (2.8).

As with the  $\varepsilon(\Gamma_3)$  component of strain, this strain does not cause a shift in the band energies at the  $\Gamma$  point to linear order, due to symmetry. Additionally, the symmetry of the  $\Gamma_4$  strain

prevents it from affecting band energies along the  $\Delta$  direction, meaning that in Si there is no shift in the conduction band minima due to this type of strain. This strain does shift energies along the  $[110]$  direction, and has an associated deformation potential  $\Xi_p$ , which has been measured to be 5.7 eV [46].

It has been noted [47, 48] that  $\Gamma_4$  strain may play a role in the valley physics of Si, despite it not shifting the energies. In particular, the introduction of  $\Gamma_4$  strain breaks the degeneracy between the  $\Delta_1$  and  $\Delta_2'$  bands at the  $X$  point in Si [47]. It has been proposed [48] this could split the degeneracy between the (+) and (-) valleys in bulk Si, although it is not yet understood what effect if any this would have in Si-based devices used for quantum computing.

## 2.4 Finite-element modeling of stress and strain

We simulate a device with an “overlapping gate” design to calculate the expected strain and conduction-band energy shift due to CTE mismatch. We find that the strain induces local maxima of the conduction band underneath the electrodes intended to accumulate quantum dots, consistent with the findings in Ref. [41].

The simulated device features the electrode layout shown in Fig. 2.1a. The electrodes are extruded in the  $z$  direction by different lengths depending on the order in which they would be fabricated: the orange, blue and pink electrodes are extruded 33 nm, 44 nm, and 55 nm, respectively. In this simulation, we have neglected the physical requirement that the

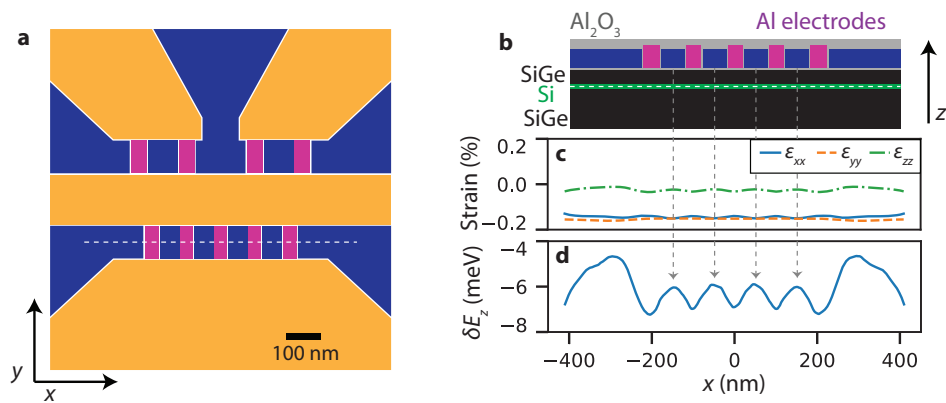


Figure 2.1: **Coefficient of thermal expansion mismatch expected to induce local maxima of conduction band underneath accumulation gates.** **a**, Electrode layout for simulated device. Electrodes are intended to accumulate (blue), deplete (pink), or screen (orange) the 2DEG in the Si layer below. **b**, Vertical slice of simulated device along white dashed line in **a**. **c,d**, Strain (**c**) and conduction band shift (**d**) in one-dimensional cut through quantum well (white dashed line in **b**). The strain is induced as the device is cooled from 293.15 K to 1 K due to the difference in coefficient of thermal expansion in the materials within the device. All three components of the strain are relatively constant over the considered region, with variations on the order of 0.01%. The conduction band of the  $z$  valley  $\delta E_z$  is expected to shift to lower energy, with local maxima underneath the accumulation gates, as indicated by the gray arrows. The local variations are on the order of 1 meV, the same order of magnitude as the quantum dot orbital splitting.



electrodes must extend out to the edges of the device, as we expect the biggest effects to come from the components of the device most proximal to the quantum dots. It would be of interest to determine in future work if imposing this restriction would affect the results of this study.

The heterostructure of the simulated device is shown in Fig. 2.1b. The device consists of (from bottom to top) a 100 nm SiGe buffer, a 9 nm Si well, a 40 nm SiGe spacer, and a 66.5 nm Al<sub>2</sub>O<sub>3</sub> layer. The Al electrodes are located within the Al<sub>2</sub>O<sub>3</sub> layer, starting 5 nm above the top of the SiGe spacer.

Material	Color	$E$ (GPa)	$\nu$	$G$ (GPa)	$\alpha$ ( $10^{-6}$ K <sup>-1</sup> )	$\Xi_d$ (eV)	$\Xi_u$ (eV)
Al	blue/pink	70	0.33	26.32	23	-	-
Al <sub>2</sub> O <sub>3</sub>	gray	300	0.22	123	5.4	-	-
SiGe	black	124.6	0.276	77.08	2.6	-	-
Si	green	130 [49]	0.28 [49]	79.6 [49]	2.6	1.1	10.5

Table 2.1: Materials parameters used in strain simulations. The ‘‘Color’’ column refers to the color of the material as shown in Fig. 2.1b.

To calculate the strain in this device due to CTE mismatch, we first find the thermal strain  $\varepsilon^{th}$  in the devices using Eq. (2.5). We assume  $T_{initial} = 293.15$  K,  $T_{final} = 1$  K, and we use the material parameters given in Table 2.1. We next solve the differential equation in Eq. 2.1 over the domains of the devices using the finite-element method as implemented in the COMSOL multiphysics package [50]. We assume that the device does not experience any external force, which leads to a Neumann boundary condition. We additionally impose Dirichlet boundary conditions on three corners in the geometry to ensure that the device does not translate or rotate.

The strain in the vicinity of the quantum dots is shown in Fig. 2.1c. The  $x$  and  $y$

components of strain  $\varepsilon_{xx}$  and  $\varepsilon_{yy}$  are both larger in magnitude than the  $z$  component of strain  $\varepsilon_{zz}$ . All three components of strain are relatively constant over this domain, with variations on the order of 0.01%.

Once the strain tensor is calculated, we use Eq. (2.11) to find the expected shift in the conduction-band minima, shown in Fig. 2.1d. We use the deformation potentials given in Table 2.1. The local variations in the  $\varepsilon_{zz}$  strain field generate local variations in the conduction band, yielding local maxima underneath the electrodes intended for quantum dot accumulation. These variations in the conduction band affect the properties of the quantum dots formed in this device.

While the strain-induced conduction-band variations lead to a decrease in the magnitude of the quantum dots' orbital splittings, a rough estimate of the expected change in energy implies that the effect is small. We divide the one-dimensional cut shown in Fig. 2.1c,d into four 100 nm segments, each centered around a gray arrow in Fig. 2.1c,d. In each of these segments, we assume that the unperturbed electrostatic potential can be approximated as a simple harmonic oscillator with an orbital splitting of  $\hbar\omega \approx 2$  meV [51]. The unperturbed potential in each segment is then  $E_z = \frac{1}{2}m^*\omega^2(x - x_0)^2 \approx (5 \times 10^{-3} \text{ meV/nm}^2)(x - x_0)^2$ , where  $m^*$  is the transverse effective mass in Si and  $x_0$  is the center of the segment. We then perform a least-squares quadratic fit of the predicted energy shift  $\delta E_z$  in each segment. In each case, we find the energy shifts to be approximately  $\delta E_z \approx -(5 \times 10^{-4} \text{ meV/nm}^2)(x - x_0)^2$ , an order of magnitude smaller than the unperturbed potential. The orbital splitting of the combined potential  $E'_z = E_z + \delta E_z$  in each segment is within 10% of 2 meV, implying

that CTE mismatch does not lead to large changes in quantum dot orbital energy splittings.

## **2.5 Acknowledgments**

The simulation presented here was performed by myself, with assistance from Dr. Evan MacQuarrie, Prof. Mark Eriksson, Dr. Mark Friesen, and Prof. Susan Coppersmith. The electrode layout was designed by Mr. Sam Neyens. Mr. Neyens additionally provided the details of the heterostructure used in this simulation.

## Chapter 3

# Compressed Optimization of Device

## Architectures for Semiconductor

### Quantum Devices

Recent advances in nanotechnology have enabled researchers to manipulate small collections of quantum mechanical objects with unprecedented accuracy. In semiconductor quantum dot qubits, this manipulation requires controlling the dot orbital energies, tunnel couplings, and the electron occupations. These properties all depend on the voltages placed on the metallic electrodes that define the device, whose positions are fixed once the device is fabricated. While there has been much success with small numbers of dots, as the number of dots grows, it will be increasingly useful to control these systems with as few electrode voltage changes as possible. Here, we introduce a protocol, which we call the Compressed Optimization

of Device Architectures (CODA), in order to both efficiently identify sparse sets of voltage changes that control quantum systems, and to introduce a metric which can be used to compare device designs. As an example of the former, we apply this method to simulated devices with up to 100 quantum dots and show that CODA automatically tunes devices more efficiently than other common nonlinear optimizers. To demonstrate the latter, we determine the optimal lateral scale for a triple quantum dot, yielding a simulated device that can be tuned with small voltage changes on a limited number of electrodes.

### 3.1 Introduction

Nanoscale systems are challenging to control in part because their size makes them susceptible to even the smallest materials defects. Quantum devices present special challenges because their energy spectra and tunnel couplings each require precise control [52, 53]. Here, we focus on quantum bits (qubits) formed in electrostatically-gated quantum dots [25]. In these systems, voltages are simultaneously tuned on many electrodes to precisely shape the electrostatic potential landscape within a device. Working with a small number of qubits, researchers have already demonstrated excellent qubit coherence and performance in devices based in GaAs [54, 55, 56, 57] and silicon [58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68], including the successful implementation of two-qubit gates [69, 70, 71] and algorithms [72]. Additionally, there has been rapid progress in systems with electrons bound to donors [73, 74, 75, 76, 77, 78, 79, 80, 81, 82], which share many of the same design challenges as

quantum dot qubits. Tuning schemes for these devices are typically determined empirically; however, there has been recent progress towards automatic tuning of quantum dots both experimentally [83, 84] and in simulated systems [85]. One challenge in controlling quantum dot devices is ensuring that the voltage changes on the electrodes remain small during tuning, a property that we refer to as *voltage moderation*. Imposing voltage moderation reduces the power required during dynamic operation and lowers the risk of instability. Another challenge is ensuring that a small number of electrodes can be used to tune quantum dots, preferably with those electrodes proximal to the relevant dots. We refer to this property as *voltage sparsity*, which eases the demands on control electronics and will be increasingly important as devices are scaled to very large numbers of quantum dots. In order to achieve these goals, it is critical to use simulations both to identify moderate and sparse tunings, and to design these features into devices before they are fabricated.

Here, we introduce the Compressed Optimization of Device Architectures (CODA) protocol, which both determines optimized ways to change the voltages in a given system to achieve a desired outcome, and provides a metric to characterize the ease with which the device can be tuned. We show that by minimizing the  $L_1$  norm of the applied voltage changes, we can simultaneously achieve voltage moderation and voltage sparsity. We minimize this norm by implementing the CODA protocol, which relies on results and methods used for compressed sensing [86, 87] in the signal processing literature. Using a simulated eight-dot device, we demonstrate that CODA yields solutions that are simultaneously sparse and moderate. Moreover, we show that by imposing voltage sparsity and moderation, we

obtain solutions that only involve gates that are proximal to the dot being manipulated – an extremely desirable property for extensibility. To further demonstrate the extensibility of the CODA protocol, we use a simple model to measure the number of device simulations required to tune devices with up to 100 quantum dots. We find that CODA requires nearly an order of magnitude fewer simulations than other commonly-used nonlinear optimization techniques. Additionally, formulating control as an optimization problem allows us to directly compare the effectiveness of different device architectures, enabling improvement of the electrode designs themselves. To demonstrate this, we use CODA to optimize the overall lateral scale of a triple quantum dot, which leads to a device that is optimally “tunable.”

## 3.2 Automatic tuning of simulated devices

For qubit applications, the main properties we wish to control are the quantum dot occupations, the orbital energy splittings, and the tunnel rates between the dots. Such properties are referred to as *quantities of interest*, which we represent as a vector  $\mathbf{q}$  in a vector space  $\mathcal{Q}$ . These values are controlled by the voltages applied to the electrodes. A given set of voltages is represented as a vector  $\mathbf{v}$  in a vector space  $\mathcal{V}$ . A physical system, or a simulation thereof, acts as a function that maps the voltages to the quantities of interest:  $\hat{S} : \mathcal{V} \rightarrow \mathcal{Q}$ , as shown in Fig. 3.1(a).

Because generically there are more electrodes than quantities of interest, many different  $\mathbf{v}$  can yield target values for the quantities of interest  $\mathbf{q}_{\text{target}}$ . The solutions are not all

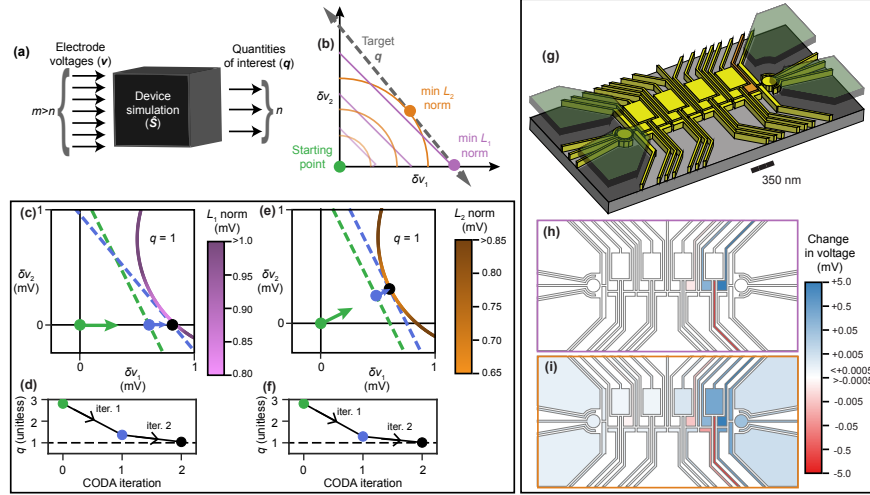


Figure 3.1: Example implementations of the CODA protocol, including comparisons of using the  $L_1$  norm (the sum of absolute values) and the  $L_2$  norm (the Euclidian length) of the voltage changes needed to achieve the desired values of the quantities of interest (dot occupations, tunnel couplings, *etc.*). (a) A simulated quantum dot device maps electrode voltages to quantities of interest. Generically, there are more electrodes in a device than quantities of interest. (b) There are many combinations of electrode voltages that result in a target system state (grey dashed line). By choosing the solution associated with the minimum  $L_1$  norm of voltage changes, described in Eq. (3.1) (purple), we simultaneously ensure that voltages are changed by small amounts on a small number of electrodes. Minimizing the  $L_2$  norm (orange) does not ensure that voltage changes will be applied to a small number of electrodes. (c)-(f) Depiction of CODA algorithm tuning voltages to obtain  $q_{\text{target}} = 1$  in a simple system. The voltages which yield the target quantity of interest are indicated with the circular segment in panels (c) and (e). Starting at  $\delta v_1 = \delta v_2 = 0$  mV (green circle), we calculate the Jacobian  $\mathbf{J}$  and find all of the solutions to the linear equation  $q_{\text{target}} - q_{\text{current}} = \mathbf{J} \cdot \delta \mathbf{v}$ , shown as a green dashed line. We then find the voltage changes on this line which minimize the  $L_1$  and  $L_2$  norm, blue circle in panel (c) and (e), respectively. Using the derivatives at this new point, we again estimate the voltage changes required to hit the target (blue dashed line), and again find the solution which minimizes the appropriate norm (black circle). The solution found here has converged on  $q_{\text{target}}$ . Additionally, we have obtained the voltage changes associated with the minimum of the appropriate norm, indicated by the purple color scale in panel (c), and the orange color scale in panel (e). (g)-(i) CODA tuning of an 8-dot device. (g) Schematic of a simulated 8-dot device, with metal electrodes colored yellow/orange (lower level) and green (upper level). Here, the objective is to increase the occupation of the right-most quantum dot (underneath the orange electrode) by one electron, keeping all other dot occupations and tunnel couplings unchanged. (h),(i) Visualization of the voltage changes obtained by the optimization protocol, plotted on a logarithmic color scale (electrodes with voltage changes less than  $0.5 \mu\text{V}$  are colored white), minimizing the (h)  $L_1$  and (i)  $L_2$  norms of the voltage changes. As expected, minimizing the  $L_1$  norm ensures that a limited number of electrode voltages are changed, whereas minimizing the  $L_2$  norm does not.



equivalently useful – it is our goal to select changes in voltages that are simultaneously moderate and sparse. In Fig. 3.1(b), we highlight the advantage of choosing the solution that satisfies the equation

$$\widetilde{\delta\mathbf{v}} = \underset{\delta\mathbf{v}}{\operatorname{argmin}} \|\delta\mathbf{v}\|_1, \text{ subject to } \hat{S}(\mathbf{v}_{\text{init}} + \delta\mathbf{v}) = \mathbf{q}_{\text{target}}, \quad (3.1)$$

where  $\mathbf{v}_{\text{init}}$  is the voltage vector at the starting point,  $\mathbf{q}_{\text{target}}$  is the vector of target quantities of interest,  $\delta\mathbf{v}$  is the voltage change from the starting point, and  $\|\cdot\|_1$  is the  $L_1$  norm, which is the sum of the absolute value of each element in the vector.

Minimizing the  $L_1$  norm of the voltage change vector ensures both the magnitude of the individual voltage changes remain small (*i.e.*, voltage moderation) and that the voltage change vector is sparse. The sparsity achieved by  $L_1$  norm minimization is a property used extensively in the field of compressed sensing [86, 87]. In contrast,  $L_2$  norm minimization (minimizing the Euclidian length of  $\delta\mathbf{v}$ ) does not guarantee voltage sparsity, and  $L_0$  pseudonorm minimization (minimizing the number of nonzero elements of  $\delta\mathbf{v}$ ) does not guarantee voltage moderation. In principle, one could achieve both moderation and sparsity by minimizing some combination of these two quantities, but this would involve making an arbitrary choice for the relative weight given to the  $L_0$  and  $L_2$  norms. Additionally, finding the solution that minimizes the  $L_0$  pseudonorm is known to be an NP-hard problem [88], whereas convex programming methods can be used to minimize the  $L_1$  norm efficiently [89]. Therefore, by selecting the changes in voltages described in Eq. (3.1), the CODA procedure

yields a device tuning in which a small number of voltages are changed by modest amounts.

To demonstrate the CODA protocol, we first perform automatic tuning of a simple toy example shown in Fig. 3.1(c)-(f). This system comprises two electrode voltages  $\delta v_1$  and  $\delta v_2$  and one quantity of interest  $q(\delta v_1, \delta v_2) = (\delta v_1/(1 \text{ mV}) - 3/2)^2 + (\delta v_2/(1 \text{ mV}) - 3/4)^2$ . During the protocol, we treat this as a black box function, as in more complicated device models we do not have access to an analytical form of the mapping from voltages to quantities of interest. Starting with  $\delta v_1 = \delta v_2 = 0 \text{ mV}$ , we wish to find the voltage changes with minimal  $L_1$  norm that are necessary to change the quantity of interest from  $q_{\text{current}} = 2.8125$  to  $q_{\text{target}} = 1$ . To achieve this, we apply an iterative algorithm to minimize the distance between the simulated quantities of interest and the target quantities of interest. At the starting point, we calculate the Jacobian  $\mathbf{J} = (\frac{\partial q}{\partial \delta v_1}, \frac{\partial q}{\partial \delta v_2})$  and find all of the solutions to the linear equation  $q_{\text{target}} - q_{\text{current}} = \mathbf{J} \cdot \delta \mathbf{v}$ , shown as a red dashed line in Fig. 3.1(c). In this example, it is easy to minimize the  $L_1$  norm along this line; however, for more general and complicated problems, we employ a matrix-free conic optimization algorithm [89] to determine the particular  $\delta \mathbf{v}$  that minimizes the  $L_1$  norm while satisfying this equation. This vector defines a search direction in voltage space; CODA moves along this direction in voltage space, stopping at the point that brings the simulation closest to the target quantities of interest. This process then repeats until the quantity of interest converges on  $q_{\text{target}}$ . In general, convergence is declared when  $\|\mathbf{q} - \mathbf{q}_{\text{target}}\|_2$  falls below some threshold. We typically choose an error threshold of  $\sim 10^{-2}$ ; in this example, we choose a threshold of  $5 \times 10^{-2}$ . We achieve convergence after two iterations, at which point  $q \approx 1.041$ , as shown in Fig. 3.1(d).

In this example, we achieved the target quantity of interest by changing only one electrode voltage.

It is instructive to contrast this solution to one obtained using an analogous procedure in which the  $L_2$  norm is minimized, rather than the  $L_1$  norm, as demonstrated in Fig. 3.1(e),(f). Here, we again declare convergence after two iterations, with  $q \approx 1.012$ . However the final solution involves changing the voltage on both electrodes to achieve the target quantity of interest, and is therefore less sparse than the solution obtained using the  $L_1$  norm. For more details on the CODA algorithm, see Appendix A.1.

In a more realistic demonstration of CODA's capabilities, we can use it to perform automatic tuning of the accumulation-mode eight-dot device shown in Fig 3.1(g). The device contains four capacitively-coupled double quantum dot qubits in a Si/SiGe heterostructure. In principle, the CODA protocol can use any underlying physical model. Here, we model the device using the semiclassical Thomas-Fermi approximation [90] to compute electron densities and potentials, and the WKB approximation [16] to calculate tunnel rates, as described in Appendices A.2-A.3. The quantities of interest are chosen to be the eight dot occupations and the four intra-qubit tunnel rates. While this semiclassical approach is appropriate for these quantities of interest, one may need a more sophisticated model to correctly capture the physics of different quantities of interest, *e.g.* valley splitting within a dot [91]. Our starting point is chosen to give dot occupations of 1 electron, and transmission coefficients between the dots of 0.01, corresponding to tunnel rates of approximately 400 MHz (see Appendices A.2-A.3 for details about the simulation parameters and methods). Our goal is to

find the optimal changes in voltages that can be applied to the device electrodes to add one electron to the right-most dot, leaving all other dot occupations and transmission coefficients unchanged.

In Fig. 3.1(h), we show the voltage changes needed to achieve convergence of the CODA procedure applied to the eight-dot device. Note that although voltages are allowed to vary on all the electrodes, CODA chooses to vary only four electrode voltages, and those electrodes are proximal to the target of interest. For comparison, we also performed an alternative optimization protocol based on  $L_2$ -minimization, with results shown in Fig. 3.1(i). Note that although the latter protocol achieves the same target quantities of interest, the solution involves voltage changes on almost all of the electrodes, indicating that this solution is neither sparse nor proximal. In both tuning protocols, the magnitudes of the voltage changes applied to the electrodes are all under 5 mV. Because minimizing the  $L_2$  norm explicitly ensures voltage moderation, the similarity in the magnitude of voltage changes applied in both cases confirms that the solution found via  $L_1$  norm minimization exhibits voltage moderation as well. Hence, we confirm that CODA is a practical tool for tuning a device, because it selects voltages that are both sparse and moderate while achieving the specified target quantities of interest.

### 3.3 Extensibility

In the previous section, we applied the CODA protocol to simulated devices with up

to eight quantum dots, showing that the method can find sparse voltage tunings which are moderate and sparse for these devices. These solutions were obtained after two iterative steps, suggesting that this procedure is efficient even for large systems. Here, we show that the CODA protocol scales efficiently with device size, and can be used to find sparse tunings for devices with 100 quantum dots.

We consider the device shown in Fig. 3.2(a), which consists of  $m$  quantum dots,  $m - 1$  tunnel rates, and  $4m - 1$  electrodes. The electrode separating the quantum dots (thin blue rectangle in Fig. 3.2(a)) is 35 nm from the center of the neighboring electrodes (red and blue squares in Fig. 3.2(a)). The centers of the square-shaped electrodes are separated by 50 nm. The quantum dots are located 20 nm below the electrodes.

There are many methods one could employ to model this device, including taking the semiclassical approach described in Appendices A.2–A.3, or self-consistently solving the Schrödinger and Poisson equations, which more accurately take into account quantum effects. Here, we use a simple model that can be regarded as phenomenological, although it is physically motivated, describing a non-linear system in which an electrode’s proximity to a quantum dot or tunnel barrier determines that electrode’s effect on the corresponding quantity of interest. Specifically, we define the occupation  $n_i$  of dot  $i$  to be

$$n_i = \sum_j \frac{(V_j/(1 \text{ mV}) + \frac{1}{10}\text{sgn}(V_j)V_j^2/(1 \text{ mV}^2))}{(\|\vec{r}_{V_j} - \vec{r}_{n_i}\|_2)^3 / (1 \text{ nm}^3)}, \quad (3.2)$$

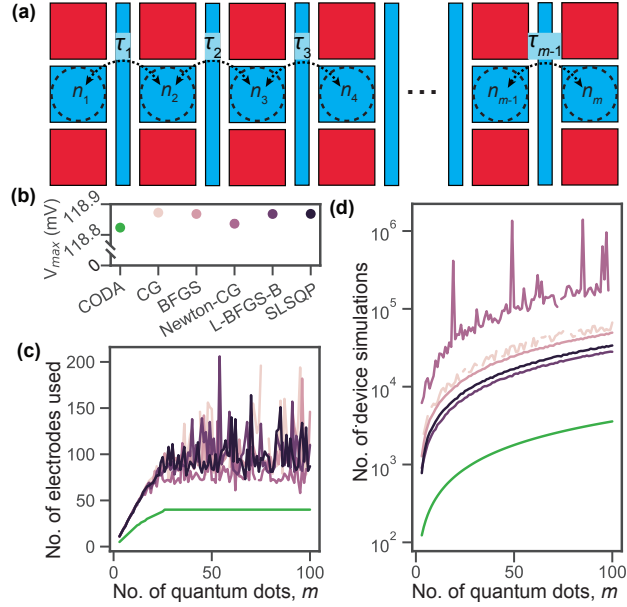


Figure 3.2: Demonstrating the extensibility of the CODA protocol using a simple model involving up to 100 quantum dots. (a) Diagram of simulated device. The quantities of interest in this device are the occupations  $n_i$  of the quantum dots (dashed circles) and tunnel rate  $\tau_i$  between the dots (dashed double arrows). Each quantum dot has three electrodes in close proximity: one located directly above the quantum dot (blue square) and two located above and to either side of the quantum dot (red squares). Additionally, there is an electrode separating each pair of dots (thin blue rectangle). The dependence of  $n_i$  and  $\tau_i$  on the electrode voltages is defined phenomenologically in Eqs. (3.2) and (3.3), respectively. (b)-(d) Given a device with  $m$  quantum dots, we use a variety of nonlinear optimizers (including CODA) to find the changes in electrode voltages which add one electron to the left-most quantum dot, keeping all other occupations and tunnel rates constant. We consider devices with  $m$  ranging from 2 to 100. The results corresponding to CODA are shown in green, and the results corresponding with the other nonlinear optimizers are shown in hues of purple. In panel (b) we show the maximum voltage change applied to any of the electrodes, averaged across all of the considered devices. The standard deviation is smaller than the points used in this plot. The uniformity of the results here indicates that all of the optimizers find solutions with similar voltage moderation. In panel (c) we show the number of electrodes used in each solution. The CODA procedure consistently requires fewer electrodes than any of the other nonlinear optimizers considered. In panel (d) we show the number of function calls used by each optimizer. CODA requires roughly an order of magnitude fewer function calls than any of the other nonlinear optimizers considered.

and the tunnel rate  $\tau_i$  between the  $i^{\text{th}}$  and  $(i + 1)^{\text{th}}$  quantum dot is given by

$$\tau_i = \frac{1}{100} \sum_j \frac{(V_j/(1 \text{ mV}) + \frac{1}{2}\text{sgn}(V_j)V_j^2/(1 \text{ mV}^2))}{(\|\vec{r}_{V_j} - \vec{r}_{n_i}\|_2)^3 / (1 \text{ nm}^3)}. \quad (3.3)$$

In these equations,  $V_j$  is the voltage applied to electrode  $j$ ,  $\text{sgn}(x)$  is the sign function of  $x$ ,  $\vec{r}_{V_j}$  is the center of electrode  $j$ ,  $\vec{r}_{n_i}$  is the center of dot  $i$ , and  $\vec{r}_{\tau_i}$  is directly between dots  $i$  and  $(i + 1)$ . The  $r^{-3}$  scaling of these quantities is the expected dependence due to screening in a 2DEG [92], while the voltage dependence was chosen assuming that these quantities scale approximately linearly with voltage, with an additional, explicitly non-linear contribution.

To study the extensibility of our approach, we employ a variety of nonlinear optimizers, including CODA, to tune the voltages in devices with  $m = 2$  to 100 dots. In each case, we begin with voltages -100 mV applied to each of the electrodes indicated with red in Fig. 3.2(a), and with positive voltages applied to the electrodes indicated with blue in Fig. 3.2(a). The latter voltages are set such that  $n_i = 1$  and  $\tau_i = 0.01$  for every  $i$ . We then find a combination of voltage changes that adds one electron to dot  $i = 1$ , keeping all other  $n_i$  and  $\tau_i$  fixed. We specifically consider the CODA protocol, the conjugate gradient (CG) algorithm, the Broyden, Fletcher, Goldfarb, and Shanno (BFGS) algorithm, the Newton conjugate gradient (Newton-CG) algorithm, the limited memory BFGS algorithm L-BFGS-B, and the Sequential Least Squares Programming (SLSQP) algorithm, as implemented in the SciPy package [93]. In all of these algorithms we minimize  $\|\mathbf{q} - \mathbf{q}_{\text{target}}\|_2$ , where  $\|\cdot\|_2$  is the  $L_2$  norm,  $\mathbf{q}$  is the vector consisting of every quantity of interest  $n_i$  and  $\tau_i$ , and  $\mathbf{q}_{\text{target}}$  is

the vector consisting of the target values for the quantities of interest. We define the system to be converged on the target when  $\|\mathbf{q} - \mathbf{q}_{\text{target}}\|_2 < 10^{-5}$ .

We assess the voltage moderation and sparsity of each of the solutions found by the nonlinear optimizers. The maximum voltage change  $V_{max}$  applied to a given device is shown in Fig. 3.2(b). The number of nonzero voltage changes found by each optimizer as a function of  $m$  is shown in Fig. 3.2(c). While all of the nonlinear solvers apply voltage changes of comparable magnitudes, CODA consistently finds solutions which require changing fewer electrode voltages than any of the other optimizers we consider. Moreover, the number of electrodes used by CODA does not change for devices with  $m \geq 26$ . In contrast, all of the other nonlinear optimizers show high variability in the number of electrodes used as a function of device size.

We find that the CODA protocol requires fewer function calls than any of the other nonlinear optimizers considered. In Fig. 3.2(d), we show the number of function calls made in each optimization. The average number of function calls made by a given nonlinear optimizer scales linearly with the number of quantum dots in the system, regardless of the optimizer used. Although the scaling is the same for each of these optimizers, the CODA protocol achieves convergence using approximately an order of magnitude fewer function calls compared with the next-most efficient optimizer, the L-BFGS-B algorithm. Since the device simulations dominate the computation time, the CODA protocol can therefore automatically tune simulated devices significantly faster than any of the other nonlinear optimizers considered. We conclude that CODA can efficiently find sparse and moderate solutions in



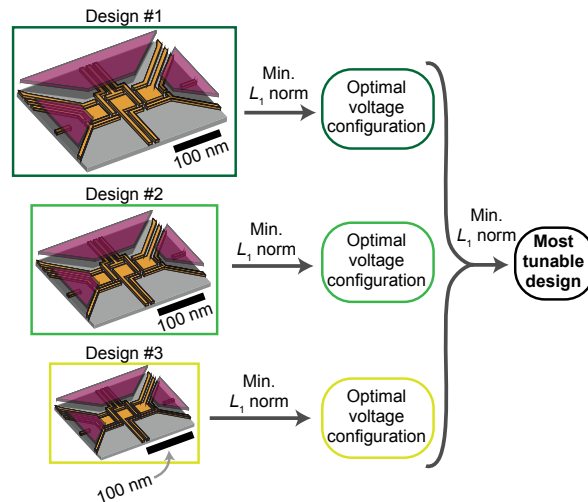


Figure 3.3: Protocol used to compare the “tunability” of device designs. Given multiple simulated devices, we use the CODA protocol to find the minimum  $L_1$  norm of voltage variations needed to induce a common change in each device (*e.g.*, change the dot occupations in device #1, device #2, *etc.*). The device with the minimum norm can simultaneously provide voltage moderation and sparsity, and should therefore be regarded as the most “tunable” device.

large devices.

### 3.4 Device design optimization

In addition to automatically tuning quantum dot devices, the CODA protocol can be used to characterize the voltage sparsity and moderation of typical gate operations, thus providing a key metric for evaluating and comparing different device designs. Here, we consider a series of triple-dot devices, shown in Fig 3.3. Each device has the same electrode layout, save for the overall lateral scale – we parameterize this scale via the width of the paddle electrode, as shown in Fig 3.4(a). All devices have an identical Si/SiGe heterostructure with a silicon quantum well a distance 30 nm below the electrodes. Optimizing device size is important, because a device with electrodes too small will lead to instability and larger

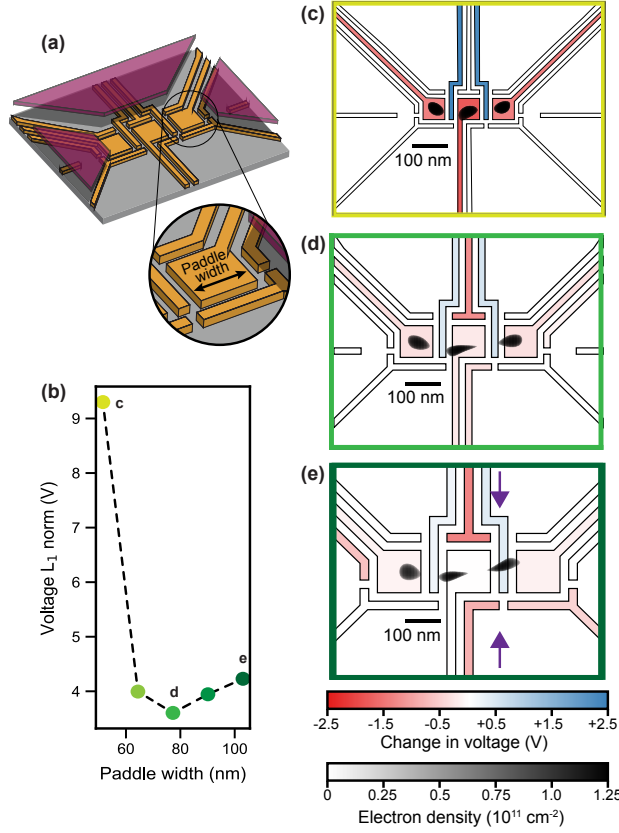


Figure 3.4: Using CODA to optimize the triple quantum dot designs shown in Fig. 3.3. (a) Every device in the series has an identical electrode layout, save for the lateral scale of the device, which we characterize in terms of the width of the paddle electrode. The Si/SiGe heterostructures for all devices considered are identical. (b) For each device in the series, we use the CODA protocol to lower all three dots from an occupation of thirty electrons to one electron, keeping the transmission coefficients fixed at 0.01. Additionally, the single-electron dots are required to have orbital excitation energies greater than 1 meV. (c)-(e) Visualization of the optimized voltage variations required to tune each device. Voltage changes are shown in red and blue, and the resulting electron density distributions are shown in black. For smaller devices (*e.g.*, (c)), fewer electrodes are required to tune the device. However larger voltage changes must be applied to those electrodes, resulting in a high voltage  $L_1$  norm. For larger devices (*e.g.*, (e)), quantum dots no longer form underneath the paddle electrodes (*e.g.*, right-most dot, indicated with purple arrows), so that many electrodes are required to tune the device. Balancing these effects leads to a local minimum in the voltage  $L_1$  norm, corresponding to device (d).

power requirements for switching, and a device with electrodes too big will not have sufficient control over the potential landscape at the required length-scales. In particular, it has been observed in experiments [69, 72, 94] that in larger devices it is often necessary to form the quantum dots away from their intended locations. We now show that CODA can be used to determine an optimal device scale.

We again use the semiclassical Thomas-Fermi approximation [90] and the WKB approximation [16] to model the devices. Since it is relatively difficult to determine the gate voltages needed to achieve single-electron occupancies in each dot, we choose a starting point for our simulations with 30 electrons in all three dots, and tunnel couplings that yield transmission coefficients of 0.01 between the left and middle dots and the middle and right dots. In each device, we then use the CODA protocol to automatically tune gate voltages to achieve single-electron occupation of each dot, while keeping the transmission coefficients constant, exploiting CODA's ability to automatically tune voltages. Additionally, we require the orbital energy splitting of each dot to be 1 meV or more, as consistent with recent experiments. Orbital energy splittings are calculated using a 2D finite-difference Schrödinger solver; see Appendix A.2 for details.

The minimized voltage  $L_1$  norms required for the auto-tuning process in each device are plotted in Fig 3.4(b). Comparing these results, we see that the voltage  $L_1$  norm is minimized for the device labeled (d), with a paddle width of approximately 80 nm, suggesting that this device is optimal from a control standpoint.

The voltage changes for the auto-tuning protocol and the resulting electron charge density

distributions are shown in Figs. 3.4(c)-3.4(e). For the optimal device, shown in Fig. 3.4(d), the number of gates with voltage changes is relatively small, indicating good voltage sparsity. For devices smaller than the optimal device, voltage sparsity is still maintained, as a benefit of small device size. However, voltage moderation is not, with large voltage changes required on multiple electrodes, as shown in Fig. 3.4(c). For devices larger than optimal, voltage moderation is maintained, but the solution is no longer sparse, as shown in Fig. 3.4(e). The constraints on dot occupation, dot energy and transmission coefficients lead to constraints on the size and relative position of the dots. In smaller devices, such as those shown in Fig. 3.4(c),(d), the dots can be formed underneath the paddle electrodes and still meet these requirements. In larger devices, such as the device in Fig. 3.4(e), to achieve the required quantities of interest, the dots can no longer form under the paddle gate electrodes, with the right-most dot forming under the nominal right tunnel-barrier electrode as indicated by the arrows. This misalignment between electrodes and dots, which has also been observed in experimental devices [69, 72, 94], leads to solutions that are less sparse than in the smaller devices.

## 3.5 Conclusion

We have introduced a protocol for the Compressed Optimization of Device Architectures, which determines the optimal voltage changes for a given device operation by minimizing their  $L_1$  norm. We have demonstrated the effectiveness of this scheme by considering its

application to semiconductor nanoelectronic quantum dot systems. As devices continue to grow in complexity, such automated control schemes will be essential for design and operation. Our protocol is computationally efficient to implement, and it provides a systematic approach for achieving local and sparse control. Through realistic semiclassical simulations of multi-dot devices, we have illustrated how the CODA scheme can be used for quantitative benchmarking and device development. While the current work focuses on quantum dot geometries, we note that the CODA protocol could also be applied to other device geometries, including donor-bound qubits, using simulation tools appropriate for those systems. This method provides a path toward the rational design and operation of extensible quantum nanodevices.

### **3.6 Funding acknowledgments**

This work was supported in part by ARO (W911NF-12-1-0607, W911NF-17-1-0274), NSF (PHY-1104660), and the Vannevar Bush Faculty Fellowship program sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering and funded by the Office of Naval Research through Grant No. N00014-15-1-0029. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office (ARO), or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

This paper describes objective technical results and analysis. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. The authors gratefully acknowledge support from the Sandia National Laboratories Truman Fellowship Program, which is funded by the Laboratory Directed Research and Development (LDRD) Program. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

## Chapter 4

# Adiabatic two-qubit gates in capacitively coupled quantum dot hybrid qubits

The ability to tune qubits to flat points in their energy dispersions (“sweet spots”) is an important tool for mitigating the effects of charge noise and dephasing in solid-state devices. However, the number of derivatives that must be simultaneously set to zero grows exponentially with the number of coupled qubits, making the task untenable for as few as two qubits. This is a particular problem for adiabatic gates, due to their slower speeds. Here, we propose an adiabatic two-qubit gate for quantum dot hybrid qubits, based on the tunable, electrostatic coupling between distinct charge configurations. We confirm the absence of a conventional sweet spot, but show that controlled-Z (CZ) gates can nonetheless be optimized

to have fidelities of  $\sim 99\%$  for a typical level of quasistatic charge noise ( $\sigma_\varepsilon \simeq 1 \mu\text{eV}$ ). We then develop the concept of a dynamical sweet spot (DSS), for which the time-averaged energy derivatives are set to zero, and identify a simple pulse sequence that achieves an approximate DSS for a CZ gate, with a  $5\times$  improvement in the fidelity. We observe that the results depend on the number of tunable parameters in the pulse sequence, and speculate that a more elaborate sequence could potentially attain a true DSS.

## 4.1 Introduction

Since their original proposal [25], semiconductor quantum dot qubits have progressed greatly, demonstrating excellent qubit coherence and performance through the use of sweet spots [54, 55, 56, 57, 58, 59, 60, 62, 63, 65, 66] and control of the spin degree of freedom [61, 64, 67, 68]. There has also been remarkable progress in systems with small numbers of donor-bound electrons [73, 74, 75, 77, 78, 79, 80, 81, 82]. Recently, two-qubit gates [69, 71] and algorithms [72] have been realized using exchange-coupled single-spin qubits. Capacitive coupling has also been employed to entangle and perform two-qubit operations between singlet-triplet qubits [70, 95], and has been proposed as the basis for two-qubit gates between resonant-exchange qubits [96] and flip-flop qubits [80]. In these experiments and proposals, the two-qubit gate times are typically measured in microseconds or hundreds of nanoseconds, which is much longer than typical single-qubit gate times. In contrast, the predicted two-qubit gate times for capacitively-coupled quantum dot hybrid qubits [97, 98] (QDHQs) are comparable



to single-qubit gates, which are of order 10 ns [63, 99]. However, the methods proposed in refs [97] and [98] rely on applying quickly varying electrical pulses, which can cause leakage from the qubit subspace [99].

In this paper we study an adiabatic entangling protocol based on capacitive couplings between QDHQs. The gate is inspired by an early proposal for entangling singlet-triplet qubits [100]. Although the necessary voltage changes are slow relative to the qubit frequencies, we show that high-fidelity adiabatic gates can be achieved in under 50 ns, which is significantly faster than those in recent singlet-triplet experiments [70, 95]. While the pulse sequences used in adiabatic protocols are more resilient against pulse errors than non-adiabatic pulses and are less susceptible to leakage errors, a potential concern is that they could be more susceptible to charge noise due to slower speeds. It is therefore crucial to study the effect of charge noise on the gate fidelities.

We begin by considering the system of two capacitively coupled QDHQs, deriving the effective couplings between the two qubits, and describing how a slowly-varying electrical pulse on the qubits can yield an entangling gate. Next, we optimize the pulse sequence for a two-qubit system to maximize the process fidelity of the resulting gate. We find that gate fidelities  $> 99\%$  are feasible, assuming quasistatic charge noise with a standard deviation of  $\sigma_\epsilon = 1 \mu\text{eV}$ , and that the infidelity scales roughly as  $\sigma_\epsilon^2$ . Finally, we show that gate fidelities can be further improved to  $\sim 99.9\%$  by modifying pulse sequences to impose a “dynamical sweet spot,” a technique similar to dynamical decoupling [101].

## 4.2 Results

The QDHQ consists of three electrons shared between two quantum dots. The minimal Hilbert space of the qubit can be defined as the spin states  $|\cdot S\rangle = |\downarrow\rangle |S\rangle$ ,  $|\cdot T\rangle = \sqrt{\frac{1}{3}} |\downarrow\rangle |T_0\rangle - \sqrt{\frac{2}{3}} |\uparrow\rangle |T_-\rangle$ , and  $|S\cdot\rangle = |S\rangle |\downarrow\rangle$  [33, 102], where  $|\cdot S\rangle$  and  $|\cdot T\rangle$  correspond to (1,2) charge configurations (one electron in the left dot, two electrons in the right),  $|S\cdot\rangle$  corresponds to a (2,1) charge configuration, and the singlet state,  $|S\rangle = 1/\sqrt{2} (|\downarrow\uparrow\rangle - |\uparrow\downarrow\rangle)$ , and triplet states,  $|T_0\rangle = 1/\sqrt{2} (|\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle)$  and  $|T_-\rangle = |\downarrow\downarrow\rangle$ , refer to the dot with two electrons. In this basis, the single-qubit Hamiltonian is

$$\mathcal{H}_{1q} = \begin{pmatrix} -\varepsilon/2 & 0 & \Delta_1 \\ 0 & -\varepsilon/2 + E_{ST} & -\Delta_2 \\ \Delta_1 & -\Delta_2 & \varepsilon/2 \end{pmatrix}, \quad (4.1)$$

where the detuning parameter,  $\varepsilon$ , corresponds to the energy separation between the quantum dots,  $\Delta_1$  is the tunnel coupling between states  $|\cdot S\rangle$  and  $|S\cdot\rangle$ ,  $\Delta_2$  is the tunnel coupling between states  $|\cdot T\rangle$  and  $|S\cdot\rangle$ , and  $E_{ST}$  is the energy splitting between the singlet-like and triplet-like basis states,  $|\cdot S\rangle$  and  $|\cdot T\rangle$ . A typical energy spectrum for  $\mathcal{H}_{1q}$  is plotted as a function of detuning in Fig. 4.1a. Here, the two lowest-energy eigenstates  $|0\rangle$  and  $|1\rangle$  form the qubit, while the remaining state  $|L\rangle$  is regarded as a leakage state.

The qubit states and charge configurations hybridize as a function of the detuning. The large detuning regime (right-most dashed line in Fig. 4.1a) is characterized by the asymptotic

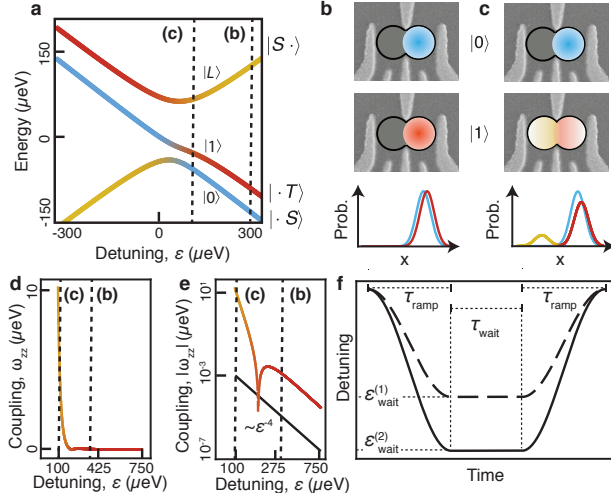


Figure 4.1: **Implementing a CZ gate between capacitively coupled QDHQs.** **a**, Energy dispersion of a single QDHQ as a function of detuning, as defined in Eq. (4.1), for typical experimental values given by  $\Delta_1 = 18.1 \mu\text{eV}$ ,  $\Delta_2 = 46.7 \mu\text{eV}$ , and  $E_{\text{ST}} = 51.7 \mu\text{eV}$  [63]. On the right-hand side of the plot, the logical states  $|0\rangle$  and  $|1\rangle$  converge to the basis states  $|S\rangle$  (indicated in blue) and  $|T\rangle$  (red), as defined in the main text, while the leakage state  $|L\rangle$  converges to  $|S\rangle$  (yellow). **b,c**, Charge distributions of a third electron added to an underlying (1,1) charge configuration at two different values of  $\varepsilon$ . **b**, In the large- $\varepsilon$  regime (right-most dashed line in **a**), the qubit states have very similar charge distributions (same color coding as **a**). **c**, For  $\varepsilon$  near the charge transition (left-most dashed line in **a**), some charge moves from the right dot to the left dot, especially for state  $|1\rangle$ , setting up a dipole moment between states  $|0\rangle$  and  $|1\rangle$ . **d,e**, Effective two-qubit coupling versus  $\varepsilon = \varepsilon^{(1)} = \varepsilon^{(2)}$ , plotted on linear-linear (**d**) and log-log (**e**) scales. When  $\varepsilon$  is large, the coupling is negligible and decreases as  $\varepsilon^{-4}$  (see Eq. (B.13)). When  $\varepsilon$  is simultaneously lowered on both qubits, their dipole moments grow, and the effective coupling increases. **f**, Detuning pulse sequences for qubits 1 and 2 (dashed and solid lines, respectively; see Appendix B.3 for details).




Figure 4.2: **Optimized infidelities of sub-50 ns adiabatic CZ gates.** **a**, Process infidelities obtained in the presence of quasistatic charge noise, with a standard deviation of  $\sigma_\epsilon=1 \mu\text{eV}$ ; other Hamiltonian parameters are given in the main text. Following the optimization procedure described in Methods, we obtain variable gate times, as indicated by the contours; however, we discard total gate times with  $\tau_{\text{total}}=2\tau_{\text{ramp}}+\tau_{\text{wait}}> 50 \text{ ns}$  (cross-hatched region) to ensure that the two-qubit gate is comparable in length to single-qubit gates. The highest fidelity pulse sequence is obtained at  $(\epsilon_{\text{wait}}^{(1)}, \epsilon_{\text{wait}}^{(2)})=(80,100) \mu\text{eV}$ . **b**, Minimum CZ gate infidelities (black squares and dashed black line) plotted as a function of the standard deviation of the charge noise,  $\sigma_\epsilon$ . The infidelity roughly falls off as  $\sigma_\epsilon^2$  (solid black line).

behavior  $|0\rangle \simeq |S\rangle$  and  $|1\rangle \simeq |T\rangle$ , for which both states have the same charge configuration, as depicted in Fig. 4.1b. Here, the information is stored entirely in the spin degree of freedom and the qubit is well protected from charge noise [63]; however the single-qubit gate speeds can be slow [103, 104]. (Below, we show the same is true for two-qubit gates.) To perform efficient gates, we must therefore lower the detuning, bringing it near the anticrossing region (left-most dashed line in Fig. 4.1a). In this regime,  $|S\rangle$  and  $|T\rangle$  begin to hybridize with  $|S\rangle$ , which has a (2,1) charge configuration, as depicted in Fig. 4.1c. Since the admixture of (2,1) is different for  $|0\rangle$  and  $|1\rangle$ , the qubit states acquire distinct dipole moments that can be used to mediate two-qubit dipole-dipole interactions, but which also couple to environmental charge noise. The goal of this work is to optimize the control parameters, to achieve high-fidelity two-qubit gates.

We can formalize the concept of a dipole moment by defining the operator

$$\hat{x} = \text{diag}\{d/2, d/2, -d/2\}, \quad (4.2)$$

describing the position of the third electron in the double dot, as depicted in Fig. 4.1b,c. Here for simplicity, we assume that states  $|\cdot S\rangle$  and  $|\cdot T\rangle$  have identical charge configurations. The dimensionless dipole operator is therefore given by  $\mathcal{P} = -\hat{x}/d = \text{diag}\{-1/2, -1/2, 1/2\}$ , which is related to Eq. (4.1) through  $\mathcal{P} = \partial\mathcal{H}_{1q}/\partial\varepsilon$ , where  $\varepsilon$  plays the role of an electric field along the axis between the dots. The two-qubit Coulomb interaction can be expressed in terms of the dipole moments  $\mathcal{P}^{(1)}$  and  $\mathcal{P}^{(2)}$  of qubits 1 and 2. We first note that the Coulomb interaction is classical, and therefore diagonal, when expressed in a charge-state basis. In analogy with charge qubits, the interaction can then be fully specified by  $b^{(0)}\mathcal{I}^{(1)} \otimes \mathcal{I}^{(2)} + b^{(1)}\mathcal{P}^{(1)} \otimes \mathcal{I}^{(2)} + b^{(2)}\mathcal{I}^{(1)} \otimes \mathcal{P}^{(2)} + b^{(3)}\mathcal{P}^{(1)} \otimes \mathcal{P}^{(2)}$ . The first term in this expression is a uniform energy shift, which can be ignored. The second and third terms can be absorbed into the detuning parameters through the transformation  $\varepsilon^{(i)} \rightarrow \varepsilon^{(i)} + b^{(i)}$  ( $i = 1, 2$ ). Finally, defining  $g$  as the change in Coulomb energy when one of the qubits flips its charge configuration, the two-qubit Hamiltonian becomes

$$\mathcal{H}_{2q} = \mathcal{H}_{1q}^{(1)} \otimes I^{(2)} + I^{(1)} \otimes \mathcal{H}_{1q}^{(2)} + g\mathcal{P}^{(1)} \otimes \mathcal{P}^{(2)}. \quad (4.3)$$

This form is generic and does not depend on qubit geometry. However, the value of  $g$  depends on the geometry, and has been found to be of order  $75 \mu\text{eV}$  for a linear dot array [94]. The

full 9D basis set for Eq. (4.3) is given by  $\{|\cdot S\rangle, |\cdot T\rangle, |S\cdot\rangle\}^{(1)} \otimes \{|\cdot S\rangle, |\cdot T\rangle, |S\cdot\rangle\}^{(2)}$ , and the corresponding matrix representation for  $\mathcal{H}_{2q}$  is given in Appendix B.1.

We first discuss qubit initialization and the implementation of single-qubit gates. In the large-detuning regime, the qubit logical states are energetically well separated from the leakage states, as shown in Fig. B.1, allowing leakage-free qubit initialization. To gain insight into gate operations, we perform a canonical transformation to decouple the logical states from the leakage states in the large-detuning limit. Additionally, we evaluate this Hamiltonian in the adiabatic basis (the basis that diagonalizes  $\mathcal{H}_{2q}$ ), yielding the effective Hamiltonian

$$\mathcal{H}_{2q,\text{eff}} \simeq \frac{-\hbar\omega_{z1}}{2}\sigma_z^{(1)} \otimes I^{(2)} + \frac{-\hbar\omega_{z2}}{2}I^{(1)} \otimes \sigma_z^{(2)} + \frac{\hbar\omega_{zz}}{2}\sigma_z^{(1)} \otimes \sigma_z^{(2)}, \quad (4.4)$$

where the leading-order contributions to the single-qubit prefactors are of order  $\hbar\omega_{zi} = E_{\text{ST}}^{(i)} + \mathcal{O}[\Delta_z^{(i)2}/\varepsilon^{(i)}]$ , and the effective two-qubit coupling  $\hbar\omega_{zz}$  is discussed below. (See Appendix B.1 for details of the calculation.) Here, the sub- (or super)-script  $i$  identifies the qubit,  $\Delta_z^{(i)2}$  is a quadratic function of the tunnel couplings, and the identity and Pauli matrices,  $I^{(i)}$  and  $\sigma_z^{(i)}$ , act on the logical subspace. Although the detuning parameters provide some control over the qubit resonant frequencies,  $\omega_{z1}$  and  $\omega_{z2}$ ,  $E_{\text{ST}}$  typically varies significantly from dot to dot, resulting in well separated resonances. Single-qubit gates thus proceed by lowering one of the detunings (say,  $\varepsilon^{(1)}$ ) from its high value to a regime where fast ac gates can be performed (e.g., the first dashed line in Fig. 4.1a). At this point, the dipole on qubit 1 is non-negligible;

however, we can operate it near a single-qubit sweet spot to minimize dephasing, as described in Appendix B.2. Since  $\varepsilon^{(2)}$  remains at a large value, there is no danger of implementing either a single-qubit gate on qubit 2, or a two-qubit gate. Superimposing an ac drive on  $\varepsilon^{(1)}$  at the resonant frequency of qubit 1 yields an additional term in Eq. (4.4) proportional to  $\cos(\omega_{z1}t) \sigma_x^{(1)} \otimes I^{(2)}$ , which induces Rabi oscillations about the  $\hat{x}$  axis of the qubit; additional modulation of the phase in  $\cos(\omega_{z1}t + \phi)$  enables rotations about an arbitrary axis in the  $x$ - $y$  plane. To suppress the coupling of the dipole moment to external charge noise, we return  $\varepsilon^{(1)}$  to its large value when the gate is finished.

Next, we consider two-qubit gate operations, which are performed adiabatically, and do not involve ac driving. The canonical transformation leading to Eq. (4.4) yields the leading order result at high detuning,  $\hbar\omega_{zz} = \mathcal{O}[g\Delta^4/\varepsilon^{(1)2}\varepsilon^{(2)2}]$ , where  $\Delta^4$  is a quartic function of the tunnel couplings in both qubits. As anticipated, the adiabatic gate speed  $|\omega_{zz}|/2\pi$  depends linearly on  $g$ , and requires  $\varepsilon^{(1)}$  and  $\varepsilon^{(2)}$  to be simultaneously reduced from their high values to initiate a two-qubit gate. The canonical transformation breaks down when  $\varepsilon^{(1)}$  and  $\varepsilon^{(2)}$  take their low values; however under adiabatic operation, the projection onto the logical subspace, Eq. (4.4), is still meaningful. We can compute  $\hbar\omega_{zz}$  at arbitrary detuning values by evaluating  $\mathcal{H}_{2q}$  in its adiabatic basis and projecting it onto the 4D logical subspace,  $\mathcal{H}_{2q} \rightarrow \mathcal{H}_{4D}$ . We then identify  $\hbar\omega_{zz} = \frac{1}{2}\text{Tr}[(\sigma_z^{(1)} \otimes \sigma_z^{(2)})\mathcal{H}_{4D}] = \frac{1}{2}(E_{00} - E_{01} - E_{10} + E_{11})$ , where  $E_{ij}$  is the energy eigenvalue corresponding to the two-qubit logical state  $|ij\rangle$ . In Fig. 4.1d,e, we plot numerical results for  $\hbar\omega_{zz}$  assuming typical qubit parameters and  $\varepsilon \equiv \varepsilon^{(1)} = \varepsilon^{(2)}$ . Here we observe the predicted asymptotic behavior  $\hbar\omega_{zz} \propto \varepsilon^{-4}$ . We also note that  $\hbar\omega_{zz}$  changes sign

when  $\varepsilon$  is of order  $g$ , in the low-detuning regime where the canonical transformation breaks down.

A simple protocol for implementing adiabatic two-qubit gates is shown in Fig. 4.1f, and can be summarized as follows. We begin with the detuning parameters  $\varepsilon^{(1)}$  and  $\varepsilon^{(2)}$  set to ‘high’ values of  $500 \mu\text{eV}$ , and smoothly lower them to the ‘low’ values  $\varepsilon_{\text{wait}}^{(1)}$  and  $\varepsilon_{\text{wait}}^{(2)}$  over a ramp time  $\tau_{\text{ramp}}$ . The detunings are held constant at these values for a waiting period  $\tau_{\text{wait}}$ , and then smoothly returned to  $\varepsilon^{(1)} = \varepsilon^{(2)} = 500 \mu\text{eV}$  over the same ramp time  $\tau_{\text{ramp}}$ . The parameters defining the pulse sequence are chosen to approximately yield a controlled-Z (CZ) gate operation. This protocol also produces incidental single-qubit  $Z^{(1)}$  and  $Z^{(2)}$  rotations, which can be eliminated, if necessary, by applying additional  $Z^{(1)}$  and  $Z^{(2)}$  gates. Explicit functional forms for  $\varepsilon^{(1)}(t)$  and  $\varepsilon^{(2)}(t)$  are given in Appendix B.3.

We now compute the two-qubit gate fidelity for this sequence including both leakage and charge noise. While we do not explicitly consider the effects of relaxation error, in Appendix B.9 we provide a rough estimate of the expected  $T_1$  times, finding them larger than the gate time. Leakage corresponds to the filling of quantum levels outside the logical subspace, and is primarily caused by non-adiabatic gate pulses. It is taken into account in our simulations by retaining the full 9D Hilbert space, comprising both logical and leakage states. In Section 4.4 we describe a method for computing the process fidelity of a CZ gate in the presence of charge noise. This procedure allows us to identify optimal values of  $\tau_{\text{ramp}}$  and  $\tau_{\text{wait}}$ , consistent with fast pulse sequences, low leakage, and high fidelity. Figure 4.2 shows the results of such fidelity calculations, for a range of  $\varepsilon_{\text{wait}}^{(1)}$  and  $\varepsilon_{\text{wait}}^{(2)}$  values, assuming



the typical quantum dot parameters  $E_{\text{ST}}^{(1)} = 52 \mu\text{eV}$ ,  $E_{\text{ST}}^{(2)} = 47 \mu\text{eV}$ ,  $g = 75 \mu\text{eV}$ , and  $\sigma_\epsilon = 1 \mu\text{eV}$ . Here, we choose  $\Delta_1^{(i)} = 0.64E_{\text{ST}}^{(i)}$  and  $\Delta_2^{(i)} = 0.58E_{\text{ST}}^{(i)}$ , which suppresses the single-qubit dephasing, as discussed in Appendix B.2. We also omit pulse sequences with total gate times  $\tau_{\text{total}} = 2\tau_{\text{ramp}} + \tau_{\text{wait}} > 50 \text{ ns}$  (the cross-hatched regions in the plot), to ensure that entangling gates are performed on a timescale comparable to the QDHQ single-qubit gates [99]. The fastest pulse sequence in the viable regime corresponds to  $\epsilon_{\text{wait}}^{(1)} = 90 \mu\text{eV}$ ,  $\epsilon_{\text{wait}}^{(2)} = 110 \mu\text{eV}$ ,  $\tau_{\text{ramp}} = 2.4 \text{ ns}$ , and  $\tau_{\text{wait}} = 2.8 \text{ ns}$  ( $\tau_{\text{total}} = 7.6 \text{ ns}$ ), and exhibits an average process infidelity of  $9.9 \times 10^{-3}$ . The highest-fidelity sequence corresponds to  $\epsilon_{\text{wait}}^{(1)} = 80 \mu\text{eV}$ ,  $\epsilon_{\text{wait}}^{(2)} = 100 \mu\text{eV}$ ,  $\tau_{\text{ramp}} = 4.0 \text{ ns}$ , and  $\tau_{\text{wait}} = 8.0 \text{ ns}$  ( $\tau_{\text{total}} = 16.0 \text{ ns}$ ), with an average process infidelity of  $9.2 \times 10^{-3}$ . This optimized value depends on the standard deviation of the charge noise,  $\sigma_\epsilon$ . In Fig. 4.2b, we plot the minimized CZ gate infidelity  $\mathcal{I}$  as a function of  $\sigma_\epsilon$ , revealing the scaling behavior  $\mathcal{I} \propto \sigma_\epsilon^2$ .

The strong dependence of infidelity on  $\sigma_\epsilon$  indicates that dephasing, rather than leakage, is the main source of gate errors. In Appendix B.7, we explain the observed behavior by assuming that charge noise is quasistatic, obtaining

$$\mathcal{I}_{\text{cn}} \approx \frac{1}{4} \sigma_\epsilon^2 \sum_{i=1,2} \sum_{j=z1,z2,zz} \left( \int \frac{\partial \omega_j}{\partial \epsilon^{(i)}} dt \right)^2, \quad (4.5)$$

as expected in the absence of a sweet spot [56]. To confirm the absence of a sweet spot, we perform an exhaustive search over the detuning ( $\epsilon^{(i)}$ ), tunnel coupling ( $\Delta_j^{(i)}$ ), and Coulomb ( $g$ ) parameters in Eq. (4.3), finding that it is impossible to simultaneously set  $\partial \omega_j / \partial \epsilon^{(i)} = 0$ ,

for all  $i$  and  $j$ , in the parameter range of interest. However, it is clear that this conventional, time-independent definition of a sweet spot is overly restrictive for ensuring that  $\mathcal{I}_{\text{cn}} \approx 0$  in Eq. (4.5).

We now introduce the concept of a dynamical sweet spot (DSS) in which the *time-averaged* derivatives in Eq. (4.5) are made to vanish, as sketched in Fig. 4.3a. Through an exhaustive search (Fig. 4.2a), we have already demonstrated that no DSS exists for the simple pulse sequence of Fig. 4.1f. Moreover, because of the monotonic dependence of  $\omega_j$  on  $\varepsilon^{(i)}$  (for example, see Fig. B.1), it appears unlikely that a more elaborate detuning pulse sequence could provide significant improvements in the fidelity. We therefore augment the detuning sequence with a tunnel-coupling sequence,  $\Delta_j^{(i)}(t)$  ( $i, j = 1, 2$ ). Our initial investigations suggest that a relatively large number of pulse parameters are needed to achieve significant improvements in the fidelity. We therefore consider the more elaborate pulse shape, shown in Fig. 4.3b, which incorporates seven parameters for each of four tunnel couplings.

Because of the large number of parameters in the combined detuning-tunnel-coupling sequence, we do not attempt an exhaustive search in this case. Instead, we maximize the CZ gate fidelity by performing a hundred separate Broyden-Fletcher-Goldfarb-Shanno [105, 106, 107, 108] (BFGS) searches using the method of ref. [93], and adopting a basin-hopping protocol with randomized initial values to help escape any local minima [109]. To simplify the calculation, we adopt the following hybrid infidelity functional:  $\mathcal{I}_{\text{total}} = \mathcal{I}_{\text{cn}} + \mathcal{I}_{\text{nf}} + \mathcal{I}_{\text{na}}$ , which treats the charge noise (cn), noise-free (nf), and non-adiabatic (na) infidelity contributions separately. Calculating  $\mathcal{I}_{\text{total}}$  is computationally efficient because all three contributions,

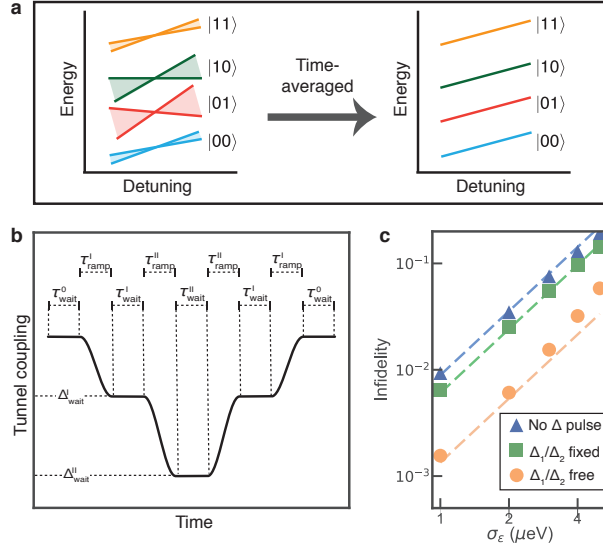


Figure 4.3: **Dynamical sweet spot (DSS) and a tunnel-coupling pulse sequence.** **a**, Two-qubit energy levels plotted schematically as a function of a single detuning variable. At any given time, it is difficult to arrange for all the energy dispersions to be parallel, as indicated on the left, leaving the qubits susceptible to dephasing. However, it may be possible to construct a pulse sequence for which the levels vary in time (shaded regions), such that their time-averaged dispersions are parallel, yielding a DSS that is more resilient to quasistatic fluctuations of the detuning. Here, we explore a DSS construction in which the detuning pulse sequence of Fig. 4.1f is augmented with the tunnel coupling pulse sequence defined in **b** and Appendix B.3. The latter is simple enough that it can be optimized using the method described in the main text, but complex enough that it provides a significant improvement in the CZ gate fidelity. **c**, Infidelities calculated for three different pulse sequences as a function of the standard deviation of the quasistatic charge noise  $\sigma_\epsilon$ . The markers correspond to full gate simulations averaged over a charge noise distribution, as described in Methods. The dashed lines correspond to the much simpler infidelity estimate of Eq. (4.5). For the blue line and triangles, the tunnel couplings are held constant, as in Fig. 4.1. For the green line and squares, the tunnel couplings are pulsed as in **c** with the ratios  $\Delta_1^{(i)}/\Delta_2^{(i)}$  held constant. For the orange line and circles, the tunnel coupling sequence is optimized with no constraints.

including the charge noise term defined in Eq. (4.5), do not require taking an average over charge noise. The noise-free term describes the CZ gate infidelity in the absence of charge noise, as described in Methods. We find that minimizing just the  $\mathcal{I}_{\text{cn}}$  and  $\mathcal{I}_{\text{nf}}$  terms (without  $\mathcal{I}_{\text{na}}$ ) yields extremely short and fast pulse sequences that first populate then depopulate the leakage state. Since  $\mathcal{I}_{\text{cn}}$  was derived assuming an adiabatic pulse, these short and fast pulses are not guaranteed to have low process infidelity. Hence we also introduce the  $\mathcal{I}_{\text{na}}$  term, as defined in Methods, to explicitly penalize non-adiabatic evolution.

We now obtain two different sets of solutions for the detuning-tunnel-coupling pulse sequence. In the first, all the tunnel coupling parameters in Fig. 4.3b, as well as  $\tau_{\text{ramp}}$  and  $\tau_{\text{wait}}$ , are varied independently, under the constraint that the detuning and tunnel coupling sequences have the same length; this sequence contains 26 free parameters. The second case is similar, except that the ratio between the tunnel couplings in each double dot is assumed to be fixed throughout the sequence, with  $\Delta_1^{(i)}(t)/\Delta_2^{(i)}(t) = 1.1034$ , as consistent with Appendix B.2; this sequence contains 14 free parameters. In both cases, we set the detuning parameters to the values giving the fastest detuning-only pulse sequence in Fig. 4.2 ( $\varepsilon_{\text{wait}}^{(1)} = 90 \mu\text{eV}$  and  $\varepsilon_{\text{wait}}^{(2)} = 110 \mu\text{eV}$ ), and we use  $\tau_{\text{ramp}} = 2.4 \text{ ns}$  and  $\tau_{\text{wait}} = 2.8 \text{ ns}$  as the starting points for our optimization procedure; initial values of the other parameters are chosen randomly, according to the basin-hopping protocol. The results of this procedure are presented in Table B.1. Using these results, we recompute the infidelity as described in Methods, performing an average over the charge noise.

Infidelity results using the tunnel coupling pulse sequence are plotted in Fig. 4.3c as a

function of charge noise,  $\sigma_\varepsilon$ . Here we observe clear improvements compared to the detuning-only sequence, with the best results obtained for the sequence with the largest number of pulse parameters. Appendix B.8 suggests that this result can largely be attributed to the suppression of the time-averaged derivatives  $\partial\omega_j/\partial\varepsilon^{(i)}$ , as consistent with a DSS. For a true sweet spot, we would expect a power-law exponent in  $\mathcal{I} \propto \sigma_\varepsilon^\alpha$ , with  $\alpha > 2$ . Although the large- $\sigma_\varepsilon$  data in Fig. 4.3c hint at such behavior, Fig. B.6 indicates that a full suppression of the time-averaged derivatives has not yet been achieved in the current pulse sequences.

### 4.3 Discussion

We have proposed a scheme for entangling capacitively coupled quantum-dot hybrid qubits by applying adiabatic pulse sequences to detuning parameters. We have optimized the sequences in the presence of quasistatic charge noise and computed the resulting process fidelities for a controlled-Z gate, obtaining fidelities approaching 99% for typical noise levels. Further improvements are obtained by simultaneously applying pulse sequences to the tunnel couplings. These results are explained by invoking the concept of a dynamical sweet spot (DSS), for which the splittings between the two-qubit energy levels are insensitive to fluctuations of the detuning parameters when averaged over the whole pulse sequence. Our analysis shows that a true DSS cannot be achieved using simple pulse sequences. However, fidelities >99% are achieved when the pulse sequences include a large number of tunable parameters. As indicated by ref [110], these fidelities can be further improved by exploring a wider range

of pulse shapes. Using the GRAPE algorithm [111], we have identified high-frequency pulses that significantly improve the expected performance under quasistatic charge noise. Despite the potential improvement in fidelity, the experimental bandwidth limitations make implementing some of these pulses impractical. We speculate that a bandwidth-limited version [112, 113] of the GRAPE algorithm [111] could be used to explore a much larger parameter space of adiabatic pulse sequences, possibly allowing us to identify a true DSS. The GRAPE algorithm also provides a means for exploring non-adiabatic pulse sequences. However the simplicity and relatively high fidelity achieved with the sequences studied here, and the robustness of adiabatic gating methods, make the current proposal attractive for two-qubit gates.

## 4.4 Methods

To study the time evolution resulting from the pulse sequences applied to capacitively coupled qubits, we numerically integrate the Hamiltonian in Eq. (4.3), for which the time-dependent control parameters  $\tilde{\varepsilon}^{(1)}(t)$ ,  $\tilde{\varepsilon}^{(2)}(t)$ , and  $\vec{\Delta}(t)$  depend on the particular pulse sequence. Here,  $\vec{\Delta}(t)$  refers to the set of four intra-qubit tunnel couplings, and we define  $\tilde{\varepsilon}^{(i)}(t) = \varepsilon^{(i)}(t) + \delta\varepsilon^{(i)}$ , where  $\varepsilon^{(i)}(t)$  is the ideal, noise-free detuning sequence for qubit  $i$ , and the (quasi-static) noise term  $\delta\varepsilon^{(i)}$  is assumed to remain constant for the duration of the sequence. The resulting

unitary operator is given by

$$U_{2q}(t) = \exp \left[ -i/\hbar \int_0^t \mathcal{H}_{2q} \left( \varepsilon^{(1)}(t'), \varepsilon^{(2)}(t'), \vec{\Delta}(t') \right) dt' \right]. \quad (4.6)$$

In most cases, we take  $t$  to be the final time in the pulse sequence, with one exception, described below.

We employ the following procedure to determine the detuning pulse parameters used in Fig. 4.2. (For additional details, see Appendix B.4.) We first choose the fastest ramp time  $\tau_{\text{ramp}}$  consistent with leakage errors  $< 0.1\%$  in the absence of charge noise. We then compute  $U_{2q}$  as a function of  $\tau_{\text{wait}}$  for a fixed level of quasi-static charge noise. (High-frequency noise can also affect the fidelity of slow QDHQ gates [114]; however we do not consider such processes here.) The process fidelity  $\mathcal{F}$  is computed, comparing  $U_{2q}$  to a perfect CZ gate, modulo single-qubit rotations, using the Choi-Jamiolkowski formalism [115], as described in Appendix B.5. We then average  $\mathcal{F}$  over charge noise, using the method described below, and choose the value of  $\tau_{\text{wait}}$  that maximizes  $\langle \mathcal{F} \rangle$ .

To optimize the detuning-tunnel-coupling pulse sequence used in Fig. 4.3, we choose pulse parameters that minimize the total infidelity function  $\mathcal{I}_{\text{total}} = \mathcal{I}_{\text{cn}} + \mathcal{I}_{\text{nf}} + \mathcal{I}_{\text{na}}$ , as discussed in the main text. Here, the noise-free term  $\mathcal{I}_{\text{nf}}$  describes the CZ gate infidelity, computed using the Choi-Jamiolkowski formalism, as described above, in the absence of charge noise. In this work, we also introduce a penalty term to suppress non-adiabatic evolution, defined as  $\mathcal{I}_{\text{na}} = \max_t \left[ 1 - \frac{1}{4} \sum |\langle ij(t) | U_{2q}(t) | ij(0) \rangle|^2 \right]$ , where the sum is taken over the logical basis

states  $(i, j) = (0, 1)$ , and the function  $\max_t$  picks out the maximal occupation of leakage states, at any point in the pulse sequence. Note that the state  $|ij(0)\rangle$  is an eigenstate of  $H_{2q}(t)$  at time  $t = 0$ , while  $|ij(t)\rangle$  is the corresponding eigenstate at time  $t$ . Under perfect adiabatic operation, the mapping  $U_{2q}(t) |ij(0)\rangle \rightarrow |ij(t)\rangle$  is exact, yielding  $\mathcal{I}_{\text{na}} = 0$ ; however for non-adiabatic operation, we obtain  $\mathcal{I}_{\text{na}} > 0$ . In practice, we find that the exact form of  $\mathcal{I}_{\text{na}}$  does not significantly affect our results.

To average the fidelity over charge noise, we assume that the noise terms  $\delta\varepsilon^{(1)}$  and  $\delta\varepsilon^{(2)}$  are uncorrelated and sample them independently at 17 values in the range between -25 and +25  $\mu\text{eV}$ , corresponding to 1089 unique pairs. We then linearly interpolate  $\mathcal{F}$  over the results and calculate its average value, assuming a gaussian probability distribution with standard deviation  $\sigma_\varepsilon$ :

$$p(\delta\varepsilon^{(1)}, \delta\varepsilon^{(2)}) = \frac{1}{2\pi\sigma_\varepsilon^2} \exp\left(-\frac{\delta\varepsilon^{(1)2} + \delta\varepsilon^{(2)2}}{2\sigma_\varepsilon^2}\right). \quad (4.7)$$

## 4.5 Funding acknowledgements

This work was supported in part by ARO (W911NF-12-1-0607, W911NF-17-1-0274) and the Vannevar Bush Faculty Fellowship program sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering and funded by the Office of Naval Research through Grant No. N00014-15-1-0029. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office (ARO), or the U.S.



Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. The authors gratefully acknowledge support from the Sandia National Laboratories Truman Fellowship Program, which is funded by the Laboratory Directed Research and Development (LDRD) Program.

## Chapter 5

# Power law scaling for the adiabatic algorithm for search engine ranking

An important method for search engine result ranking works by finding the principal eigenvector of the “Google matrix.” Recently, a quantum algorithm for generating this eigenvector as a quantum state was presented, with evidence of an exponential speedup of this process for some scale-free networks. Here, we show that the run-time depends on features of the graphs other than the degree distribution, and can be altered sufficiently to rule out a general exponential speedup, consistent with the findings of Garnerone, Zanardi, and Lidar. According to our simulations, for a sample of graphs with degree distributions that are scale-free, with parameters thought to closely resemble the Web, the proposed algorithm for eigenvector preparation does not appear to run exponentially faster than the classical case.

## 5.1 Introduction

Quantum algorithms, which run on quantum computers, are known to be able to outperform classical algorithms for certain computational problems [2, 116]. Thus, finding a new algorithm that exhibits a quantum speedup, in particular an exponential speedup, is of great interest [117]. An extremely important problem in computer science is calculating ranking for search engine results. PageRank, first proposed by Brin and Page [118] underlies the success of the Google search engine [119]. In this algorithm, websites are represented as nodes on a network graph, connected by directed edges that represent links. The matrix of network connections is constructed, and the PageRank vector is its principal eigenvector. Currently, computing the PageRank vector requires a time  $O(n)$ , where  $n$  is the number of websites in the network considered (e.g. the World Wide Web) [36]. Obtaining a quantum algorithm for PageRank that runs exponentially faster than the classical algorithm would be of great interest.

Recently, Garnerone, Zanardi, and Lidar (GZL) proposed an adiabatic quantum algorithm [17] to prepare the PageRank vector for a given network [36]. Remarkably, GZL present evidence that this algorithm can prepare the PageRank vector in time  $O[\text{polylog}(n)]$ , exponentially faster than classical algorithms for certain networks. It is important to note that the algorithm only offers an exponential speedup in the preparation of the PageRank vector; the information must then be extracted from the quantum state, which ultimately yields a polynomial quantum speedup. This polylogarithmic runtime is due to the apparent logarithmic scaling of the gap between the two smallest eigenvalues of the Hamiltonian

used in the algorithm (the energy gap). This scaling emerged on graphs constructed using adapted versions of two established methods of network construction: the preferential attachment model [120] and the copying model [121]. Both of these models yield graphs that are similar to the connectivity of the World Wide Web in that they are sparse (the total number of edges scales at most proportionally to the number of nodes) and scale-free (the probability of finding a node with a specified in- or out-degree scales as a power law in those degrees). These features lead to networks that exhibit large-scale structure similar to that of the internet, such as being small-world [122] and loosely hierarchical [123]. GZL studied sets of networks that exhibited both logarithmic scaling and polynomial scaling of the gap in the system size. However, they did not demonstrate that the networks with the favorable logarithmic gap scaling are scale-free over the region studied numerically.

Here, we study the scaling of the GZL algorithm for graphs with degree distributions consistent with the internet. A realistic network model of the World Wide Web must be scale-free in both the in- and the out-degree [124, 125]. We consider a broad variety of scale-free networks constructed by different methods. Choosing three well-known models for constructing random, scale-free networks, we control for both the mean degree and the exponent of the power-law governing the degree distribution. We find that graphs with the same degree distribution can have different energy gap and run-time behaviors. Finally, we focus on degree distributions described by power laws consistent with those measured for the Web, both for the in-degree and the out-degree. We find that the relevant energy gap scales as a power of the system size, rather than logarithmically. These results demonstrate that for

scale-free graphs with parameters thought to resemble the Web, the GZL adiabatic algorithm does not appear to yield an exponential quantum speedup for preparing the PageRank vector compared to current classical algorithms.

## 5.2 Network growth models

We generate samples of graphs with prescribed degree distributions using three different network growth models. GZL [36] use modified versions of two network construction algorithms: the preferential attachment model [120] and the copying model [121]. In addition to these two models, here we include also the more complex  $\alpha$ -preferential attachment model described by Bollobás *et al.* [124, 126]. All three models grow random networks using probabilistic rules at discrete construction steps, which are detailed in Fig. 5.1.

All three of these models produce sparse, scale-free directed networks, in which the probability of the in-degree (the number of incoming edges) and out-degree (the number of outgoing edges) of node  $i$  being equal to  $k$  are each proportional to a power law:

$$P(d_{in}(i) = k) \sim k^{-\gamma_{in}} \tag{5.1}$$

$$P(d_{out}(i) = k) \sim k^{-\gamma_{out}}, \tag{5.2}$$

where  $d_{in}(i)$  and  $d_{out}(i)$  are the in- and out-degrees of node  $i$ , respectively, and the exponents  $\gamma_{in}$  and  $\gamma_{out}$  are typically between 2 and 4 [120]. The GZL versions [36] of the preferential attachment and copying models [120, 127] produce networks that are scale-free in the limit

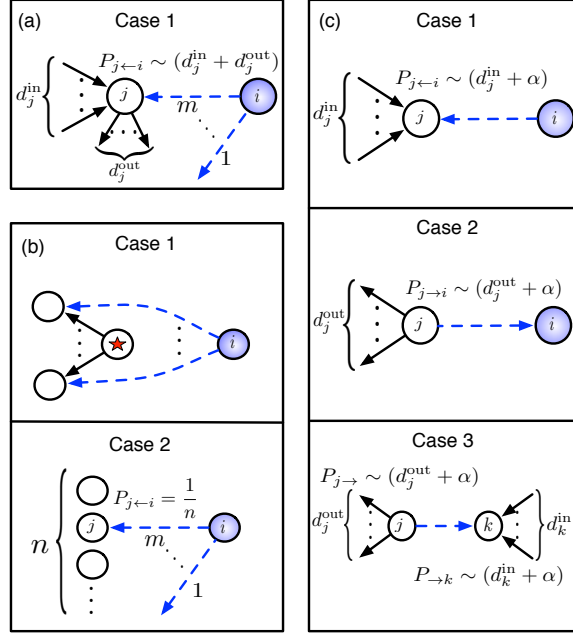


Figure 5.1: Illustrations of the three network generation models used. (a): GZL [36] preferential attachment, (b): GZL copying, and (c):  $\alpha$ -preferential attachment [124, 126]. In all three models, a network is constructed by adding vertices and edges sequentially. (a): At each time step a new vertex  $i$  is added with  $m$  outgoing edges. The probability that one of these edges connects to a node  $j$  is proportional to the total degree of  $j$ . (b): At each time step there are two possible actions. With probability  $(1 - p)$ , the new vertex points to all of the same vertices as the “star vertex,” which is a pre-existing vertex chosen uniformly at random at each time step. With probability  $p$ ,  $m$  outgoing edges are added to the new vertex, each pointing to vertices chosen uniformly at random. (c): There are three possible actions at each time step. With probability  $p_1$ , a new vertex is added with a single outgoing edge, pointing to a node  $j$  with probability proportional to the in-degree of  $j$  plus a parameter  $\alpha$ . With probability  $p_2$ , a new vertex is added with a single incoming edge, pointing from a node  $j$  with probability proportional to the out-degree of  $j$  plus  $\alpha$ . With probability  $(1 - p_1 - p_2)$ , no vertex, only an edge, is added. Its ending and starting points are determined as in cases 1 and 2, respectively. In all panels, the newly-added edges are indicated by dashed lines.

of large graph size. However, due to the addition procedure described below, the networks are not necessarily scale-free for the sizes of graphs studied numerically here and in Ref. [36]. To achieve networks that are scale-free in the out-degree, GZL suggest to construct two networks,  $X$  and  $Y$ , independently.  $X$  and  $Y$  are each generated as in Fig. 5.1, except that for  $Y$  the direction of the edges added is reversed. The networks can then be added together, and the weights and loops discarded [36, 128]. The resulting composite network is scale-free in both in-degree and out-degree, provided  $X$  and  $Y$  have the same number of edges per node. (See Appendix C.1 for details.) In contrast to Ref. [36], the graphs studied here are all constrained in this way. However, the graphs exhibiting logarithmic scaling in [36] are not so constrained [128], and so they do not exhibit truly scale-free degree distributions over the numerically studied region. On the other hand, the  $\alpha$ -preferential attachment model (considered here but not in [36]) constructs a network which is scale-free in both in- and out-degrees without requiring an additional combination step. As with the GZL preferential attachment model, all weights and loops are removed from the final  $\alpha$ -preferential attachment network.

The exponents  $\gamma_{in}$  (Eq. 5.1) and  $\gamma_{out}$  (Eq. 5.2) of the degree distribution are model-dependent. In the GZL preferential attachment model the number of edges added at each construction step controls the sparsity, and it is always the case that  $\gamma_{in} = \gamma_{out} = 3$  [120]. Both the GZL copying model and  $\alpha$ -preferential attachment allow for independently tunable exponents and mean degree. (See Appendix C.1 for details.) This flexibility enables us to create three ensembles of model networks that have nearly identical degree distributions for

$\gamma_{in} = \gamma_{out} = 3$ . Further, the last two models can be set with the exponents estimated for the World Wide Web [121, 124], namely  $\gamma_{in} = 2.1$  and  $\gamma_{out} = 2.72$  [125].

### 5.3 Algorithm description

The Google matrix is constructed by taking as input an unweighted, simple network with  $n$  nodes [118], and representing it as an adjacency matrix  $A$ , where  $A(i, j) = 1$  if a directed edge points from node  $i$  to node  $j$ , and 0 otherwise. From this, one defines the matrix  $P$ :

$$P(i, j) = \begin{cases} 1/d_{out}(i) & \text{if } A(i, j) = 1 \\ 1/n & \text{if } \forall j, A(i, j) = 0 \\ 0 & \text{otherwise} \end{cases} \quad \begin{array}{l} (5.3a) \\ (5.3b) \\ (5.3c) \end{array}$$

The matrix  $P$  is stochastic because  $\sum_j P(i, j) = 1$  for all  $i$ .  $P$  can be thought of as a random walk (i.e. a web-surfer), where the walker follows the network with equal likelihood of traversing all allowed links. If the walker ever reaches a dangling node (a node with  $d_{out} = 0$ ), Eq. 5.3b implies that it can randomly hop to any vertex with equal probability. To prevent the walker from becoming trapped in an isolated portion of the network (a sink), the probability  $(1 - \alpha_g)$  of moving to a node uniformly at random (including the possibility of staying still) is included, where  $0 < \alpha_g < 1$ ; Google uses  $\alpha_g = 0.85$ , which we also use here [36]. The Google matrix  $G$  is defined as the transpose of this resulting transition matrix:

$$G = \alpha_g P^T + (1 - \alpha_g) J, \quad (5.4)$$



where  $J$  is the matrix of all ones. The PageRank vector  $\vec{p}$  is the unique eigenvector associated with the largest eigenvalue of  $G$ , which is 1. The runtime of the best classical algorithm, which calculates the PageRank vector via power iteration, is  $O(n)$  [36, 118].

To formulate an adiabatic quantum algorithm, GZL construct the Hamiltonian  $h(G)$ :

$$h(G) = (\mathbb{I} - G)^\dagger (\mathbb{I} - G), \quad (5.5)$$

which is Hermitian, even though  $G$  is not. The ground state of this Hamiltonian is the normalized PageRank vector. The adiabatic algorithm is completely defined by the interpolation Hamiltonian  $H(s) = sh(G) + (1 - s)h(G_c)$ , where  $s \in [0, 1]$ , and  $G_c$  is the Google matrix for the complete network (including loops), whose ground state is a uniform superposition. The adiabatic theorem guarantees that if we initialize our system in the ground state of  $h(G_c)$  and change  $s$  from 0 to 1 sufficiently slowly, the system remains in the ground state [17]. Since the PageRank vector is the ground state of  $H(1) = h(G)$ , the PageRank vector is obtained when  $s = 1$ . The required slowness is also determined by the adiabatic theorem: as long as  $s(t)$  is a smooth function of the time  $t$  with  $0 \leq t \leq T$ , the runtime  $T \sim \delta^{-b}$ , where  $b$  is  $O(1)$  and  $\delta$  is the energy gap between the ground and first excited state of  $H(s)$ , minimized over  $s$  [17]. Thus, an exponential speedup over the classical case is possible if  $\delta^{-1}$  is  $O[\log(n)]$ , since then  $T$  is  $O[\text{polylog}(n)]$ .

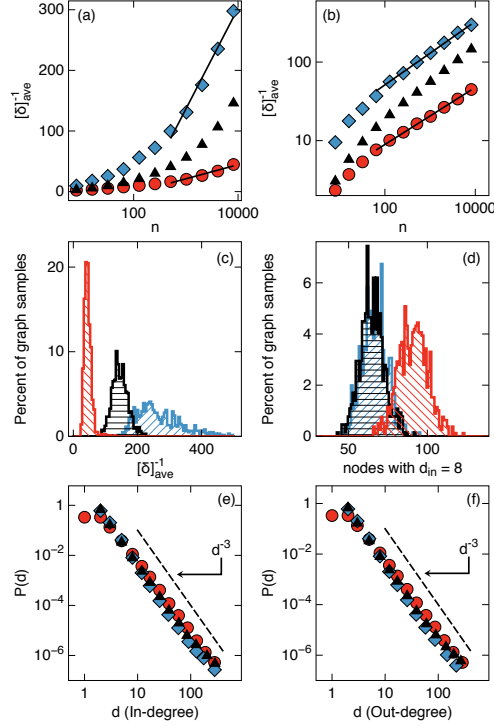


Figure 5.2: Comparison of the scaling of the inverse energy gap  $\delta^{-1}$  for the GZL [36] preferential attachment model (triangles, horizontal hatching), GZL copying model (diamonds, upward-sloping hatching), and  $\alpha$ -preferential attachment model [124] (circles, downward-sloping hatching), shown on (a): Semilog and (b): Log-Log scales, demonstrating that  $\delta^{-1}$  is not proportional to  $\log(n)$  for these models. Results are averaged over 1000 random instances for  $n < 8192$ , and over 500 random instances at  $n = 8192$ . The fitting lines showed in (a) are  $72.2 \cdot \ln(n) - 363$  for the copying model and  $10.1 \cdot \ln(n) - 48.8$  for the  $\alpha$ -preferential attachment model. In (b), the fits shown are  $8.0 \cdot n^{0.4}$  for the copying model and  $1.7 \cdot n^{0.4}$  for the  $\alpha$ -preferential attachment model. If we fit the data instead to a power of a logarithm (not shown), we obtain  $0.56 \cdot \ln^{2.9}(n)$  for the copying model and  $0.18 \cdot \ln^{2.5}(n)$  for the  $\alpha$ -preferential attachment model. (c): Histogram of the inverse energy gaps for the data shown in panels (a)-(b) at  $n = 8192$ . (d): Histogram showing the distribution of number of vertices with in-degree  $d_{in} = 8$  for  $n = 8192$ . (e)-(f): Degree-distributions of the three models, demonstrating scale-free behavior and indicating that  $\gamma_{in} = \gamma_{out} = 3$ . Adaptive binning was used, as described in Appendix C.3. In all cases, both the mean in- and out-degree of each graph are 2 edges per node. These results demonstrate that  $\delta^{-1}$  differs significantly for the different graph construction methods, while the degree distributions are very similar.

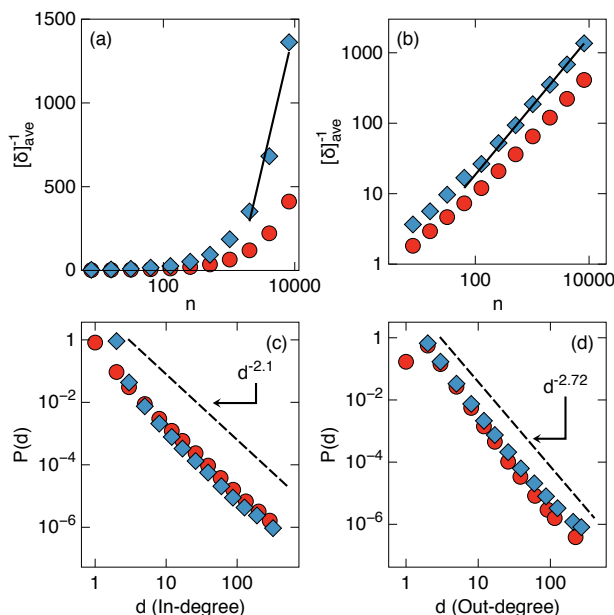


Figure 5.3: Inverse energy gap scaling for GZL [36] copying model (diamonds), and  $\alpha$ -preferential attachment model [124] (circles) of WWW-like networks, shown on (a): Semilog and (b): Log-Log scales. Results are averaged over 1000 random instances for  $n < 8192$ , and over 500 random instances at  $n = 8192$ . In (a), the line fit shown is  $730 \cdot \ln(n) - 5300$ , while in (b) the line fit is  $0.2 \cdot n^{0.97}$ . If we fit the data to a power of a logarithm (not shown), for the copying model we obtain  $3 \times 10^{-5} \cdot \ln^{8.0}(n)$ . Because of the large power of the logarithm required for the polylogarithmic fit, the power-law dependence on  $n$  appears more natural and plausible. (c)-(d): Degree-distributions of the two models, histogrammed using adaptive binning (see Appendix C.3), indicating that  $\gamma_{\text{in}} = 2.1$  and  $\gamma_{\text{out}} = 2.72$ , corresponding to the estimates for the degree distribution of the World Wide Web [125]. In all cases, the mean in- and out-degree of each network were each 2 edges per node.

## 5.4 Numerical results

To study the scaling of the minimum energy gap  $\delta$  with the network size  $n$ , we compute  $\delta$  for the GZL Hamiltonian  $H(s)$ , averaging the results over many network realizations (typically 1000). Specifically, we calculate the minimum value of  $\delta$  over  $s \in [0, 1]$  using the Nelder-Mead method [129], where each objective function call calculates directly the eigenvalue spectrum of  $H(s)$ . We find that for most, but not all, network choices the minimum gap occurs when  $s = 1$ . Since  $H(s)$  is a dense matrix, this process is computationally intensive. We use the University of Wisconsin-Madison Center for High Throughput Computing and Open Science Grid to perform the simulations.

To assess whether the inverse energy gap  $\delta^{-1}$  scales logarithmically or as a power-law in  $n$ , we plot in Fig. 5.2  $\delta^{-1}$  versus the network size on both log-linear and log-log scales, with data for the GZL preferential attachment, GZL copying, and  $\alpha$ -preferential attachment models. The model parameters are tuned (see Appendix C.1) so that all three have  $\gamma_{in} = \gamma_{out} = 3$  and have an average of 2 in- and 2 out-edges per node. Despite having nearly identical degree distributions (shown in Figs. 5.2(e) and 5.2(f)), the scaling of  $\delta^{-1}$  depends significantly on the method used to construct the graphs when viewed in Fig. 5.2(a). In Fig. 5.2(c), we show the distribution corresponding to the final data points in Fig. 5.2(a), where we see that the distributions are well-separated and hence the construction models give different values of  $\delta^{-1}$ . By contrast, the degree distributions are difficult to distinguish, as shown in Fig. 5.2(d). Finally, we conclude that for all three models, the data are more consistent with  $\delta^{-1}$  scaling as a power law or a high-order polylogarithm, rather than a logarithm, as consistent with

the data presented by GZL in the supplemental information of Ref. [36].

We next perform a similar analysis for degree distributions more closely related to the network of primary interest, the World Wide Web, for which a realistic set of degree parameters is given by  $\gamma_{in} = 2.1$  and  $\gamma_{out} = 2.72$  [125]. As mentioned above, the preferential attachment model cannot be tuned to obtain degree parameters other than 3. However, the other two network models can be adjusted to match these values [121, 124]. More details on this are discussed in Appendix C.1. As before, we set the mean degree to be 2 in- and 2 out-edges per node.

Fig. 5.3 presents the results of these simulations, clearly indicating that  $\delta^{-1}$  scales at least as a power of  $n$ . In particular, we note that the prefactor of the logarithmic fit is over 700 and the power of the logarithm in the polylogarithmic fit is 8, while the power law fit exponent is close to one. The results do not change substantially when the mean degree is varied and the degree distributions exponents are fixed. These data indicate that for graphs with degree distributions similar to those measured for the World Wide Web, the GZL adiabatic algorithm for PageRank vector preparation is unlikely to provide an exponential speedup over the classical case.

## 5.5 Discussion

We have investigated the recently proposed adiabatic quantum algorithm for preparing the PageRank vector using an adiabatic quantum algorithm [36]. We find that the eigenvalue

gap that determines the algorithm runtime depends on the method of construction of the network, even when the feature believed to be critical for large-scale network structure, the degree distribution, is held fixed. The exponent governing the variation of the gap with graph size does not vary significantly with the method of construction only if power-law scaling of the gap with size is assumed. For networks that are scale-free in their in- and out-degree distributions, and particularly when the degree distributions similar to those measured for the World Wide Web, our numerical results indicate strongly that the GZL adiabatic algorithm for PageRank vector preparation does not offer an exponential speedup over current classical algorithms.

Consistent with the findings of GZL, we observe that altering the degree distribution changes the scaling properties for preparation of the PageRank state. As our research shows that the degree distribution and the average number of edges per node do not fully determine the algorithm's performance; it is not currently known if additional network features contribute to the polylogarithmic gap scaling observed by GZL. Answering this question would be an interesting topic for future work.

## 5.6 Funding acknowledgements

This work was supported in part by ARO, DOD (W911NF-09-1-0439) and NSF (CCR-0635355, DMR 0906951). AF acknowledges support from the NSF REU program (PHY-PIF-1104660). We thank S. Garnerone, D. A. Lidar, and D. Bradley for useful discussions.

We also thank the HEP, Condor, and CHTC groups at University of Wisconsin-Madison for computational support.

# Chapter 6

## Conclusion

In this thesis, we have described work concerning the performance of semiconductor-based devices intended for quantum computation. We began by focusing on the materials used to fabricate semiconductor-based qubit devices. We studied the strain induced on the semiconductor domain of a simulated device due to the cooling of Al electrodes in the vicinity of the studied domain. Specifically, we presented the results of a simulation modeling the expected strain in a simulated device, and found the corresponding shift in the conduction band of the region intended to contain the quantum dots.

We next considered the design of semiconductor-based qubit devices, introducing a protocol, which we call the Compressed Optimization of Device Architectures (CODA). This protocol both efficiently identifies sparse sets of voltage changes that control quantum systems, and introduces a metric which can be used to compare device designs. As an example of the former, we applied this method to simulated devices with up to 100 quantum dots



and showed that CODA automatically tunes devices more efficiently than other common nonlinear optimizers. To demonstrate the latter, we determined the optimal lateral scale for a triple quantum dot, yielding a simulated device that can be tuned with small voltage changes on a limited number of electrodes.

Finally, we studied the operation of semiconductor-based qubit devices. We considered a system comprised of two capacitively-coupled quantum dot hybrid qubits, and developed a simple electronic pulse sequence that achieves a high-fidelity two-qubit gate under realistic levels of quasistatic charge noise. We also considered the performance of a quantum algorithm for PageRank which takes as input a graph of the World Wide Web. We showed that the runtime of this algorithm depends on features of the graphs other than the degree distribution, and can be altered sufficiently to rule out a general exponential speedup.

# Appendix A

## Supplemental information for Ch. 3

### A.1 CODA protocol

Here, we provide further details about the CODA protocol described in the main text. The simulated device is considered to be a nonlinear function  $\hat{S} : \mathcal{V} \rightarrow \mathcal{Q}$ , where  $\mathcal{V}$  is the space of electrode voltages and  $\mathcal{Q}$  is the space of quantities of interest (*e.g.*, dot occupations, dot energies, transmission coefficients). Suppose that we have  $n$  quantities of interest and  $m$  electrode voltages, and that  $m > n$ , so the system is underconstrained. We first identify a starting point of experimental interest  $(\mathbf{v}_{\text{op}}^0, \mathbf{q}_{\text{op}}^0)$  such that  $\hat{S}(\mathbf{v}_{\text{op}}^0) = \mathbf{q}_{\text{op}}^0$ , and a target quantity of interest  $\mathbf{q}_{\text{target}}$ . It is our goal to find the vector  $\delta\mathbf{v}_{\text{tot}}$  with the minimum  $L_1$  norm which satisfies  $\hat{S}(\mathbf{v}_{\text{op}}^0 + \delta\mathbf{v}_{\text{tot}}) = \mathbf{q}_{\text{target}}$ .

The following are the steps taken at the  $(i + 1)^{\text{th}}$  iteration of CODA.

1. Given a working point  $(\mathbf{v}_{\text{op}}^i, \mathbf{q}_{\text{op}}^i)$ , consider a set of linearly independent, small voltage

variations  $\{\boldsymbol{\epsilon}_1, \boldsymbol{\epsilon}_2, \dots, \boldsymbol{\epsilon}_m\}$  ( $\boldsymbol{\epsilon}_j \in \mathcal{V}$ ) about the current working point. In the simulations described in the main text, we assumed took  $\boldsymbol{\epsilon}_j$  to correspond to a voltage change of 0.1 mV on the  $j^{\text{th}}$  electrode. Perform  $m$  simulations  $\hat{S}(\mathbf{v}_{\text{op}}^i + \boldsymbol{\epsilon}_j)$  to obtain the resulting  $m$   $\mathbf{q}_{\text{op}}^i + \delta\mathbf{q}_j$ . From the collection of  $\{\boldsymbol{\epsilon}_1, \boldsymbol{\epsilon}_2, \dots, \boldsymbol{\epsilon}_m\}$  and the associated  $\{\delta\mathbf{q}_1, \delta\mathbf{q}_2, \dots, \delta\mathbf{q}_m\}$ , construct the Jacobian matrix  $\mathbf{J}_i$  using the method of least squares. For small  $\delta\mathbf{v}$ ,

$$\hat{S}(\mathbf{v}_{\text{op}}^i + \delta\mathbf{v}) \approx \mathbf{q}_{\text{op}}^i + \mathbf{J}_i \cdot \delta\mathbf{v}.$$

2. Using a convex program (such as the matrix-free conic optimization implemented in the CVXPY package [89]), find the  $\delta\mathbf{v}_i$  that minimizes  $\|\mathbf{v}_{\text{op}}^i - \mathbf{v}_{\text{op}}^0 + \delta\mathbf{v}_i\|_1$  subject to the constraint  $\mathbf{q}_{\text{target}} = \mathbf{q}_{\text{op}}^i + \mathbf{J}_i \cdot \delta\mathbf{v}_i$ . The vector  $\mathbf{v}_{\text{op}}^i - \mathbf{v}_{\text{op}}^0 + \delta\mathbf{v}_i$  is the total change in voltage from the initial working point ( $\mathbf{v}_{\text{op}}^0$ ).
3. The voltage change vector  $\delta\mathbf{v}_i$  defines a search direction, similar to the gradient used in nonlinear gradient descent. Evaluate  $\hat{S}(\mathbf{v}_{\text{op}}^i + \delta\mathbf{v}_i)$ ; if the quantities of interest move closer to the target, *i.e.*,  $\|\hat{S}(\mathbf{v}_{\text{op}}^i + \delta\mathbf{v}_i) - \mathbf{q}_{\text{target}}\|_2 < \|\mathbf{q}_{\text{op}}^i - \mathbf{q}_{\text{target}}\|_2$ , then define  $\mathbf{v}_{\text{op}}^{i+1} = \mathbf{v}_{\text{op}}^i + \delta\mathbf{v}_i$ . If not, then replace  $\delta\mathbf{v}_i$  with  $\delta\mathbf{v}_i/2$ , and repeat this step. Continue until the quantities of interest move closer to the target.

We repeat this process until  $\|\mathbf{q}_{\text{op}}^i - \mathbf{q}_{\text{target}}\|_2$  is below a specified threshold value. For the simulations described in the main text, we assumed a threshold value of 0.01. For a detailed explanation of the units of this vector, see the following section.

To decrease the total number of device simulations in CODA, one can replace step 1 in the protocol described above by Broyden's method [130]. This method finds an approximate

Jacobian matrix  $\mathbf{J}_{i+1}$  by combining the Jacobian obtained in the previous iteration  $\mathbf{J}_i$  and the nonlinear error from the previous step:  $\hat{S}(\mathbf{v}_{\text{op}}^i + \delta\mathbf{v}) - (\mathbf{q}_{\text{op}}^i + \mathbf{J}_i \cdot \delta\mathbf{v})$ . Although this eliminates the need for explicit Jacobian formation, using Broyden’s method often does increase the number of iterations required for convergence. In many cases, using this method leads to an overall speedup. However, in cases where sparse control cannot be achieved, we find that the number of iterations required for convergence increases dramatically, which negates any potential speed-up.

As with all “hill-climbing” nonlinear optimization algorithms, there is no guarantee that the local optimum found by the CODA protocol is a global optimum. However, because CODA is a regularized optimization protocol, the solution with the globally minimal  $L_1$  norm is by definition “close” in control space to the starting point, and therefore it is likely that the solution found by CODA is the global minimum. While it is certainly possible to devise systems in which the CODA protocol does not find the global minimum, all indications are that the solutions found in the main text are indeed global minima. In principle, one could better ensure global optimization for these systems by implementing a version of CODA which uses a basin-hopping protocol to sample across several local minima.

## A.2 Simulation details

We perform semi-classical Thomas-Fermi calculations [90] using the COMSOL Multiphysics software package to solve a nonlinear Poisson equation in three dimensions. We use zero-field

boundary conditions on all sides of the simulated domain, with the exception of the bottom of the SiGe buffer, which is grounded. We assume the following heterostructure profile for all the modeled devices. This profile is consistent with the accumulation-mode devices described in Refs. [131] and [132]: 200 nm of Si<sub>0.7</sub>Ge<sub>0.3</sub> (with dielectric constant  $\varepsilon = 13.19$ ), a 10 nm Si quantum well ( $\varepsilon = 11.7$ ), 30 nm of Si<sub>0.7</sub>Ge<sub>0.3</sub>, 10 nm of Al<sub>2</sub>O<sub>3</sub> ( $\varepsilon = 9.0$ ), a 10 nm layer of metallic electrodes embedded in the dielectric, 80 nm of Al<sub>2</sub>O<sub>3</sub>, and a second 10 nm layer of metallic electrodes, followed by vacuum. Midway within the Si quantum well, we define a plane of charge with the charge density given by

$$\sigma_{2D}(x, y) = -2 \times 2 \times \frac{em_{\text{eff}}(U(x, y) + E_F)}{2\pi\hbar^2} \times \theta(U(x, y) + E_F), \quad (\text{A.1})$$

where  $e$  is the charge of an electron,  $m_{\text{eff}} = 0.19 m_{\text{electron}}$  is the transverse effective mass of a conduction electron in silicon,  $U(x, y)$  is the strength of the electrostatic potential energy as a function of position,  $E_F$  is the Fermi energy (which we take to be at ground), and  $\theta(x)$  is the step function. The two prefactors account for the spin and valley degeneracies.

The dot occupations are calculated via integrating the charge density found with the Thomas-Fermi approximation. Transmission coefficients between dots are calculated by finding the center of charge of each dot, and applying the WKB approximation [16] across a straight line connecting the two charge centers. Orbital dot energies are calculated via a 2D finite-difference Schrödinger solver in the plane of charge, using the electrostatic confinement potential obtained from the Thomas-Fermi analysis, and the transverse effective mass of a

conduction electron in silicon.

Our CODA protocol requires all the components of the voltage vector to have the same units (and comparable magnitudes, for numerical stability). The quantities of interest considered in our simulations were electron occupations and tunnel barrier heights. When populating our vectors in the space  $\mathcal{Q}$ , we use the units of electron number for dot occupation, meV for dot energy and we take the logarithm of the transmission coefficient, divided by 1000, since tunnel couplings can vary by orders of magnitude as a function of gate voltage. These units were chosen to ensure rapid convergence.

### A.3 Tunnel rates and transmission coefficients

Following Ref. [133] one can derive that the tunnel coupling  $\Delta$  between two one-dimensional simple harmonic oscillators with frequencies  $\omega_1$  and  $\omega_2$  is approximately

$$\Delta \approx \frac{1}{2\pi} T_{WKB} \hbar \sqrt{\omega_1 \omega_2}, \quad (\text{A.2})$$

where  $T_{WKB}$  is the transmission coefficient calculated via the WKB approximation. For quantum dots with orbital energies of  $\sim 1$  meV, a transmission coefficient of 0.01 corresponds with a tunnel coupling of  $\Delta \approx 1.6 \mu\text{eV}$ , or a tunnel rate of  $\Delta/h \approx 400$  MHz.

# Appendix B

## Supplemental information for Ch. 4

### B.1 Derivation of an effective Hamiltonian for two coupled QDHQS

In this section, we derive an effective 4D Hamiltonian describing the logical subspace of two capacitively coupled QDHQs. In the basis  $\{|\cdot S\rangle, |\cdot T\rangle, |S\cdot\rangle\}^{(1)} \otimes \{|\cdot S\rangle, |\cdot T\rangle, |S\cdot\rangle\}^{(2)}$ , the full

9D Hamiltonian, Eq. (4.3) of the main text, is given by

$$\mathcal{H}_{2q} = \left( \begin{array}{cccc|cccc|c} E_0 & 0 & 0 & 0 & \Delta_1^{(2)} & \Delta_1^{(1)} & 0 & 0 & 0 \\ 0 & E_1 & 0 & 0 & -\Delta_2^{(2)} & 0 & 0 & \Delta_1^{(1)} & 0 \\ 0 & 0 & E_2 & 0 & 0 & -\Delta_2^{(1)} & \Delta_1^{(2)} & 0 & 0 \\ 0 & 0 & 0 & E_3 & 0 & 0 & -\Delta_2^{(2)} & -\Delta_2^{(1)} & 0 \\ \hline \Delta_1^{(2)} & -\Delta_2^{(2)} & 0 & 0 & E_4 & 0 & 0 & 0 & \Delta_1^{(1)} \\ \Delta_1^{(1)} & 0 & -\Delta_2^{(1)} & 0 & 0 & E_5 & 0 & 0 & \Delta_1^{(2)} \\ 0 & 0 & \Delta_1^{(2)} & -\Delta_2^{(2)} & 0 & 0 & E_6 & 0 & -\Delta_2^{(1)} \\ 0 & \Delta_1^{(1)} & 0 & -\Delta_2^{(1)} & 0 & 0 & 0 & E_7 & -\Delta_2^{(2)} \\ \hline 0 & 0 & 0 & 0 & \Delta_1^{(1)} & \Delta_1^{(2)} & -\Delta_2^{(1)} & -\Delta_2^{(2)} & E_8 \end{array} \right), \quad (\text{B.1})$$



where we adopt the same notation as the main text, and define

$$E_0 = -\frac{\varepsilon^{(1)}}{2} + \frac{g}{4} - \frac{\varepsilon^{(2)}}{2}, \quad (\text{B.2})$$

$$E_1 = -\frac{\varepsilon^{(1)}}{2} + \frac{g}{4} + E_{\text{ST}}^{(2)} - \frac{\varepsilon^{(2)}}{2}, \quad (\text{B.3})$$

$$E_2 = -\frac{\varepsilon^{(1)}}{2} + \frac{g}{4} + E_{\text{ST}}^{(1)} - \frac{\varepsilon^{(2)}}{2}, \quad (\text{B.4})$$

$$E_3 = -\frac{\varepsilon^{(1)}}{2} + \frac{g}{4} + E_{\text{ST}}^{(1)} + E_{\text{ST}}^{(2)} - \frac{\varepsilon^{(2)}}{2}, \quad (\text{B.5})$$

$$E_4 = -\frac{\varepsilon^{(1)}}{2} + \frac{\varepsilon^{(2)}}{2} - \frac{g}{4}, \quad (\text{B.6})$$

$$E_5 = \frac{\varepsilon^{(1)}}{2} - \frac{\varepsilon^{(2)}}{2} - \frac{g}{4}, \quad (\text{B.7})$$

$$E_6 = -\frac{\varepsilon^{(1)}}{2} + \frac{\varepsilon^{(2)}}{2} + E_{\text{ST}}^{(1)} - \frac{g}{4}, \quad (\text{B.8})$$

$$E_7 = \frac{\varepsilon^{(1)}}{2} + E_{\text{ST}}^{(2)} - \frac{\varepsilon^{(2)}}{2} - \frac{g}{4}, \quad (\text{B.9})$$

$$E_8 = \frac{\varepsilon^{(1)}}{2} + \frac{\varepsilon^{(2)}}{2} + \frac{g}{4}. \quad (\text{B.10})$$

In Eq. (B.1), the solid lines delineate three distinct subspaces.  $E_0$  through  $E_3$  represent the logical manifold, in which the energy levels decrease with  $\varepsilon^{(1)}$  and  $\varepsilon^{(2)}$ .  $E_4$  through  $E_7$  represent a leakage manifold in which  $\varepsilon^{(1)}$  and  $\varepsilon^{(2)}$  have opposite effects on the energy.  $E_8$  is an additional leakage state for which the energy increases with  $\varepsilon^{(1)}$  or  $\varepsilon^{(2)}$ . Some typical eigenvalues of Eq. (B.1) are plotted in Fig. B.1, where we have set  $\varepsilon^{(1)} = \varepsilon^{(2)}$  and added  $\varepsilon^{(1)}$  to all the eigenstates. At large detunings, we observe a large energy splitting between the logical and leakage manifolds. To gain insight into the gate operations, we can therefore perform a Schrieffer-Wolff [134] decomposition to adiabatically eliminate the leakage states. Working to fourth order, and further diagonalizing the resulting 4D Hamiltonian, we obtain

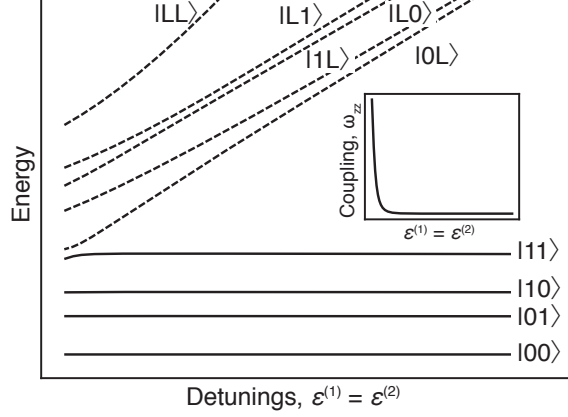


Figure B.1: **Energy spectrum of two capacitively-coupled QDHQs as a function of qubit detunings.** In the limit of large detuning, the low-energy logical subspace is well separated from the leakage states. For smaller detuning values, the mixing between the leakage states and the low-energy states grows, giving rise to an effective Coulomb interaction between the qubits (inset – inter-qubit coupling  $\omega_{zz}$  of four-dimensional effective Hamiltonian, see Eq. 4.4 of main text). The results shown here assume  $E_{\text{ST}}^{(1)} = 52 \mu\text{eV}$ ,  $E_{\text{ST}}^{(2)} = 32 \mu\text{eV}$ ,  $\Delta_1^{(i)} = 0.64E_{\text{ST}}^{(i)}$ ,  $\Delta_2^{(i)} = 0.58E_{\text{ST}}^{(i)}$ , and  $g = -75 \mu\text{eV}$ . (See Section S6, below, for discussion about the sign of  $g$ .)

the leading-order contributions to the prefactors in Eq. (4.4):

$$\hbar\omega_{z1} = E_{\text{ST}}^{(1)} + \frac{\Delta_2^{(1)2} - \Delta_1^{(1)2}}{\varepsilon^{(1)} - g/2} + \mathcal{O}(\varepsilon^{-2}), \quad (\text{B.11})$$

$$\hbar\omega_{z2} = E_{\text{ST}}^{(2)} + \frac{\Delta_2^{(2)2} - \Delta_1^{(2)2}}{\varepsilon^{(2)} - g/2} + \mathcal{O}(\varepsilon^{-2}), \quad (\text{B.12})$$

$$\hbar\omega_{zz} = \frac{8g(g - \varepsilon^{(1)} - \varepsilon^{(2)}) \left( \Delta_1^{(1)2} - \Delta_2^{(1)2} \right) \left( \Delta_1^{(2)2} - \Delta_2^{(2)2} \right)}{(g - 2\varepsilon^{(1)})^2 (g - 2\varepsilon^{(2)})^2 (\varepsilon^{(1)} + \varepsilon^{(2)})}. \quad (\text{B.13})$$

Here, we have also assumed that  $E_{\text{ST}}^{(i)} \ll \varepsilon^{(i)} - g/2$ .

## B.2 Optimizing single qubit dispersions

To achieve high-fidelity two-qubit gates, one must suppress both single- and two-qubit errors. Dephasing due to charge noise is a significant source of single-qubit errors in semiconductor-based quantum dot qubits, but can be mitigated by tuning the qubit near a sweet spot where the splitting between energy levels of the logical states,  $\hbar\omega$ , is insensitive to small fluctuations of the detuning, i.e.,  $\partial\omega/\partial\varepsilon = 0$ . Since the pure dephasing time  $T_2^*$  is inversely proportional to  $|\partial\omega/\partial\varepsilon|$  [55, 63], it can increase significantly near such a sweet spot. Here, we identify a working regime for a QDHQ where approximate sweet spot behavior can be achieved over a wide range of  $\varepsilon$ .

For the QDHQ, the qubit energy dispersion can be made extremely flat at large detunings by choosing the special values  $\Delta_1 = \Delta_2 = E_{ST}/\sqrt{2}$ . However, faster two-qubit gates can be achieved by working at somewhat lower detuning values, where the qubits have large dipole moments. In Fig. B.2 we plot  $\partial\omega/\partial\varepsilon$  for a range of  $\Delta_1$  and  $\Delta_2$ . Here we note that large dipole moments occur in the vicinity of the dips in the dispersion. The pulse sequences shown in Fig. 4.1f require spending time at very large detuning values, and then transitioning adiabatically (i.e., slowly) to a region with such fast gates.

To reduce the dephasing effects caused by charge noise, we therefore choose values of  $\Delta_1$  and  $\Delta_2$  that provide relatively flat dispersions over a wide range of epsilon. In Fig. B.2 this corresponds to tunnel couplings given by  $\Delta_1 = 0.64E_{ST}$  and  $\Delta_2 = 0.58E_{ST}$ , and  $\Delta_1/\Delta_2 \simeq 1.103$ . For these values, the qubit can be operated near the charge transition point ( $\varepsilon \simeq 0$ ), where gates are fast, while still maintaining long single-qubit coherence times.




Figure B.2: **Adjusting tunnel couplings to optimize single qubit dispersions.** The derivative of the QDHQ frequency  $\omega$  with respect to the detuning  $\varepsilon$ , plotted as a function of  $\varepsilon$ . **a**, Here we consider a fixed value of  $\Delta_2 = 0.58E_{ST}$  and several values of  $\Delta_1$ . From top to bottom, in units of  $E_{ST}$ :  $\Delta_1 = 0.6, 0.62, 0.64, 0.66$ , and  $0.68$ . **b**, Here we consider a fixed value of  $\Delta_1 = 0.64E_{ST}$  and several values of  $\Delta_2$ . From top to bottom, in units of  $E_{ST}$ :  $\Delta_2 = 0.54, 0.56, 0.58, 0.6$ , and  $0.62$ . The combination of  $\Delta_1 = 0.64E_{ST}$  and  $\Delta_2 = 0.58E_{ST}$  yields a dispersion for which  $\hbar|\partial\omega/\partial\varepsilon| < 0.002$  for all detunings  $\varepsilon > 1.48E_{ST}$ .

### B.3 Pulse sequences

The simple detuning-only entangling pulse sequence considered in this work is shown in Fig. 1f of the main text. It is parameterized by the four quantities  $\varepsilon_{\text{wait}}^{(1)}$ ,  $\varepsilon_{\text{wait}}^{(2)}$ ,  $\tau_{\text{ramp}}$ , and  $\tau_{\text{wait}}$ . Here we provide the explicit sequence used in our simulations, including the ramp functions:

$$\varepsilon^{(i)} = \begin{cases} \varepsilon_{\text{init}}^{(i)} & t < 0, \\ \varepsilon_{\text{init}}^{(i)} + \left( \varepsilon_{\text{wait}}^{(i)} - \varepsilon_{\text{init}}^{(i)} \right) \sin^2 \left( \frac{\pi t}{2\tau_{\text{ramp}}} \right) & 0 < t < \tau_{\text{ramp}}, \\ \varepsilon_{\text{wait}}^{(i)} & \tau_{\text{ramp}} < t < \tau_{\text{ramp}} + \tau_{\text{wait}}, \\ \varepsilon_{\text{init}}^{(i)} - \left( \varepsilon_{\text{wait}}^{(i)} - \varepsilon_{\text{init}}^{(i)} \right) \sin^2 \left( \frac{\pi(t - \tau_{\text{ramp}} - \tau_{\text{wait}})}{2\tau_{\text{ramp}}} \right) & \tau_{\text{ramp}} + \tau_{\text{wait}} < t < 2\tau_{\text{ramp}} + \tau_{\text{wait}}, \\ \varepsilon_{\text{init}}^{(i)} & 2\tau_{\text{ramp}} + \tau_{\text{wait}} < t. \end{cases} \quad (\text{B.14})$$

In our simulations we adopt the initial detuning values  $\varepsilon_{\text{init}}^{(1)} = \varepsilon_{\text{init}}^{(2)} = 500 \mu\text{eV}$ , for which effective one- and two-qubit couplings are negligible. These values are then varied during the optimization procedure.

To explore the possibility of a dynamical sweet spot, we also consider the tunnel-coupling pulse sequence shown in Fig. 4.3b, which is parameterized by the variables listed in Table B.1. Again we provide the explicit sequence used in our simulations, including the ramp functions:

where  $t_1 = \tau_{\text{wait}}^0$ ,  $t_2 = t_1 + \tau_{\text{ramp}}^{\text{I}}$ ,  $t_3 = t_2 + \tau_{\text{wait}}^{\text{I}}$ ,  $t_4 = t_3 + \tau_{\text{ramp}}^{\text{II}}$ , and  $t_5 = t_4 + \tau_{\text{wait}}^{\text{II}}/2$ . For

$$\Delta(t) = \begin{cases} \Delta_{\text{init}} & t \leq t_1 \\ \Delta_{\text{init}} + (\Delta_{\text{wait}}^{\text{I}} - \Delta_{\text{init}}) \sin^2 \left( \frac{\pi(t-t_1)}{\tau_{\text{ramp}}^{\text{I}}} \right) & t_1 \leq t < t_2 \\ \Delta_{\text{wait}}^{\text{I}} & t_2 \leq t < t_3 \\ \Delta_{\text{wait}}^{\text{I}} + (\Delta_{\text{wait}}^{\text{II}} - \Delta_{\text{wait}}^{\text{I}}) \sin^2 \left( \frac{\pi(t-t_3)}{\tau_{\text{ramp}}^{\text{II}}} \right) & t_3 \leq t < t_4 \\ \Delta_{\text{wait}}^{\text{II}} & t_4 \leq t < t_5, \end{cases} \quad (\text{B.15})$$

		$\tau_{\text{wait}}^0$	$\tau_{\text{ramp}}^{\text{I}}$	$\tau_{\text{wait}}^{\text{I}}$	$\tau_{\text{ramp}}^{\text{II}}$	$\tau_{\text{wait}}^{\text{II}}$	$\Delta_{\text{wait}}^{\text{I}}$	$\Delta_{\text{wait}}^{\text{II}}$	$\tau_{\text{ramp}}$	$\tau_{\text{wait}}$
$\Delta_1^{(i)}/\Delta_2^{(i)}$ =1.1034	$\Delta_1^{(1)}$	1.10	0.45	0.00	1.81	2.10	33.263	32.752	2.847	2.587
	$\Delta_1^{(2)}$	1.14	0.31	1.40	0.76	1.83	35.245	28.9555		
$\Delta_1^{(i)}/\Delta_2^{(i)}$ free	$\Delta_1^{(1)}$	0.31	1.09	1.27	2.21	1.01	11.200	46.332	4.185	1.705
	$\Delta_2^{(1)}$	0.74	0.58	0.38	1.40	2.79	3.977	31.341		
	$\Delta_1^{(2)}$	1.06	2.01	0.00	0.69	2.21	20.974	20.828		
	$\Delta_2^{(2)}$	0.25	1.28	1.03	2.46	0.86	34.976	67.408		

Table B.1: Parameters used in the tunnel coupling pulse sequence shown in Fig. 4.3b, and accompanying detuning pulse sequence shown in Fig. 4.1f, obtained using the optimization procedure described in the Ch. 4. The specific ratio,  $\Delta_1^{(i)}/\Delta_2^{(i)} = 1.1034$  used in the first set of solutions is consistent with the discussion in Section B.2, above. Times are given in ns, and tunnel couplings are given in  $\mu\text{eV}$ .

$t > t_5$ ,

$$\Delta(t) = \Delta(t_5 - t). \quad (\text{B.16})$$

For brevity here, we have dropped the superscripts and subscripts on the tunnel coupling parameters. In this case, we adopt the initial values  $\Delta_{1,\text{init}}^{(1)} = 33.28 \mu\text{eV}$ ,  $\Delta_{2,\text{init}}^{(1)} = 30.16 \mu\text{eV}$ ,  $\Delta_{1,\text{init}}^{(2)} = 28.8 \mu\text{eV}$ , and  $\Delta_{2,\text{init}}^{(2)} = 26.1 \mu\text{eV}$ , which were chosen according to the considerations of Section B.2, above. For this sequence, we also assumed the fixed values  $\varepsilon_{\text{wait}}^{(1)} = 90 \mu\text{eV}$  and  $\varepsilon_{\text{wait}}^{(2)} = 110 \mu\text{eV}$ , as discussed in the main text. The optimized values obtained for the other parameters in the sequence are listed in Table B.1.

## B.4 Optimization of detuning pulse parameters

The simple two-qubit detuning-only pulse sequence used in this work is shown in Fig. 4.1f, and is completely defined by the parameters  $\varepsilon_{\text{wait}}^{(1)}$ ,  $\varepsilon_{\text{wait}}^{(2)}$ ,  $\tau_{\text{ramp}}$ , and  $\tau_{\text{wait}}$ , as described in Section B.3, above. In the current section, we explain how these parameters are chosen in our analysis, while constraining the total gate time  $2\tau_{\text{ramp}} + \tau_{\text{wait}} < 50$  ns. The procedure is summarized as follows.

1. **Choose specific values for  $\varepsilon_{\text{wait}}^{(1)}$  and  $\varepsilon_{\text{wait}}^{(2)}$ .**
2. **Determine  $\tau_{\text{ramp}}$ .** Consider a “ramp-only” detuning sequence with  $\tau_{\text{wait}} = 0$ . Non-adiabatic effects such as leakage occur *only* during these ramp steps, and we determine  $\tau_{\text{ramp}}$  by ensuring that it satisfies a “fast-adiabatic” criterion, defined as follows. We first define the leakage fidelity as  $\mathcal{F}_{\text{leak}} = \frac{1}{4} \sum |\langle ij | U_{2q} | ij \rangle|^2$ , where the sum is taken over the logical states,  $(i, j) = (0, 1)$ , defined at time  $t = 0$  in the pulse sequence, and  $U_{2q}$  is the 9D unitary operator derived from  $\mathcal{H}_{2q}$ , and computed here in the absence of charge noise. (Note that  $\mathcal{F}_{\text{leak}}$  is defined similarly to  $\mathcal{I}_{\text{na}}$  in Methods, except that here, the logical states on either side of the matrix element are evaluated at the initial time, and are the same as the logical states at the final time of the pulse sequence.) We then choose a  $\tau_{\text{ramp}}$  that corresponds to the shortest ramp time giving  $\mathcal{F}_{\text{leak}} > 99.9\%$ . Finally, we omit any solutions with  $2\tau_{\text{ramp}} > 50$  ns from the rest of the analysis; these correspond to the cross-hatched regions of Fig. B.3a. As discussed in the main text, the fidelity of the results shown in Fig. 4.3c is limited by charge noise, not leakage,

indicating that the particular choice of  $\mathcal{F}_{\text{leak}} > 99.9\%$  does not affect our final results.

3. **Determine the angles  $\phi_1$  and  $\phi_2$** , defined such that  $U_{4\text{D}} \simeq U_{\text{ideal}}$ , where  $U_{\text{ideal}} = Z^{(1)}(\phi_1)Z^{(2)}(\phi_2)\text{CZ}$ . Here,  $U_{4\text{D}}$  is the projection of  $U_{2q}$  onto the 4D logical subspace; in this step,  $U_{2q}$  is computed as a function of  $\tau_{\text{wait}}$  in the absence of charge noise (but including leakage), using the values of  $\varepsilon_{\text{wait}}^{(1)}$ ,  $\varepsilon_{\text{wait}}^{(2)}$ , and  $\tau_{\text{ramp}}$  chosen above. CZ is defined as  $\text{diag}\{1, 1, 1, -1\}$ . The explicit method used to determine  $\phi_1$  and  $\phi_2$  as a function of  $\tau_{\text{wait}}$  is given as follows. First, we adjust the overall phase of  $U_{2q}$  such that  $\langle 00 | U_{2q} | 00 \rangle$  is real and approximately equal to 1. Then we define  $\phi_1 = \text{angle}[\langle 01 | U_{2q} | 01 \rangle] + \pi/2$  and  $\phi_2 = \text{angle}[\langle 10 | U_{2q} | 10 \rangle] + \pi/2$ , where the function  $\text{angle}[u]$  gives the complex phase of  $u$ .
4. **Determine  $\tau_{\text{wait}}$** . Recompute  $U_{2q}$  as a function of  $\tau_{\text{wait}}$ , now including charge noise. Compute the process fidelity  $\mathcal{F}$ , as described in Section B.5 below, for each value of  $\tau_{\text{wait}}$ , where the actual  $\chi$  matrix is obtained from  $U_{2q}$ , while the ideal  $\chi$  matrix is derived from  $U_{\text{ideal}}$ . Perform an average of  $\mathcal{F}$  over charge noise configurations, as described in Section 4.4. Maximize this  $\langle \mathcal{F} \rangle$  with respect to  $\tau_{\text{wait}}$ , omitting any results for which  $\tau_{\text{total}} = 2\tau_{\text{ramp}} + \tau_{\text{wait}} > 50$  ns.
5. Finally, **choose the optimal values of  $\varepsilon_{\text{wait}}^{(1)}$  and  $\varepsilon_{\text{wait}}^{(2)}$** , as described in the main text, by determining the maximum fidelity shown in Fig. 4.2.






Figure B.3: **Key steps in the detuning-pulse optimization procedure, as described in Section B.4.** **a**, Step 2: determine  $\tau_{\text{ramp}}$ . We simulate the ramp portion of the pulse sequence in the absence of noise and compute the leakage fidelity  $\mathcal{F}_{\text{leak}}$  that quantifies the adiabaticity of the pulse, defined as the probability that a system initialized into a two-qubit logical state remains in that state after the pulse is applied. The plot shows the shortest values of  $\tau_{\text{ramp}}$  giving  $\mathcal{F}_{\text{leak}} > 99.9\%$ . Results with  $2\tau_{\text{ramp}} > 50$  ns are omitted from the rest of the analysis because they are much slower than single-qubit gates (cross-hatched region). **b**, Step 4: determine  $\tau_{\text{wait}}$ . We compute the noise-averaged fidelity  $\langle \mathcal{F} \rangle$  of a full CZ pulse sequence, modulo single-qubit gate operations. The plot shows the value of  $\tau_{\text{wait}}$  that maximizes  $\langle \mathcal{F} \rangle$ . Solutions with  $\tau_{\text{total}} = 2\tau_{\text{ramp}} + \tau_{\text{wait}} > 50$  ns are now omitted.

## B.5 Process fidelity

In Figs. 4.2 and 4.3, we report fidelities that are averaged over a noise distribution. To compute the fidelity for a given instance of noise, we first solve Eq. (4.6) for the appropriate pulse sequence to obtain the corresponding unitary operator  $U_{2q}$ . The process fidelity is defined as  $\mathcal{F} = \text{Tr}(\chi_{\text{ideal}}\chi)$ , where  $\chi$  is the actual process matrix, including noise and leakage effects, and  $\chi_{\text{ideal}}$  is the ideal process matrix derived from  $U_{\text{ideal}}$ . In this case, we use the 4D  $U_{\text{ideal}}$  defined above, in Section B.4, which is then embedded in the full 9D Hilbert space.

We obtain the process matrix  $\chi$  from  $U_{2q}$  using the Choi-Jamiolkowski isomorphism [115], in which the process matrix is given by  $\chi = d\rho$ , where  $d$  is the dimensionality of the logical system (in this case,  $d = 4$ ), and  $\rho$  is given by

$$\rho = [I \otimes U_{2q}] (|\Phi\rangle\langle\Phi|). \quad (\text{B.17})$$

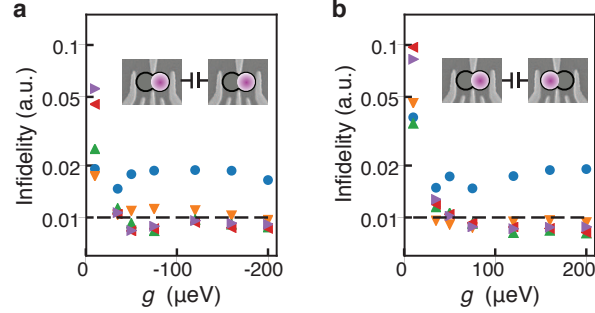


Figure B.4: **Optimized process infidelities of adiabatic CZ gates as a function of  $\Delta E_{\text{ST}}$  and  $g$ .** Results are obtained using the procedure described in Section B.4. **a**,  $g < 0$ . **b**,  $g > 0$ . In both cases, the markers correspond to  $\Delta E_{\text{ST}} = 1 \mu\text{eV}$  (blue circles),  $\Delta E_{\text{ST}} = 3 \mu\text{eV}$  (orange downwards-facing triangles),  $\Delta E_{\text{ST}} = 5 \mu\text{eV}$  (green upwards-facing triangles),  $\Delta E_{\text{ST}} = 7 \mu\text{eV}$  (red left-facing triangles), and  $\Delta E_{\text{ST}} = 9 \mu\text{eV}$  (purple right-facing triangles). The threshold for acceptable infidelity is chosen to be 1% (black dashed lines). Insets: **a**, Aligned dipoles ( $g < 0$ ); **b**, Anti-aligned dipoles ( $g > 0$ ).

Here,  $I$  is the identity matrix of the 9D Hilbert space and the Jamiolkowski state  $|\Phi\rangle$  is defined as

$$|\Phi\rangle = \frac{1}{2} \sum |ij\rangle |ij\rangle, \quad (\text{B.18})$$

where the sum is taken over the logical eigenstates  $(i, j) = (0, 1)$  of the two-qubit Hamiltonian described in Eq. (4.3) of the main text. Note that the Jamiolkowski state only includes four states despite the full system having nine states.

## B.6 Dependence of the fidelity on $g$ and $\Delta E_{\text{ST}}$

In the CZ gate analyses presented in Figs. 4.2 and 4.3,  $g$  was chosen to be  $75 \mu\text{eV}$ , as this was an experimentally measured value [94]. Additionally, the singlet-triplet splittings  $E_{\text{ST}}^{(1)}$  and  $E_{\text{ST}}^{(2)}$  were chosen to be  $52 \mu\text{eV}$  and  $47 \mu\text{eV}$ , respectively, inspired by experimentally measured

values [63]. However, both  $g$  and the singlet-triplet splittings will vary from experiment to experiment; indeed, capacitive couplings as large as  $200 \mu\text{eV}$  have been measured [51]. In this section, we determine whether either  $g$  or the difference between the singlet-triplet splittings,  $\Delta E_{\text{ST}} = E_{\text{ST}}^{(1)} - E_{\text{ST}}^{(2)}$ , can be used to further optimize the fidelity, finding no significant improvements.

We first note that it is possible to change the sign of the Coulomb interaction  $g$  by reversing the alignment of the charge dipole of one of the qubits (say,  $i$ ), which amounts to changing the sign of  $\varepsilon^{(i)}$ . Here, we adopt the convention in Eq. (4.3) of the main text that  $g < 0$  corresponds to the dipoles being aligned in the limit of large detunings, as indicated in the inset of Fig. B.4a, while  $g > 0$  corresponds to the dipoles being anti-aligned, as indicated in the inset of Fig. B.4b. Clearly the sign of  $g$  also affects the qubits' tendency to align or anti-align as the detunings are varied, and we therefore expect our results to depend on this sign.

In Fig. B.4, we plot the noise-averaged infidelities obtained for a range of  $g$  (both positive and negative values) and  $\Delta E_{\text{ST}}$ , assuming a constant level of quasistatic charge noise,  $\sigma_\varepsilon = 1 \mu\text{eV}$ . The results are obtained using the procedure described in Section B.4 to obtain the absolute minimum infidelity for the detuning-only pulse sequence. We find that the infidelities generally fall below a threshold criterion of 0.01 (i.e., 1%), except when  $-35 \leq g \leq 50 \mu\text{eV}$ , or when  $\Delta E_{\text{ST}} \leq 3 \mu\text{eV}$  (with  $g < 0$ ), or  $\Delta E_{\text{ST}} \leq 1 \mu\text{eV}$  (with  $g > 0$ ).

We can understand the behaviors observed in Fig. B.4 as follows. The infidelity decreases with  $|g|$  because the entanglement frequency  $\omega_{zz}$  is roughly proportional to  $g$ , as observed




Figure B.5: **Adiabatic ramp time diverges for small  $\Delta E_{ST}$ .** Ramp time  $\tau_{\text{ramp}}$  obtained by optimizing the process fidelity of the detuning pulse sequence is plotted as a function of  $g$  and  $\Delta E_{ST} = E_{ST}^{(1)} - E_{ST}^{(2)}$ . When  $\Delta E_{ST} \rightarrow 0$ , the resulting energy level degeneracy causes  $\tau_{\text{ramp}}$  to diverge, to preserve adiabaticity, thus increasing the exposure to charge noise and the infidelity.

in Eq (B.13). When  $g$  is small, the entangling gate speed is therefore slow. This must be compensated by reducing  $\varepsilon^{(1)}$  and  $\varepsilon^{(2)}$ ; however this also increases the susceptibility to charge noise, and the infidelity.

The dependence on  $\Delta E_{ST}$  in Fig. B.4 can be understood by noting that the limit  $\Delta E_{ST} \rightarrow 0$  corresponds to the degeneracy of logical states  $|01\rangle$  and  $|10\rangle$  in the limit of large detunings. Degenerate energy levels cause problems for adiabatic operation, which can only be solved by reducing the ramp speed. This is demonstrated in Fig. B.5 where we plot the optimized value of  $\tau_{\text{ramp}}$  as a function of  $g$  and  $\Delta E_{ST}$ . Here we observe little dependence on  $g$ , but for small  $\Delta E_{ST}$ ,  $\tau_{\text{ramp}}$  and therefore  $\tau_{\text{total}}$  increase significantly. The longer gates are more exposed to charge noise, resulting in lower fidelity. On the other hand, for  $\Delta E_{ST} > 3 \mu\text{eV}$ , the total gate time is dominated by the waiting time  $\tau_{\text{wait}}$ , so the further reduction of  $\tau_{\text{ramp}}$  has a marginal effect on the infidelity.

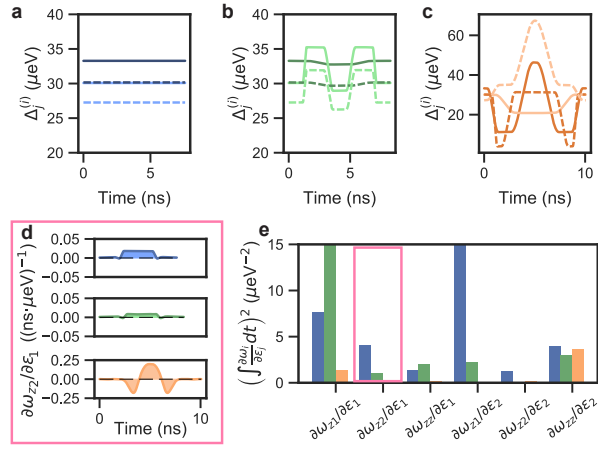


Figure B.6: **Pulse sequences involving both the tunnel coupling and detuning parameters are able to improve the fidelity of a CZ gate.** The tunnel coupling pulses shown in Fig. 4.3 improve the process fidelity by lowering the time-averaged value of the derivatives  $\partial\omega_j/\partial\varepsilon_i$ . As described in the main text, optimized results are obtained for three different tunnel coupling pulse sequences when **a**, tunnel couplings are held fixed for the duration of the sequence, **b**, only the ratios  $\Delta_1^{(1)}/\Delta_2^{(1)} = \Delta_1^{(2)}/\Delta_2^{(2)}$  are held fixed throughout the sequence, and **c**, no constraints are placed on the tunnel coupling sequence parameters. Here, the dark solid lines correspond to  $\Delta_1^{(1)}$ , the dark dashed lines correspond to  $\Delta_2^{(1)}$ , the light solid lines correspond to  $\Delta_1^{(2)}$ , and light dashed lines correspond to  $\Delta_2^{(2)}$ . The color codings are the same as in Fig. 4.3c. **d**, The derivative  $\partial\omega_{z2}/\partial\varepsilon_1$  is plotted as a function of time, showing a change of sign due to the application of the tunnel coupling pulse sequence. **e**, The squared, time-averaged derivatives are shown for all the different qubit frequencies. For a perfect dynamical sweet spot, these integrals would all vanish. Here, the pink box corresponds to the results shown in **d**.

## B.7 Approximate formula for the infidelity

In this section, we derive an approximate analytical expression for the charge-noise induced infidelity,  $\mathcal{I}_{\text{cn}}$ , to more efficiently identify tunnel coupling pulse sequences that improve the gate fidelity.

In the absence of non-adiabatic processes, we can evaluate the effective two-qubit Hamiltonian  $\mathcal{H}_{2q}$  in its adiabatic basis, giving Eq. (4.4). Since this Hamiltonian is strictly diagonal, it is trivial to compute the resulting unitary operator for the logical subspace:

$$\begin{aligned}
 U_{4\text{D}} = \text{diag}(\exp\{i(-\theta_{z1} - \theta_{z2} + \theta_{zz})/2\}, \\
 \exp\{i(\theta_{z1} - \theta_{z2} - \theta_{zz})/2\}, \\
 \exp\{i(-\theta_{z1} + \theta_{z2} - \theta_{zz})/2\}, \\
 \exp\{i(\theta_{z1} + \theta_{z2} + \theta_{zz})/2\}),
 \end{aligned}$$

where  $\theta_i = \int \omega_i dt$ . Quasistatic charge noise causes the phases to evolve with errors defined as  $\Delta\theta_i = \theta_i - \theta_i^{\text{ideal}}$ . However, the resulting time evolution is unitary, and the methods of Section B.5 easily give an expression for the process infidelity:

$$\begin{aligned}
 \mathcal{I} = 1 - \frac{1}{8}[2 + \cos(\Delta\theta_{z1} + \Delta\theta_{z2}) + \cos(\Delta\theta_{z1} - \Delta\theta_{z2}) \\
 + 2[\cos(\Delta\theta_{z1}) + \cos(\Delta\theta_{z2})]\cos(\Delta\theta_{zz})].
 \end{aligned} \tag{B.19}$$

Again for quasistatic charge noise, we can approximate

$$\Delta\theta_i \approx \delta\varepsilon^{(1)} \int \frac{\partial\omega_i}{\partial\varepsilon^{(1)}} dt + \delta\varepsilon^{(2)} \int \frac{\partial\omega_i}{\partial\varepsilon^{(2)}} dt, \quad (\text{B.20})$$

where  $\delta\varepsilon^{(j)}$  is the noise on  $\varepsilon^{(j)}$ , which is assumed to be constant over the duration of the pulse sequence. Substituting these definitions into Eq. (B.19), expanding in small  $\Delta\theta_i$ , and averaging over the noise distribution as described in Section 4.4 yields Eq (4.5).

## B.8 Dynamical sweet spot analysis for the tunnel coupling pulse sequence

In Fig. 4.3, we showed optimized infidelity results for three different tunnel coupling pulse sequences. The optimized pulse sequences are shown in Fig. B.6a-c with the same color coding as Fig. 4.3. In Fig. B.6a the tunnel couplings are held fixed for the duration of the detuning pulse. Figure B.6b shows the non-constant tunnel coupling pulse sequence obtained under the constraint that the ratios  $\Delta_1^{(1)}/\Delta_2^{(1)} = \Delta_1^{(2)}/\Delta_2^{(2)} = 1.1034$  are held fixed throughout the pulse sequence. Figure B.6c shows the pulse sequence obtained when the tunnel couplings are allowed to vary without constraint.

A dynamical sweet spot (DSS) is formed when the time-averaged derivatives of the qubit frequencies go to zero, as described below Eq. (4.5) of the main text. We plot these time averages in Fig. B.6e for each of the different pulse sequences, with details of the time

dependence for one of the qubit frequencies shown in Fig. B.6d. As seen here, certain choices for the pulse sequence cause the derivatives to change sign as a function of time, leading to an overall suppression of the time average and the infidelity. We note that the effect is especially pronounced for the pulse sequence with the largest number of tuning parameters. However, we also note that the derivative of the  $\omega_{zz}$  qubit frequency is particularly difficult to suppress.

## B.9 Relaxation time estimate

The hybridization of the spin degree of freedom with the charge degree of freedom in a QDHQ causes the qubit to be affected by the environmental charge noise. In addition to the dephasing considered in the main text, this charge noise could also lead to increased qubit relaxation, as indicated by short  $T_1$  times measured in charge qubits [59, 135]. In addition, phonons can also induce qubit relaxation. The relaxation rates measured in Refs. [135] and [59] decrease as the detuning is increased, which is consistent with a charge-noise-induced [136] as opposed to a phonon-induced [137] relaxation mechanism.

Here, we perform a simple estimate of the expected  $T_1$  relaxation times for the pulses considered in Fig. 4.2, finding the estimated infidelity due to relaxation to be less than the infidelity due to dephasing. To estimate the relaxation time, we use Fermi's golden rule to



calculate the relaxation rate assuming a qubit-boson interaction:

$$1/T_1 = \frac{2\pi}{\hbar} \sum_{\mathbf{Q}} |\langle 0| H_{q-b} |1\rangle|^2 \delta(\hbar\omega_{\mathbf{Q}} - \hbar\omega), \quad (\text{B.21})$$

where  $\mathbf{Q}$  and  $\omega_{\mathbf{Q}}$  are the wavevector and angular frequency of the boson respectively,  $|0\rangle$  and  $|1\rangle$  are the ground state and first excited state of the qubit,  $H_{q-b}$  is the qubit-boson coupling term of the Hamiltonian, and  $\omega$  is the frequency of the qubit. Eq. (B.21) applies whether the mechanism leading to qubit relaxation are phonons or classical charge fluctuators.

We first consider classical charge fluctuators. In this case, the charge noise couples to the qubits mainly through the dipole  $\mathcal{P} = -\hat{x}/d$  as described in the main text. In the  $\{|\cdot S\rangle, |\cdot T\rangle, |S\cdot\rangle\}$  basis, we take qubit-fluctuator coupling term  $H_{q-cn} = \text{diag}\{-P, -P, P\}$ . Using these approximations for the perturbation, we can estimate the expected relaxation time. First note that  $\langle 0| H_{q-cn} |1\rangle = \alpha P$ , where  $\alpha = \langle 0|S\cdot\rangle \langle S\cdot|1\rangle - \langle 0|\cdot S\rangle \langle \cdot S|1\rangle - \langle 0|\cdot T\rangle \langle \cdot T|1\rangle$  is a mixing parameter between zero and one that reflects the composition of the qubit's eigenstates; it can be thought of as a measure of how “spin-like” ( $\alpha = 0$ ) or “charge-like” ( $\alpha = 1$ ) a qubit is. In this proposal, qubit gates are performed by tuning the qubits from  $\alpha \approx 0$ , to states with much larger  $\alpha$  values; the  $\alpha$  parameter therefore varies considerably. In contrast, the qubit frequency  $\omega$  is essentially constant throughout the detuning pulses considered here (see Fig. B.1). Thus, we assume that the density of charge fluctuator states  $\delta(\hbar\omega_{\mathbf{Q}} - \hbar\omega)$  remains approximately constant. From Eq. (B.21), this allows us to estimate




Figure B.7: **Estimated error due to relaxation is smaller than error due to dephasing.** **a**, The estimated relaxation time  $T_1$  for the system of coupled QDHQs described in the main text. While a typical charge-noise-induced relaxation time for a charge qubit is 10 ns, the expected  $T_1$  time for the coupled hybrid system as the qubits become entangled is over an order of magnitude larger. **b**, The ratio between the estimated infidelity due to relaxation  $I_{relax}$  and the infidelity due to dephasing  $I_{dephase}$ .  $I_{relax}$  is smaller than  $I_{dephase}$  for every pair of qubit detunings considered here, in most cases by over a factor of two. We expect phonon-mediated relaxation times to be significantly longer than and proportional to the estimated charge-noise-induced relaxation times.

the relaxation time of a QDHQ as

$$T_1 \approx T_1^{min,charge} / \alpha^2, \quad (\text{B.22})$$

where  $T_1^{min,charge}$  is the relaxation time of a pure charge qubit at its avoided crossing.

Taking  $T_1^{min,charge} = 10$  ns, as in recent Si charge-qubit experiments [59, 135], in Fig. B.7a we use Eq. (B.22) to estimate the relaxation time for the capacitively-coupled QDHQ system with the parameters considered in Fig. 2 of the main text. Because the charge noise relaxation mechanism will be maximal when the qubit is the most “charge-like” for the present calculation, we simplify the pulse shown in Fig. 1f of the main text, considering instead that the system is simply held constant at the detuning values  $\varepsilon^{(1)} = \varepsilon_{wait}^{(1)}$ ,  $\varepsilon^{(2)} = \varepsilon_{wait}^{(2)}$  for a time  $\tau_{wait}$ . We find that for  $\varepsilon_{wait}^{(1)} > 115 \mu\text{eV}$  and  $\varepsilon_{wait}^{(2)} > 110 \mu\text{eV}$  the anticipated  $T_1$  time is over 100 ns, an order of magnitude larger than the expected minimum  $T_1$  time for a charge qubit.

Lowering the charge noise will further increase the expected  $T_1$  time.

In Fig. B.7b we compare the expected infidelity due to relaxation  $I_{relax}$  with the infidelity due to dephasing under quasistatic charge noise  $I_{dephase}$ . We estimate the infidelity due to relaxation as

$$I_{relax} \approx 1 - \exp(-\tau_{wait}/T_1) \quad (\text{B.23})$$

where the optimized time  $\tau_{wait}$  is calculated as described in Section S4.  $I_{dephase}$  is calculated as described in the Methods section of the main text. In order to accurately compare these two error channels, we assume a quasistatic noise level of  $\sigma_\varepsilon = 4.5 \mu\text{eV}$ , which has been measured [63] in a device for which  $T_1^{min,charge} \approx 10 \text{ ns}$  [59]. We find that  $I_{dephase} > I_{relax}$  for all of the pulses we consider here, in most cases by more than a factor of two. For the 10 ns detuning pulse considered in Fig. 3 of the main text, the expected infidelity due to relaxation is approximately a factor of four smaller than the calculated infidelity due to dephasing.

We expect the relaxation due to charge noise to dominate over phonon-mediated noise under experimental conditions. However, as we show now, we expect the two effects to scale proportionally. The qubit-phonon coupling term  $H_{q-pn}$  in the  $\{|\cdot S\rangle, |\cdot T\rangle, |S\cdot\rangle\}$  basis can be approximated as

$$H_{q-pn} = \begin{pmatrix} H_{RR} & 0 & H_{RL} \\ 0 & H_{RR} & H_{RL} \\ H_{RL}^* & H_{RL}^* & H_{LL} \end{pmatrix}, \quad (\text{B.24})$$

where  $R$  and  $L$  refer to the third electron being in the right or left dot, respectively. Here we assume that the intra-dot relaxation is negligible [63]. Since  $|H_{LL} - H_{RR}| \gg |H_{RL}|$  for

the majority of  $\mathbf{Q}$ -space [135, 137], we further approximate  $H_{q-pn} \approx (H_{LL} - H_{RR})\mathcal{P}$ . Thus the qubit-phonon coupling term is approximately proportional to the charge noise coupling term, which by Eq. (B.21) means that the phonon-induced relaxation rate is proportional to the charge-noise-induced relaxation rate. The  $T_1^{min,charge}$  parameter employed previously in Eq. (B.22) encompasses both relaxation mechanisms simultaneously.

In summary, using rough estimates of the expected  $T_1$  time, we have shown that the under realistic charge noise levels and without sophisticated pulse optimization, charge-noise-induced dephasing rather than relaxation is the dominant source of error for the proposed two-qubit gates.

# Appendix C

## Supplemental information for Ch. 5

### C.1 Parameters of web graph models

In implementing the models used in this paper, the relationship between the parameters of the network generation algorithms and the generated networks themselves is not always obvious, so in the following section we explain it in detail.

#### C.1.1 GZL preferential attachment

The method of graph construction in the GZL Preferential Attachment Model [36] consists of two phases, each with its own parameter. First, a graph  $X$  (with adjacency matrix  $A_X$ ) is created by adding a new vertex at each time step, where each vertex is created with  $m_X$  out-going edges. Next, a second graph  $Y$  (with adjacency matrix  $A_Y$ ) is created in the same fashion, only with each new vertex having  $m_Y$  in-coming edges.  $A_X$  and  $A_Y$  are then added

together, with loops and weights discarded, forming the adjacency matrix of the desired network.  $m_X$  and  $m_Y$  are the two parameters to consider in this algorithm.

In order for a graph to be scale-free,  $\Pr(d_{in} = k)$  and  $\Pr(d_{out} = k)$ , the probabilities that the in-degree  $d_{in}$  and the out-degree  $d_{out}$  of a random node have the value  $k$ , must satisfy

$$\Pr(d_{in} = k) \sim k^{-\gamma_{in}}, \quad (\text{C.1})$$

$$\Pr(d_{out} = k) \sim k^{-\gamma_{out}},$$

where  $\gamma_{in}$  and  $\gamma_{out}$  are positive real numbers, and it is understood that  $\Pr(d_{in} = k) = 0$  when  $k < m_X$  and  $\Pr(d_{out} = k) = 0$  when  $k < m_Y$ . To compute  $\gamma_{in}$  and  $\gamma_{out}$ , one starts from the undirected version from Ref. [138]. This result is then combined with a constant offset, since each vertex of  $X$  has  $m_x$  outgoing edges and each vertex of  $Y$  has  $m_Y$  incoming edges. The resulting composite probability distributions follow

$$\Pr(d_{in} = k) \sim (k + m_X - m_Y)^{-3}, \quad (\text{C.2})$$

$$\Pr(d_{out} = k) \sim (k - m_X + m_Y)^{-3}.$$

Thus, for sufficiently large  $k$ , these distributions are scale-free. However, for a large range of intermediate  $k$ , we expect substantial deviation from the power law dependence of Eq. (C.1). According to GZL [128], the parameters used to generate Fig. 2 in their paper [36], which provides the main evidence for logarithmic scaling of the gap, follow  $m_Y \gg m_X$ . In Fig. C.1,




Figure C.1: Degree distributions for the GZL preferential attachment model with  $m_X = 1$  and  $m_Y = 15$ , taken at graph size  $n = 8196$  and averaged over approximately 1000 random graph realizations. Both the in-degree (blue circles) and out-degree (red squares) distributions are shown. For reference, the in-degree distribution for  $m_X = 1$  and  $m_Y = 1$  (duplicated from Fig. 5.2 is shown (black diamonds). The dashed line is the expected power law scaling of  $d^{-3}$ , which is applicable for large  $d$ . As predicted by Eq. (C.2), shown as fitting curves, the  $m_X = 1$  and  $m_Y = 15$  distributions exhibit non-scale-free behavior over a wide region of  $d$ .

we show the degree distributions for such a network, where we set  $m_X = 1$  and  $m_Y = 15$ . There, we see that the degree distributions are well-described by Eq. (C.2), and that the addition process does indeed distort the degree distributions. By requiring  $m_X = m_Y$ , as we have done in this paper (and GZL did for a portion of their supplemental material [36]),  $\gamma_{in} = \gamma_{out} = 3$  for all  $k$ , meaning that the in-degrees and out-degrees both follow the desired power law behavior.

The asymptotic (large number of nodes) value of average edges per node for the composite graph is also determined by the parameters  $m_X$  and  $m_Y$ . Because  $m_X$  is the number of outgoing edges per vertex in graph  $X$ , it is also the average number of edges per vertex in  $X$ . The same logic holds for  $m_Y$  and graph  $Y$ . Thus, when constructing the composite graph, the asymptotic average edges per node would be simply  $m_X + m_Y$ . Although loops are then

eliminated from the composite graph, the expected number of loops is much less than  $n$  in the large- $n$  case, so this has little effect on the average edges per node. To produce a graph with  $\gamma_{in} = \gamma_{out} = 3$  and average in- and out-edges per node of 2 (as in Fig. 5.2), we use this model with  $m_X = m_Y = 1$ .

### C.1.2 GZL copying model

The parameters of the GZL Copying Model [36] are similar to the GZL Preferential Attachment, as they both involve the adding of two graphs to form a composite graph. We again have the parameters  $m_X$  and  $m_Y$ , which again indicate the number of out-going edges per node in one component graph and the number of in-coming edges per node in the other.

This model has two new parameters,  $p_X$  and  $p_Y$ , which are the probabilities of a new node connecting to nodes chosen uniformly at random at a given time step during the construction of  $X$  and  $Y$ , respectively. We follow Ref. [127] and add a constant offset (just as in the preferential attachment case). Doing so, we again obtain the result that the graphs are scale-free only for  $m_X = m_Y$ . Assuming this constraint, the composite graph follows

$$\gamma_{in} = \frac{2 - p_X}{1 - p_X}, \quad (\text{C.3})$$

$$\gamma_{out} = \frac{2 - p_Y}{1 - p_Y}. \quad (\text{C.4})$$

For the data in Fig. 2 of the main text, we used the parameters  $p_X = p_Y = 0.5$  and  $m_X = m_Y = 1$ . In Fig. 3 of the main text, we used  $p_X = 1/11$  and  $p_Y = 35/86$  and



$$m_X = m_Y = 1.$$

### C.1.3 $\alpha$ -preferential attachment

Just as in the GZL Copying Model, there are multiple possible actions at each time step in the  $\alpha$ -Preferential Attachment Model [124], and each of these steps has an associated probability.  $p_1$  is the probability of adding a new vertex with a single out-going edge,  $p_2$  is the probability of adding a new vertex with a single in-coming edge, and  $1 - p_1 - p_2$  is the probability of an edge being added to the existing network without the addition of a new vertex.  $\alpha$ , the third parameter, measures how far the generated network deviates from the GZL preferential attachment model.

As laid out in Ref. [124], the relationship between these 3 parameters and the exponents is

$$\gamma_{in} = \frac{2 + (p_1 + p_2)\alpha - p_2}{1 - p_2}, \quad (\text{C.5})$$

$$\gamma_{out} = \frac{2 + (p_1 + p_2)\alpha - p_1}{1 - p_1}. \quad (\text{C.6})$$

The connection between these parameters and the average number of directed edges per node in the graph is clear when one considers that the probability that a new node will be added at a given time step is  $p_1 + p_2$ , and a new edge is added at each step.

Using these constraints, we can find appropriate values for the parameters for both

Fig. 5.2 and Fig. 5.3. In Fig. 5.2, we used  $p_1 = p_2 = 0.25$ , and  $\alpha = 1$ , and in Fig. 5.3, we used  $p_1 = 0.415$ ,  $p_2 = 0.0851$ , and  $\alpha = 0.0128$ . These choices in parameters keep  $\gamma_{in}$  and  $\gamma_{out}$  fixed at our desired values, while simultaneously keeping the graph at an average of 2 in- and 2 out-edges per node.

## C.2 Initial conditions

For each of these models, it is necessary to specify an initial graph to seed the network growth. In our simulations we used a complete graph (including loops) with  $m + 1$  vertices, where  $m$  is the number of edges added per vertex (in the  $\alpha$ -Preferential Attachment Model, we used  $m = 1$ ).

## C.3 Adaptive binning

In the plots of the degree distributions (Figs. 5.2(e)-(f), Figs. 5.3(c)-(d), and Fig. C.1), numerical noise caused by few high-degree vertices leads to data which are difficult to interpret. In order to combat this, we use adaptive binning, which functions as follows. First, some sampling threshold  $s_t$  is set, which we take to be 200 in our analysis. If a given data point, corresponding to a degree, contains at least  $s_t$  samples, then it is included. If the data point instead has fewer than  $s_t$  samples, it is combined with nearby points until the aggregated samples total at least  $s_t$ . The weighted average degree and probability are then recorded.

# Bibliography

- [1] Rivest, R. L., Shamir, A. & Adleman, L. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* **21**, 120–126 (1978).
- [2] Shor, P. W. Algorithms for quantum computation: discrete logarithms and factoring. In *FOCS '94: Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, 124–134 (IEEE Computer Society, Washington, DC, USA, 1994).
- [3] Hamdi, S. M., Zuhori, S. T., Mahmud, F. & Pal, B. A compare between Shor's quantum factoring algorithm and general number field sieve. In *2014 International Conference on Electrical Engineering and Information Communication Technology*, 1–6 (2014).
- [4] Kaye, P., Laflamme, R. & Mosca, M. *An introduction to quantum computing* (Oxford University Press, 2007).
- [5] Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, New York, NY, USA, 2011), 10th edn.
- [6] Albash, T. & Lidar, D. A. Adiabatic quantum computation. *Rev. Mod. Phys.* **90**, 015002 (2018).
- [7] Ekert, A. & Jozsa, R. Quantum computation and shor's factoring algorithm. *Rev. Mod. Phys.* **68**, 733–753 (1996).
- [8] Horodecki, R., Horodecki, P., Horodecki, M. & Horodecki, K. Quantum entanglement. *Rev. Mod. Phys.* **81**, 865–942 (2009).
- [9] Barenco, A. *et al.* Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457–3467 (1995).
- [10] Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A* **52**, R2493–R2496 (1995).
- [11] Steane, A. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences.* **452** (1996).

- [12] Bombin, H. & Martin-Delgado, M. A. Topological quantum distillation. *Phys. Rev. Lett.* **97**, 180501 (2006).
- [13] Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A* **86**, 032324 (2012).
- [14] Knill, E., Laflamme, R. & Zurek, W. Threshold accuracy for quantum computation. *arXiv e-prints* arXiv:quant-ph/9610011 (1996).
- [15] Aharonov, D. & Ben-Or, M. Fault-tolerant quantum computation with constant error rate. *arXiv e-prints* arXiv:quant-ph/9906129 (1999).
- [16] Shankar, R. *Principles of Quantum Mechanics* (Springer, 1994), 2 edn.
- [17] Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum computation by adiabatic evolution. *arXiv e-prints* arXiv:quant-ph/0001106v1 (2000).
- [18] Jansen, S., Ruskai, M.-B. & Seiler, R. Bounds for the adiabatic approximation with applications to quantum computation. *Journal of Mathematical Physics* **48**, 102111 (2007).
- [19] Young, K. C., Sarovar, M. & Blume-Kohout, R. Error suppression and error correction in adiabatic quantum computation: Techniques and challenges. *Phys. Rev. X* **3**, 041013 (2013).
- [20] Pudenz, K. L., Albash, T. & Lidar, D. A. Error-corrected quantum annealing with hundreds of qubits. *Nature Communications* **5**, 3243 (2014).
- [21] Aharonov, D. *et al.* Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM review* **50**, 755–787 (2008).
- [22] Devoret, M. H. & Schoelkopf, R. J. Superconducting circuits for quantum information: An outlook. *Science* **339**, 1169–1174 (2013).
- [23] Monroe, C. & Kim, J. Scaling the ion trap quantum processor. *Science* **339**, 1164–1169 (2013).
- [24] Mourik, V. *et al.* Signatures of Majorana Fermions in Hybrid Superconductor-Semiconductor Nanowire Devices. *Science* **336**, 1003–1007 (2012).
- [25] Loss, D. & DiVincenzo, D. P. Quantum computation with quantum dots. *Phys. Rev. A* **57**, 120–126 (1998).
- [26] Friis, N. *et al.* Observation of entangled states of a fully controlled 20-qubit system. *Phys. Rev. X* **8**, 021012 (2018).

- [27] Neill, C. *et al.* A blueprint for demonstrating quantum supremacy with superconducting qubits. *Science* **360**, 195–199 (2018).
- [28] Kane, B. E. A silicon-based nuclear spin quantum computer. *Nature* **393**, 133–137 (1998).
- [29] Saraiva, A. L., Calderón, M. J., Hu, X., Das Sarma, S. & Koiller, B. Physical mechanisms of interface-mediated intervalley coupling in si. *Phys. Rev. B* **80**, 081305 (2009).
- [30] Koppens, F. *et al.* Control and detection of singlet-triplet mixing in a random nuclear field. *Science* **309**, 1346–1350 (2005).
- [31] Petta, J. R. *et al.* Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science* **309**, 2180 (2005).
- [32] Laird, E. A. *et al.* Coherent spin manipulation in an exchange-only qubit. *Phys. Rev. B* **82**, 075403 (2010).
- [33] Koh, T. S., Gamble, J. K., Friesen, M., Eriksson, M. A. & Coppersmith, S. N. Pulse-gated quantum-dot hybrid qubit. *Phys. Rev. Lett.* **109**, 250503 (2012).
- [34] Friesen, M., Ghosh, J., Eriksson, M. A. & Coppersmith, S. N. A decoherence-free subspace in a charge quadrupole qubit. *Nature Communications* **8**, 15923 (2017).
- [35] Russ, M., Petta, J. R. & Burkard, G. Quadrupolar exchange-only spin qubit. *Phys. Rev. Lett.* **121**, 177701 (2018).
- [36] Garnerone, S., Zanardi, P. & Lidar, D. A. Adiabatic quantum algorithm for search engine ranking. *Phys. Rev. Lett.* **108**, 230506 (2012).
- [37] Frees, A. *et al.* Compressed optimization of device architectures for semiconductor quantum devices. *Phys. Rev. Applied* **11**, 024063 (2019).
- [38] Frees, A., Mehl, S., King Gamble, J., Friesen, M. & Coppersmith, S. N. Adiabatic two-qubit gates in capacitively coupled quantum dot hybrid qubits. *arXiv e-prints* arXiv:1812.03177 (2018).
- [39] Frees, A. *et al.* Power-law scaling for the adiabatic algorithm for search-engine ranking. *Phys. Rev. A* **88**, 032307 (2013).
- [40] Yu, P. & Cardona, M. *Fundamentals of Semiconductors: Physics and Materials Properties* (Springer, 2010), 4th edn.
- [41] Thorbeck, T. & Zimmerman, N. M. Formation of strain-induced quantum dots in gated semiconductor nanostructures. *AIP Advances* **5**, 087107 (2015).

- [42] Adachi, S. *Physical Properties of III-V Semiconductor Compounds: InP, InAs, GaAs, GaP, InGaAs, and InGaAsP* (John Wiley and Sons, Inc., 1992).
- [43] Ando, T., Fowler, A. B. & Stern, F. Electronic properties of two-dimensional systems. *Rev. Mod. Phys.* **54**, 437–672 (1982).
- [44] Herring, C. & Vogt, E. Transport and deformation-potential theory for many-valley semiconductors with anisotropic scattering. *Phys. Rev.* **101**, 944–961 (1956).
- [45] Schaffler, F., Tobben, D., Herzog, H. J., Abstreiter, G. & Hollander, B. High-electron-mobility si/SiGe heterostructures: influence of the relaxed SiGe buffer layer. *Semiconductor Science and Technology* **7**, 260–266 (1992).
- [46] Hensel, J. C., Hasegawa, H. & Nakayama, M. Cyclotron resonance in uniaxially stressed silicon. ii. nature of the covalent bond. *Phys. Rev.* **138**, A225–A238 (1965).
- [47] Ohkawa, F. J. & Uemura, Y. Theory of valley splitting in an N-channel (100) inversion layer of Si II. electric break through. *Journal of the Physical Society of Japan* **43**, 917–924 (1977).
- [48] Sverdlov, V. & Selberherr, S. Electron subband structure and controlled valley splitting in silicon thin-body SOI FETs: Two-band k·p theory and beyond. *Solid-State Electronics* **52**, 1861 – 1866 (2008).
- [49] Hopcroft, M. A., Nix, W. D. & Kenny, T. W. What is the Young’s modulus of silicon? *Journal of Microelectromechanical Systems* **19**, 229–238 (2010).
- [50] COMSOL Multiphysics. <http://www.comsol.com>.
- [51] Zajac, D. M., Hazard, T. M., Mi, X., Nielsen, E. & Petta, J. R. Scalable gate architecture for a one-dimensional array of semiconductor spin qubits. *Physical Review Applied* **6**, 054013 (2016).
- [52] Hanson, R., Kouwenhoven, L. P., Petta, J. R., Tarucha, S. & Vandersypen, L. M. K. Spins in few-electron quantum dots. *Rev. Mod. Phys.* **79**, 1217–1265 (2007).
- [53] Zwanenburg, F. A. *et al.* Silicon quantum electronics. *Rev. Mod. Phys.* **85**, 961 (2013).
- [54] Petersson, K. D., Petta, J. R., Lu, H. & Gossard, A. C. Quantum coherence in a one-electron semiconductor charge qubit. *Phys. Rev. Lett.* **105**, 246804 (2010).
- [55] Dial, O. E. *et al.* Charge noise spectroscopy using coherent exchange oscillations in a singlet-triplet qubit. *Phys. Rev. Lett.* **110**, 146804 (2013).
- [56] Cao, G. *et al.* Tunable hybrid qubit in a GaAs double quantum dot. *Phys. Rev. Lett.* **116**, 086801 (2016).

- [57] Martins, F. *et al.* Noise suppression using symmetric exchange gates in spin qubits. *Phys. Rev. Lett.* **116**, 116801 (2016).
- [58] Kim, D. *et al.* Quantum control and process tomography of a semiconductor quantum dot hybrid qubit. *Nature* **511**, 70–74 (2014).
- [59] Kim, D. *et al.* Microwave-driven coherent operation of a semiconductor quantum dot charge qubit. *Nature Nano.* **10**, 243–247 (2015).
- [60] Reed, M. D. *et al.* Reduced sensitivity to charge noise in semiconductor spin qubits via symmetric operation. *Phys. Rev. Lett.* **116**, 110402 (2016).
- [61] Kawakami, E. *et al.* Gate fidelity and coherence of an electron spin in an Si/SiGe quantum dot with micromagnet. *P. N. A. S.* **113**, 11738–11743 (2016).
- [62] Schoenfield, J. S., Freeman, B. M. & Jiang, H. Coherent manipulation of valley states at multiple charge configurations of a silicon quantum dot device. *Nature Communications* **8**, 64 (2017).
- [63] Thorgrimsson, B. *et al.* Extending the coherence of a quantum dot hybrid qubit. *npj Quantum Information* **3**, 32 (2017).
- [64] Yoneda, J. *et al.* A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%. *Nature Nano.* **13**, 102–106 (2018).
- [65] Samkharadze, N. *et al.* Strong spin-photon coupling in silicon. *Science* **359**, 1123–1127 (2018).
- [66] Mi, X. *et al.* A coherent spin–photon interface in silicon. *Nature* **555**, 599 (2018).
- [67] Crippa, A. *et al.* Electrical spin driving by  $g$ -matrix modulation in spin-orbit qubits. *Phys. Rev. Lett.* **120**, 137702 (2018).
- [68] Jock, R. M. *et al.* A silicon metal-oxide-semiconductor electron spin-orbit qubit. *Nature Communications* **9**, 1768 (2018).
- [69] Veldhorst, M. *et al.* A two-qubit logic gate in silicon. *Nature* **526**, 410–414 (2015).
- [70] Nichol, J. M. *et al.* High-fidelity entangling gate for double-quantum-dot spin qubits. *npj Quantum Information* **3**, 3 (2017).
- [71] Zajac, D. M. *et al.* Resonantly driven CNOT gate for electron spins. *Science* **359**, 439–442 (2018).
- [72] Watson, T. F. *et al.* A programmable two-qubit quantum processor in silicon. *Nature* **555**, 633 (2018).

- [73] Pla, J. J. *et al.* A single-atom electron spin qubit in silicon. *Nature* **489**, 541–545 (2012).
- [74] Fuechsle, M. *et al.* A single-atom transistor. *Nature Nano.* **7**, 242–246 (2012).
- [75] Pla, J. J. *et al.* High-fidelity readout and control of a nuclear spin qubit in silicon. *Nature* **496**, 334–338 (2013).
- [76] Foote, R. H. *et al.* Transport through an impurity tunnel coupled to a Si/SiGe quantum dot. *Applied Physics Letters* **107**, 103112 (2015).
- [77] Watson, T. F., Weber, B., House, M. G., Büch, H. & Simmons, M. Y. High-fidelity rapid initialization and read-out of an electron spin via the single donor  $D^-$  charge state. *Phys. Rev. Lett.* **115**, 166806 (2015).
- [78] Laucht, A. *et al.* A dressed spin qubit in silicon. *Nature Nanotechnology* **12**, 61 (2016).
- [79] Watson, T. F. *et al.* Atomically engineered electron spin lifetimes of 30 s in silicon. *Science Advances* **3**, e1602811 (2017).
- [80] Tosi, G. *et al.* Silicon quantum processor with robust long-distance qubit couplings. *Nature Communications* **8**, 450 (2017).
- [81] Harvey-Collard, P. *et al.* Coherent coupling between a quantum dot and a donor in silicon. *Nature Communications* **8**, 1029 (2017).
- [82] Broome, M. A. *et al.* Two-electron spin correlations in precision placed donors in silicon. *Nature Communications* **9**, 980 (2018).
- [83] Baart, T. A., Eendebak, P. T., Reichl, C., Wegscheider, W. & Vandersypen, L. M. K. Computer-automated tuning of semiconductor double quantum dots into the single-electron regime. *Applied Physics Letters* **108**, 213104 (2016).
- [84] van Diepen, C. J. *et al.* Automated tuning of inter-dot tunnel coupling in double quantum dots. *Applied Physics Letters* **113**, 033101 (2018).
- [85] Zwolak, J. P., Kalantre, S. S., Wu, X., Ragole, S. & Taylor, J. M. Qflow lite dataset: A machine-learning approach to the charge states in quantum dot experiments. *PLOS ONE* **13**, 1–17 (2018).
- [86] Candès, E. J., Romberg, J. K. & Tao, T. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics* **59**, 1207–1223 (2006).
- [87] Donoho, D. Compressed sensing. *IEEE Transactions on Information Theory* **52**, 1289–1306 (2006).



- [88] Natarajan, B. K. Sparse approximate solutions to linear systems. *SIAM Journal on Computing* **24**, 227–234 (1995).
- [89] Diamond, S. & Boyd, S. Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* **17**, 1–5 (2016).
- [90] Stopa, M. Quantum dot self-consistent electronic structure and the coulomb blockade. *Phys. Rev. B* **54**, 13767–13783 (1996).
- [91] Boykin, T. B. *et al.* Valley splitting in low-density quantum-confined heterostructures studied using tight-binding models. *Phys. Rev. B* **70**, 165325 (2004).
- [92] Davies, J. H. *The Physics of Low-dimensional Semiconductors: An Introduction* (Cambridge University Press, 1997).
- [93] Jones, E., Oliphant, T., Peterson, P. *et al.* SciPy: Open source scientific tools for Python (2001–). [Online; accessed 18 May 2018].
- [94] Ward, D. R. *et al.* State-conditional coherent charge qubit oscillations in a Si/SiGe quadruple quantum dot. *npj Quantum Information* **2**, 16032 (2016).
- [95] Shulman, M. D. *et al.* Demonstration of entanglement of electrostatically coupled singlet-triplet qubits. *Science* **336**, 202–205 (2012).
- [96] Taylor, J. M., Srinivasa, V. & Medford, J. Electrically protected resonant exchange qubits in triple quantum dots. *Phys. Rev. Lett.* **111**, 050502 (2013).
- [97] Mehl, S. Quantum computation with three-electron double quantum dots at an optimal operation point. *arXiv e-prints* arXiv:1507.03425 (2015).
- [98] Ferraro, E., Fanciulli, M. & De Michielis, M. Controlled-NOT gate sequences for mixed spin qubit architectures in a noisy environment. *Quantum Inf Process* **16** (2017).
- [99] Kim, D. *et al.* High-fidelity resonant gating of a silicon-based quantum dot hybrid qubit. *Npj Quantum Information* **1**, 15004 (2015).
- [100] Taylor, J. M. *et al.* Fault-tolerant architecture for quantum computation using electrically controlled semiconductor spins. *Nature Physics* **1**, 177 (2005).
- [101] Viola, L., Knill, E. & Lloyd, S. Dynamical decoupling of open quantum systems. *Phys. Rev. Lett.* **82**, 2417–2421 (1999).
- [102] Shi, Z. *et al.* Fast hybrid silicon double-quantum-dot qubit. *Phys. Rev. Lett.* **108**, 140503 (2012).
- [103] Wong, C. H. High-fidelity ac gate operations of a three-electron double quantum dot qubit. *Phys. Rev. B* **93**, 035409 (2016).

- [104] Yang, Y.-C., Coppersmith, S. N. & Friesen, M. Achieving high-fidelity single-qubit gates in a strongly driven silicon-quantum-dot hybrid qubit. *Phys. Rev. A* **95**, 062321 (2017).
- [105] Broyden, C. G. The convergence of a class of double-rank minimization algorithms 1. general considerations. *IMA Journal of Applied Mathematics* **6**, 76–90 (1970).
- [106] Fletcher, R. A new approach to variable metric algorithms. *The Computer Journal* **13**, 317–322 (1970).
- [107] Goldfarb, D. A family of variable-metric methods derived by variational means. *Mathematics of Computation* **24**, 23–26 (1970).
- [108] Shanno, D. F. Conditioning of quasi-newton methods for function minimization. *Mathematics of Computation* **24**, 647–656 (1970).
- [109] Leary, R. H. Global optimization on funneling landscapes. *Journal of Global Optimization* **18**, 367–383 (2000).
- [110] Setser, A. A. & Kestner, J. P. Rapid adiabatic gating for capacitively coupled quantum dot hybrid qubits without barrier control. *arXiv e-prints* arXiv:1812.04681 (2018).
- [111] Khaneja, N., Reiss, T., Kehlet, C., Schulte-Herbrüggen, T. & Glaser, S. J. Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance* **172**, 296 – 305 (2005).
- [112] Lucarelli, D. Quantum optimal control via gradient ascent in function space and the time-bandwidth quantum speed limit. *Phys. Rev. A* **97**, 062346 (2018).
- [113] Frey, V. M. *et al.* Application of optimal band-limited control protocols to quantum noise sensing. *Nature Communications* **8**, 2189 (2017).
- [114] Gamble, J. K., Friesen, M., Coppersmith, S. N. & Hu, X. Two-electron dephasing in single Si and GaAs quantum dots. *Phys. Rev. B* **86**, 035302 (2012).
- [115] Gilchrist, A., Langford, N. K. & Nielsen, M. A. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A* **71**, 062310 (2005).
- [116] Grover, L. K. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, STOC '96, 212–219 (ACM, New York, NY, USA, 1996).
- [117] Bacon, D. & van Dam, W. Recent progress in quantum algorithms. *Commun. ACM* **53**, 84–93 (2010).
- [118] Brin, S. & Page, L. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems* **30**, 107 – 117 (1998).

- [119] Berkhin, P. A survey on pagerank computing. *Internet Math* **2**, 73–120 (2005).
- [120] Barabasi, A. L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
- [121] Bollobás, B., Riordan, O., Spencer, J. & Tusnády, G. The degree sequence of a scale-free random graph process. *Random Structures and Algorithms* **18**, 279–290 (2001).
- [122] Cohen, R. & Havlin, S. Scale-free networks are ultrasmall. *Phys. Rev. Lett.* **90**, 058701 (2003).
- [123] Tangmunarunkit, H., Govindan, R., Jamin, S., Shenker, S. & Willinger, W. Network topology generators: degree-based vs. structural. In *SIGCOMM '02: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, 147–159 (ACM, New York, NY, USA, 2002).
- [124] Bollobás, B., Borgs, C., Chayes, J. & Riordan, O. Directed scale-free graphs. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms, SODA '03*, 132–139 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003).
- [125] Broder, A. *et al.* Graph structure in the web. *Comput. Netw.* **33**, 309–320 (2000).
- [126] Chung, F. & Lu, L. *Complex Graphs and Networks (Cbms Regional Conference Series in Mathematics)*. 107 (American Mathematical Society, Boston, MA, USA, 2006).
- [127] Kleinberg, J. M., Kumar, R., Raghavan, P., Rajagopalan, S. & Tomkins, A. S. The web as a graph: measurements, models, and methods. In *Proceedings of the 5th annual international conference on Computing and combinatorics, COCOON'99*, 1–17 (Springer-Verlag, Berlin, Heidelberg, 1999).
- [128] Garnerone, S. (2012). Private communication.
- [129] Nelder, J. & Mead, R. A simplex method for function minimization. *Computer Journal* **7**, 308–313 (1965).
- [130] Broyden, C. G. A class of methods for solving nonlinear simultaneous equations. *Math. Comp.* **19**, 577–593 (1965).
- [131] Wu, X. *et al.* Two-axis control of a singlet–triplet qubit with an integrated micromagnet. *P. N. A. S.* **111**, 11938–11942 (2014).
- [132] Kawakami, E. *et al.* Electrical control of a long-lived spin qubit in a Si/SiGe quantum dot. *Nature Nano.* **9**, 666–670 (2014).
- [133] Landau, L. & Lifshitz, E. *Course of Theoretical Physics. Vol. 3: Quantum Mechanics. Non-relativistic Theory* (London, 1958), third edn.

- [134] Winkler, R. *Spin-orbit Coupling Effects in Two-Dimensional Electron and Hole Systems* (Springer-Verlag Berlin Heidelberg, 2003), 1 edn.
- [135] Wang, K., Payette, C., Dovzhenko, Y., Deelman, P. W. & Petta, J. R. Charge relaxation in a single-electron Si/SiGe double quantum dot. *Phys. Rev. Lett.* **111**, 046801 (2013).
- [136] Schoelkopf, R. J., Clerk, A. A., Girvin, S. M., Lehnert, K. W. & Devoret, M. H. Noise and measurement backaction in superconducting circuits: qubits as spectrometers of quantum noise. In *Proc. SPIE 5115, Noise and Information in Nanoelectronics, Sensors, and Standards* (2003).
- [137] Nguyen, K. T. *et al.* Spontaneous emission in a silicon charge qubit. *arXiv e-prints* arXiv:1403.3704 (2014).
- [138] Barabási, A., Albert, R. & Jeong, H. Mean-field theory for scale-free random networks. *Physica A* **272**, 173–187 (1999).